Artificial Intelligence in the News

How AI Retools, Rationalizes, and Reshapes Journalism and the Public Arena

Felix M. Simon

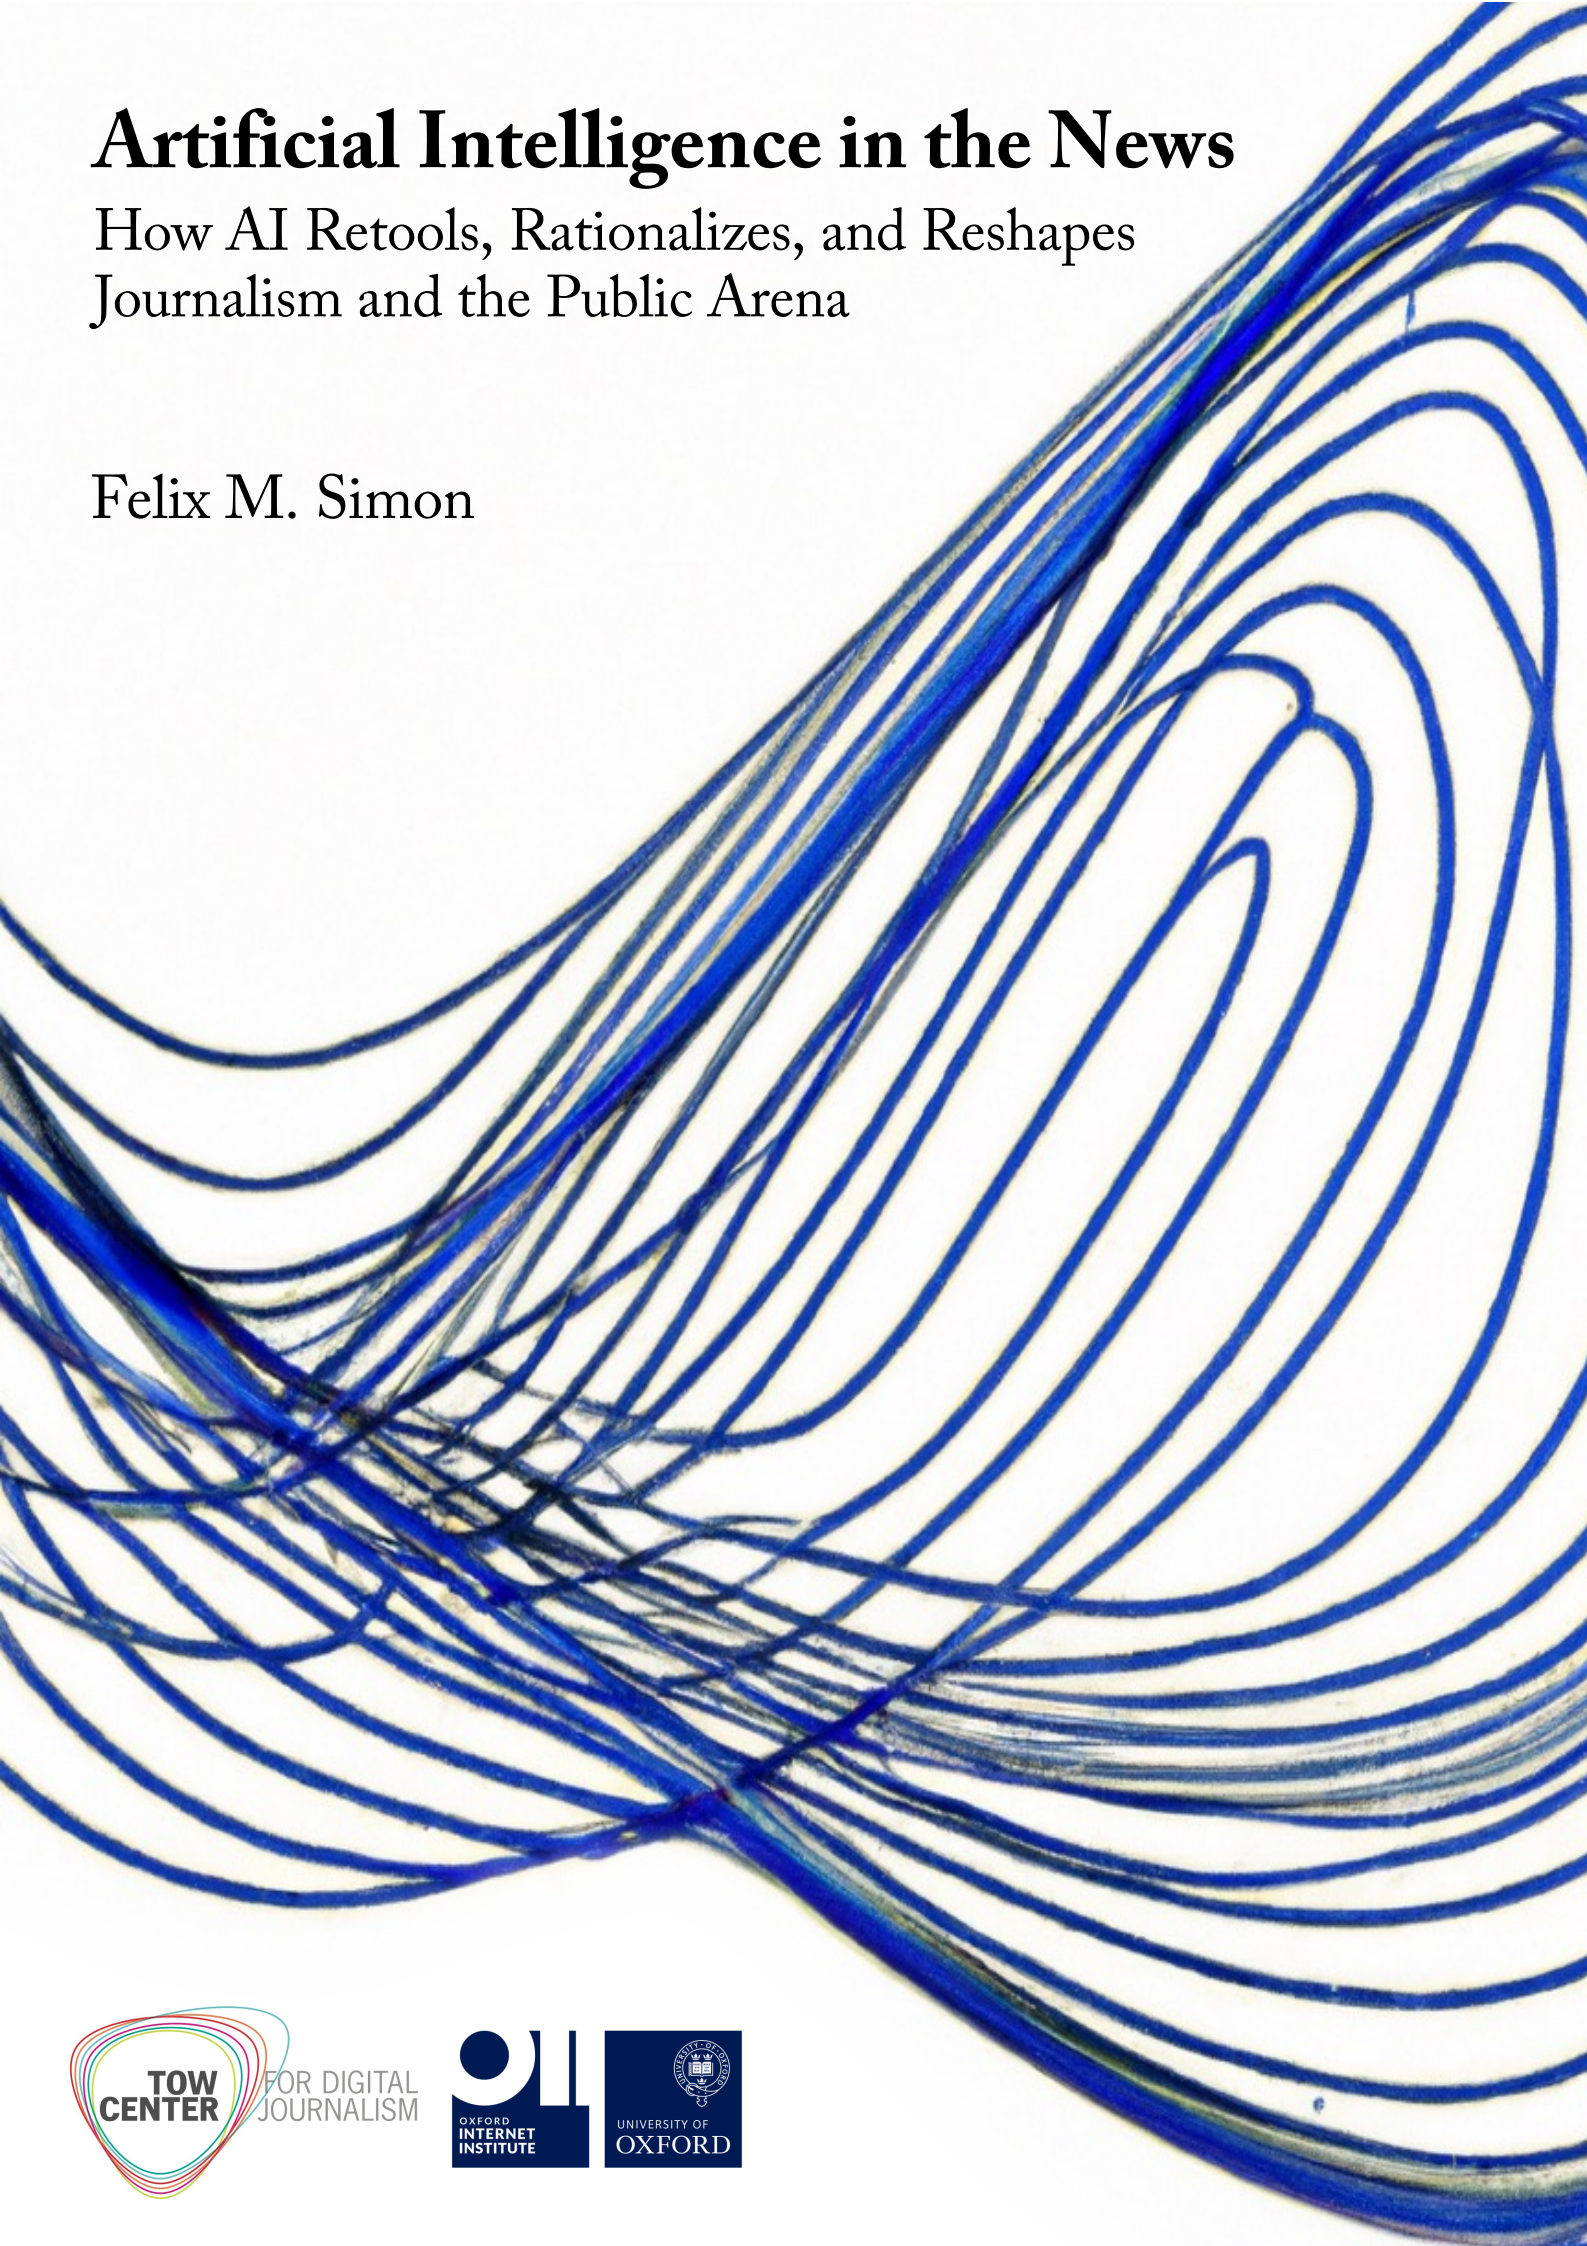

TOW CENTER FOR DIGITAL JOURNALISM
OXFORD INTERNET INSTITUTE
UNIVERSITY OF OXFORD

Table of Contents

Keywords: Artificial Intelligence, AI, Democracy, Journalism, Generative AI, LLMs, News, News Industry, Platform companies

This report was funded by the Tow Center for Digital Journalism at Columbia University and received support from Balliol College and the Oxford Internet Institute at the University of Oxford.

Executive Summary

Despite growing interest, the effects of AI on the news industry and our information environment — the public arena — remain poorly understood. Insufficient attention has also been paid to the implications of the news industry's dependence on technology companies for AI. Drawing on 134 interviews with news workers at 35 news organizations in the United States, the United Kingdom, and Germany — including outlets such as *The Guardian, Bayerischer Rundfunk,* the *Washington Post, The Sun*, and the *Financial Times* — and 36 international experts from industry, academia, technology, and policy, this report examines the use of AI across editorial, commercial, and technological domains with an eye to the structural implications of AI in news organizations for the public arena. In a second step, it considers how a retooling of the news through AI stands to reinforce news organizations' existing dependency on the technology sector and the implications of this.

Chapter 1 is broken down into three parts, exploring (i) news organizations' motives for introducing AI into their businesses; (ii) the ways in which AI is currently being used for the production and distribution of journalism; and (iii) the expectations being placed on AI's scope to deliver efficiency.

- In terms of motivations, news organizations have adopted AI as a result of recent technological advancements, market pressures stemming partially from the industry's financial challenges, competitive dynamics with a focus on innovation, and the pervasive sense of uncertainty, hype, and hope surrounding AI.
- AI is now applied across an ever greater range of tasks in the production and distribution of news. Contrary to some assertions, many of the most beneficial applications of AI in news are relatively mundane, and AI has often not proved to be a silver bullet in many cases.
- AI's potential to increase efficiency in news organizations is a central motivator for its adoption. Various examples demonstrate that efficiency and productivity gains have been achieved, including dynamic paywalls, automated transcription, and data analysis tools in news production.
- Such efficiency gains are task- and context-dependent. Potential efficiency gains can be curtailed by factors such as the unreliability of AI outputs, concerns about reputational damage resulting from inaccurate AI outputs, and the difficulty of automating certain tasks.

Reflecting on the extent to which AI has impacted news organizations, I argue that it presents 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 that is imbued with ideas of rationality, efficiency, and speed — and that does indeed provide greater efficiency and effectiveness *in some contexts.* However, the effects of AI in the news are subject to contextual

factors, with professional norms, resistance from news workers, regulations, audience preferences, and existing technological infrastructures all acting as constraints.

Chapter 2 explores the questions of how and why news organizations rely on technology companies for AI. Again, it is broken down into three parts, analyzing (i) the contexts in which publishers rely on AI and AI infrastructure from platform companies; (ii) the reasons for this reliance; and (iii) the implications of this relationship. Key takeaways include:

- News organizations make extensive use of AI products and infrastructure from major tech companies like Google, Amazon, and Microsoft across various aspects of their operations.
- Larger, better resourced news organizations are more likely to engage in in-house AI development. The majority of other publishers, especially smaller ones, opt for third-party solutions from platform companies because of the high costs associated with custom AI.
- Publishers turn to platform companies' AI offerings due to the costs and challenges associated with independent development, including the need for extensive computing power, competition for tech talent, and the scarcity of large datasets. The convenience, scalability, and cost-effectiveness of platform offerings make them attractive, allowing publishers to leverage AI capabilities without the financial burden of in-house development.
- Despite reservations in some quarters of the news industry, the adoption of "platform AI" is largely viewed as a pragmatic choice driven by economic challenges and the competitive landscape for tech talent.
- 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 organizations' autonomy and renders them vulnerable to price hikes or the shifting priorities of technology companies that may not align with their own.
- The lack of transparency in AI systems raises worries about biases or errors creeping into journalistic output, especially as generative AI models gain prominence. There is also a risk that the use of AI undercuts journalists' autonomy by limiting their discretionary decision-making abilities.

The growing use of AI in news work tilts the balance of power toward technology companies, raising concerns about "rent" extraction and potential threats to publishers' autonomy business models, particularly those reliant on search-driven traffic. As platforms prioritize AI-enhanced search experiences, publishers fear a shift where users opt for short answers, impacting audience engagement and highlighting the increasing control exerted by platform companies over the information ecosystem.

Bringing all this together, **Chapter 3** interrogates the question of whose interests are being served by the increasing adoption of AI in the news and how this shift stands to reshape the public arena — our information ecosystem. In this chapter I argue:

- Currently, AI aids news workers rather than replaces them, but there are no guarantees this will remain the case. AI is sufficiently mature to enable the replacement of at least some journalism jobs, either directly or because fewer workers are needed.
- It is not a given that AI will free up news workers to do deeper or better journalism. It is just as likely that any time savings will immediately be filled with new or additional demands.
- AI's effects on the news and the public arena will largely be determined by the decisions news organizations and managers make about when, where, and how the technology is used. The use of AI will not automatically improve journalism or the quality of information available to the public; this will only be achieved if the technology is used for this purpose.
- The increasing use of AI will likely reinforce existing inequalities among news organizations, with well-resourced, international publishers getting a head start. Local news organizations and publishers in the Global South are often an afterthought in the current conversations around AI in the news.
- On a macro level, news organizations are a vital component of the public arena. They act as gatekeepers for the common attention space most of us inhabit. As news organizations change through AI, so does the makeup of the broader system that they constitute and shape.
- The adoption of AI is shifting newswork, and the public arena, further toward the technical and the logics of platform companies, e.g. prioritizing greater rationalization and calculability (on the audience side in particular), and efficiencies and productivity in journalistic work. But this approach may not necessarily prioritize the welfare of journalism or the needs of audiences. Not every problem the news faces can be addressed with technological solutions.
- Publishers' use of platforms' AI for their own services, and their growing dependence on technology companies for AI more generally, could further weaken the news industry. The visibility of news content could shrink as AI user experiences become more popular.
- At times, publishers' use of AI helps improve the AI systems of major technology giants. This provides a pathway for platform companies to build better general-purpose AI products and services, further cementing their control over information, and potentially enabling them to take over tasks that were once central to journalism.

Finally, the **conclusion** summarizes my overall insights into how AI shapes and reshapes the news and the public arena.

- For now, I argue, AI mostly constitutes a retooling of the news rather than a fundamental change in the needs and motives of news organizations. It does not impact the fundamental need to access and gather information, to process and package that information into "news," to reach existing and new audiences, and to make money.
- AI will play a transformative role in reshaping news work, from editorial to the business side. We are witnessing — to a degree — a further rationalization of news work through

AI. It is important to recognize that the extent of this reshaping will be context- and task-dependent, and will also be influenced by institutional incentives and decisions.

- Winners and losers will emerge. In fact, they already have. News organizations that have been able to invest in research and development, devote staff time, attract and retain talent, and build infrastructure already have something of a head start. These "winners" are also in a stronger position to demand better terms when negotiating with platforms and technology companies.
- As news organizations get reshaped by AI, so too will the public arena that is so vital to democracy and for which news organizations play a vital gatekeeper role. The way this takes shape will depend on decisions made by two sets of actors: one that wields direct control over the conditions of news work (executives and managers, journalists) and, increasingly, one that does not (technology companies, regulatory bodies, and the public).
- AI will be far from the only thing that shapes the news and the public arena in the coming years. Journalism is not fundamentally altered by a single technology: It interacts with institutions and other forces in society and the economy.
- Productivity gains from AI in the news will not be straightforward. The benefits of AI to the news will be staggered. They will incur costs in the early stages and necessitate changes at the organizational and strategic level.
- The adoption of AI in news organizations will not be frictionless. Regulation, resistance from news workers, audience preferences, and incompatible technological infrastructure are just some of the variables that will shape the speed at which news organizations adopt AI, and, by extension, the rate at which tangible effects on the news come into focus.
- AI will not be a panacea for the many deep-seated problems and challenges facing journalism and the public arena. Technology alone cannot fix intractable political, social, and economic ills. News organizations will continue to be forced to make a case for why they still matter in the modern news environment — and why they deserve audiences' attention and money.
- The concentration of control over AI by a small handful of major technology companies must — and will — remain a key area of scrutiny. Control over infrastructure confers power.
- Developing frameworks to balance innovation — which is bound to continue — through AI in the news with concerns around issues like copyright and various forms of harms will remain a difficult and imperfect but necessary task.
- As with any new technology entering the news, the effects of AI will neither be as dire as the doomsayers predict, nor as utopian as the enthusiasts hope.

Introduction

Recent developments in the field of large language models (LLMs) have supercharged the news industry's thinking and experimentation with artificial intelligence (AI). What had previously been a relatively slow-burning development — once, that is, the initial, pre-pandemic hype cycle had tailed off — is now the talk of the town. As seen in the frenzied discourse that followed OpenAI's launch of ChatGPT in November 2022,[1] many news industry leaders have high hopes for AI to not just be the next big thing, but to be the "big thing" that delivers for their industry. It is for that reason that news organizations around the world are scrambling to come up with AI strategies or expand existing initiatives.

In recent years, many scholars and journalists have investigated the news industry's growing interest in AI. However, much of this earlier work has largely ignored (i) the ways in which news organizations' use of AI could reshape our information ecosystem (which I refer to as the public arena); (ii) the ways in which this race to "retool" the news industry could deepen news organizations' reliance on technology companies; and (iii) the possible consequences for the public arena. Drawing on more than four years of research, this report sets out to start filling these gaps.

This report is presented in two parts. I begin by considering the structural implications that the integration of AI into news organizations could have for the public arena. To do this, I explore the motivations behind news organizations' current use of AI across editorial, commercial, and technological domains, paying particular attention to what has emerged as one of the technology's key promises to the news industry and one of the news industry's big hopes for the technology: AI's capacity to increase efficiency and facilitate the production of more high-quality journalism.

Next, I analyze how a retooling of the news through AI may reinforce news organizations' existing dependencies on the technology sector and the potential implications for the public arena. Technology companies, especially large platform companies, are central players in AI, and their AI technologies are already embedded in many news organizations. This growing dependence — which will carry a sense of déjà vu for many — forces us to consider whether these moves to retool journalism through AI are forcing news organizations to cede (more) control to technology companies that have minimal interest in their craft, and the ramifications of this seemingly imbalanced power dynamic.

Method

This research draws upon 170 semi-structured interviews — 134 with news workers[2] and 36 with experts — I conducted between July 2021 and September 2023. My sample of news workers was

[1] Bartholomew & Mehta, 2023.

[2] I favor the term news worker over journalist, as it includes those working for news organizations in fields such as data science, product management, and research and development.

drawn from 35 news organizations in the United States, the United Kingdom, and Germany. The supplementary interviews included industry, technology, and policy experts from across Europe and the United States. In an effort to produce a suitably varied dataset, I included publishers operating in various media markets and media systems that have both a commercial or public service mission.[3] A full list of organizations is provided below.

Research ethical approval for the study was granted by the University of Oxford's Central University Research Ethics Committee.[4] Having started with subjects who have worked with or on AI in the broadest sense (as indicated by their job descriptions, press reports, or descriptions provided by colleagues), a mixture of purposive and snowball sampling was used to recruit interview participants from across their respective organizations.

Most interviews were conducted between July 2021 and December 2022. However, further interviews took place between January and September 2023 to capture more recent developments around large language models and generative AI.[5] Most interviews were conducted through Microsoft Teams and Zoom. All interviews were transcribed and thematically coded using a mixture of inductive and deductive coding. All interviews were anonymized; participants are identified only by broad descriptions of their roles and locations to maintain their anonymity, and all quotes and examples were carefully checked to ensure that they do not reveal participants' identity.

Like all research, this study has limitations. For example, it is limited to three countries and does not include local news organizations. As a qualitative study, no claim is made to the generalizability of findings to any population or context. Instead, the goal was to produce data rich enough to explain and interpret the phenomena. Interviews were conducted until I had reached saturation.

Name of organization	Country	Organization type
ARD	Germany	Broadcaster (public service media)
Bayerischer Rundfunk (BR)	Germany	Broadcaster (public service media)
Der Spiegel	Germany	Magazine
Deutsche Presse Agentur (dpa)	Germany	News agency
Deutsche Welle	Germany	Broadcaster (public service media)
Die Welt (Axel Springer)	Germany	Upmarket newspaper
Frankfurter Allgemeine Zeitung (F.A.Z.)	Germany	Upmarket newspaper

[3] Färdigh, 2010; García Avilés et al., 2004.

[4] CUREC, Approval Reference: SSH_OII_CIA_20_71

[5] Wolfram, 2023.

ProSieben/Sat.1	Germany	Broadcaster (commercial)
Rundfunk Berlin-Brandenburg (rbb)	Germany	Broadcaster (public service media)
Süddeutsche Zeitung (SZ)	Germany	Upmarket newspaper
Westdeutscher Rundfunk (WDR)	Germany	Broadcaster (public service media)
ZDF	Germany	Broadcaster (public service media)
Die Zeit	Germany	Upmarket weekly newspaper
BBC	United Kingdom	Broadcaster (public service media)
DMG Media Group (Daily Mail)	United Kingdom	Midmarket newspaper
Financial Times	United Kingdom	Upmarket newspaper
FullFact	United Kingdom	Digital-born outlet (public service oriented)
New Statesman	United Kingdom	Upmarket magazine
Press Agency (PA)/RADAR	United Kingdom	News agency
Reuters	United Kingdom	News agency
Sky News	United Kingdom	Broadcaster (commercial)
The Daily Telegraph	United Kingdom	Upmarket newspaper
The Economist	United Kingdom	Upmarket magazine
The Guardian	United Kingdom	Upmarket newspaper
The Sun (News UK)	United Kingdom	Tabloid newspaper
The Times (News UK)	United Kingdom	Upmarket newspaper
Associated Press	United States	News agency
Bloomberg	United States	News agency/digital-born outlet
The International Consortium of Investigative Journalists (ICIJ)	United States	Digital-born outlet/consortium (public service oriented)
NPR	United States	Broadcaster (public service media)
The Markup	United States	Digital-born outlet (public service oriented)
The New York Times	United States	Upmarket newspaper
The Wall Street Journal	United States	Upmarket newspaper
The Washington Post	United States	Upmarket newspaper

Table 1: List of news organizations

Definition of AI

There is much debate about how artificial intelligence should be defined and what should and should not count as true AI. While it is beyond the scope of this report to explore this debate in detail, it cannot be ignored. It is fair to say there is no consensus about what constitutes AI, nor is there a generally accepted definition of AI. There is, however, agreement around what AI *is not*: namely, a conscious, general intelligence that understands and works across domains.[6] Some, like sociologist Elena Esposito, have argued that the focus on recreating intelligence may not even be the point: "What algorithms [and AI] are reproducing is not the intelligence of people but the informativity of communication." What we can observe in interactions with algorithms, and especially chatbots, she writes, "is not necessarily an artificial form of intelligence, but rather an artificial form of communication" for which the question of whether the system is actually "intelligent" (whatever that means[7]) is mostly irrelevant.[8] Recent developments around generative AI — that is, AI systems capable of generating new, realistic forms of data such as text, images, and audio — have somewhat reinforced Esposito's point. In many ways, it matters little if AI is truly intelligent. Of more importance is that we treat it as such and acknowledge that its capabilities are edging closer to facilitating outputs previously perceived to be uniquely human.[9]

What, then, is AI in practice? The forms seen most commonly in the wild can be broadly categorized as "narrow" and "weak."[10] These include a diverse range of applications and techniques with different levels of complexity, autonomy, and abstraction, chipping away at various fairly narrowly defined tasks and problems.[11] These systems and programs are unable to operate beyond the "frontier of [their] own design"[12] — a point that even remains true for large language models such as GPT-4, which are able to operate across multiple text-based domains but have no consciousness[13] or comprehensive model of the world.[14] Examples of narrow forms of AI include applications of machine learning (ML) and its subfield, deep learning, as well as various forms of natural language processing (NLP) that often build on ML approaches. What these have in common is that a computer program 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. This training — or learning — happens on a scale between supervised and unsupervised, differs among AI systems and the approaches they use, and can also include further steps such as reinforcement learning from human feedback.

---

[6] Mitchell, 2019; Newport, 2023.

[7] Cave, 2020.

[8] Esposito, 2022, pp. 2, 16.

[9] Siegele, 2023.

[10] Broussard, 2018.

[11] Lewis & Simon, 2023.

[12] Diakopoulos, 2019, p. 243; Esposito, 2022.

[13] Defined here as the "ability to maintain a constantly updated conception of itself as a distinct entity interacting with a model of the external world." (Newport, 2023, p. 7.)

[14] Newport, 2023.

In the news industry, AI is largely used as an umbrella term to communicate with colleagues, partners, or the public about a set of technologies, with technological definitions mostly focusing on the techniques mentioned above.[15] For the purpose of this report, I define AI as:

> *The act of computationally simulating human activities and skills in narrowly defined domains, typically the application of machine learning approaches through which machines learn from data and/or their own performance.*

Chapter I: AI in the News: Reshaping Production and Distribution?

The purpose of this chapter is to begin exploring the broad question of how news organizations' adoption of AI could change our information environment. To do this, I present an overview of three relevant themes from my interview data: (a) news organizations' motives for using AI, (b) current applications of AI for news production and distribution, and (c) the extent to which these fulfill AI's key promise to improve efficiency and aid the production of better-quality news.

1a: A Question of Motives

Broader technological developments,[16] institutional factors, and sociocultural conditions all play a role in news organizations' adoption of new technologies, and so it is with the current drive toward AI.

The primary motivations for adopting AI cited by my interviewees can be grouped into four broad categories: (i) *technological developments*, (ii) *market pressures*, (iii) *industry dynamics*, and (iv) *uncertainty, hype, and hope*.

In terms of *technological developments*, many news organizations recognize advancements in AI and the extent to which it is being used by other industries over the past decade. As one digital executive in the United Kingdom put it:

> So it originates from the fact that the technology has been improving over the years, to the point where it is now something that is, to some extent, accessible […] and widely used by other industries [like finance].

This usage has intensified of late, owing to the rapid developments around generative AI and the larger rollout of these tools.

[15] *E.g.* Beckett, 2019.
[16] Westlund et al., 2021.

Market pressures also play a role. With the news industry still reeling from the collapse of its traditional business model,[17] many publishers are hoping AI will help to fight off this existential threat. In the words of a former audience analyst and manager from Germany:

> Well, listen: I think one of the truths about the media industry is that it is an industry that is under a certain obvious strain for cash, for new business models, figuring out what their future is. Basically this "What's going to save us?" question is all out there.

Recent surveys of news media executives corroborate this view. Hopes for financial profit and new and improved business models are key drivers for the adoption of the technology at many news organizations.[18] To cite a media manager in the United States:

> Revenue — […] how can I use these AI technologies to increase my audience, to increase my subscriber base, to increase the time that people are spending on the page and scrolling and viewing my gorgeous ads that are alongside it? That's a motivation for us.

These hopes for profit stem mainly from the promise that AI can deliver meaningful gains in efficiency and productivity by, for example, speeding up existing workflows (although, as I show in later sections, these promises should be taken with a grain of salt) or improving product experiences.

Competitive *industry dynamics* also play an important role. As Lucy Kueng has shown,[19] news organizations often anxiously watch their competitors, plagued by concerns that their own innovations have historically lagged behind those of their peers. My interviewees frequently cited this as a strong motivating factor for introducing AI to their organizations, suggesting that this dynamic is repeating itself.

A final theme to emerge was that *uncertainty* — together with *hype* and *hope* — around the future potential of AI as a set of technologies has engulfed publishers, driving investments, experimentation, and early adoption:

> I think it's the same with anything that there's hype around […] people that read any of the kind of reporting in tech are kind of aware of AI being something that has all this promise. So I think that quite naturally leads itself to people thinking, "Well, could we use it to do something," you know? (*Journalist, U.K.*)

[17] Nielsen, 2018.

[18] Newman, 2019; 2023.

[19] Kueng, 2017.

Collectively, these interlocking dynamics — *technological developments*, *market pressures*, *industry dynamics,* and *uncertainty, hype, and hope* — capture the main motivations driving the uptake of AI across the news organizations in my sample. Underpinning all of them is the hope that this technology will deliver *greater efficiencies* and unlock possibilities previously thought impossible.

1b: News Organizations' Use of AI in Production and Distribution

The current hype around generative AI somewhat masks the fact that AI is not entirely new terrain for publishers. As Charlie Beckett and his colleagues at the London School of Economics have shown, the technology has gradually moved into various aspects of news production and distribution in recent years, often in ways that audiences (and journalists) may not necessarily notice. While listing every single application of AI in the news is beyond the scope of this report,[20] the tasks for which news organizations currently turn to AI can be grouped into a small number of broad — albeit imperfect — categories. *(See Table 2.)*

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

[20] Anna Hansen and colleagues provide a very good overview in one of their white papers (Hansen et al., 2023). The Journalism AI project at the LSE has a growing list of case studies. *Case studies: Exploring the intersection of AI and journalism*. https://www.lse.ac.uk/media-and-communications/polis/JournalismAI/Case-studies

Many of these use cases can bring value to news organizations by, for example, aiding the creation of new products or features. One product manager in the United States described the implementation of a text-to-speech feature they felt enabled their organization to kill two birds with one stone:

> We have this product feature on our app where if you are, I don't know, commuting and you don't want to read the article I can have it read to me — that is an AI system that we use, that translates the article text into an audio file that you can read, and that has real advantages for accessibility.

For some publishers, this technology quietly contributes to the recommendation and curation of content on their owned and operated platforms. For example, a German data scientist described how machine learning informs article recommendations on their outlet's website and app:

> [These] article recommendations … come from a tool of ours […] which we then use. This does a classic look-alike model in the background, i.e., simply a segment based on the behavior of what users read, so a similar group is formed.

Another example comes from an editor at a U.K. publication where AI is used to improve recommendations offered to readers in newsletters:

> Some of the items in our newsletters for subscribers are automated with machine learning. Most of the newsletter is still curated by our editors, of course, but some parts are fully automated, yes.

The rise of generative AI has spurred a degree of creativity in this domain, although at the time of writing many organizations are still experimenting. However, contrary to some claims, AI is far from a silver bullet for many news-related tasks and often brings significant limitations. In fact, it may be a surprise to some that many of the tasks for which AI has so far proved most beneficial to news organizations are relatively mundane — or rather they may appear to be to anyone who has been taken in by the recent hype around LLMs.

1c: The Promise of Efficiency

AI's potential to increase efficiency in journalistic work is a topic of great debate. It was therefore unsurprising that this emerged as one of the core motivations cited by interviewees for adopting this technology in their organizations. This sentiment was particularly strong in relation to news production. As one U.S.-based executive described it:

> The strategic question is: With the limited amount of time and resources, how could we make the most use of our journalistic talent?

Another manager in Germany put it less poetically:

> What's the big driver behind the use of this technology? In its simplest form, it's automation in the pursuit of efficiency and productivity gains.

The efficiency debate can be broadly divided into two camps. On the one hand, there are experts and practitioners who believe that AI will significantly free up journalists, allowing them to focus on more creative and strategic tasks while the technology takes care of the grunt work. Others are more skeptical, arguing that AI's impact on productivity is likely to be more limited.

In the absence of solid empirical evidence to support either argument, it helps to look at concrete examples, guided by two overarching questions: (i) In what context are we talking about efficiency? (ii) What kind of AI do we mean?

Beginning with the *distribution or business side*, one example comes from a publisher that uses AI to enrich its podcast experience, providing users with additional information and recommendations based on their listening history. In this instance, various NLP and machine learning approaches were used to extract and analyze (meta-)data from existing podcasts, which was then combined with user data.

A manager at this organization described this as a combination of AI-induced efficiency and effectiveness:

> So, that was using AI to do something that, you know […] at a scale that wouldn't be [otherwise] possible, a level of detail that wouldn't be possible.

While the manager could not provide specifics, the results of this approach were deemed successful enough for the tool to be fully implemented. In this instance, AI helped to achieve a desired result (effectiveness) — a better experience for audiences and greater user retention at scale — and did so by providing efficiency gains in terms of processing large amounts of data in a reasonable amount of time and with minimal effort (efficiency).

Another salient example in this category is the use of dynamic paywalls. These increasingly popular[21] systems draw upon a vast trove of data points pertaining to individuals' behavior while

[21] Zaffarano, F. (2019, March 7). How Neue Zürcher Zeitung increased its conversion rate up to five times with dynamic "paygates." Journalism.co.uk. https://www.journalism.co.uk/news/how-nzz-increased-its-conversion-rate-up-to-five-times-with-dynamic-paygates-/s2/a735623/ *and* Supekar, R. (2022, Aug. 10). How The New York Times Uses Machine Learning To Make Its Paywall Smarter. *NYT Open*. https://open.nytimes.com/how-the-new-york-times-uses-machine-learning-to-make-its-paywall-smarter-e5771d5f46f8

using a website — time and duration of visit, device used, content consumed, time spent consuming content — to predict the likelihood of converting them to paying subscribers and adapting paywall access accordingly. Various machine learning approaches help to "essentially decide when […] you should see our journalism [which is] essentially all behind the paywall," as one data scientist at a U.S. organization explained it. They continued:

> [There is all this] math that goes into determining your propensity to subscribe. How likely are you to actually click the subscribe button? A lot of that computational prowess is essentially [about] trying to predict what is the right thing to show, [and] what is the right thing to hide.

Although publishers tend to be reticent about these systems, the data that flows into them, and their effectiveness at increasing conversion rates — my interviewees reported estimates ranging from 2 percent to 10 percent in comparison to a random policy, although these numbers are difficult to verify — they are becoming increasingly popular across commercial news outlets. A well-implemented and carefully fine-tuned paywall can be very impactful for a news organization's business, as Rohit Supekar, a data scientist at the *New York Times*, has described:

> The *Times* achieved its goal of 10 million subscriptions and set a new target of 15 million subscribers by the end of 2027. This success has been possible in part due to continuous improvements in the paywall strategy over the years.[22]

We may not know how big a difference AI makes in this context, but we can say with some confidence that it is significant enough that many leading publishers — including many in my sample — have decided to implement, keep, and improve these systems, suggesting they deliver meaningful efficiencies on the business side.

To a data scientist at a U.S. publisher, AI holds great appeal in this context because of its ability to help solve a complex optimization problem more efficiently:

> If we reduce free articles, like [reducing the number to] three for example, this will help subscriptions because users come to our site, they like our articles, but they can only read three and they want to read more, so they subscribe. But at the same time, it'll reduce overall visitors to the site, or overall impressions on the site, which will impact negatively on advertisements. So that's a grand optimization problem, and machine learning helps with that.

On the *production side*, many interviewees pinpointed AI tools (or approaches) that allow journalists to find connections in large datasets. One investigative reporter in the U.K., for example, identified fuzzy matching — a machine learning technique that can identify similar but not necessarily

[22] Supekar, 2022.

identical elements within a dataset — as a tool they frequently turned to when using large document sets to investigate subjects such as corruption and tax evasion:

> I think the fuzzy matching just speeds things up, and you quite quickly find out whether there's a match or not in the data and it saves you spending ages painstakingly going through documents looking for … looking for what you're looking for.

For this interviewee, it was not just that fuzzy matching made part of their work *more efficient* (although they were unable to specify exact time savings) — it actually made significant aspects of the work *possible in the first place.* Thus, AI ultimately *made them more effective* in their reporting because it enabled them to cover more stories than otherwise would have been possible. A similar experience was described by a team leader at a U.S. outlet:

> Well, [one of] the benefits [of AI] is that we can often look at data and look at data sources that we wouldn't usually get a sense for. […] It also lets us tell stories in very specific ways. I feel like our election forecasting model is exactly that. We're able to describe uncertainty and different possibilities of what might happen in a very visceral sense. Instead of just describing, we're able to show exactly what the outcomes may be. That just was not possible before.

Another news outlet in my sample has sped up production of its finance reporting by developing a system that combines machine learning and natural language processing to automate the process of analyzing and extracting key points from financial statements. An editor involved in building this system, which now operates largely autonomously, said:

> It gives our journalists the time to actually look at, say, contextual information, for example for a surprise announcement of a company. Say, the CEO being done for sexual harassment or whatever it happens to be, you know. It — it's freed up a lot of our journalists' time.

Similarly, a journalist at a U.K. news organization explained how an AI-assisted archive system has proved particularly valuable for its ability to streamline workflows during high-pressure breaking news situations:

> One of our toughest moments on the news desk is when something happens like a celebrity death, et cetera, right, and we quickly need to find archive material of an event or a person — so it's really good for those kinds of situations.

However, the main area in which AI appeared to deliver tangible improvements in efficiency was transcription. In fact, almost all interviewees brought up transcription as the foremost area where

AI makes a significant, measurable difference to their work — primarily in the form of significant time savings. One German editor put it thus:

> For my interviews, I must transcribe them like everyone else [laughs]. And that really takes time if I do it manually. Like, an hour-long interview would usually take me three or four hours to type up, [although it] kind of depends on how much I need, of course. With AI, that easily comes down to 15 minutes.

In this instance, the time saving is about 9 percent, although we should not lose sight of the fact that manual transcription allows journalists to develop a sense of the underlying material, which may speed up subsequent tasks.[23] A journalist at a U.K. publisher also explained that, in addition to time savings, AI transcription technology has benefited their work in other ways:

> I think it has made it more … I think it's probably made me more confident if I've missed something, I can always go back and have a quick read over the transcript. And if I need to I can check the recording.

Contrary to some of the early hype about AI, however, my research suggests that its ability to improve the efficiency of journalistic work — and the work of news organizations more broadly — is not as straightforward as it might be assumed. First, there is no one singular journalistic process that can be neatly separated and measured for efficiency gains (let alone automated with AI), just as there is no one single "AI" whose effect could be studied across the board.[24] This means that the impact of AI on journalistic work varies depending on the specific tasks being automated. In some cases, it may, in fact, *decrease* efficiency, e.g., if something produced by AI ends up needing to be laboriously checked by a human, or if its output cannot be fully trusted. These considerations can also limit the scope to scale up certain products or processes that use automation. As one U.K.-based newsroom manager explained:

> We pride ourselves on putting out trustworthy and reliable news. It's kind of in our statutes that things have to be reliable. So we have to have a handbrake on some systems, actually. Some things you cannot scale.

This is particularly true for large language models. While LLMs have become increasingly popular for a range of tasks — including summarization, translation, transcription, data processing, extractive and abstractive summarization of unstructured texts, creating article drafts, and

[23] The task used to take 4 hours, which is equal to 240 minutes. Now it takes 15 minutes, saving 225 minutes. If we define time savings as "(original time - new time) / original time," then in this case, the time savings is equal to (240 minutes - 15 minutes) / 240 minutes = 93.75 percent.

[24] In some cases, this will be clearer; in some cases our knowledge will likely remain hazy. Interviews and surveys are unable to paint an accurate picture as they rely on people's self-reports, which may not be accurate depending on what we are looking at specifically. Experimental studies would go further here, but these, too, are of limited use as they often do not capture the everyday reality of work — or, if so, only to a limited degree.

simplifying complicated writing[25] — they can be prone to producing unreliable results that hinder the journalistic process more than they help. As one U.K. interviewee put it:

> AI summarization can be wobbly. Depending on the length, it is really actually not very good, I find. I tried it a lot and, well, checking sometimes takes longer than writing a summary myself. Also, the story ideas it gives me are very homogenous. So, yes, it will get better, but I am not sure if this technology is the great flex people think it is.

Indeed, concerns about these kinds of limitations have led some newsroom leaders to conclude that LLMs' ability to deliver short-term efficiencies may currently be outweighed by their potential to cause longer-term reputational damage. For example, one editor at a U.K. organization said:

> Our newsroom is … actually [a] very conservative place because we've got to get things right. We've got to be very, very careful. We've got to think of just the normal editing process of, you know, how an editor commissions something, how a reporter goes out and reports it, they go through fact-checking for everything before it goes. I don't want to be BuzzFeed or CNET, just putting out sort of, you know, junk.

Far from liberating news workers, AI technology could introduce new demands to an already stressful profession. For example, one journalist described how automated transcription has allowed them to, in their words, "sort of be in two places at once," insofar as they can use the time they previously would have spent transcribing audio to watch something else or write up another story. While this is a relatively positive assessment, it raises a lingering question of whether — as optimists hope — efficiencies resulting from AI will enable journalists to do better and/or more in-depth reporting, or whether, in a journalistic version of the Jevons Paradox,[26] they will simply be expected to use the time savings to churn out more content. In other words, it's a question of whether AI will facilitate an increase in quality or quantity. A response from one U.K. editor who touched on this implies they expect the latter:

> It's freed up our journalists' time. But for any of those journalists who thought, "Oh my gosh, that's gonna take away my job." Oh no, don't worry. […] We've got more journalism for you to do.

Sociologist Randall Collins has a straightforward reply to this pursuit of efficiency. Maximal efficiency, he argues, is effectively a pipe dream. One can undoubtedly strive for improvements, but, he writes, as a wealth of studies from organizational theory show, "There is no such thing as a pure optimal solution to a situation of great complexity. […] If you try to optimize one thing, you

[25] Nishal & Diakopoulos, 2023, p. 1ff.

[26] "Any increase in efficiency in resource use will generate an increase in resource consumption."

sacrifice something else, [and] many of these processes involve uncertainties that you simply cannot control in advance."[27] A U.S.-based manager whom I interviewed agrees:

> I believe fairly strongly that the most effective and efficient AI tools I know about today are ones that are very much a hybrid system, where the machine is not deciding but the system is making a recommendation and a human is deciding. And I think that both helps with ethical concerns, but also just makes AI tools more efficient and more effective.

Reflection: AI in the News, A Difference That Makes a Difference?

What emerges here is a complex picture. The answer to the question: "Does AI fundamentally make a difference to the production and distribution of news?" must be both yes and no. The available evidence shows that AI has been — or can be — employed in a variety of settings to improve (or partially replace) a variety of tasks. Ultimately, though, I submit that what 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 that is imbued with ideas of rationality, efficiency, and speed — and which does indeed provide greater efficiency and effectiveness *in some contexts*.

Yet, given that the production and distribution of journalism is a complex sociotechnical system, it is inevitable that any attempt to disrupt the status quo by introducing automation and/or AI will encounter some form of resistance:

> There is a natural buffer against the adoption of this technology. Some of it is human and organizational. Some of it is technical. Some news organizations don't even have modern IT infrastructure, or they have CMS [content management systems] that are very old. There are so many things that they have to sort out first before they can even think about AI. (*Senior editor, U.K.*)

Resistance from news workers, adverse public opinion (or just the anticipation thereof), legislative conditions,[28] a lack of skills, insufficient data or technical infrastructure,[29] or a combination of these and other factors can and will act as bottlenecks.

Additionally, the news-making process is, to a large extent, only Taylorist in principle. While there are some standard procedures, the production and distribution of news is not an assembly line of neatly defined components that can be automated with AI. This is particularly true of news production, which is often a messy and unpredictable process that makes markedly different

[27] Collins, 1992, p. 86.

[28] *See, e.g.* Jerome, 1934, p. 19.

[29] Varian, 2019.

demands of journalists depending on the story, project, or deadline. The decidedly unscientific nature of this work is precisely what makes so much of it unsuited to automation. Take, for example, investigative journalism involving large datasets: While machine learning can help streamline certain tasks like detecting patterns or translating documents, much of the process currently remains beyond the technology's reach.

> Even people who have worked with AI on big datasets still say there are stories in there that we haven't found. And you still have to fact-check the information, too. There are so many steps that you cannot automate easily. (*Investigative journalist, Germany*)

A purely Taylorist view also underestimates the complexity of journalistic work, some of which will always defy automation because it either rests upon a large repertoire of embodied experiences and knowledge[30] or simply does not follow a standard procedure. For example, building a network of trusted sources or convincing sources to share their secrets — the remit of every good reporter — is not something AI will be able to achieve any time soon. Consequently, the integration of AI into journalists' day-to-day work is more complex than simply replacing human tasks with automated processes. As a senior U.K. editor quipped regarding LLMs:

> The job of journalism is to find stuff not on the internet already. Artificial intelligence won't be able to do that.

A German investigative reporter was also skeptical:

> How will the exclusive stuff we want to find out be in any kind of AI? It isn't. That's not where you find information that ultimately gives us exclusives.

A product manager at a U.S. organization shared a similar view with respect to AI more broadly:

> I don't see how you could really write some of the investigative stories where you're asking a specific question and … like, it only knows [as much as] it knows. So you do have to have someone who you know can think outside that box.

For applications of AI in news organizations, this means two things. First, far from being uniform, AI is used for a variety of tasks across the variety of settings in which aspects of journalistic work take place, such as news outlets' content management systems, reporters' mobile devices like phones or cameras, and the software used to create and distribute news. Second, the complex realities of publishing often constrain how and when the technology can be put to practical use, limiting the extent to which some of the more eye-catching capabilities showcased in controlled experiments or anecdotal accounts — often framed to foreground their supposed ability to achieve greater

[30] *See also* Scott, 2020, p. 329, where he refers to this as Mētis.

efficiency — can feasibly translate to a professional setting. While efficiency and productivity gains are real, they do not apply across the board. In a twist on the famous adage that "Culture eats strategy for breakfast," one could argue that "Workplace reality eats outsized expectations of AI for breakfast," given that applications of AI in news organizations are so often messy, varied, and idiosyncratic.

Chapter II: Platform Companies and AI in the News

In this section we turn our attention to the second piece of the puzzle: the role of platform companies in journalism. These companies have a long history of framing their products as efficient solutions to news organizations' problems — and efficiencies are again central to technology companies' pitch to news organizations about AI.

Over the last decade, technology companies such as Facebook (Meta), Google (Alphabet), Twitter (now X), Apple, and TikTok (ByteDance) have become influential actors in the news.[31] For example, they provide access to audiences through their platforms, and direct readers toward news content via search engines.[32] This not only makes news organizations partially dependent on these platforms for distribution,[33] it also allows these companies to shape the flow of attention online.[34] Additionally, these technology companies provide important services to the news industry, offering business-to-business products including cloud storage and computing, audience analytics, app developments, advertising exchanges, and revenue-sharing agreements.[35] Some platforms also fund journalism projects and research, with Google in particular standing out as the largest (and, as of the time of writing, still active) international funder of such schemes.[36]

At the same time, platform companies are also leaders in the development and application of artificial intelligence, as recent research by Nur Ahmed and colleagues shows.[37] Platform companies boast large in-house teams of computer scientists whose work covers every aspect of AI, and continue to invest heavily in the expansion of their AI capabilities. Many have acquired or invested in companies that are innovating in this space. Examples here include Google's acquisition of DeepMind in 2014 and Microsoft's $16 billion purchase of Nuance Communications, a company working on conversational AI and ambient intelligence across different domains, in 2021. Combined with a skillful exploitation of their existing infrastructure (e.g. servers and computing facilities, custom software, and the organizational structures that maintain and develop them), this has positioned platform businesses as important nodes and intermediaries in the AI field.

[31] Bell et al., 2017; Rashidian et al., 2019.
[32] Newman et al., 2022.
[33] Ananny, 2018; Nielsen & Ganter, 2022; Rashidian et al., 2019.
[34] Diakopoulos, 2019, p. 179; Newman et al., 2020, p. 23.
[35] Nielsen & Ganter, 2022, p. 69.
[36] Fanta & Dachwitz, 2020; Nechushtai, 2018; Papaevangelou, 2023.
[37] Ahmed et al., 2023.

Consequently, they now act as providers of AI services, tools and models, and infrastructure — all of which are increasingly required to build functioning and cost-effective AI applications — across industries.[38]

This has accelerated since the rise of "generative AI," an emerging field that swiftly muscled itself to the center of platform companies' long-term strategies when OpenAI launched ChatGPT in November 2022. In response to OpenAI's success, Google's CEO, Sundar Pichai, declared a "code red." Google has since unveiled new products — including its own chatbot, Bard, and its own family of multimodal LLMs, Gemini — and is building an advanced search engine that provides AI-generated answers to user queries.[39] Microsoft, in turn, has announced a multibillion-dollar investment in AI, betting that AI systems will have the power to transform the tech giant's business model and products and allow it to stay competitive.[40] It has also struck a deal to integrate OpenAI's technology into a range of its software products and rolled out Bing Chat, a new search engine feature that builds on OpenAI's GPT-4 system and, among other things, answers user queries with AI-generated replies. Given the nature of the technology — which has a high barrier to entry due to the vast data and computational power requirements — platform companies are likely to dominate this space for the foreseeable future.[41]

While it has been argued that the open-sourcing of AI models — which can provide powerful functionalities at low cost — and the emergence of a smaller crop of new firms such as Hugging Face and OpenAI will help counter this platform hegemony, claims that this signals a democratization of AI should be treated with caution. Technology companies of all sizes have numerous commercial incentives for "democratizing" their AI, from influencing market competitions and shaping standards to improving their corporate brands and hiring highly sought-after technical talent.[42] While open-sourcing broadens access to AI, it does not necessarily democratize resources, decision-making, or the creation of new AI.[43] Even when AI models are not developed by large companies, they often rely on their architectures and require substantial computing resources that often have to be licensed from these firms.[44]

In this context, the question that arises for us is obvious: Given platform companies' centrality to the AI space and their uneasy relationship with the news industry, how significant is their role in

---

[38] Simon, 2022.

[39] Grant, N. (2023, Jan. 20). Google Calls In Help From Larry Page and Sergey Brin for A.I. Fight. *New York Times*. https://www.nytimes.com/2023/01/20/technology/google-chatgpt-artificial-intelligence.html

[40] Bradshaw, T. and Criddle, C. (2023, Jan. 23). Microsoft confirms "multibillion-dollar investment" in ChatGPT maker OpenAI. *Financial Times*. https://on.ft.com/3H4lQgg

[41] Vipra & Myers West, 2023.

[42] *See also* Elizabeth Seger, What Do We Mean When We Talk About "AI Democratisation"? The Centre for the Governance of AI. Feb. 7, 2023. https://www.governance.ai/post/what-do-we-mean-when-we-talk-about-ai-democratisation

[43] Heaven, 2023.

[44] Whittaker, M. May 19, 2023. https://twitter.com/mer__edith/status/1659566944377241602?s=20

shaping the use of AI in the news industry? And what, if anything, does this mean for the public arena and the news we get to see?

2a: (For) Everything, Everywhere, All at Once? Where Publishers Use Platforms' AI

A primary objective of this research was to understand where and why the news organizations in my cross-national sample are using platform companies' proprietary AI products. The short answer to this is: Almost everywhere. As one manager from a big U.K. news organization told me, "You can't use AI without using these companies in some way."

News publishers in the United Kingdom, United States, and Germany use AI and related infrastructure provided by companies like Google, Amazon, and Microsoft to automate various tasks — including many of those described in the previous chapter.

> We encounter again and again situations where we end up using the tech giants' [AI] infrastructures. We might end up using their algorithms, or they provide us a service. We might use their cloud hosting systems, but we're going to build our own versions of these things. But they are always in there somewhere. (*IT manager, U.K.*)

As noted in the previous section, automated transcription is one of the foremost ways in which AI has been integrated into news production. This was also an area where many interviewees described a reliance on platforms' AI tools, citing their use of products like Amazon Transcribe[45] and Google WaveNet[46] to transcribe interviews, create automated subtitles, or generate audio for articles. Many news organizations also use platforms' pre-trained AI models to help investigate large documents or images. For example, the *Washington Post* uses Amazon Textract for advanced optical character recognition (OCR) when digitizing documents for investigative work. According to a public testimonial from Jeremy Bowers, the *Post*'s former director of engineering, this allows the paper's journalists to "study records of public interest" and extract "structured data that is found in newsworthy documents," a sentiment that was echoed by my interviewees.[47] Also popular are Google's Vision services, Amazon's Rekognition Image, and Microsoft's Azure AI Vision, which many interviewees described using to label and classify images, detect objects within images, and for OCR.

Examples also abound on the distribution side. The *Financial Times*' use of unsupervised machine learning for consistent article labeling relies on infrastructure provided by Amazon Web Services (AWS) and Google.[48] The German newspaper *Frankfurter Allgemeine Zeitung* (FAZ) has migrated

[45] Speech to Text — Amazon Transcribe. https://aws.amazon.com/transcribe/

[46] Text-to-Speech AI. https://cloud.google.com/text-to-speech

[47] Amazon Textract customers. https://aws.amazon.com/textract/customers/

[48] Gajtkowski, A. (2021, Feb. 9). FT Article Clustering. Medium. https://medium.com/ft-product-technology/ft-article-clustering-ffce1e8e32d0

its online offerings to Microsoft Azure, and partially uses Azure's machine learning to improve personalization.[49] The newspaper also employs Google's AI services for a machine learning tool that provides editors with predictions about which articles will work best behind the paywall.[50]

The publishers I interviewed mostly use these companies for business and distribution tasks. But in the future, as generative AI applications become commonplace, they will increasingly be used for creating and producing news. Most publishers in my survey depend on more than one of these big technology companies for infrastructure and services, with Google, Amazon, and Microsoft currently the most popular. While I observed no significant differences between commercial and public service organizations, it was notable that larger, better resourced news organizations were far more likely to do at least some in-house AI development than their smaller counterparts. The reason for this is fairly simple: The larger and wealthier an organization, the greater the likelihood it can dedicate time and resources to the development of custom-made applications, AI teams, and R&D resources. While there can, of course, be exceptions — flatter hierarchies and more nimble organizational structures may enable some smaller organizations to innovate at speed — my findings suggest that the high cost of custom AI development puts it out of reach for all but the best-resourced news organizations. For everyone else, the most viable solutions are third-party offerings from platform companies and the like.

2b: Cheaper, Quicker, Better: Why Publishers Rely on Platform Companies' AI

When asking publishers to describe their reasons for relying on platform companies' proprietary AI infrastructure, services, and applications, I quickly became accustomed to hearing variations on these statement from interviewees in Germany:

> We can't do everything ourselves. And if you want to do it, if you want to stay on an island, then you have so little data and so few resources that you're not getting anywhere. *(Data scientist, Germany)*

> We host most of our [AI] work in a Microsoft Azure environment. And there is so much out of the box, especially when it comes to kickstarting processes or building pipelines and AI applications. As a publishing house, we wouldn't sit down and build everything ourselves. It's presumptuous, and that's why you license or use their stuff. *(Manager, Germany)*

[49] https://customers.microsoft.com/cs-cz/story/faz-media-azure-de-germany *(link not archived)* and https://web.archive.org/web/20231215110328/https://www.jambit.com/kompetenzen/innovationstories/migration-und-weiterentwicklung-des-nachrichtenportals-faz.net/

[50] Rabenstein, G. (2021, Jun. 8). Using AI to predict what should go behind a paywall. Google News Initiative. https://blog.google/outreach-initiatives/google-news-initiative/using-ai-predict-what-should-go-behind-paywall/

Notably, though, while interviewees frequently pointed to the ease, convenience, and scale of platform companies' integrated offerings, many of them admitted they had misgivings. Their newfound reliance on the platform companies with whom so many have had a rocky recent relationship was born more out of necessity than choice. Describing the choice between using platform-provided AI or having no AI, one U.S.-based manager admitted:

> It's a real challenge because, you know, you're damned if you do and damned if you don't, right? It's really, really problematic because the industry is so challenged [economically].

As interviewees in all three countries explained, the independent development and implementation of AI solutions is not just prohibitively expensive, but sometimes almost impossible. The computing power required to train very large models is expensive, as is the hiring and retention of skilled personnel — computer scientists, software engineers, data scientists — so fierce is the competition for their talent. Here's how a team leader in the United States put it:

> Google and Facebook and Apple are our competitors […] in the space of tech talent. There's no way we can pay the amount of money that the big tech jobs can. That creates a problem.

A U.K.-based developer described the same issue:

> We want to do things. We want to experiment, but how do you[?] ... Where does the talent come from? We simply can't pay that kind of money.

News organizations also lack the vast amount of data required in many instances. Platform companies, by contrast, have been able to expand their hardware, network infrastructure, and software simultaneously, scaling up their operations with great efficiency and achieving ever greater economies of scale. Their resulting structural advantages in the AI space allow them to innovate at a scale and pace that makes it difficult for most other industries, including the news industry, to keep up:

> They're like the landlords who offer the computing power, cloud storage, and then they have these tenants, smaller AI startups. … Everything leads to Big Tech. [Even] all the smaller AI companies are dependent on Big Tech computing power. (*Journalist, U.K.*)

However, for all that some news workers lament the hold platform companies have in this arena, others embrace it because they do not see it as a news organization's mission to develop AI solutions for themselves. As one U.K. executive put it:

> If they have the best technology on the market, why should we not make use of that? We can't build everything from scratch … I mean, I don't think we should either.

This view of platform companies as AI service providers akin to utility providers was particularly pronounced on the business side, where it is often assessed through a cost-benefit lens. These interviewees argued that it was cheaper and more effective to rely on these companies because it lowered their financial risk. Many praised platform companies for their AI products' ease of use, stability, and scalability. As one product developer in the United States explained:

> It's easy to have success quickly. You upload your training data, click something together, even without being able to program very much, and you can build quite nice things.

Emphasizing the affordability of platform offerings, a U.K. data manager said:

> I think [the current situation] … means a pretty good position for news organizations […], because [AI] is much cheaper than it used to be, widely available, commoditized. And you know, I don't mind applying technology that is owned by one company as long as the price is right and the competition works.

What we are witnessing here is something that one might call the "Hotel California Paradox." Up in the distance, publishers see a shimmering light: *AI is the future. We have to be part of it*. But pursuing that distant light risks making them prisoners of their own devising. Or, as the song goes: "You can check out any time you like. But you can never leave."

2c: Relying on Platforms for AI: Does It Matter?

As we have seen, publishers already use AI tools provided by platform companies in a variety of ways across every part of their operations. The level of their dependence — both reluctant and welcome — described in the previous section leaves two final questions: *To what extent does this latest shift in control matter for the news?* And *What difference does it make to the public arena?*

A fruitful way to think about this shift is to look at the autonomy of the media. Broadly defined, autonomy refers to the absence of external control[51] and the ability for agents to act and make decisions according to their own logic.[52] The opposite of autonomy in a news context is media capture, where a news organization is under the influence of another agent, such as a government

[51] *See also* Philip Pettit's political theory of freedom or autonomy as the absence of domination (Pettit, 1999).

[52] Following Haveman and Gualtieri, I use logic here in the sense of institutional logics: "systems of cultural elements (values, beliefs, and normative expectations) by which people, groups, and organizations make sense of and evaluate their everyday activities, and organize those activities in time and space" (Haveman and Gualtieri, 2017).

or business, and loses some or all of its autonomy in relation to it. A specific aspect of this is infrastructure capture, where a news organization is dependent on the physical or digital resources and services provided by an external actor, thereby ceding some of its autonomy.[53]

In the context of platform companies and AI, the complexity of AI increases platform companies' control, creating lock-in effects that risk keeping news organizations tethered to the platforms and their products. This risk of *vendor lock-in* — reinforced by the high switching costs to which many of my interviewees alluded — undercuts publishers' autonomy on a macro level and leaves them vulnerable to price hikes and other whims of the vendor. An I.T. manager in Germany admitted:

> Lock-in effects and such really bother me […]. From my experience, switching tech providers isn't a casual affair. Think of the costs. It's not like moving a box from A to B. So you build a solid relationship with a specific provider, but of course that comes at a risk.

Low costs and stable pricing models are crucial for news organizations, particularly those whose inability to build or maintain their own tools and systems leaves them dependent on outside vendors and off-the-shelf solutions.[54] Platforms also possess *artifactual and contractual control over their AI*, giving them carte blanche to dictate what activities are permitted or restricted. This creates a familiar power imbalance between platforms and publishers whereby the latter are largely at the mercy of the former. In this instance, the platforms not only get to determine the overall conditions of use, they also have control over more granular terms, such as the extent to which they permit publishers to customize AI applications built on top of their technology — a dynamic that could end up restricting the tools or systems publishers can build, or affecting existing applications in unforeseen but problematic ways.

Indeed, some interviewees have already experienced the fallout of becoming overly dependent on a third-party AI service. Central to one cautionary tale is Graphiq, a U.S. company that provided publishers with, among other things, AI-informed search and interactive data-driven infographics,[55] before things took a turn in July 2017. As one U.S. journalist recalled:

> The AP was using it. The large papers … the *LA Times* was using it, and plenty of other major news organizations were using it. Nobody actually is using it anymore, because that company was bought by Amazon a few years ago — and Amazon decided to discontinue that service for newsrooms.

[53] Nechushtai, 2018; Simon, 2022; 2023.

[54] Nishal & Diakopoulos, 2023; Rinehart & Kung, 2022.

[55] Graphiq. Retrieved Dec. 17, 2023, from https://en.wikipedia.org/wiki/Graphiq

Graphiq paid lip service to the news industry in a statement ("We greatly enjoyed working with publishers over the last few years to help them tell the news and look forward to continuing to use our technology in other exciting areas"),[56] but publishers who had been using it were left hanging.

This story is far from unique to Graphiq, Amazon, or AI tools, recalling as it does Apple's purchase and shuttering of social analytics service Topsy,[57] Google's treatment of Freebase Gridworks (now OpenRefine) after its acquisition of Metaweb,[58] and current concerns that Meta is phasing out CrowdTangle, the widely used analytics tool it acquired in 2016. While Graphiq is obviously not a platform company, its story illustrates the risks publishers face when they place too many eggs in a third party's basket: If priorities or business interests change, news organizations can easily be cut off. To quote the same U.S. journalist again:

> A service simply vanishing is ... It's a total waste of time for newsrooms to have gone through all that effort.

A final recurring theme in my interviews was frustration at the opaqueness of these services' inner workings, a situation that forces news organizations to either place absolute faith in the platform companies that provide them, or expend valuable resources conducting laborious manual tests to try and peer inside the *black box* of these systems. This particularly matters on an individual, micro level. Many of my interviewees expressed concern that AI systems from external providers could undercut their autonomy by limiting discretionary decision-making abilities and journalistic values more broadly in subtle, unforeseeable ways, by structuring their view of what is newsworthy in ways that make it hard for them to think about counterfactuals or alternatives, or by introducing bias into their output. As one German data journalist at a broadcaster argued:

> If I send some images [for analysis] to a Google API and it's supposed to tell me what's there, then I don't know what it was trained with and what bias it might have. And that of course has an influence on what kind of story I might tell.

A U.K.-based journalist, who likened a platform company's AI tool he uses for investigations to an unreliable calculator, expressed similar skepticism:

> I think my main concern is: Is the tool missing something, is it a bad tool, is it misinterpreting what I want? And I think if you keep not finding stuff that you expect to, you can do, you know, more manual tests.

---

[56] Dave, P. (2017, Jul. 20). Amazon acquires Santa Barbara start-up Graphiq to try to bolster Alexa. *Los Angeles Times*. https://www.latimes.com/business/technology/la-fi-tn-graphiq-amazon-20170719-story.html

[57] Moon, A., and Fares, M. (2015, Dec. 16) Two years after acquisition, Apple shuts social analytics platform Topsy. Reuters. https://www.reuters.com/article/us-apple-topsy-idUSKBN0TZ2NV20151216/

[58] From Freebase Gridworks to Google Refine and now OpenRefine. RefinePro. *Retrieved Dec. 16, 2023, from* https://kb.refinepro.com/2012/10/from-freebase-gridworks-to-google.html

The number of news workers feeling unease about the opaqueness of these black boxes will surely grow as more are drawn toward the tools emanating from the boom in generative AI, the biggest of which are developed by platform companies — such as Google's Genesis, an experimental product to help produce news stories[59] — or are dependent on their financial or technological support, such as OpenAI's ChatGPT.[60]

> I think the models [like GPT-3 and 4] are too complicated, and I think we're going to be too reliant on these big companies that make them accessible to us. (*Data scientist, U.S.)*

Follow-up interviews conducted in spring and summer 2023 revealed that anxieties were already starting to emerge about the impact of these newer systems on news workers' personal autonomy. Interviewees expressed concerns about errors and bias, privacy and data protection, and being implicitly steered away from what they see as core values of their work.

But it isn't only individual journalists who are concerned. People representing news organizations at the institutional level are anxious, too, albeit on a more macro level. On the one hand, there are concerns about infrastructure capture and the conditions by which transformer models were trained, including with publishers' content:

> They're all dominating in AI … like, models, infrastructure, right? They provide so much and keep on growing with these models. That's terrifying. And, again, from a journalistic perspective, they're using all of our content. We're getting less for it, but it makes [platform companies'] systems better. (*Manager, U.S.*)

Further, as these powerful AI models become increasingly essential to everyday news work, the (im)balance of power will tip ever more toward the technology companies that provide access to them — at the expense of the news organizations whose journalism has been used to train and improve them. In economists' parlance, platform companies can extract "rent": payments that far exceed what is economically necessary to provide the service and make a profit.

Some publishers also worry about the effects on their business models, which partially depend on audiences reaching them via search or platforms — something that might not be a given in the future, especially as platforms consider AI-enhanced search experiences. As one U.K. editor put it:

> Yes, the technology has accelerated, and that is the driver for a lot of the adoption we are seeing. And of course, they [the platforms] say to publishers, "There is nothing to worry about, we care about quality news," and so on. But a lot of the decisions made by the platforms seem to have nothing to do with an improvement

[59] Mullin & Grant, 2023.

[60] Heaven, 2023; Lehdonvirta, 2023.

> to some of the things that matter for us. Platforms are absolutely driven by their own corporate interests. And I think, like, Google Search is a massive, massive issue.

Others echoed these concerns, including this German manager:

> Roughly two-thirds of our online audience come from search. And 90 percent through Google. That's a big risk for us, if clicks to our content become optional because Google has decided to go all-in on AI-enhanced search where users just get short answers.

Pondering the same existential threat, one U.S. journalist asked:

> Why would people still come to our website and read a story if they can get something [via AI-enhanced search] that is tailored to their interests? Something that's short and doesn't mean they have to make another click? And people will consume this information; I mean, we already do. It's convenient.

Time will tell whether platform companies' AI products and services become integral to news organizations' collective future. A crystal ball will be required to know how the rapidly changing information environment will impact the business of news. What is already clear is that platform companies exert a certain degree of control over the technological conditions under which news organizations operate — and that control will only grow as AI becomes more widely adopted.

Chapter III: Drawing the Loose Ends Together

Whose Interests Are Being Maximized?

A pivotal question regarding the integration of AI into journalism and the information ecosystem more generally is: Whose interests are being served? The answer to this question will arguably go a long way toward determining which logics come to dominate and, by extension, the extent to which AI (re)shapes news organizations and the public square. Differing priorities and expectations — particularly as they pertain to overpromised and underdelivered *efficiencies* — mean news organizations should brace for battles both internally (e.g. where motives and priorities differ between managers and employees, or between the business side and the newsroom) and externally (e.g. with platform companies, or audiences/the public).

Managers versus News Workers

The notion — particularly recurrent among interviewees in senior roles — that news organizations should wholeheartedly embrace AI and reap the (in their view, inevitable) rewards was at times

reminiscent of what political scientist and anthropologist James C. Scott terms "high-modernism": a "bold self-confidence about scientific and technical progress" and a "sweeping vision of how the benefits of technical and scientific progress might be applied."[61] However, as Scott has demonstrated, such belief systems rarely take into account that things might go awry — and that such benefits may not be equally distributed.

For now, AI systems mostly aid, rather than replace, journalists, product managers, or audience analysts. Consequently, there is a precarious balance within news organizations between top-down wishes and bottom-up interests as far as the adoption and use of AI is concerned. How long this will remain the case is hard to say. One could easily imagine more advanced LLMs replacing copy editors or illustrators, particularly at news organizations with limited resources. In such a scenario, it would be difficult to argue that the technology is doing anything more than maximizing the interests of those who call the shots at the expense of those "making the news." Despite public proclamations to the contrary, some managers I interviewed tacitly admitted that AI could replace certain jobs in the middle to long term.

Even if we assume for a moment that the technology remains mostly augmentative, we can again ask whose interests and which logics will win out. Joque has demonstrated the link between statistical systems, such as AI, and capitalist logics of increasing marginal utilities.[62] One cannot be disentangled from the other. We rarely talk about how AI systems could make journalistic work more creative, imaginative, or interesting. Instead, disproportionate emphasis is placed on the technology's potential to deliver increased efficiencies and productivity — and all in the hope that these gains will be deemed satisfactory and won't just lead to a shifting of goalposts whereby time savings are immediately filled with new or additional demands.

Putting job losses to the side for a moment, the use of AI will not automatically improve journalism and, by extension, the quality of information available to the public, if news executives make decisions that mean this is not what AI gets used for. A core part of what is often conceived of as "good" journalism is the work of reporting. While the crucial work of public service reporting can be *aided* by AI, this technology cannot entirely *replace* it or make it vastly more cost-effective. No AI can convey the horrors of war by going into a war zone and talking to a mother of starving children; nor can it gain the trust of a whistleblower that leads to a story that uncovers massive corruption.

Second, a considerable proportion of modern journalism already consists of desk-based work that, at its worst, is merely a regurgitation of existing material with a dusting of additional reporting. Rather than broadening audiences' horizons, this arguably delivers a narrower view of the world. Depending on the decisions managers make in the short to medium term, the use of AI could end

[61] Scott, 2020, pp. 4, 20.
[62] Joque, 2022.

up bolstering the latter version of journalism at the expense of the former — with knock-on effects for the quality of information in the public arena.

Both examples accentuate an often unacknowledged truth: No matter the shaping power of the technology, AI's effects on the news and the public arena will largely be determined by the decisions news organizations and managers make about when, where, and how it will get used. The technology might enable some of these uses, but it does not ultimately call the shots.

Platform Companies versus Publishers

Looking beyond intra-organizational dynamics, we can apply the same lens to the relationship between platform companies and publishers. Platforms' business priorities determine the algorithmic systems that are underpinning their products as well as their "objective functions," the weighted goals on which they are supposed to maximize (e.g. "engagement" for social media companies).[63] Unsurprisingly, platform and technology companies' development and deployment of AI follows the same logic. AI is a technology to drive rationality, efficiency, and speed, and is therefore utilized to make the operations of platform companies more efficient by providing better service quality, developing new products, and offering customization across their various business offerings.[64] Their bets on AI here have already paid off in some areas, cutting electricity costs in data centers or providing users with better experiences in search.[65]

None of this ends at platform companies' front doors, of course. Instead, it extends into the settings where their AI systems come to bear, which includes the news industry. As Papa and Kouros argue, the news industry has already adopted the Silicon Valley approach of solving problems through technology (see the heavy reliance on Big Data to address a raft of industry challenges, from revenue shortfalls to reaching and connecting with audiences). This also comes through in some of the journalism-facing products and formats developed by technology companies that incentivize the creation of content primed to circulate widely on social media (e.g. the so-called "pivot to video," the portrait "Story" format) and/or "solutions" that sell news organizations on the promise that the proprietary publishing product du jour somehow offers a viable route to sustainability (e.g. Facebook Instant Articles, Google AMP).[66]

AI doesn't just continue this dynamic: it intensifies it. It shifts newswork even further toward the technical and the logics of platform companies: prioritizing greater rationalization and calculability (on the audience side in particular), and efficiencies and productivity (where journalistic work is concerned). But the prevailing logic may not necessarily prioritize the welfare of journalism or the needs of audiences. Prevailing logic dictates that the reduction of human beings to a series of data

[63] Farrell & Fourcade, 2023, p. 231.
[64] Barwise & Watkins, 2018, p. 29
[65] Big tech and the pursuit of AI dominance. (2023), p.5. Hindman, 2018.
[66] Papa & Kouros, 2023.

points that can be quantified and controlled is the key to understanding a news organization's audience. However, one can — and should — question whether this approach truly produces a deeper understanding than, for instance, those that recognize audience members as complex individuals with diverse backgrounds and perspectives.

The Public Arena in the Age of AI

News organizations are a vital component of the public arena, acting as gatekeepers for the common attention space most of us share. As news organizations change through technology, so does the makeup of the broader system that they constitute and shape.

How will the increasing use of AI play out in this context? One answer is that it will reinforce and perhaps even exacerbate existing inequalities among publishers. Organizations that have been able to make early investments in AI are far more likely to reap the rewards of the technology than those that have been unable or unwilling to embrace it. Early indications of this are already showing, with the usual suspects — typically well-resourced, international publishers — gaining an edge over their competitors. Local news organizations and publishers in the Global South are often an afterthought in the current conversations around AI, despite a slew of studies demonstrating that local news plays a vital democratic role within smaller communities and drives various forms of accountability on the micro-level of democracy. That said, the prospect of certain news organizations getting a head start need not be bad for the public arena if those "winners" use their newfound powers for good and double down on providing quality journalism to a plurality of audiences. There are, however, no guarantees that this will happen, not least because decisions about how AI gets used are made by executives whose primary concerns may differ from communication scholars who bang on about the importance of strong and well-resourced newsrooms.

Leaving aside the ways in which AI may end up strengthening journalistic work through greater rationalization in various areas, the public arena will also be reshaped in accordance with how the balance of control shakes out between platforms and publishers. While traditional news organizations continue to hold a great deal of control over what does and does not end up as news, their control has been greatly diminished in recent years, thanks to the rise of digital media — which has significantly lowered the cost of producing, publishing, sharing, and consuming information — and the emergence of platforms as central information intermediaries. The use of AI by platform companies may well end up further weakening this structural role of the news. In many ways, it already has, with platform companies shaping what audiences see online, not least through their use of AI to rank, curate, filter, and increasingly create and display information on their social media platforms and search engines. News organizations have played — and continue to play — roles in all of this, often as providers of information that can be found through search or shared on social media more generally.

But these things are not set in stone. The 2023 edition of the *Reuters Institute Digital News Report*, alongside various other studies, reveals a significant decline in news outlets' direct access to audiences, due to audiences' increased usage of third-party platforms and aggregators for news content. This trend is particularly pronounced among younger users, who perceive social media platforms as accessible and engaging, and gravitate toward interactive formats that prioritize personalities.[67]

At the same time, traffic from social media platforms such as Facebook is declining, partly as a result of parent company Meta's pivot away from news, the impact of which was felt particularly hard by smaller publishers.[68] It is entirely conceivable that the visibility of news on platforms will diminish, depending on how generative AIs are integrated into search engines and other products. Some organizations fear losing as much as half of the audience reach they currently get from search. The consequences could be dire. As John Herrman writes, the informal deal with publishers that has sustained them for years was effectively "You make content; we send traffic." This, in turn, offered publishers the prospect of advertising revenue, subscription conversions, and/or e-commerce revenue.[69]

It is far from a given that platform companies and especially search engines like Google will continue to afford visibility to news content and send valuable traffic to publishers' sites. Crucially, this will depend on strategic choices made by a set of powerful actors over whom the news industry has little control, but whose decisions could have severe ramifications for publishers — both in terms of their financial position and in their ability to reach audiences. As one senior manager at a U.K. publisher put it:

> Current tests [of AI from platform companies] are very hard to judge, but from what we've seen there are grave risks to referral [traffic] — and also reputational risk in it citing us against content that may be inaccurate or libelous.

The introduction of generative AI tools like the Search Generative Experience (SGE) at Google, which provides AI-powered overviews combining relevant information for user searches, offers early clues as to where the journey could go. Now expanded to more than 120 countries and territories with support for a range of languages, "SGE allows for easier follow-up questions, AI-powered translation assistance, and more definitions for various topics."[70] Products such as this could lead to a shift in the way users interact with search engines, potentially affecting the amount of traffic directed toward publishers' sites.

---

[67] Newman et al., 2023.

[68] Majid, A. (2023, May 4). As Reach warns of traffic slowdown: How Facebook referrals to publishers have plummeted. *Press Gazette*. https://pressgazette.co.uk/platforms/how-far-facebook-referral-traffic-to-news-sites-has-plummeted/

[69] Herrman, 2023.

[70] Budaraju, H. (2023, Nov. 8). Generative AI in Search expands to more than 120 new countries and territories. The Keyword. https://blog.google/products/search/google-search-generative-ai-international-expansion/

Perhaps the cruelest irony of all is that in using platform companies' AI services, news organizations are playing a key role in improving the very AIs that may ultimately pose an existential threat to their business models and position as gatekeepers. The stock of high-quality language data available on the internet has already been used extensively to train LLMs — and further data required for training "is locked away in small amounts in corporate databases or on personal devices, inaccessible at the scale and low cost that Common Crawl allows."[71] Whenever news organizations (i) provide access to their own structured data (as is now the case for the Associated Press and Axel Springer, who struck individual deals with OpenAI giving the company access to their archives as well as new content), (ii) allow platform companies to scrape their content, or (iii) use platform companies' AI products on their own data (particularly where options to decline data sharing are impossible or impractical), they only end up improving these systems. This risk is particularly pronounced with easy-to-use off-the-shelf tools. For example, AWS's AI services such as Amazon Rekognition, Amazon CodeWhisperer, or Amazon Transcribe are by default using users' data to train the company's own models—as it specifies in its terms of service: "[We] might store and use customer content processed by those services for the development and continuous improvement of other AWS services."[72] While opting out is possible, it is not a straightforward process and many newsrooms will not necessarily be aware of this issue. Given that continuous learning is central to AI, this could provide a pathway for platform companies to not only build better general-purpose AI products and services — which would reinforce their hegemony in the AI space, thereby further cementing their control over information[73] — but also potentially enable them to take on tasks that were once central to the news, such as providing their audiences with vital information about public affairs, political positions, and the like. Whether this would be beneficial to the public arena and broader publics is anyone's guess.

Conclusion

It is easy to assume that new technology is destined to make a vast difference to our lives or to certain industries, especially when the hype machine is in full flow. AI and journalism are no exception in this regard. That brings us to the central question of this report: What impact will AI have on the future of news and the public arena? As things stand, the only reasonable response has to be: It depends. That, I concede, is unlikely to be a popular answer. But context and nuance matter. Valid answers depend on it, as the sociologist Charles Tilly once put it.[74]

AI, I argue, for now mostly constitutes a "retooling" of the news rather than a fundamental change in the needs and motives of news organizations. It does not impact the fundamental need to access and

[71] Hodson, 2023.

[72] AI services opt-out policies. AWS. Retrieved Dec. 10, 2023, from https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_ai-opt-out.html

[73] Dolata, 2018.

[74] Tilly & Goodin, 2006.

gather information, to process it into "news," to reach existing and new audiences, and to make money. The ways in which news organizations go about pursuing these needs has already been changed by digital technologies — and they will change further with the arrival and implementation of AI.

That said, I am in no way dismissive of the shaping power of AI. Based on available evidence, it seems increasingly clear that *AI will play a transformative role in reshaping news work, from editorial to the business side.* 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. It is important to recognize that the extent of this reshaping will vary based on the specific context and task at hand, and will also be influenced by institutional incentives and decisions.

In this context, *winners and losers will emerge. In fact, they already have*. News organizations that have been able to invest in research and development, devote staff time, attract and retain talent, and build infrastructure already have something of a head start when it comes to adopting new AI technologies and developing new products and services in meaningful ways. These "winners" are also in a stronger position to demand better terms when negotiating with platforms and technology companies, e.g. regarding the release of news content to train AI technology. While major media outlets or publishing groups like News Corp, Axel Springer, or *The New York Times* can engage in direct negotiations with the likes of OpenAI, Google, or Microsoft, *The Philadelphia Inquirer*, *Offenbach Post*, or the *Oxford Mail* might not be so lucky.

As news organizations get *reshaped by AI, so too will the public arena that is so vital to democracy and for which news organizations play a gatekeeper role.* Depending on how it is used, AI has the potential to structurally strengthen news organizations' position as gatekeepers to an information environment that provides "people with relatively accurate, accessible, diverse, relevant, and timely independently produced information about public affairs" which they can use to make decisions about their lives.[75] For this to be achieved, *news organizations must use AI to help them (i) strengthen their business operations (thereby improving the conditions that make journalism viable and sustainable in the first place) and/or (ii) improve the quality of their output and the manner in which they serve their audiences (i.e. strengthen reporting and the provision of quality news).* This, however, is not a foregone conclusion. Instead, it will depend on decisions made by the set of actors who wield control over the conditions of news work — executives, managers, and journalists, but also increasingly technology companies, regulatory bodies, and the public.

[75] Nielsen, 2017, p. 1251.

Coda: A Few Final Thoughts About the Future

Most of the research for this report took place before the explosion of hype around ChatGPT and other large language and transformer models in the winter of 2022. (As noted earlier, additional interviews were conducted to capture stakeholders' thoughts about the impact of ChatGPT et al.) These newly prominent forms of AI have generated much speculation about these models' capacity to produce news content, the accessibility and reliability of the data sources and techniques they use to generate text and images, and the potential for these sources to provide misleading information. They have also been discussed in terms of copyright issues, liability, and the existential risks they may pose.

These debates, in turn, have affected discussions of news: What if AI is used to write news? Will journalists be laid off en masse? How are we to tell whether a human or AI wrote a story? Should we be able to tell, and in which context? And what will audiences think? These and similar questions can and perhaps will be taken up in future Tow reports. For now, it is important to maintain a sense of perspective. This report deals with developments that are already ongoing — many of which are instructive for these newer forms of AI — and once there is more evidence of their implications, these more speculative questions may be analyzed in more depth.

Part of this perspective comes from looking at the past: a past that the future might not repeat, but one with which it often sings in tune.

AI will be far from the only thing that shapes the news and the public arena in the coming years. Journalism does not change only through a single technology. To quote economic historian Carl Benedikt Frey, "Technology is not a soloist but part of an ensemble. It interacts with institutions and other forces in society and the economy."[76]

Productivity gains from the use of AI in the news will not be straightforward. Technology often improves productivity, but only after long delays. As economist Robert Solow once quipped, "You can see the computer age everywhere but in the productivity statistics." The benefits of AI to the news will be staggered. Its use will incur costs in the early stages and require organizational and strategic changes.[77]

The adoption of AI in news organizations will not be frictionless. Regulation, resistance from news workers, audience preferences, and incompatible technological infrastructure are just some of the variables that will shape the speed at which news organizations adopt AI, and, by extension, the rate at which AI's tangible effects on news creation come into focus. The speed of adoption should not be expected to move evenly across domains and applications — first, because some areas will be easier than others to automate with AI, and also because some organizations will have an easier

[76] Frey, 2019, p. 22.
[77] Frey, *ibid.*, p. 326.

time adopting AI than others. This is another reason winners and losers will emerge, another factor that will shape the composure of the public arena.

AI will not be a panacea for the many deep-seated problems and challenges facing journalism and the public arena. Technology alone cannot fix intractable political, social, and economic ills. Political attacks will not stop because news organizations use AI. Audience habits and consumption patterns will not revert to those of a bygone era. Instead, news organizations will still be forced to make a case for why they still matter in this modern news environment — and why they are still deserving of audiences' attention and money. The use of AI can help address some of these issues, but only the most deluded of Silicon Valley acolytes would believe that AI can miraculously solve them overnight.

The concentration of control over AI by a small handful of major technology companies will remain a key area of scrutiny. Neither established platform companies nor the fledgling start-ups developing (generative) AI necessarily care much for the concerns of publishers, or indeed the concerns of the public. They are large firms interested in concentrating information and making revenue by seeking efficiency gains and new business opportunities. But decisions these platforms make — about how AI gets used across the communication structures they control, who gets access to the technology, and the conditions under which that access is granted — will matter greatly. Control over infrastructure confers power. Structural dependencies around AI will likely chip away at news organizations' autonomy — potentially undermining their business models and thus their long-term viability — leading many to reconfigure themselves in ways that bring them yet closer to the logics of the technology sector and platform companies. At the same time, a tightening of the technology sector's stranglehold on (i) people's attention and (ii) information — as well as their increased capacity to manage, analyze, process, and serve that information — will further reshape the makeup of the public arena.

Developing frameworks to balance innovation through AI in the news — which is bound to continue — with concerns around issues like copyright and various forms of harms will remain a difficult and imperfect, but necessary task. As the sociologist Alondra Nelson puts it, "There are always harms that we can't foresee or that we can't anticipate, use cases that we might have thought about but didn't consider quite in the right way."[78] However, these technologies and their use can be shaped, and their risks can be assessed and mitigated. Not all of this work can or should be done by publishers, but they must not shirk their responsibilities in this regard. Luckily, as has been evidenced by the push to establish AI guidelines and develop responsible forms of AI, a growing number of publishers are already taking these risks seriously.[79]

[78] Transcript: Ezra Klein Interviews Alondra Nelson. (2023, Apr. 11). *New York Times.* https://www.nytimes.com/2023/04/11/podcasts/ezra-klein-podcast-transcript-alondra-nelson.html

[79] Becker et al., 2023.

As with any new technology entering the news, the effects of AI will neither be as dire as the doomsayers predict nor as utopian as the enthusiasts hope. AI's power to shape society and institutions such as journalism will be subject to the contexts in which it is used. It will be limited by professional norms and resistance to the technology itself, as well as technical and organizational bottlenecks or "reverse salients" that for now hold back its technological momentum. But it would be wrong to assume that it is a passing fad. AI has already had an impact on journalism, the news industry, and by extension the public arena. This impact will only increase. But its true size and significance will only become clear with time.

Acknowledgments

First, my sincerest thanks go out to my interviewees for their enthusiasm in participating in this research project and for answering all the (follow-up) questions I had. This report would not have been possible without all the news workers at various organizations who were willing to lend me their time. The same is true for the experts who talked to me. They all know who they are, and I hope I have done them justice in reflecting their work and industry.

I am also indebted to Michelle Disser and my PhD supervisors at the Oxford Internet Institute and the University of Cambridge, Ralph Schroeder, Gina Neff, and Ekaterina Hertog, for helping me coordinate this work with my overall PhD thesis and for providing crucial input at various stages of this project. A particularly big note of thanks goes to the Tow Center at Columbia University for funding and supporting this research, and especially to its research director, Pete Brown, for his patience, encouragement, and valuable input in cutting down the behemoth of a manuscript I submitted to something more readable. Hana Joy and Katie Johnston I am grateful to for their administrative support. Vicky Walker did a wonderful job of copy-editing the final version. A special thanks also to Emily Bell, who has kindly acted as a sounding board for me on some of the themes discussed herein, and Andreas Jungherr, who kindly provided feedback on the final draft.

I would also like to express my sincerest thanks to my colleagues and friends at Balliol College, but especially the Reuters Institute and the Oxford Internet Institute for their unwavering support and valuable contributions to my work. There are too many to list them all but they, too, know who they are. Additionally, I am profoundly appreciative of all the individuals I have had the privilege and pleasure of encountering over the years, who dedicated their time and efforts to studying, writing, and contemplating this topic. As it is impossible to acknowledge each of you individually within these pages, let me just say that you all played a role in shaping my understanding and perspective — something for which I am tremendously grateful.

Finally, it goes without saying that all mistakes are mine, and mine alone.

Felix M. Simon, Oxford, December 2023

Biography

Felix M. Simon is a communication researcher and doctoral student at the Oxford Internet Institute (OII) and Balliol College at the University of Oxford, where he has been studying the effects of AI in journalism and the news industry since 2019. His research seeks to understand the structural implications of AI—including forms of generative AI—for news organizations' production and distribution processes as well as the public sphere. Felix has published and presented at a number of leading academic journals and conferences and has co-authored various research reports and papers on topics ranging from innovation in the media to COVID-19 misinformation.

His research and commentary have appeared, among others, in *The Guardian*, *The Washington Post*, *Politico*, and the *Financial Times* and he has given evidence to inquiries of the UK House of Lords and House of Commons, press regulator IMPRESS, and the United Nations, among others. In May 2023, he was awarded the Hans Bausch Media Prize by German public broadcaster SWR in cooperation with the Institute for Media Studies at the University of Tübingen for his work on AI, news, and platform companies. Felix is a Knight News Innovation Fellow at Columbia University's Tow Center for Digital Journalism, and an affiliate at the Center for Information, Technology, and Public Life (CITAP) at the University of North Carolina at Chapel Hill. He also works as a research assistant at the Reuters Institute for the Study of Journalism (RISJ).

He holds a BA in Film and Media Studies as well as English Studies from Goethe-University Frankfurt and an MSc in Social Science of the Internet from the OII. He is currently a fellow at the Salzburg Global Seminar and an Associate Fellow of the UK Higher Education Academy and sits on the AI and Local News Steering Committee of Partnership on AI.

He can be found on Twitter, BlueSky and LinkedIn.

References

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

Ananny, M. (2018). *The partnership press: Lessons for platform-publisher collaborations as Facebook and news outlets team to fight misinformation.* (Tow Center for Digital Journalism Publications). Tow Center for Digital Journalism, Columbia University. https://www.cjr.org/tow_center_reports/partnership-press-facebook-news-outlets-team-fight-misinformation.php/

Bartholomew, J., & Mehta, D. (2023, May 26). *How the media is covering ChatGPT*. Columbia Journalism Review. https://www.cjr.org/tow_center/media-coverage-chatgpt.php

Barwise, P., & Watkins, L. (2018). The evolution of digital dominance: How and why we got to GAFA. In *Digital Dominance: The Power of Google, Amazon, Facebook, and Apple* (pp. 21–49). Oxford University Press. https://fdslive.oup.com/www.oup.com/academic/pdf/openaccess/9780190845124.pdf

Becker, K. B., Simon, F. M., & Crum, C. (2023). *Policies in Parallel? A Comparative Study of Journalistic AI Policies in 52 Global News organizations*. SocArXiv. https://doi.org/10.31235/osf.io/c4af9

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

Bell, E. J., Owen, T., Brown, P. D., Hauka, C., & Rashidian, N. (2017). *The Platform Press: How Silicon Valley Reengineered Journalism* (Tow Center for Digital Journalism Publications). Tow Center for Digital Journalism, Columbia University. https://doi.org/10.7916/D8R216ZZ

Big tech and the pursuit of AI dominance. (2023, March 26). *The Economist*. https://www.economist.com/business/2023/03/26/big-tech-and-the-pursuit-of-ai-dominance

Broussard, M. (2018). *Artificial Unintelligence. How Computers Misunderstand the World.* (1st ed.). MIT Press. https://mitpress.mit.edu/9780262537018/artificial-unintelligence/

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

Collins, R. (1992). *Sociological Insight: An Introduction to Non-Obvious Sociology.* (2nd ed.). Oxford University Press. https://search.worldcat.org/title/268794703

Diakopoulos, N. (2019). *Automating the News: How Algorithms Are Rewriting the Media*. Harvard University Press. https://www.hup.harvard.edu/books/9780674976986

Dolata, U. (2018). Internet Companies: Market Concentration, Competition and Power. In *Collectivity and Power on the Internet* (pp. 73-93). SpringerBriefs in Sociology. Cham: Springer. https://doi.org/10.1007/978-3-319-78414-4_5

Esposito, E. (2022). *Artificial Communication: How Algorithms Produce Social Intelligence*. The MIT Press. https://doi.org/10.7551/mitpress/14189.001.0001

Fanta, A., & Dachwitz, I. (2020). *Google, the media patron. How the digital giant ensnares journalism.* (103; OBS-Arbeitsheft). Otto Brenner Stiftung. https://doi.org/10.31235/osf.io/3qbp9

Färdigh, M. (2010). *Comparing Media Systems in Europe: Identifying Comparable Country-level Dimensions of Media Systems* (QoG WORKING PAPER SERIES 2010: 2). The Quality of Government Institute. https://www.gu.se/sites/default/files/2020-05/2010_2_Fardigh.pdf

Farrell, H., & Fourcade, M. (2023). The Moral Economy of High-Tech Modernism. *Daedalus*, *152*(1), 225–235. https://doi.org/10.1162/daed_a_01982

Frey, C. B. (2019). *Technology Trap: Capital, Labor, and Power in the Age of Automation*. Princeton University Press. https://press.princeton.edu/books/hardcover/9780691172798/the-technology-trap

García Avilés, J. A., León, B., Sanders, K., & Harrison, J. (2004). Journalists at digital television newsrooms in Britain and Spain: Workflow and multi-skilling in a competitive environment. *Journalism Studies*, *5*(1), 87–100. https://doi.org/10.1080/1461670032000174765

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

Haveman, H. A., & Gualtieri, G. (2017). *Institutional Logics*. Oxford: Oxford University Press. https://oxfordre.com/business/view/10.1093/acrefore/9780190224851.001.0001/acrefore-9780190224851-e-137

Heaven, W. D. (2023, December 5). *The open-source AI boom is built on Big Tech's handouts. How long will it last?* MIT Technology Review. https://www.technologyreview.com/2023/05/12/1072950/open-source-ai-google-openai-eleuther-meta/

Herrman, J. (2023, June 1). *Will Google's AI Plans Destroy the Media?* Intelligencer. https://nymag.com/intelligencer/2023/06/will-ai-powered-google-eat-the-publishing-industry.html

Hindman, M. (2018). *The Internet Trap: How the Digital Economy Builds Monopolies and Undermines Democracy*. Princeton University Press. https://press.princeton.edu/books/hardcover/9780691159263/the-internet-trap

Hodson, H. (2023, April 22). Large, creative AI models will transform lives and labour markets. *The Economist*. https://www.economist.com/interactive/science-and-technology/2023/04/22/large-creative-ai-models-will-transform-how-we-live-and-work

Jerome, H. (1934). Changes in Mechanization in Selected Manufacturing Industries. In Mechanization in Industry (pp. 55-119). National Bureau of Economic Research, p. 19. https://www.nber.org/system/files/chapters/c5242/c5242.pdf

Joque, J. (2022). *Revolutionary Mathematics: Artificial Intelligence, Statistics and the Logic of Capitalism* (1st ed.). Verso. https://www.versobooks.com/products/816-revolutionary-mathematics

Kueng, L. (2017). *Going Digital. A Roadmap for Organisational Transformation* (Reuters Institute Report, p. 48). Reuters Institute for the Study of Journalism. https://reutersinstitute.politics.ox.ac.uk/sites/default/files/2017-11/Going%20Digital.pdf

Lehdonvirta, V. (2023, March 17). *Behind AI, a massive infrastructure is changing geopolitics*. Oxford Internet Institute. https://www.oii.ox.ac.uk/news-events/news/behind-ai-a-massive-infrastructure-is-changing-geopolitics

Lewis, S. C., & Simon, F. M. (2023). Why human-machine communication matters for the study of artificial intelligence in journalism. *In* A. L. Guzman, R. McEwen, & S. Jones (Eds.; 1st ed.), *The SAGE Handbook of Human-Machine Communication* (pp. 516–523). SAGE Publishing. https://ora.ox.ac.uk/objects/uuid:39a45af3-086a-42d6-85be-d11208f6b531

Mitchell, M. (2019). *Artificial Intelligence: A Guide for Thinking Humans*. Macmillan. https://us.macmillan.com/books/9780374715236/artificialintelligence

Mullin, B., & Grant, N. (2023, July 20). Google Tests A.I. Tool That Is Able to Write News Articles. *The New York Times*. https://www.nytimes.com/2023/07/19/business/google-artificial-intelligence-news-articles.html

Nechushtai, E. (2018). Could digital platforms capture the media through infrastructure? *Journalism*, *19*(8), 1043–1058. https://doi.org/10.1177/1464884917725163

Newman, N. (2019). *Journalism, Media, and Technology Trends and Predictions 2019* (Reuters Institute Report, p. 48). Reuters Institute for the Study of Journalism. https://reutersinstitute.politics.ox.ac.uk/our-research/journalism-media-and-technology-trends-and-predictions-2019

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

Newman, N., Fletcher, R., Robertson, C. T., Eddy, K., & Nielsen, R. K. (2022). *Reuters Institute Digital News Report 2022* (Digital News Report). Reuters Institute for the Study of Journalism. https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2022

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

Newman, N., Fletcher, R., Schulz, A., Andi, S., & Nielsen, R. K. (2020). *Reuters Institute Digital News Report 2020* (Digital News Report, p. 112). Reuters Institute for the Study of Journalism. https://www.digitalnewsreport.org/survey/2020/

Newport, C. (2023, April 13). What Kind of Mind Does ChatGPT Have? *The New Yorker*. https://www.newyorker.com/science/annals-of-artificial-intelligence/what-kind-of-mind-does-chatgpt-have

Nielsen, R. K. (2012). *Ground Wars: Personalized Communication in Political Campaigns*. Princeton University Press. https://press.princeton.edu/books/hardcover/9780691153049/ground-wars

Nielsen, R. K. (2017). The One Thing Journalism Just Might Do for Democracy. *Journalism Studies*, *18*(10), 1251–1262. https://doi.org/10.1080/1461670X.2017.1338152

Nielsen, R. K. (2018). The Changing Economic Contexts of Journalism. *In* T. Hanitzsch & K. Wahl-Jorgensen (Eds.), *Handbook of Journalism Studies* (2nd ed.). Routledge. https://rasmuskleisnielsen.files.wordpress.com/2018/05/nielsen-the-changing-economic-contexts-of-journalism-v2.pdf

Nielsen, R. K., & Ganter, S. A. (2022). *The Power of Platforms: Shaping Media and Society*. Oxford University Press. https://global.oup.com/academic/product/the-power-of-platforms-9780190908867

Nishal, S., & Diakopoulos, N. (2023). *Envisioning the Applications and Implications of Generative AI for News Media*. 7. https://nishalsach.github.io/pdfs/2023-genaihci-chi.pdf

Papa, V., & Kouros, T. (2023). Do Facebook and Google Care about Journalism? Mapping the Relationship between Affordances of GNI and FJP Tools and Journalistic Norms. *Digital Journalism, 11*(8), 1475-1498. https://doi.org/10.1080/21670811.2023.2211626

Papaevangelou, C. (2023). Funding Intermediaries: Google and Facebook's Strategy to Capture Journalism. *Digital Journalism*, *0*(0), 1–22. https://doi.org/10.1080/21670811.2022.2155206

Pettit, P. (1999). Liberty as Non-Domination. In *Republicanism: A Theory of Freedom and Government* (pp. 43-73). Oxford Academic. Retrieved from https://doi.org/10.1093/0198296428.003.0003 on December 16, 2023.

Rashidian, N., Tsiveriotis, G., Brown, P., Bell, E., & Hartstone, A. (2019). *Platforms and Publishers. The End of an Era* (Tow Center for Digital Journalism Publications). Tow Center for Digital Journalism, Columbia University. https://academiccommons.columbia.edu/doi/10.7916/d8-sc1s-2j58

Rinehart, A., & Kung, E. (2022). *Artificial Intelligence in Local News: A survey of US newsrooms' AI readiness* (p. 56). The Associated Press. https://www.ap.org/assets/files/ap_local_news_ai_report_march_2022.pdf

Scott, J. C. (2020). *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. Yale University Press. https://yalebooks.yale.edu/book/9780300078152/

Siegele, L. (2023). How AI could change computing, culture and the course of history. *The Economist*. Retrieved April 25, 2023, from https://www.economist.com/essay/2023/04/20/how-ai-could-change-computing-culture-and-the-course-of-history

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

Simon, F. M. (2023). Escape Me If You Can: How AI Reshapes News Organisations' Dependency on Platform Companies. *Digital Journalism*, *0*(0), 1–22. https://doi.org/10.1080/21670811.2023.2287464

Tilly, C., & Goodin, R. E. (2006). It Depends. *In* R. E. Goodin & C. Tilly (Eds.), *The Oxford Handbook of Contextual Political Analysis* (pp. 3–32). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199270439.001.0001

Varian, H. (2019). Artificial Intelligence, Economics, and Industrial Organization. *In* A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The Economics of Artificial Intelligence: An Agenda* (pp. 211-234). University of Chicago Press. https://doi.org/10.7208/chicago/9780226613475.003.0016

Vipra, J., & Myers West, S. (2023). *Computational Power and AI* (p. 32). AI Now Institute. https://ainowinstitute.org/wp-content/uploads/2023/09/AI-Now_Computational-Power-an-AI.pdf

Westlund, O., Krumsvik, A. H., & Lewis, S. C. (2021). Competition, Change, and Coordination and Collaboration: Tracing News Executives' Perceptions About Participation in Media Innovation. *Journalism Studies*, *22*(1), 1–21. https://doi.org/10.1080/1461670X.2020.1835526

Wolfram, S. (2023, February 14). *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/