

SETH C. LEWIS, RODRIGO ZAMITH, TOMÁS DODDS

JOURNALISM IN THE AGE OF



FROM ACCELERATION
TO REIMAGINATION

Praise for *Journalism in the Age of AI*

“An invitation to see AI not only as a tool, a threat, or yet another instantiation of technological disruption driven by sometimes reckless and always self-interested competing corporations, but also as an opportunity to offer fresh answers to the questions of what journalism is, ought to be, and who it is for.”

Rasmus Kleis Nielsen, University of Copenhagen

“This book provides a much-needed calm, conceptual analysis of AI in contemporary journalism that cuts through the dystopian and utopian hype. It makes a compelling case that journalism is facing a fundamental challenge and needs to respond with public interest at the heart of its future.”

Charlie Beckett, LSE JournalismAI Project

“With *Journalism in the Age of AI*, we have a comprehensive yet highly readable account of how AI is disrupting news organizations and journalism, when they are already reeling from 25 years of digital transformations. A must-read for journalists, news workers, and those who teach and study in the field.”

Terry Flew, University of Sydney

“The authors do not buy into convenient narratives of inevitable AI – instead they map the powers, politics, and professional forces that shape journalism’s future with AI. This book is a timely call for editors, journalists, and policy-makers: the risks are real, but so are the opportunities for journalism’s democratic role – if we are willing to make courageous choices.”

Natali Helberger, University of Amsterdam

“This book’s great achievement is to show that AI is not just rewriting the news – it is reshaping how all of us come to know the world and tell truth from fabrication. Lewis, Zamith, and Dodds make a compelling case that reimagining journalism means rethinking the very infrastructure on which our shared sense of reality depends.”

Andreas Hepp, ZeMKI, University of Bremen

“Journalists’ stories play a key role in society, when they are told by humans to other humans. AI risks disrupting that, not because it is necessarily bad, but because its adoption driven only by economic values risks undermining the social meaning of journalism as a practice. This sharp, well-focused book highlights those dangers, but also the future promise, if media workers – journalists themselves – are empowered to take control of how AI is used and how it isn’t.”

Nick Couldry, author of *Predatory AI*

“This book sets the stage perfectly: AI is a choice point. Lewis, Zamith, and Dodds make a compelling case that the hamster wheel isn’t the path. Using AI to elevate journalism is. Its concrete examples make this volume one I will continue to recommend.”

Florent Daudens, CEO, Mizal

“Journalism in the Age of AI arrives at exactly the right moment and it does something that’s very much needed. It looks beyond the newest tools and models to map the structural reordering of the information ecosystem itself – where power is moving, how value is being redefined, and what happens when journalism becomes just one layer in a much larger, AI-mediated system of knowledge creation and distribution. If you’ve been trying to make sense of where all this is heading, this book won’t give you easy answers. But it will help you ask better questions, and that may be exactly what this moment demands.”

Ezra Eeman, Strategy and Innovation Director at NPO

“This well-researched, intellectually honest, and ultimately hopeful book deserves the attention of policymakers, educators, and technologists who believe that the quality of our public knowledge determines the quality of our shared future.”

Daniel Erasmus, Head of AI for Internet Archive Europe

“AI has forced a reckoning with what matters most about journalism; which parts belong to humans, machines, or some combination of both. *Journalism in the Age of AI* gives readers the language to confront what has been, until now, only a felt sense among those of us working in news and innovation.”

Aimee Rinehart, Founder of FrontierCollective.io and former Senior Product Manager for AI Strategy at the Associated Press

“A masterful synthesis of what really counts in journalism and AI that cuts through the hype with a grounded vision forward.”

Nick Diakopoulos, author of *Automating the News*

“Journalism in the Age of AI provides a rare service: a clear-eyed and careful assessment of AI’s significance as a disrupter of news, combined with a plausible argument for developing new relevance for journalism in an AI-mediated age – one that may eventually offer journalists a catharsis for the frantic and unsustainable practices of the digital era. Lewis, Zamith, and Dodds don’t sugarcoat the challenges ahead, but they provide far more than just anxious criticism and they gently invite journalists to imagine what journalism might become.”

David Caswell, Co-founder of Signals at Scale and former Executive Product Manager at BBC News Labs

“As AI both accelerates and complicates the structural shifts transforming journalism today, urgent questions arise about the future of media production and democratic society. Fortunately, *Journalism in the Age of AI* has arrived at the perfect moment. Written by leading scholars of journalism studies, the book expertly assembles key insights on journalistic labor, platform power, and democracy, while showing how AI not only is disrupting how we share and consume information, but also provides a rare opportunity to reimagine what good journalism could be.”

Victor Pickard, University of Pennsylvania and author of *Democracy without Journalism?*

Journalism in the Age of AI

AI & Society Series

David J. Gunkel, *Understanding Artificial Intelligence*
Seth C. Lewis, Rodrigo Zamith, and Tomás Dodds, *Journalism
in the Age of AI*

Journalism in the Age of AI

From Acceleration to Reimagination

Seth C. Lewis
Rodrigo Zamith
Tomás Dodds

polity

Copyright © Seth C. Lewis, Rodrigo Zamith, and Tomás Dodds 2026

The right of Seth C. Lewis, Rodrigo Zamith, and Tomás Dodds to be identified as Author of this Work has been asserted in accordance with the UK Copyright, Designs and Patents Act 1988.

An online version of this work is available under the terms of a Creative Commons Attribution – Non Commercial – No Derivatives 4.0 International license (CC-BY-NC-ND 4.0). To view a copy of this license, visit <https://creativecommons.org/licenses/by-nc-nd/4.0/>. This license permits re-use, distribution and reproduction in any medium for non-commercial purposes providing appropriate credit to the original work is given. You must not produce any derivative works or make use of this work for commercial purposes.

First published in 2026 by Polity Press Ltd.

Polity Press Ltd.
65 Bridge Street
Cambridge CB2 1UR, UK

Polity Press Ltd.
111 River Street
Hoboken, NJ 07030, USA

All rights reserved. Except for the quotation of short passages for the purpose of criticism and review, no part of this publication may be reproduced, stored in a retrieval system or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior permission of the publisher.

ISBN-13: 978-1-5095-7069-0
ISBN-13: 978-1-5095-7070-6 (pb)

A catalogue record for this book is available from the British Library.

Library of Congress Control Number: 2026939776

Typeset in 10.5/12pt Sabon LT Pro
by Cheshire Typesetting Ltd, Cuddington, Cheshire
Printed and bound in Great Britain by Ashford Colour Press

The publisher has used its best endeavours to ensure that the URLs for external websites referred to in this book are correct and active at the time of going to press. However, the publisher has no responsibility for the websites and can make no guarantee that a site will remain live or that the content is or will remain appropriate.

Every effort has been made to trace all copyright holders, but if any have been overlooked the publisher will be pleased to include any necessary credits in any subsequent reprint or edition.

For further information on Polity, visit our website:
politybooks.com

Contents

<i>Prologue</i>	vi
1 Journalism's Constitutive Moment	1
2 The Newsroom Accelerates	31
3 The New Architecture of Media Power	63
4 Designing AI for Democracy	92
5 Learning to Think in the Age of AI	119
6 Reimagining Journalism for the Age of AI	145
<i>Acknowledgments</i>	158
<i>Notes</i>	164
<i>Index</i>	192

Prologue

This is a book about journalism in the age of artificial intelligence – but that’s only part of the story. It’s also a book about a much larger phenomenon: the rapidly expanding scope and influence of AI, which is transforming how people around the world create, access, and make sense of information.

All of this has profound implications for journalism, as we show in the chapters that follow. But if you look closely, you will see that the challenges and opportunities we explore – for journalists and their audiences, educators and their students, policy-makers and citizens alike – have deep resonance across public life.

Questions about creativity and authorship, platform power and infrastructural capture, public knowledge and democratic accountability, the fate of education and the future of personal empowerment – these extend well beyond journalism. They matter for the broader question of societal wellbeing and human self-determination in an age when both are increasingly mediated by machines.

In this book, we argue that the status quo of journalism and AI development is only accelerating the hamster wheel that has characterized – and diminished – news work in the twenty-first century.

But it doesn’t have to be this way.

AI can serve as a catalyst for improving journalism – a constitutive moment that offers an off-ramp to escape from that

hamster wheel. Taking advantage of that requires making deliberate, and difficult, choices, ones that can allow us to reimagine what we want from journalism and what it means to practice and engage with it.

This book is an effort not only to explain how we arrived at this point and why it matters, but also to chart what it could mean to reinvent journalism as an institution – and news as a public good – for the age of AI.

1

Journalism's Constitutive Moment

In early 2025, *The New York Times* announced it would integrate artificial intelligence tools throughout its newsroom.¹ Reporters at the “newspaper of record” would now have access to Echo, an in-house AI summarization tool, alongside GitHub Copilot, Google’s NotebookLM, and, with special approval, even OpenAI’s API. “Generative AI can assist our journalists in uncovering the truth and helping more people understand the world,” the *Times* declared in editorial guidelines shared with staff.

At the same time, the *Times* was actively suing OpenAI and Microsoft for billions of dollars over the alleged theft of its content to train those very AI systems.² The irony wasn’t subtle. Here was one of the world’s most prestigious news organizations simultaneously fighting AI companies in court and welcoming their technology into the newsroom.

Just weeks before the *Times*’ announcement, researchers at BBC News released findings from an experiment testing whether AI chatbots could accurately summarize news.³ The results were alarming. Even when the AI systems were given access to the BBC News website – one of the most trusted news sources globally – and were prompted to use articles from it when possible, the chatbots stumbled badly. Google’s Gemini bungled National Health Service guidance on quitting smoking. OpenAI’s ChatGPT and Microsoft’s Copilot claimed that politicians Rishi Sunak and Nicola Sturgeon remained in office months after

they'd left. When asked about the death of TV doctor Michael Mosley, Perplexity reported he had disappeared in October and that his body was found in November – when he had actually died the previous June.

More than half of all AI-generated responses the BBC looked at contained significant issues. Nearly one in five stories introduced outright factual errors. Thirteen percent included quotes that had been altered or seemingly fabricated entirely.

“We live in troubled times,” said Deborah Turness, CEO of BBC News, when the study was released.⁴ “[H]ow long will it be before an AI-distorted headline causes significant real-world harm?”

Meanwhile, *The Washington Post* was charting a different course. In April 2025, the newspaper announced a strategic partnership with OpenAI.⁵ Under the deal, ChatGPT's hundreds of millions of weekly users would see *Post* journalism appear more frequently in AI-generated responses – complete with quotes, summaries, and links. “We’re all in on meeting our audiences where they are,” explained Peter Elkins-Williams, the *Post*'s head of global partnerships.

But where, exactly, were those audiences? By late 2025, the answer was becoming increasingly clear: they were staying on platforms, and no longer clicking through to news sites. One analysis found that Google traffic from organic search (i.e., unpaid listings) to over 2,500 websites declined by 33 percent globally between November 2024 and November 2025, with publishers expecting further declines in the years to come.⁶

The suspected culprit was Google's quiet shift from being a search engine to an answer engine. Users no longer needed to visit news websites to get information; the AI summary at the top of search results was sufficient. For publishers, the implications were catastrophic. Without clicks, there could be no advertising revenue, no opportunities to convert casual readers into paying subscribers, and no relationship-building with audiences. “This is a serious threat to journalism that should not be underestimated,” warned then-*Washington Post* publisher William Lewis.⁷

Then came a scandal that summed up the sorry state for publishers. In the summer of 2025, *Wired* published what appeared to be an endearing portrait about couples getting married in virtual worlds like Minecraft and Roblox.⁸ The story included

interviews, quotes from a Chicago-based “ordained officiant” who specialized in “digital celebrations,” and heartwarming anecdotes about love in digital spaces. Mashable praised it as “a charming feature.”

But the story had vanished from *Wired*'s website a few months later. So had similar articles at *Business Insider* and several other outlets. It turned out the married couples didn't exist. Neither did the Chicago officiant. Nor, it turned out, did the journalist – at least not in any conventional sense.

The freelance writer who wrote the pieces, working under the pseudonym “Margaux Blanchard,” had pitched to editors and filed their stories on time. But the names, quotes, and details in those captivating stories turned out to have been elaborate AI-generated fiction, sophisticated enough to deceive editors at respected and tech-savvy publications. “If anyone should be able to catch an AI scammer, it's *Wired*,” the magazine reflected in an editor's note. “Unfortunately, one got through.”⁹

How were audiences making sense of all this chaos? In fall 2025, researchers at the Reuters Institute for the Study of Journalism surveyed 12,000 people across six countries.¹⁰ While respondents reported growing comfort with AI in their daily lives – weekly ChatGPT usage had doubled year-over-year in some cases – they remained deeply skeptical about it in news production. Only 12 percent said they were comfortable with news made entirely by AI. With human oversight added, that figure rose to just 21 percent. Even with human journalists leading production and AI relegated to marginal assistance, only 43 percent found it acceptable. News made entirely by human journalists, by contrast, had 62 percent support.

Perhaps most revealing was the perception gap. About half of respondents believed journalists were already using AI regularly in their work. A quarter thought newsrooms were using AI to generate images when photographs weren't available – an idea that would horrify photojournalists, for whom realism is an ethical cornerstone. Only one-third believed that newsrooms were bothering to verify AI outputs before publication. And audiences overwhelmingly felt that AI would “primarily benefit publishers rather than them as users” – making news cheaper and faster to produce, yes, but also less trustworthy and less transparent.

By the end of 2025, it was becoming increasingly hard to ignore the fact that AI was helping reshape journalism and, more

broadly, the enterprises of creative production and knowledge work.¹¹ These examples weren't isolated incidents but a constellation of upheavals: prestigious newsrooms simultaneously suing AI companies and adopting their technologies; AI systems that confidently presented fabricated information and attributed it to trusted news organizations; publishers scrambling for lifelines as search traffic collapsed; editors fooled by AI-generated fakery; and audiences expressing profound distrust of the very AI-powered journalism that industry leaders saw as central to the future of the profession.

How do we make sense of journalism in the age of artificial intelligence? This book offers a window into what is happening, why, and what it means for journalists, citizens, educators, and the future of knowledge in democratic societies.

What's Different about the Age of AI?

There have been many stories of technological salvation in journalism's lore. The web would democratize access to information. Social media would connect journalists to communities in a virtuous circle of pro-social reciprocity. Mobile apps would allow news organizations to meet people where they spend most of their time – on their phones. Each successive transition carried the hope of revitalizing journalism through greater reach and audience engagement, and the promise of being a help rather than a hindrance to the quality and impact of journalistic work.

Instead, each wave of technological change made journalists' jobs harder or more complex. Business models collapsed, production became frantic, and reporters struggled to reach audiences who preferred streaming, gaming, and other online entertainment over news. A couple of decades into the twenty-first century, many journalists have found themselves working faster and harder, but struggling to keep up.

Generative AI (GenAI) arrives as both the latest iteration of this cycle and something potentially different – indeed, possibly so profoundly different that it requires a fresh evaluation of the present and future of journalism. Yes, the rise of easy-to-use AI echoes past disruptions. It's another technology laden with hopes for “saving” journalism through efficiency and innovation, and

another that carries an undercurrent of concern: “If journalists don’t adapt,” people seem to say, “they will be left behind.” Sound familiar?

But AI also introduces dynamics that journalism hasn’t confronted before. GenAI tools like ChatGPT, Gemini, and Claude can mimic human communication, creativity, and authorship to a degree far beyond anything seen before.¹² Writing stories, creating compelling graphics and videos, producing code – all these and more can be produced with limited human involvement. Agentic systems go well beyond chatbots to proactively pursue goals and execute multi-step tasks on their own. Think of them less like a search engine you query and more like a tireless colleague you can hand a complex assignment – “Dig into these financial records and tell me what looks suspicious” – and walk away from. That kind of autonomous, end-to-end capability is what makes agentic AI so disruptive: where earlier waves of automation came for physical and routine work, this one is aimed squarely at the tasks that define professional knowledge work. These capabilities force journalism to confront fundamental questions about what it means to create, verify, and communicate in an age when machines can do many tasks once thought to be uniquely human.

In finance, the phrase “this time is different” typically signals irrational overconfidence – the belief that new conditions have rewritten old rules – and it tends to arrive on the scene at the height of inflated expectations, just before a market crash. And yet, it’s also true that *every time is different*, at least to some degree. History doesn’t repeat, as the saying goes, but it does tend to rhyme. For journalism and AI, both sides of the paradox are in play: the story of digital acceleration is familiar, but the specific dimensions of this moment are not – particularly because AI challenges longstanding assumptions about humans, machines, and communication.¹³

A Constitutive Moment

We argue in this book that the rapid integration of AI – by news professionals and the public alike – represents not a peripheral shift but a constitutive moment: a critical juncture that unsettles how journalism understands its roles, routines, and reasons for

being. This is a veritable shock to the field, forcing journalists to engage with deep questions about the division of labor between humans and machines, about media production when synthetic content can be produced at scale, and about journalism's claims to special authority when audiences increasingly encounter information through algorithmic intermediaries.¹⁴

Constitutive moments are not necessarily destructive ones. None of this means that AI will "replace" journalism. Rather, it suggests that a fundamental renegotiation is underway – one that calls for fresh examination of who does what, how work is organized, and why journalism matters. At stake is not simply how journalists use AI but also what journalism becomes and in relation to whom. Who counts as an author when machines draft articles? Or when they just contribute – as Anthropic's Claude did for us with the anecdotes that opened this chapter? What constitutes original reporting when AI synthesizes existing information? How is trust established when synthetic content can convincingly impersonate journalism itself?

The reconfiguration of journalism in the age of AI illuminates transformations that go beyond news production. This moment reveals wider changes in how society forges a sense of knowledge, meaning, and truth in a digital environment where machines act as creators, collaborators, and communicators. Journalism, with its professional commitments to accuracy and accountability, becomes a crucial site for understanding what happens when machine creations permeate information ecosystems. The issues that journalism encounters – about authenticity, authority, and human judgment – matter also for the future of any endeavor involving the making and sharing of knowledge: from education and science to law and governance. These issues matter for how all of us – as consumers and citizens – make sense of a digital world increasingly defined by AI's influence.

Journalism's Hamster Wheel

To grasp why this moment feels so consequential, consider the hamster wheel – that tiny spinning contraption that allows small pets to stay active while inside a cage. In journalism and media, the hamster wheel has emerged as a defining metaphor for digital-era work: relentless motion but too little progress to show for it.

Dean Starkman, writing in the *Columbia Journalism Review* in 2010, crystallized the problem: “The Hamster Wheel isn’t speed; it’s motion for motion’s sake. The Hamster Wheel is volume without thought. It is news panic, a lack of discipline, an inability to say no.”¹⁵ The metaphor captured what had been happening in journalism in the early twenty-first century: an almost pathological push by newsrooms to do more with less. At *The Wall Street Journal*, Starkman noted, staff produced 22,000 stories in 2000; by 2008, that figure had swelled to 38,000, even as the newsroom shrank by 13 percent. Output per journalist had jumped 69 percent in less than a decade.

But the picture Starkman painted in 2010 wasn’t the end of the acceleration story. The pile-on of increasing demands – more platforms to manage, more social media feeds to tend, more audiences to reach – has persisted and left many of today’s journalists feeling trapped and exhausted.

From Web to Social to Mobile

Each wave of technological change has only added rotational fuel to the wheel’s velocity. The rise of the World Wide Web in the late 1990s forced many newspaper journalists to learn new content management systems and find ways to repurpose their work for the internet while maintaining their print duties. In the 2000s, blogging introduced expectations of frequent updates in addition to traditional reporting.¹⁶ Then social media layered on demands to engage with audiences everywhere, participate in multiple conversations as quickly as possible, and even make the news dance to new trends on platforms like TikTok.¹⁷ As a journalist lamented to one of us in 2015, their manager had berated them for not posting at least 10 tweets per day. The journalist, who had a paltry following, couldn’t understand why they were being saddled with extra work that had no clear benefit. They have been proven right for the most part: a decade later, hamster wheel posting on social media has largely yielded little tangible benefit to journalism’s financial sustainability or its quality.

Mobile technology accelerated the wheel further. Smartphones untethered journalism from the office, creating expectations of always-on availability and real-time responsiveness.¹⁸ Push notifications and pull-to-refresh feeds trained audiences to expect

immediate updates, which trained journalists to provide them. The 24-hour news cycle, already exhausting in the cable era, became even more frenetic.

To be sure, these pressures toward overproduction weren't limited to journalism. Productivity expert Cal Newport describes how the "hyperactive hive mind" of emails and texts has drowned most knowledge workers in messaging mania, making deep work – focused labor on a single significant task – vanishingly rare.¹⁹ But for journalists, each technological layer compounded demands even as the pool of people to help produce news was, for most organizations, shrinking. By the mid-2010s, as layoffs left more work for fewer employees, the hamster wheel symbolized the seemingly unceasing acceleration of media labor, as platform-driven pressures replaced a slower, more deliberate rhythm from a bygone era.²⁰ The human cost was significant: burnout, depression, and a profession quietly hemorrhaging the people it could least afford to lose.²¹

The labor implications extend well beyond journalists themselves. Behind every AI system integrated into a newsroom lies a largely invisible workforce: the data labelers, content moderators, and annotation workers – often in the Global South – whose low-wage labor makes these technologies possible in the first place.²² When a newsroom deploys an AI transcription tool or content recommendation system, it benefits from a global supply chain of human labor that rarely appears in discussions of journalistic automation.²³

Why GenAI Poses a Breaking Point

Through all these transformations, a fundamental assumption remained: the hamster wheel was dependent on humans running it. New technologies might change the pace of work, but there was no question about the engine: the human journalist served as the irreplaceable driver of content creation. The wheel spun faster as journalists were expected to live-blog, tweet, and engage across platforms, but it was still a person – however exhausted – running that wheel.

GenAI complicates that assumption. It doesn't simply remove humans from media work; it also reshapes what the work of journalism entails – and with it, how journalists understand their value and purpose. With tools like ChatGPT, Gemini, and

Claude, certain aspects of journalism production can now be automated or augmented in unprecedented ways.

But here's the crucial distinction: AI's capacity to write doesn't eliminate the need for journalism any more than it eliminates the need for humans; instead, it forces a reckoning with what matters most about journalism, and which of those things are best served by humans, machines, or some combination thereof. For example, writing – a task large language models can now execute better than many journalists – has long been central to journalism, but it has never been *all* of journalism. The work has always involved shoe-leather reporting, reconciling conflicting accounts, relationship-building, and ethical decision-making – things requiring sensibilities that are still best accomplished by humans.

The age of AI thus marks a breaking point – not because machines are replacing journalists wholesale, but because the profession must confront what happens when its core activities become partially automatable. If writing becomes less central to what journalists do day-to-day, does shoe-leather reporting become more so? Do journalists spend less time drafting prose and more time cultivating sources, verifying claims, and exercising editorial judgment? Thus, whereas previous technologies asked journalists to do more of what they already did – faster, across more platforms, in more formats – AI asks something different: which parts of journalism are best supported by humans and which parts by machines?

When journalists are given dedicated time to work with AI technologies – removed from regular production duties – researchers observe something remarkable. “[Journalists] see it as a break, a moment that they can reflect and engage in things – not just on the hamster wheel, if you like, of production,” notes Bronwyn Jones, research lead at the Responsible Innovation Centre for Public Media Futures.²⁴ Whether AI creates breathing room and novel work or simply raises expectations depends on how newsrooms respond, including through managerial priorities. If bosses care solely about “optimization,” “efficiency,” and “cost-cutting,” the wheel will just keep spinning. If they care about making their journalism as valuable as possible, an off-ramp might just emerge.

Reimagining News

These questions about AI's impact have consequences not only inside the newsroom but also for citizens and educators – people who have been, in a sense, running on hamster wheels of their own. For example, there is growing evidence that people feel bombarded by “AI slop,” or synthetic content often lacking effort, originality, or substance.²⁵ So ubiquitous was this content by 2025 that *The Economist* named “slop” its word of the year.²⁶ This is more than just an inconvenience. The junk adds to existing information avalanches and crowds out quality content – essentially, making it harder for citizens to get the same benefits they once did from more easily organized information ecosystems.

To be sure, AI tools are likely to help citizens in this filtering process, too. AI can sift and summarize with ease. But this tug of war – between tools that automate the production of content and those that make information ecosystems more navigable – adds instability and demands new literacies. After all, what's scarce is no longer content but human capacities for judgment and sense-making: for establishing validity, provenance, context, and social authority. In other words, for many citizens, AI's proliferation will only further complicate their ability to puzzle through what's happening and exacerbate problems of “burnout” with news, with even more people experiencing frustration as they consume journalism or avoid it altogether.

The breaking of the hamster wheel's fundamental premise – that humans must run it – creates space for reimagining what journalism could be if not organized around acceleration as its governing logic. If machines can manage a growing share of production work, what uniquely human contributions will become correspondingly valuable? If audiences can no longer assume human creation in the news they encounter, what signals of credibility and authenticity will emerge, and with what consequences for trust in news and knowledge experts broadly?

The hamster wheel might finally come apart because journalists and audiences learn to say no to the rat race of digital demands – and say yes to a long-overdue reckoning about what journalism is for, and to much-needed reflection about whether, and how, the rise of GenAI might help or hinder the future of news. As we will see throughout this book, AI's accelerating tendencies are not inherently beneficial or detrimental to the pro-

duction, distribution, and consumption of quality information that advances the public good. The impact depends far more on how these tools are leveraged, toward what ends, and with what normative aims in mind.

To explore those possibilities further, we first need to answer three questions that set up the contextual landscape for this book: What exactly is AI? How did we arrive at this moment of AI and journalism? And how does the current technology work?

What Is AI?

The story of artificial intelligence begins not with a line of code, but with a philosophical question. In a foundational 1950 paper, British mathematician Alan Turing sidestepped the intractable question “Can machines think?” by proposing something more practical: the “Imitation Game.”²⁷ Rather than wrestling with poorly defined terms like “machine” and “think,” Turing reframed the debate from abstract philosophy to a concrete behavioral challenge – an engineering puzzle to be solved through code.

Turing's work didn't emerge in a vacuum. It was the culmination of a transatlantic intellectual debate about “the mind as a machine” in the late 1940s and early 1950s. Norbert Wiener's cybernetics suggested machines could “learn” like animals.²⁸ Claude Shannon's information theory argued that all communication could be expressed mathematically as ones and zeros.²⁹ John von Neumann's “stored-program” concept laid the groundwork for how computers would fetch instructions, process them, and display results.³⁰ Together, these developments created an ecosystem that would define how people imagined a “thinking machine.”

This momentum culminated in 1956 at Dartmouth College, where John McCarthy introduced the term “artificial intelligence.” It was a clever bit of branding: an evocative phrase designed to attract donor money and a vibes-driven conjecture and over-promise about building a synthetic “mind” that could “think.” It's the kind of language that should sound familiar to anyone following AI marketing today.³¹

Toward a Definition of AI

Despite decades of development, defining AI remains contentious. Classic definitions focus on replicating human intelligence. Marvin Minsky, a pioneer in the field, framed AI as building computer programs that could perform tasks that were believed to be, at that moment, carried out better by humans because they demanded higher-level mental processes, such as perception, learning, memory organization, and critical reasoning.³² Many decades later, the seventh edition of Oxford's *A Dictionary of Computer Science* echoes Minsky by defining AI as "the building of computer programs that perform tasks requiring intelligence when done by humans."³³

These human-centered definitions face three critiques.³⁴ First, intelligence itself is contested, making it easy to subjectively label systems "intelligent" or not. Second, tying machine intelligence to human mimicry risks overlooking capabilities that diverge from human cognition. Third, defining AI relative to what humans currently do makes it a moving target – as soon as machines match a benchmark, the goalposts shift.

This shifting process even has a name: the "AI effect."³⁵ Defeating a world chess champion was once seen as the arrival of genuine AI. When IBM's Deep Blue did it in 1997, critics dismissed it as brute force rather than true intelligence. The bar moved to the game Go, with its complex strategy. When Google DeepMind's AlphaGo beat top players in 2016, definitions shifted again toward broader reasoning. When ChatGPT launched in late 2022, it was dismissed as fancy autocomplete. Each milestone invites reimagining that makes AI seem perpetually just out of reach.

To sidestep debates over "intelligence," many computer scientists describe AI as a suite of computational techniques accomplishing complex tasks, including automated reasoning, computer vision, knowledge representation, natural language understanding, and speech recognition.³⁶ This centers computational feats rather than simulating human thought.

Applied industries, including journalism, tend to define AI less by its ability to simulate human action and more by its purpose. The JournalismAI initiative at the London School of Economics uses AI "as an umbrella term to refer to the use of algorithms and automation by news organisations, usually to make journal-

ists' work more efficient or to deliver more relevant content to audiences."³⁷ The Associated Press describes AI as "a range of technologies that are human-designed to automate, accelerate or extend the human work required for specific tasks," noting that "some applications involve simple automation" while "others are more advanced, helping journalists with tasks that they couldn't reasonably accomplish with human effort alone."³⁸

We adopt the following definition to guide our thinking throughout this book:

Artificial intelligence refers to computational systems that learn from data to recognize patterns, make predictions, and generate content – automating or extending tasks at scales and speeds beyond practical human effort – though with capabilities and limitations that vary unevenly across tasks.

What Are News and Journalism?

Like AI, "news" and "journalism" are terms that seem self-evident but carry multifaceted meanings worth unpacking.

Simply put, news refers to novel information about recent affairs.³⁹ It's not a natural phenomenon that just exists, though. It's invariably constructed through production processes involving selection, assemblage, and expression. News can absorb both intentional properties (e.g., the appearance of neutrality) and unintentional ones (e.g., journalists' subconscious biases) based on how it's organized and produced. For example, what counts as "newsworthy" reflects judgments about relevance, timeliness, prominence, and public interest – judgments made by those who craft news, whether they are professionals or not.

Journalism, on the other hand, is the institutional framework through which news is often produced and distributed. It encompasses "the systematic gathering, filtering, and circulating of information deemed to be news and in the public interest."⁴⁰ In most Western democracies, journalism is organized around principles like truth-telling, fact-checking, holding power to account, minimizing harm, and informing the public. Journalists and news organizations remain the most visible and influential actors in the field of journalism, though a rapidly shifting information environment has created space for new actors and activities to claim "journalistic" status.⁴¹

The boundaries between news and journalism have always been porous. Not all news comes from journalists – citizens' social media posts about breaking events or corporate press releases can be vehicles for news. And not everything journalists produce is “news” in the traditional sense, such as opinion columns and editorial cartoons. But news and journalism remain deeply intertwined: news is typically the precursor to and product of journalism, and journalism is the primary means of producing news in democratic societies.

A Brief History of Automation and AI in Journalism

On the night of November 4, 1952, a seven-ton machine called UNIVAC I made a prediction that its human operators refused to believe.⁴² With only a fraction of votes counted in the US presidential election, the “electronic brain” forecast a landslide victory for Dwight D. Eisenhower – defying human pollsters, who were still cautious after their embarrassing miscalls four years earlier. CBS News, which had enlisted the computer for a piece of technological theater, hesitated to broadcast something so implausible. By the end of the night, though, UNIVAC I had been vindicated. The machine captured a sense of wonder not because it was truly “smart,” but because it performed a feat of calculation that seemed almost magical.

This moment, broadcast to millions on the new medium of television, established a pattern that persists today: public perception of AI is shaped as much by spectacle and imagination as by technical reality.⁴³ It also marked journalism's first dramatic encounter with computational prediction – an encounter that began not with ChatGPT's viral arrival in 2022, nor with the first “robot reporters” of the early 2010s, but rather in that postwar moment when thinking machines seized the public imagination.

From the Pressroom to the Newsroom

Journalism's relationship with automation extends far beyond these moments of spectacle, though. Long before algorithms could write stories, technology's first major foray into the news business was fundamentally industrial.⁴⁴ For much of the twentieth

century, newspapers were often likened to factories. Automation at that time targeted the pressroom: the spaces where you would find printing presses, typesetters, and the physical labor of manufacturing a daily product. Skilled human compositors gave way to machine operators as technology made increasing portions of their work obsolete. But newsrooms – the spaces where stories were reported and written – seemed to stay largely intact. The discourse of the time centered on industrial efficiency and blue-collar displacement, not on the nature of the journalistic enterprise itself.⁴⁵

But a quieter transformation was being seeded inside the newsroom. In 1967, Philip Meyer conducted a groundbreaking analysis of Detroit riot survey data using mainframe computers, establishing what he termed “precision journalism.”⁴⁶ This was a novel form of journalism that applied social science methods to reporting. His systematic, data-driven approach revealed truths that traditional shoe-leather reporting hadn’t fully captured, contributing to a Pulitzer Prize for the *Detroit Free Press*. Meyer’s work signaled the birth of computer-assisted reporting, a precursor to modern data journalism and computational journalism – methods that rely on datasets and code to tell journalistic stories.⁴⁷

Digitalization accelerated over the next three decades. By the late 1990s, technologies like speech-to-text transcription and optical character recognition were automating previously laborious tasks that would have ordinarily been performed by news assistants inside the newsroom – if not the journalists themselves.⁴⁸ Content management systems and digital design tools likewise standardized some of the work performed by editors and other creative workers. But these developments remained firmly in the realm of assistance: tools that helped those inside the newsroom process information more efficiently, without entirely impinging on their creative or cognitive work.

The transformation that began around 2010 was qualitatively different. For the first time, algorithms would not merely assist with journalism but would create the news itself – standing in for human writers, even if initially on a limited basis. What had once seemed a distant industrial concern had now arrived at the journalist’s own desk. When automation was confined to blue-collar pressroom jobs, journalists could maintain a professional distance. The printers’ struggles represented an industrial dispute,

after all, which was separate from the cognitive work of storytelling. But as software began automating tasks closer to journalists' own work – the actual writing of news – anxiety became personal and existential.

Understanding this progression from twentieth-century automation of manual labor to twenty-first-century encroachment on knowledge work is essential for grasping why contemporary fears about AI in journalism feel both familiar and fundamentally new. The past 15 years can be understood through three distinct waves, each defined by two elements: the crest – periods of peak hype, speculative discourse, and myth-making dominated by breathless headlines and binary narratives of salvation or destruction; and the trough – the slower, messier reality of newsroom implementation, where grand promises collide with budgets, workflows, and professional cultures, and where doubts and resistance emerge more clearly.

Wave One: Foundation (2010–15)

The first wave of automation in the newsroom began in 2010 and was characterized by caution, experimentation, and managed narratives. The technology – primarily template-based natural language generation – could automate only the most routine, data-centric reporting: sports results, corporate earnings, basic event summaries, and the like. The innovation was not in sophistication but in producing simple reporting at scales and speeds humans couldn't match.⁴⁹

Two companies led the charge. Narrative Science, founded in 2010, commercialized Quill, a descendant of an academic prototype called StatsMonkey that automatically generated baseball recaps from statistics. Meanwhile, Automated Insights developed Wordsmith, which was also capable of automatically producing stories but for a different set of domains. This tool would go on to power one of the era's most significant experiments.

In 2014, the Associated Press partnered with Automated Insights to transform its coverage of corporate earnings.⁵⁰ Previously, AP journalists manually produced roughly 300 earnings stories each quarter that were in some ways just glorified data transcriptions. By feeding structured financial data into Wordsmith, the AP lifted its output to some 3,700 earnings reports per quarter – a twelve-fold increase – while freeing up an

estimated 20 percent of the time its reporters had spent on earnings coverage. Crucially, the AP framed this as augmentation, not replacement. The goal, according to business editor Lou Ferrara, was to “free journalists to spend more time on things like beat reporting and source development.”

The innovation wasn't just coming from outside the newsroom, though. Some news organizations experimented with smaller-scale initiatives that addressed niche needs and opportunities. For example, *Los Angeles Times* developer Ken Schwencke created Quakebot, an algorithm that monitored earthquake data feeds and automatically generated basic stories when tremors exceeded a certain threshold.⁵¹ On March 17, 2014, a 4.4-magnitude earthquake struck at 6:25 a.m. Schwencke woke to find Quakebot's draft waiting and published it within three minutes, beating his competitors by a significant margin. Like Ferrara, technology reporter Will Oremus saw the potential for augmentation: “Helpers like Quakebot might save a few journalists' jobs by freeing them to focus on the type of work that can only be done by a local reporter, on the ground, with a brain.”

Despite these carefully framed narratives, the crest of hype was formidable. Headlines about “robots in the newsroom” abounded, often illustrated with humanoids wearing fedoras while typing. One prediction claimed that 90 percent of journalism would be automated by 2030.⁵² In hindsight, such fears were massively overblown. Very little news became automated in this phase, and few journalists lost jobs to algorithms. The technology served as an additive resource for routine, templated narratives – the “boring” stories journalists were happy to offload.

But the anxiety was real enough to spark academic and industry interest. Scholars examined how journalists engaged in “boundary work,” articulating qualities like judgment, empathy, and creativity that separated human from synthetic journalism.⁵³ Research on audience perceptions revealed interesting findings, with some studies observing that readers could not reliably distinguish AI-generated sports stories from human-written ones even as they perceived the automated text as “more boring, but also more accurate, trustworthy, and objective.”⁵⁴ To them, machine-generated content could seem more neutral and credible precisely because it lacked the human touch that made writing engaging.

Wave Two: Expansion (2016–21)

By 2016, a new period of strategic integration took hold, particularly at well-resourced organizations. Industry leaders began moving from partnering with outside companies to developing their own automation systems that could be tailored to their journalistic needs, all the while expanding applications from simple story generation to data mining, investigative support, and content personalization.

The Washington Post's Heliograf illustrated this evolution.⁵⁵ Developed internally, it debuted during the 2016 Olympics, generating approximately 300 reports and alerts. Later, during that year's presidential election, Heliograf produced around 500 articles covering congressional and gubernatorial races – coverage the *Post* said it lacked staff to produce manually. In its first year, the system published roughly 850 articles. BBC News similarly used automation to publish a news story for every constituency that declared election results on the night of the United Kingdom's 2019 General Election – nearly 700 stories in all.⁵⁶

Reuters took a different tack with Lynx Insight, which it launched in 2018.⁵⁷ Rather than generating complete stories, it focused on mining vast datasets to identify trends, anomalies, and newsworthy items, and then alerting relevant journalists. “We’ve recognized that the key thing is not to make a machine write in-depth stories,” a top Reuters editor told a reporter, “but to marry the strength of machines with the strengths of humans.”

The immense market power of tech companies like Google and Facebook only continued to grow during this time, though.⁵⁸ That dynamic, detailed in Chapter 3, meant that even as newsrooms were trying to drive innovation from within, they were becoming increasingly dependent on the platforms managed by those tech companies – an asymmetrical dependence that persists today. In other words, the resources needed to build the technologies news organizations increasingly hoped would help them survive could only be sourced from outside their newsrooms.

As this wave matured, ethical concerns moved to the forefront. Algorithmic bias became a pressing worry as more people understood that AI systems trained on historical data could inherit and amplify human prejudices. A 2019 report noted journalists’ chief concern that “bad use of data could lead to editorial mistakes such as inaccuracy or distortion and even discrimina-

tion against certain social groups or views.”⁵⁹ Meanwhile, the emergence of deepfakes – synthetic videos, images, or audio convincingly depicting people doing or saying things they never did – transformed misuses of AI from a subtle concern to a clear threat.⁶⁰ Amid heightened worries about “fake news” and disinformation, journalism faced the difficult task of asking the public to trust its use of “good AI” while combating the “bad AI” of malicious actors, all at a time when faith in institutions overall was eroding.⁶¹

Wave Three: Generation (2022–Present)

ChatGPT's public release on November 30, 2022, marked a definitive inflection point. Its simple interface and uncanny ability to generate human-sounding responses made it an overnight sensation: ChatGPT had 1 million users within days, 100 million within two months, and nearly a billion per week within three years.⁶² It was merely the first in a cascade of chatbots – Gemini, Copilot, Claude, Grok, DeepSeek, Perplexity, and others – that rapidly defined the generative era. The diffusion of these tools triggered chaotic, bottom-up experimentation in newsrooms, far outpacing the development of institutional policies, and provoking both enthusiasm and alarm.

The fundamental difference was the shift from task-specific tools to general-purpose technology. First-wave systems could convert specific structured data into formulaic text. GenAI, by contrast, could engage in conversation, summarization, creative writing, and apparent reasoning across seemingly infinite topics. This capability triggered more intense and widespread anxiety – no longer about displacement of those writing routine reports, but about fundamental deskilling and devaluation of the entire profession.⁶³

The rush to leverage GenAI quickly produced high-profile failures. In early 2023, CNET was found to have quietly published multiple AI-generated articles containing factual errors without clear disclosure.⁶⁴ LedeAI, a service used by Gannett, produced error-riddled high school sports coverage.⁶⁵ *Sports Illustrated* published articles attributed to fake authors with AI-generated headshots – prompting the magazine's union to declare themselves “horrificed” and demand management commit to “not publishing computer-written stories by fake people.”⁶⁶

These scandals catalyzed rapid development of formal ethics policies at some organizations, which we discuss in Chapter 2.⁶⁷ But beyond questions of governance, the generative era triggered confrontation over journalism's economic foundation. AI systems had been trained on vast archives of copyrighted news content without compensation. When these systems began threatening publisher traffic through AI-powered search summaries that eliminated the need to click through to original sources, there was profound existential anxiety and a scramble to contain the fallout. *The New York Times'* landmark December 2023 lawsuit against OpenAI and Microsoft, alleging mass copyright infringement, directly challenged the GenAI business model.⁶⁸ OpenAI thus opted to sign content-licensing deals with several major publishers, including Axel Springer, the Associated Press, and *The Atlantic* – creating new revenue streams for those willing to partner rather than litigate.⁶⁹

Perhaps the most significant shift in this wave has been to journalism's position within the broader information ecosystem. In the first wave, news organizations were customers purchasing AI services to solve the problems they defined. By the third wave, they were increasingly becoming suppliers of a raw commodity – high-quality information, both generated and curated by humans – to new AI-driven intermediaries positioned between journalists and their audiences.

Finally, the pace of current transformation deserves particular emphasis. While earlier waves of automation unfolded more slowly, GenAI capabilities have advanced with striking speed. This velocity creates a distinctive challenge: newsrooms that might have had years to adapt to previous technological shifts now face the possibility that their current workflows could be fundamentally altered before institutional responses can take shape.

Overall, the restructurings required by the current wave carry profound implications for the future of news.⁷⁰ The history of automation and AI in journalism thus culminates in a reckoning: a need to clarify what makes humans distinctively valuable in an age when machines can capably do more and more of the work we associate with news production.⁷¹

How Does AI Work?

To better understand what AI can do in the present – and what it seems poised to accomplish in the foreseeable future – it is helpful to examine briefly how it works.

At its core, computing follows an input–algorithm–output model.⁷² A computer receives inputs (code, text, images, sound), runs them through an algorithm (a method or set of instructions), and produces outputs. You can think of an algorithm as a recipe: a series of step-by-step instructions to transform different ingredients into a tasty meal. When expressed in machine-readable form, an algorithm becomes software.

Traditional software stipulates precisely how a machine should respond to inputs under specific conditions. Because computers simply follow instructions, programmers must write detailed commands and anticipate many potential scenarios. If the software encounters an unanticipated situation, it generally crashes or falls back to predefined error-handling routines. Because of this, most software doesn't do everything on its own. Instead, it typically incorporates tools and resources built by others to handle some routine tasks.

This simplification helps us highlight two important considerations.⁷³ First, most software is a mixture of algorithms created by different programmers with different goals, tied together by an overarching logic that reconciles diverse philosophies and constraints. Scholars call these “algorithmic assemblages,” and they require trade-offs and acceptance of occasionally suboptimal results. Second, even the most complex assemblages aren't entirely synthetic. Trace any system back far enough and you find human choices shaping the software's behavior.

What makes AI different? It doesn't have to be its level of sophistication, at least not under our conceptualization. Amazon's Alexa and Apple's Siri were initially touted as intelligent because they handled varied phrasings, but they fundamentally relied on exhaustive phrase banks and logic structures compiled with human help that triggered relatively basic algorithms to produce useful responses. Over time, such systems have grown more sophisticated – but it's not the sophistication that sets them apart as AI.

AI and Machine Learning

What distinguishes modern AI from its predecessors is its superior machine learning capability. This enables software that improves *itself* by studying material and evaluating past performance, without *needing* additional programming. The progress stems from work in a subfield of computer science that develops algorithms to identify patterns in data. These patterns build models – mathematical representations of systems, processes, or knowledge – used to perform and refine tasks autonomously.⁷⁴

Three machine learning approaches merit special attention here.⁷⁵

First, supervised learning trains models on labeled data. A human “teacher” provides examples that they’ve labeled (e.g., images tagged as either “apple” or “carrot”) or have annotated relevant features of (e.g., color, shape, texture). The algorithm learns statistical patterns distinguishing those labels and features, enabling it to classify new images correctly – at least some of the time. Supervised learning offers control over what the model learns and can be highly effective when the training data are large, diverse, and high quality. The drawback to this approach is the typically immense manual effort needed to label and curate high-quality data.

Second, unsupervised learning unleashes algorithms to find statistical relationships in data independently, without human-identified labels or features. For instance, a clustering approach might lead the algorithm to partition a dataset into groups based on similarity. But while the algorithm creates clusters, it can’t really interpret them; they might reflect types of produce, colors, or weird combinations that are meaningless to humans. The advantage of this method is scalability – an algorithm can analyze far larger datasets than humans can label – and the ability to identify complex patterns a human teacher might miss. The risk is that it can develop internal logics too opaque for humans to interpret and latch onto spurious correlations that produce meaningless or incorrect algorithmic decisions.

Finally, reinforcement learning focuses on iteration and rewards. The algorithm interacts with an environment, takes actions, observes outcomes, and adjusts itself to maximize rewards over many trials. A model might start via supervised or unsupervised methods, then self-evaluate on labeled data or

receive human feedback on whether its choices were right or wrong. Correct classifications earn rewards for the algorithm; incorrect ones earn penalties. After each trial, the algorithm adjusts to improve future performance, and sometimes even introduces random variations to explore new optimization strategies. When you mark a chatbot's response as helpful or unhelpful, you are often contributing to its learning.

Important Limitations

Machine learning is of course far more complex than this, and not all AI relies on machine learning. Nevertheless, these simplifications allow us to highlight three additional considerations.

First, human involvement is always present.⁷⁶ Humans select algorithms, set parameters, define objectives, and evaluate performance. Even self-teaching models reflect initial human-designed logic and constraints.

Second, models reflect their training data.⁷⁷ Overrepresented examples skew sensitivities; underrepresented ones produce flawed learning. Biased training data therefore yield biased outputs.

Third, dominant learning approaches rely on statistical modeling focused on averages and are sensitive to complex, subtle patterns in large datasets.⁷⁸ They often produce useful, predictable results but can also return incoherent and erroneous outputs.

More broadly, these insights help us appreciate that AI "hallucinations" – instances when a model unintentionally produces incorrect or fabricated information and presents it as fact – are not just programmatic quirks. They are fundamental flaws of the machine learning methods currently powering popular GenAI agents.⁷⁹ This is a crucial consideration for a profession that stakes its professional authority on accuracy and truth-seeking.

Despite the rapid progress, important constraints remain for the future development of AI. First, even with ever-increasing computing power, AI relies on computational shortcuts to reduce costs or complexity at the expense of precision. The goal for mass-market systems will likely remain producing results that are "good enough" much of the time, and not perfection.

Second, data availability remains patchy and uneven. Much of what journalists care about occurs outside digital enclosures. Digitized knowledge varies by region, social group, and domain.

Efforts to digitize information into machine-readable data are time-consuming and error-prone. These gaps lead AI to especially struggle with some tasks or contexts.

Third, the proliferation of “AI slop” is creating challenges for model development as AI agents increasingly learn from AI regurgitations. This has already led to decreases in model performance at some tasks, but it has also made high-quality, human-generated data – including the material professional journalists produce – more valuable.

Finally, several AI systems are best characterized as having *narrow* intelligence, competent at specific tasks or domains but not the *general* intelligence that requires strong reasoning and high performance across several domains. The recent rise of agentic AI – systems that aim to break down complex tasks and automatically deploy the right tools to complete them – points to a future in which AI can take on more multifunctional roles. For example, asked to “plan our company’s product launch,” an AI agent might gather and analyze market data, draft a news release, identify journalist contacts, and schedule outreach, coordinating each step with minimal human direction.

This shift to agents rather than reactive chatbots expands the breadth of what AI can accomplish. It also shows how rapidly the landscape is changing, amid growing questions about the widescale disruption to white-collar jobs that could be on the horizon as agentic systems become more widely adopted in workplaces.⁸⁰ Nevertheless, for all the hand-wringing about the future, most researchers still believe that artificial general intelligence (AGI), also called superintelligence, remains far from imminent. Understanding these technical realities – what AI is, how it learns, where humans shape it, and what limits it – grounds our analysis of how these technologies are helping reshape journalism.

AI Is More Than Code

Understanding how artificial intelligence is helping reshape journalism also requires us to acknowledge that what AI can actually do is not the same thing as what AI means to people.

The anxiety rippling through newsrooms today and the lofty promises made by industry leaders are not entirely (or even largely) anchored to the technology’s technical capabilities, as the previous waves of AI have shown. They emerge from something

deeper – what media scholar Taina Bucher calls the “algorithmic imaginary,” or the way people imagine, perceive, and experience algorithms.⁸¹ This concept draws on a core insight from science and technology studies: technologies are never just their code.⁸² They also embody meanings constructed through collective uses, public discourse, historical precedents, and accumulated mythologies. In other words, things that have nothing to do with the code itself – and what it can do, or is designed to do – profoundly shape how people think about, talk about, and ultimately use AI.

The Power and Peril of Hype

AI has been hyped up dramatically in recent years, spawning an imaginary associated with both utopian promise and apocalyptic threat. An open letter from the Center for AI Safety, co-signed by hundreds of AI executives, researchers, and engineers, declared that “mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks, such as pandemics and nuclear war.”⁸³ At the same time, countless news reports paint AI as a savior – capable of curing cancer, advancing social justice, and freeing us for more leisure time. Journalists, of course, both amplify and contest such narratives.⁸⁴

For all its seeming unseriousness, technological hype actually serves five distinct functions.⁸⁵ Attentional hype makes a technology appear especially important and deserving of close scrutiny, and some journalists and outlets exploit it strategically by positioning themselves as either progressive innovators or guardians of tradition. Orientational hype steers that attention toward particular goals, products, or practices, often redirecting limited newsroom resources in the process. Signaling hype shapes how technologies are presented – both in reporting and in organizational adoption – such as through strategic messaging that reinforces narratives of innovation and inevitability. Mobilizing hype drives tangible action, including partnerships with tech vendors, policy changes, professional ethics codes, and union-led safeguards against AI misuse. Finally, reflexive hype prompts journalists to fundamentally rethink their roles, values, and professional identity – particularly when GenAI begins performing complex, creative tasks once considered distinctly human.

The news industry is indeed awash in both genuine interest and considerable hype around AI. As one journalist observed in

a 2019 survey by the London School of Economics, which was conducted well before ChatGPT sent news organizations scrambling: “Adopting AI is not a choice, it is a marathon that every organization needs to start running if they haven’t already.”⁸⁶ By late 2024, a survey of more than 1,000 journalists in the United Kingdom found that more than half were using AI professionally at least once a week and the overwhelming majority believed their organizations would only increase AI integration efforts.⁸⁷ That survey also found that the most active users of AI were those in positions to make strategic, long-term decisions on behalf of their news organization. As Liz Corbin, director of news at the European Broadcasting Union, put it: “As a newsroom leader, there is nothing more daunting in this age of generative AI than the need to be at the cutting edge of technology while also being responsible for protecting the integrity of our journalism.”⁸⁸

As this drama unfolds, newsrooms are traversing what the consulting firm Gartner calls the “hype cycle.”⁸⁹ This framework tracks how technologies tend to evolve from an initial breakthrough that leads to a peak of inflated expectations, followed by a steep decline into the trough of disillusionment when implementations fail to deliver, then gradually rebound up a slope of enlightenment as improved iterations begin meeting some promises, and finally reach a plateau of productivity where mainstream adoption takes hold – provided the technology reaches that point at all. At the time of writing, journalism broadly remains near the peak of inflated expectations, though some AI applications – automated transcription as the clearest example – have already reached widespread adoption even as some skepticism has gained steam.

Many in the news industry will eventually confront a sobering reality: AI will not match all its hype. Nevertheless, they will also likely see it proven useful and, in specific contexts, genuinely transformative. This ongoing “technological drama” will continue reshaping the myths we tell about AI.⁹⁰

Sociomateriality and Technological Flexibility

We should not limit our understanding of AI to its hype and technical capability. It is important to also recognize that when people actually use technologies, they often do so in unforeseen and novel ways.⁹¹ Twitter (now X), after all, was designed as

a personal status-update tool – its original prompt was literally “what are you doing?” – yet journalists and eyewitnesses remixed it into a global breaking-news infrastructure its designers never imagined.⁹² SMS was considered a trivial byproduct of cellular networks by the engineers who built it; it became one of the most consequential communication tools in history.⁹³ In other words, users’ behaviors are structured but not determined by designers’ intentions.

The social construction of technology framework offers a useful lens for making sense of emerging technology.⁹⁴ It highlights the flexibility of technology: both the purposes and features of technologies are remixed when users put them into practice and as they are filtered through relevant social groups like professional communities. Because of their generative and increasingly general-purpose nature, AI tools will surely be used selectively or in unexpected ways. Journalists are already using general-purpose AI tools not necessarily as writing assistants – their most marketed application – but as document-interrogation tools, thinking partners for story structure, or translation aids for reaching non-English-speaking communities. The results of such widespread experimentation will likely surprise even the most savvy experts making prognostications about the future. Such unpredictable forms of remixing become especially important to consider when a technology like AI is being developed primarily by actors outside journalism, even as it proves deeply consequential for the future of news.⁹⁵

This sociomaterial perspective reveals why we must look beyond code and technical capabilities when examining AI’s impact on journalism. Technologies are not deterministic; they do not unfold through their own internal logic with inevitable consequences. Instead, technological deployment is simultaneously shaped by and shapes an array of factors: social expectations, institutional and economic pressures, professional values and routines, regulatory frameworks, and technological literacy among practitioners.⁹⁶

This is ultimately good news for journalism. While the field is being reshaped in part by AI, it is not entirely beholden to visions others have for the technology. How journalists and news organizations understand, adopt, resist, and remix AI will matter as much as the technology’s capabilities. The social imaginary around AI – its hype, its mythology, its contested meanings – is

not mere noise obscuring the “real” technological story. It is part of that story, shaping material practices and institutional futures. Recognizing this opens space for journalism to actively shape its AI-infused future rather than passively accept it.

A Roadmap for This Book

Making sense of journalism in the age of AI requires us to look both at the journalism profession and beyond it – to understand the antecedent forces, the potential consequences, and what must be done to steer the future of news toward desired ends. The chapters that follow trace a path from the newsroom outward: examining how AI is reshaping news production, who controls the infrastructure on which journalism increasingly depends, what this means for democratic life, and how we might cultivate, through education and literacy, the human capacities necessary to navigate what lies ahead.

Throughout the book, we return to a central fork in the road: Will AI be implemented primarily in ways that accelerate the hamster wheels, with the same work being performed faster and cheaper? Or will it be harnessed in ways that disrupt them, enabling forms of journalism – and forms of human flourishing – that might have been previously unimaginable? This is not a one-time choice; it will be made continually in countless decisions by individual journalists experimenting with new tools, by editors setting workflow policies, by executives negotiating licensing deals, by policy-makers drafting regulations, by citizens deciding which information sources to trust, and by educators designing curricula. The accumulation of these decisions will determine whether the constitutive moment we are living through leads to journalism's revitalization or its further erosion – and, beyond journalism, whether AI strengthens or undermines the epistemic foundations on which democratic societies depend.

In Chapter 2, we examine a pivotal tension in contemporary journalism: whether AI merely accelerates established workflows or enables a more fundamental reimagination of news work. While AI currently automates routine tasks and enhances efficiency, most applications appear to remain augmentative rather than transformative – reflecting journalism's historical tendency to normalize new technologies within existing practices. We chal-

lenge the prevailing discourse centered on replacement, however, arguing that this framework inadequately captures AI's potential. Instead, we advocate for an approach that focuses on complementarity: asking what journalistic work AI makes newly feasible, not merely what it replicates. We further point to some emerging applications that are more likely to prove transformative. AI thus emerges as a possible catalyst for reinvigorating journalism, contingent upon intentional institutional choices.

In Chapter 3, we examine the transformation of media power in the age of AI. Where twentieth-century debates centered on ownership concentration – who controlled newspapers and broadcast channels – contemporary struggles reflect a more profound shift toward infrastructure-based concentration. The current conflict among publishers, platforms, and AI companies is not merely economic but also architectural: it concerns who controls the algorithmic systems mediating news production, distribution, and consumption. We trace this shift through the emergence of the “Great Houses of AI”: the handful of technology giants that now control the cloud computing, foundational models, and distribution systems on which journalism increasingly depends. This transformation repositions power away from content ownership toward control over the infrastructural conditions that enable media work and audience engagement. It demands reconceptualizing key journalistic principles – autonomy and professional authority chief among them – and raises urgent questions about journalism's sustainability when its operational infrastructure resides in external hands.

In Chapter 4, we contend that AI's proliferation will fundamentally reshape how individuals engage with information and participate in democratic life. We examine these implications across three dimensions essential to democracy: public understanding (how citizens make sense of the world), deliberation (how they reason together about shared concerns), and participation (how they act on those convictions). AI development is currently driven by competing capitalist and authoritarian logics (chiefly from the United States and China), which threaten democratic principles by incentivizing systems that prioritize control, isolation, and the confirmation of existing predispositions over genuine deliberation and truth-seeking. Some democratic benefits will likely emerge from AI developed under these logics, but they will be unevenly distributed while existing power struc-

tures are reinforced. These trajectories are not inevitable, though. Democratic outcomes remain possible, we argue, if stakeholders deliberately reconfigure incentive structures and infrastructures to direct the development and implementation of AI toward democratic goals – a challenge that journalists can be a part of addressing, but one that extends well beyond journalism alone.

In Chapter 5, we argue that journalism education offers a particularly revealing site for our reckoning with AI's implications for learning and pedagogy. Here, at the intersection of craft and cognition, the challenges become apparent: journalism students must learn to identify newsworthy ideas, evaluate sources, and construct compelling narratives – but these are capacities AI can increasingly approximate. While universities understandably worry about academic dishonesty, we argue that framing AI primarily as a policing problem misconstrues what is at stake. The more consequential question, we argue, is whether journalism education can serve as a site of genuine empowerment, developing students' literacy and discernment to work thoughtfully with AI tools rather than become diminished by them. This is a challenge that requires pedagogical reimagining rather than just new policies.

We conclude with Chapter 6, where we argue that the status quo of journalism and AI development appears likely to result in the acceleration of the hamster wheel's demands. We point to four interlocking dynamics covered in the preceding chapters: *deprofessionalization*, *infrastructural capture*, *democratic corrosion*, and *cognitive atrophy* – underscoring that the throughline uniting them is the notion of *epistemic thinness*, or a growing gap between *feeling* informed and *actually being* informed. But the constitutive moment we are living through remains open. We make a case for what a reimagined journalism in the age of AI might look like and offer some orientation toward that end. The hamster wheel is not destiny, and decisions being made right now – by journalists, news organizations, educators, policy-makers, and citizens alike – will shape what happens next.

2

The Newsroom Accelerates

Freelance journalist Andrew Quemere stumbled upon something unsettling in October 2025.¹ Scrolling through MassLive, a regional news website serving Massachusetts, he began noticing a small disclosure tucked at the bottom of some articles: “AI was used to draft this story based on data provided by . . .”

Sometimes the source was the IRS. Other times, a product recall notice. Once, it was even Walmart – the very company being reported on. Most stories added a reassuring coda, though: “It was reviewed and edited by MassLive.”

Quemere made his disdain public on the social media network Bluesky. This caught the attention of Jonathan Gerhardson, a fellow local journalist at *The Shoestring*, who kept digging. Using basic Google searches, he identified more than 200 stories with similar disclosures. Many seemed innocuous: where to watch a game, which state fair food was best – the kind of “soft news” that has always filled the gaps between reporting on major issues. But one story left him feeling uneasy.

Among the AI-generated articles was a story about a Massachusetts woman who had allegedly threatened to kill US Immigration and Customs Enforcement agents.² The story was published during a period of intense public debate over ICE’s enforcement actions. It carried no reporter’s byline, only a generic “MassLive.com Staff” attribution. An algorithm, fed data from federal charging documents, had drafted the article. A human

editor at MassLive – a subsidiary of the century-old media conglomerate Advance Publications – had approved it for publication.

This didn't sit right with Gerhardson. Surely a story about such a sensitive topic of national interest deserved to be written by one of the organization's human journalists.

MassLive editor-in-chief Ronnie Ramos did not fully share Gerhardson's concerns.³ He insisted the outlet remained "steadfastly committed to journalistic integrity" and pointed to the AI disclosures as proof. As for the missing bylines? "I'm not going to have a reporter put their byline on a story that they didn't write," Ramos explained.

Fair enough, Gerhardson conceded. At least they weren't fabricating synthetic journalists or forcing human reporters to claim authorship of machine-generated copy. But the logic felt hollow.

"Given the inherent risks that journalistic scrutiny imposes upon its subjects," Gerhardson argued, "the practitioner has an ethical imperative to assume full accountability for their work. One must be willing to stake their professional reputation on the integrity and veracity of the content produced. Seeking to obscure this responsibility through the opaque veil of artificial intelligence represents an abdication of journalistic duty – and doesn't help the crisis of flagging trust in media, either."⁴

Gerhardson was no AI skeptic. He had used AI to help him code an interactive map and scrape a government website and saw that organizations like ProPublica in the United States and iTromsø in Norway were deploying the technology to seemingly enhance their work. Indeed, in 2024 and 2025, AI played important roles in investigative journalism that had been recognized as some of the finest work in the profession, garnering Pulitzer Prizes.⁵

But this felt different. In Gerhardson's view, MassLive didn't seem to be using AI to empower its journalists to pursue deeper investigations or improve their work. It was largely using AI to churn out cheap content – journalism's own version of "AI slop" – without so much as a clear byline.

The MassLive case exemplifies a crucial question confronting the news industry: Will AI make journalism's hamster wheel spin faster, or might it help reporters break free from the grind?

The answer hinges on how newsrooms choose to integrate the technology. AI is already reshaping every stage of news

production – from suggesting story ideas to personalizing distribution. It automates routine tasks, helps journalists brainstorm ideas, writes headlines and whole stories, highlights relevant multimedia, and empowers copy editors to catch errors faster.⁶ More ambitiously, it offers opportunities to fundamentally reimagine news products and audience engagement through transformations like conversational journalism – a shift toward a more dialogical format that may include the use of chatbots and the like to directly address individual users’ stated questions in real time and incorporate additional context as needed.⁷

For journalists and news organizations, AI crystallizes pressing questions about identity, ethics, and institutional purpose.⁸ And the most pressing question may well be: What does it mean to be a journalist when machines can perform large portions of your work? This isn’t just about technology. It’s about how individuals and organizations adjust their thinking and practices to preserve – if not enhance – the ethical commitments, creative processes, and societal functions that underpin journalism’s authority and societal value in the face of, or with the help of, a potentially transformational technology.

But David Caswell, who previously led news product innovation at the BBC, argues that most current applications of AI have been augmentative rather than transformative.⁹ They increase efficiency within existing workflows but have rarely challenged how newsrooms fundamentally think and work. This aligns with longstanding observations from journalism scholars that new technologies are typically “normalized” to fit existing practices, at least initially.¹⁰ Indeed, transformation, if it ever emerges, usually does so gradually through competing responses of adoption, adaptation, and refusal.

Against this backdrop, much discourse around AI in journalism has focused on whether machines can replicate or replace human work. This is understandable, but it’s a limited framework that misses the technology’s broader potential.¹¹ Evaluating AI solely on its ability to mimic existing work creates, in essence, an unhelpful barometer – misleading in what it measures. Current AI systems excel at tasks like pattern recognition while struggling with others that involve, say, reasoning, because their underlying architecture differs fundamentally from human cognition.¹²

A more productive approach emphasizes complementarity and synergy. Rather than asking what AI can replace, we should

ask: How can AI enable journalistic work that previously would have been impractical, infeasible, or even impossible?¹³ This represents a rare opportunity for journalists to escape the hamster wheel they've long lamented and craft new roles, processes, and arrangements that restore meaning to their work.

Realizing this potential requires engaging with fundamental questions about which elements of journalism are necessarily human and reimagining how news work should be organized and around which ethical principles. We view AI not as a replacement for journalism but as a tool that can reinvigorate the practice for a new age – provided journalists and news organizations make deliberate choices to do just that.

Journalism Is an Evolving Practice

Journalism has never been a static institution. Its values, practices, and processes have shifted in response to new technologies, changing audience relationships, and evolving social, political, and economic contexts. Media scholar Wiebke Loosen and colleagues capture some of this dynamism through their “X journalism” framework, which treats the proliferation of new labels – robot journalism, citizen journalism, solutions journalism – as empirical evidence of the field’s ongoing adaptation.¹⁴

This raises an important question: Will AI give rise to yet another specialized variation called “AI journalism”? Or will artificial intelligence become so thoroughly integrated into everyday practice that we stop distinguishing it altogether – much as we no longer bother saying “online journalism”?

The answer depends on how journalists respond. Research across multiple disciplines reveals three distinct patterns of engagement with new technologies: adoption, adaptation, and refusal.

For example, the technology acceptance model, theory of planned behavior, and unified theory of acceptance and use of technology explain factors shaping adoption, such as an individual’s existing attitudes, how approachable the technology is perceived to be, and the expected benefits. Adaptive structuration theory, domestication theory, and social influence theory demonstrate how journalists actively *adapt* – and do not just passively *adopt* – technology through formal and informal rules (e.g., codes

of ethics), structures (e.g., training programs and informal experimentation), and reference networks (e.g., opinion leaders), often by focusing on specific features of a technology or by repurposing it in entirely unintended ways. Status quo bias theory, innovation resistance theory, and technology resistance theory help explain why some journalists may actively opt to refuse using AI in their work altogether, including their discomfort with changing routines, concerns about job security and autonomy, and perceived conflicts with professional values.

These theories collectively show that technological transformation depends as much on human factors and workplace arrangements as on the technology's material capabilities. With this in mind, we turn to how AI is already helping reshape journalistic practices across five stages of news production: coming up with story ideas, sourcing information, verifying content, telling stories, and distributing news.

AI and Story Ideation

Journalists are in the business of hunting for stories. Traditionally, this hunt involves attending public meetings and press conferences, mining records, reviewing competitors' coverage for ideas, and cultivating sources who provide tips.¹⁵ During these activities, journalists assess potential stories against familiar newsworthiness criteria:¹⁶ Is it timely? Does it involve recognizable people? Will audiences care enough to share it? And, crucially, they must manage risk – newsrooms need reliable output, so some stories must be “sure things.”¹⁷

This practical reality produces templates that make news production efficient and predictable. Journalists learn to recognize certain patterns – the political scandal, the human-interest feature, the controversial quote – and actively seek ideas that can be fitted into the proven molds. Templates offer efficiency, but they also create tension. Journalism, after all, is supposed to be a creative enterprise.¹⁸ Editors and reporters constantly seek to add something *extra* to their stories: a fresh angle on a familiar event, an unexpected voice, a compelling visual treatment. This drive stems from both intrinsic motivation (breaking the monotony) and market pressure (standing out in a crowded information landscape).

Fitting Templates and Sparking Thoughts

Some newsrooms already deploy AI to identify stories that fit established templates. These systems scan competitors' articles, press releases, social media feeds, and search trends to bring story ideas to light, essentially automating the hunt for low-risk, newsworthy topics. Sisi Wei, Chief Impact Officer at CalMatters, says that "AI can be an incredible tool for good" and that her organization uses it to keep up with government officials.¹⁹ The technology, Wei writes, "tracks pretty much everything we can track about California's state legislators, parses that fire hose of data, and then generates story ideas for our reporters." Some AI-generated tips, she notes, "would have taken a data-savvy political reporter weeks, if not months, to find on their own." The resulting stories have already influenced legislative action.

But AI doesn't have to be restricted to the familiar; it can spur creativity.²⁰ Journalists increasingly use large language models (LLMs) as sounding boards: tools for developing nascent ideas, working through story angles, and overcoming creative blocks. Meghan Murphy, director of the Online News Association's AI in Journalism Initiative, describes AI as "an interesting thought starter."²¹ While it often "spits back what I already know or what I wrote anyway," Murphy finds it sometimes highlights angles worth exploring further. The key, as data journalist Paul Bradshaw suggests, lies in well-constructed prompts that describe the situation, specify constraints, and provide examples.²²

A Serendipity Engine

As AI becomes further embedded in browsers, operating systems, and information gateways, it is poised to reshape how journalists encounter ideas without requiring direct prompting. AI agents could be configured to interrupt journalists' information consumption routines, functioning as a "second brain" that highlights unexpected connections or reveals valuable information outside a journalist's typical information-seeking patterns.²³

This vision aligns with cognitive models showing that creativity often emerges from collisions between seemingly unrelated information.²⁴ Human creativity requires both access to diverse inputs and the capacity to recognize non-obvious connections.²⁵ AI systems designed as serendipity engines could deliberately

introduce productive friction into routine behaviors, helping journalists overcome cognitive fixations and confirmation biases. Such systems would blend structured support with strategic surprise – providing scaffolding for creative work while systematically introducing the unexpected.²⁶

Any AI system that generates story ideas must be trained on what counts as newsworthy, though, which requires translating editorial values into code.²⁷ To maximize creative potential, these systems would need to point journalists toward novel, worthwhile paths rather than merely replicate their existing judgment. This is no easy task given the pattern-focused design of most current AI systems.

Journalistic creativity involves ethical and empathetic judgments about which stories matter and how to tell them in ways that respect and resonate with human experience.²⁸ The future, in our view, likely lies neither in pure human ideation nor in fully automated story generation, but in a hybrid approach where human intuition works alongside both visible and invisible digital assistants to collectively decide what stories need telling and how to best tell them.

AI and Sourcing

Journalists depend on sources to provide context, interpretation, access, and insider knowledge for stories they cannot directly observe. These sources – ranging from academic experts and government officials to advocacy organizations and ordinary citizens – help journalists understand complex issues, gain access to restricted spaces, and illustrate broader social concerns through personal narratives.²⁹ But sourcing practices have never been neutral: they reflect deep structural inequalities that AI both promises to address and threatens to reinforce.

Two key factors shape which voices appear in news coverage: authority and availability.³⁰ Individuals and organizations within established power structures – government representatives, corporate leaders, institutional experts – are treated as especially newsworthy because their decisions shape public life and their institutional standing signals legitimacy. But even diligent and well-intentioned journalists face practical constraints. Time pressures push them toward sources who are easily reached and

responsive, such as the contacts already in their address books and the spokespeople who return calls promptly.

These structural patterns are reinforced by interpersonal dynamics. Journalists and sources are subject to homophily: the human tendency to associate with people similar to ourselves. For example, journalists are more likely to contact, and receive cooperation from, sources who share their gender and race.³¹ As a result, female experts and ethnic minority community members have historically been underrepresented in coverage, and their perspectives have consequently been marginalized in important public debates.

A Sourcing Auditor

AI systems have been deployed to identify systematic sourcing biases and suggest alternatives beyond journalists' existing networks. Several organizations have experimented with automated tools that review draft stories and flag imbalanced sourcing patterns. UNHEARD, a project led by investigative journalist Bette Dam, "aims to help news organizations reveal potentially overlooked narratives by using AI to audit who is quoted in their articles."³² The ultimate goal is to allow "news organizations to analyze reporting, identify potential imbalances, and suggest alternative sets of sources to improve accuracy as a story unfolds."

This isn't a pipe dream. Research demonstrates that automated systems can identify sources with high accuracy, enabling them to predict when articles would benefit from additional perspectives and even recommend specific individuals to expand representation.³³

However, AI-powered auditing doesn't necessarily eliminate sourcing hierarchies. Journalists' disproportionate reliance on government sources doesn't stem from limited imagination but from the perceived authority those voices confer. And scientific sources tend to be underutilized not because journalists forget about them, but because academic timelines are often incompatible with journalistic deadlines – not to mention scientists' tendency toward cautious, jargon-laden responses. While AI systems can help journalists more quickly identify more diverse voices – an improvement, no doubt – they cannot resolve the deeper professional imperatives.

An Interview Assistant or Interviewee

AI is also helping journalists conduct preliminary research, learn about unfamiliar topics, and generate interview questions tailored to specific sources and story angles.³⁴ While these questions typically need refinement, they can help journalists break away from formulaic templates, provided that they use them with such intent.

In extraordinary cases, AI has even been used to automate elements of the interview process itself. For example, the Swedish tech company United Robots offers tools that, with limited human involvement, analyze sports recaps, generate questions for post-game interviews with coaches, send those questions via text messages, and insert the responses into match reports.³⁵

But interviews are rarely mechanical transactions. As Ruby Voge, Opinion co-editor at the *Daily Free Press*, contends: “ChatGPT can’t interpret the nuance of a fleeting facial expression or a subtle vocal inflection. It can’t wait outside the mayor’s office for an early-morning interview or embed itself into the life of a profile subject. It can’t interview disaster survivors with empathy and understanding or hold world leaders accountable with skepticism and perseverance.”³⁶

Even more complicated is when an AI agent becomes the source itself. In April 2023, the German magazine *Die Aktuelle* published what it claimed to be the “first interview” with Michael Schumacher, the Formula 1 legend who suffered a near-fatal brain injury in 2013 and hadn’t spoken publicly since.³⁷ But the interview wasn’t real; it was an AI-generated confabulation of what the magazine imagined he *would* say a decade after his accident. There was immediate public backlash, leading the magazine’s publisher to fire its editor, apologize for a “tasteless and misleading article,” and reach a compensatory settlement with the Schumacher family.³⁸

Other incidents of AI stand-ins have followed, raising concerns about authenticity, consent, and the blurring of lines – and whether such practices constitute legitimate journalism or exploitation of emotions like grief.³⁹ While AI can facilitate novel forms of connection, and is being used increasingly by the “digital afterlife industry” to simulate deceased loved ones through so-called “deadbots,” it also risks commodifying human tragedy, desensitizing people, and trivializing death.⁴⁰

A Rapid Transcription Tool

Audio transcription represents an area where AI has already become indispensable. Journalists routinely use AI to automatically transcribe recorded interviews, identify noteworthy segments, make audio searchable, and cross-reference sources' current statements with past remarks.⁴¹ This seemingly basic automation saves hours of work weekly – though it can be subject to mistakes, such as when popular agents like Whisper occasionally “hallucinate” and transcribe something very different from what was said.⁴² Similarly, AI can analyze large volumes of transcribed documents, extracting and summarizing key information.⁴³ The detection doesn't need to be perfect – if an AI assistant directs a journalist toward a smaller set of likely relevant documents, it has already proven valuable.

Michigan Public Radio in the United States, for example, created an AI-powered tool called “Minutes” that automatically scrapes and transcribes videos of city council meetings, allowing reporters to monitor public proceedings they cannot attend in person.⁴⁴ Reporter Dustin Dwyer finds the tool to be both a time-saver and a way to diversify his reporting: “I use it to skim transcripts to get a general idea of the meeting, without having to spend hours sitting through the whole thing. And the new keyword alert system allows me to track issues and find sources in communities that I normally don't cover.” The tool enables the station to continue performing monitorial journalism even as its resources shrink.

These examples show that AI-powered tools can improve sourcing practices in multiple ways, but journalists need to exercise care. Current tools remain best suited to identifying readily accessible, digitally present sources rather than marginalized voices who may lack online profiles or institutional affiliations, which can reinforce rather than challenge established hierarchies. Furthermore, AI-generated interview questions can inadvertently perpetuate problematic assumptions about whose perspectives matter and what questions are worth asking because of how their models were trained.⁴⁵ Perhaps most significantly, AI systems struggle to replicate the human capacity for developing rapport with sources, interpreting nonverbal cues, or approaching vulnerable subjects with the empathy their stories demand – the very kind of high-touch interactions that need to be fostered for

journalists to rebuild trust in communities.⁴⁶ A growing number of people, though, report establishing relationships with chatbots that they perceive to be meaningful, suggesting this may yet be possible. Nevertheless, the path of hybridity again strikes us not only as the most fruitful but also the most prudent given the state of the technology.

AI, Verification, and Fact-Checking

Accuracy and truth-telling have become more than professional values in journalism. As audiences drown in an ocean of information sources, verification and fact-checking practices serve as both credibility markers and a way for outlets to stand out.⁴⁷ News organizations are not just competing for clicks; they are competing for audience trust in the hope that it translates into repeat visits and revenue.⁴⁸

A fact-checking boom began around 2016, when Donald Trump's election and the Brexit referendum ushered in what some called a "post-truth" era.⁴⁹ For example, fact-checkers worldwide created ClaimReview, a standardized tagging system launched in 2015 that helps search engines and social platforms highlight verified information in places like their search results and news feeds.⁵⁰ Google initially embraced the initiative, which now includes hundreds of thousands of fact-checks from organizations like Agence France-Presse, *The Washington Post*, and Rappler. Fact-check snippets on Google products were viewed over 120 million times in the first half of 2024 in the European Union alone.⁵¹

This matters because fact-checking works. Research consistently shows that exposure to fact-checks reduces belief in misinformation, with corrective effects persisting over time.⁵² But just as fact-checking gained prominence, it became exponentially harder.

Traditional Verification and Fact-Checking

Three forces have converged to challenge traditional fact-checking.⁵³ First, the democratization of publishing tools empowered anyone to create professional-looking content and distribute claims widely online, producing a surge of misleading and fabricated information. Second, budget-strapped newsrooms began

facing even harder choices between fact-checking erroneous claims and pursuing new stories, and as dedicated fact-checkers and copy editors were increasingly laid off, those who remained were further overextended by being forced to take on the additional verification and reviewing tasks. Third, fact-checking itself became politicized, prompting major platforms, including Google and Meta, to retreat from verification partnerships. Furthermore, studies have found that initial falsehoods circulate far more widely than the fact-checks that follow, making fact-checks only partly successful most of the time.⁵⁴

When generative AI (GenAI) arrived, it turbo-charged the challenge. The speed and scale at which AI can produce and spread false information have rendered real-time, human-scale verification nearly impossible. As investigative researcher Henk van Ess observes: “Traditional fact-checking takes hours or days. AI misinformation generation takes minutes.”⁵⁵

Professional fact-checking typically follows several stages.⁵⁶ Journalists first select claims worth verifying within their story or by scanning social media, news reports, and online comments to identify statements being made elsewhere that warrant verification. Selection criteria include the claim’s verifiability, potential impact, and whether multiple sources raise the same concern. Next comes assessment and contextualization, which recognizes that technically accurate statements can mislead without proper context. Journalists trace claims to their origins, search for primary documents, review previous reporting, identify authoritative data, and assemble evidence. Finally, they decide whether to include, refute, omit, or contextualize the claim. Dedicated fact-checking organizations typically conduct internal deliberations before publishing separate fact-check stories, followed by editorial review for accuracy, clarity, and neutrality.

Automating and Augmenting Verification

Recognizing they can’t outpace AI, researchers and journalists have developed tools to automate some parts of that process while supporting others. These systems detect claims in speeches, social posts, and articles; identify synthetic media like deepfakes; and spot disinformation campaigns through pattern analysis.⁵⁷

Real-world implementations reveal both promise and pragmatism. Full Fact, Maldita.es, and the European Fact-Checking

Standards Network built a multilingual system monitoring short-form video from YouTube Shorts, TikTok, and Instagram Reels.⁵⁸ It extracts claims using speech-to-text and text recognition and then clusters them into narratives, helping journalists *prebunk* false information before it spreads further.

“AI tools help us find, filter, and follow claims at scale,” reflects Full Fact Senior Product Manager Kate Wilkinson.⁵⁹ “Human fact checkers provide the expertise, judgment and accountability. Together, they give us a fighting chance to make public debate more accurate and fair.”

Full Fact’s hybrid approach acknowledges another challenge: truth rarely fits binary categories. Journalists routinely confront ambiguous, unverifiable claims or those dependent on future events. More critically, many information sources needed for verification don’t exist digitally and thus don’t appear in AI training data.

A comparative study of fact-checking in Brazil and Germany found that journalists regularly encounter cases where reputable sources of information cannot be located or inquiries go unanswered, leaving claims unresolved despite the journalists’ rigorous efforts.⁶⁰ AI excels at verifying simple, measurable claims involving publicly documented information or previously fact-checked claims.⁶¹ It cannot navigate the messy reality of missing documents and reluctant sources.

Moreover, the propensity of LLMs to “hallucinate,” or to unintentionally produce incorrect or fabricated information and present it as fact, compounds the problem.⁶² Even when AI provides technically accurate information, it may lack contextual grounding for proper interpretation. Journalists must therefore verify what their verification tools tell them – a frustrating irony. Consequently, some journalists remain skeptical about AI’s efficiency gains. A study of 22 fact-checkers found existing tools often failed to fit their workflows and produced erroneous or irrelevant results that increased rather than reduced their workload.⁶³

Risks to Credibility

News organizations also face a new threat that goes beyond their own use of AI: having their credibility damaged by AI systems they don’t control. In December 2024, the BBC publicly criticized Apple Intelligence after it created a misleading notification that

falsely implied BBC News had reported that Luigi Mangione – the man accused of murdering healthcare executive Brian Thompson – had shot himself.⁶⁴ The BBC never published this claim, yet its credibility suffered among those who saw the error. Worse, the BBC had no way to know who was exposed to the erroneous notification or on what scale. While Apple paused that specific feature, it merely added disclaimers about potential inaccuracies elsewhere in its operating system.⁶⁵

This represents a troubling inversion: rather than using AI to combat misinformation, news organizations find their own brands weaponized or placed at risk of being tarnished by AI-generated falsehoods.

The lesson emerging from newsrooms is that while AI can help journalists rapidly process certain types of information and spotlight overlooked contentious claims, the broad practice of verification still requires a degree of judgment, contextual awareness, and accountability that algorithms currently appear to be incapable of exercising. The implication is clear: human fact-checkers will need to learn to work effectively with AI technologies to defend truth in an environment where information, misinformation, and disinformation all operate at machine speed.

AI and Storytelling

Journalists use storytelling as both craft and strategy, carefully selecting, sequencing, and framing information to engage with audiences while maintaining professional standards. Compared to other forms of storytelling, journalism tends to be more formulaic, especially in genres like breaking news.

For example, the inverted pyramid, which front-loads the most important information, has dominated breaking news in Western journalism for over a century. Of course, other structures also exist – the “martini glass” blends summary leads with chronological narration and the “kebab” bookends analysis with humanizing anecdotes – and they serve specific purposes, from accommodating time-pressed readers to adding emotional resonance to feature stories.⁶⁶

Yet even though these conventions provide consistency and efficiency, critics have long derided news writing as overly rigid. In recent years, facing fierce competition from digital upstarts

and social media influencers who favor formats like bullet points, Q&As, and visual illustrations, traditional news organizations have become more open-minded about storytelling structures and styles. Indeed, research suggests that audiences increasingly gravitate toward emotionally resonant content that breaks from formulaic norms.⁶⁷ Now, AI is accelerating this evolution – in ways that are sometimes promising and other times disastrous.

Early Automation of Newswriting

One of the more visible early applications of AI in newsrooms – even though some observers objected to calling it that – was the automation of newswriting in certain genres, including finance and sports.⁶⁸ In 2014, the Associated Press started automating the production of thousands of stories about corporate earnings reports and sports recaps by extracting information from databases or standardized reports, applying simple logic to determine newsworthiness, and inserting the information into pre-written templates. Proponents celebrated this as liberation from drudgery – few journalists relish writing yet another earnings brief, after all.

But critics charged that those automated stories were remarkably dull and risked alienating audiences, who would scroll past or, worse, develop negative associations with the news organization.⁶⁹ This wasn't merely content with little value; it potentially carried negative value by crowding out higher-quality content and reinforcing perceptions of journalism as increasingly catering to the lowest common denominator.

New Possibilities with Generative AI

The emergence of GenAI has begun to address some limitations of early automation while introducing fresh concerns. It is now routinely used to generate multiple headline variants for editors to choose from, create useful article summaries that are appended to stories, and personalize content in increasingly sophisticated ways, such as by localizing data-driven stories more naturally to make them more relevant and interesting to specific communities. These applications streamline workflows and free editorial staff to focus on higher-value tasks.

GenAI also enables some fundamentally novel storytelling opportunities, such as the use of chatbots to create individualized

dialogic news experiences. For example, *TIME* magazine launched the TIME AI Agent in November 2025, which integrates language understanding, voice synthesis, translation, and search capabilities. Readers can ask questions, request summaries, generate audio versions of stories, and translate reports – all through a single interface that can, for instance, produce an audio summary of recent interviews with world leaders. Jason Droege, the CEO of Scale AI, which powers the tool, described it as “a blueprint for how publishers can use AI agents to create a more meaningful relationship between their audience and their content.”⁷⁰

The emergence of TIME’s AI Agent – along with similar tools like *The Washington Post*’s Ask the Post AI and the *Financial Times*’s Ask FT, both launched with fanfare in 2024⁷¹ – represents the materialization of a long-imagined future: one in which people get their news from conversations with AI-powered news agents rather than discrete articles. This marks a fundamental reimagining of the journalist–audience relationship, simultaneously bringing audiences closer to a synthetic intermediary while distancing them from the human journalists whose work feeds these systems. Some tech companies are pursuing similar visions, potentially sidelining news organizations from developing direct relationships with their audiences altogether.

Prototyping New Ideas

Another key development is the emergence of vibe coding: an informal approach to software programming where a person, coder or not, uses plain language to describe the “vibe” of what they want to create, and the AI tool mostly does the rest in generating the code. The human focuses on the goal, and the AI handles the line-by-line writing. This approach, further discussed in Chapter 6, has empowered journalists to prototype nascent ideas, experiment with specialized back-end tools, and create public-facing products without having to rely on technical specialists.⁷²

Joe Amditis, associate director of operations at the Center for Cooperative Media, exemplifies this shift.⁷³ Despite describing himself as not “a real software developer,” he has used LLMs to build dashboards, create interactive features, process data dumps, and develop web tools for local news organizations across New

Jersey. According to Amditis, “the most valuable skills for vibe coding are the skills of a good project manager,” such as identifying requirements, clearly articulating what you want, providing direct and clear feedback, and organizing ideas and documentation. He adds: “These skills are now more immediately valuable than the skills of a traditional software developer. You’re a *manager* now, delegating the coding to an LLM.”

This isn’t just making some work easier. It’s a form of upskilling that allows journalists without previous technical training to do new things they otherwise would not have been able to, such as developing interactive graphics, data visualization tools, and small applications that enhance their storytelling capabilities.

When Automation Goes Wrong

The promise of automated storytelling comes with serious risks, particularly when organizations move too quickly or fail to maintain adequate oversight. AI’s tendency toward hallucination becomes even more vexing when false content reaches audiences.

In late 2023, *Sports Illustrated* faced a firestorm when other journalists discovered the magazine had published articles under AI-generated author names with fabricated biographies.⁷⁴ Around the same time, then-CNET Connie Guglielmo acknowledged that 77 machine-generated stories appeared on its website, several requiring corrections.⁷⁵ The scandal damaged the publication’s credibility and cost several executives their jobs.

Smaller organizations have also stumbled. In November 2025, reporters at Suncoast Searchlight, a Florida-based nonprofit news organization, sent a letter to their board of directors after discovering their editor-in-chief’s undisclosed use of AI editing tools – which included instances where the AI inserted hallucinated quotes into reporters’ drafts.⁷⁶ Some staff members were unaware that AI was being deployed on their work at all, raising important questions about transparency, editorial control, and ethical boundaries.

These uses of AI illuminate the emerging complications around authorship, accountability, and transparency – questions that grow more vexing as AI becomes so integrated into workflows that some people no longer consider it worth mentioning. While the need for disclosure is often underscored in debates about AI-generated news, the thresholds become murkier as AI permeates

routine support activities and sometimes even directly modifies journalists' work without their knowledge or consent. The line between tool and author therefore continues to blur, challenging foundational assumptions about who is responsible when storytelling goes wrong.

AI and Distribution

Media distribution fundamentally shapes who can access journalism and under what conditions.⁷⁷ For decades, news organizations had considerable direct control over their distribution pathways. However, the rise of digital intermediaries – search engines, social media platforms, news aggregators – has dramatically altered this landscape. Today, algorithmic systems increasingly determine which stories reach which audiences and provide the marketplace for monetizing them, forcing news organizations to adapt their editorial practices to accommodate the demands of algorithmic distribution. At the same time, AI empowers news organizations to rapidly remix their own content for multiple platforms and to expand accessibility.

Optimized Distribution

Media publishers have traditionally owned large segments of their distribution channels.⁷⁸ They often owned the printing presses and the delivery trucks. They owned the cameras and microphones to record the news, as well as the antennas to beam their stories across the country. With the advent of the internet, newsrooms also had significant control over their own websites. As social media platforms started to gain popularity in the 2000s, though, media publishers gradually became more dependent on platform distribution by tech companies like Google and Facebook, including their social media feeds, search engine results, and news aggregators.⁷⁹

Thus, news organizations have had to manage their own proprietary distribution channels while optimizing their content for algorithmic visibility. Some newsrooms even required reporters to build personal social media followings, effectively transforming individual journalists into distribution channels.⁸⁰ Meanwhile, audiences increasingly encountered news through algorithmic

feeds, where platform logic, not editorial judgment, dictated the final presentation.

Several news organizations have adapted by integrating AI into their distribution through two main strategies.⁸¹ Algorithmic optimization uses AI to decode platforms' shifting preferences, advising journalists on ideal formats, optimal publishing times, and strategic keywords to maximize visibility and reach across platforms. Targeted delivery harnesses AI for precision audience segmentation, enabling organizations to serve personalized content to specific demographic or interest groups.

BBC News exemplifies this approach. In March 2025, the broadcaster announced a new department dedicated to AI-powered personalization.⁸² As then-CEO Deborah Turness wrote to staff: "We must become ruthlessly focused on understanding our audience needs, on delivering the kind of journalism and content they want, in the places they want it, designed and produced in the shape that they enjoy it." The path forward, she argued, requires deploying "AI to support, enable and accelerate our innovation and growth."

It isn't just major organizations doing this, either. The Rural News Network, a collective of more than 500 local news non-profits across the United States, created Text RURAL to solve a distribution challenge.⁸³ With member organizations producing nearly 5,000 stories monthly, the AI-powered tool automatically selects, aggregates, and summarizes the most relevant content for individual users based on their location and interests. It then delivers personalized weekly news roundups via text message. It's a practical solution that makes local journalism more accessible to rural communities underserved by traditional distribution channels. However, it is an automated solution that bears risk, as even major publishers have discovered in their experiments with AI summaries.⁸⁴

Automating Content Adaptation

News organizations are also using GenAI for automated adaptation, or the rapid generation of derivative products tailored to different platforms and accessibility needs. AI systems can automatically produce social media posts, video summaries, audio versions, and translations from content originally created by human journalists. While these derivatives are typically reviewed

before publication, they nevertheless reduce the time, cost, and labor required to meet audiences across multiple platforms while expanding accessibility.

Dow Jones Newswires launched a custom AI language service in 2024 targeting this exact opportunity. The system produces what it describes as “fluent” translations of hundreds of English-language stories daily into Japanese, Korean, and French. “It gives us an ability to very quickly, almost instantaneously, translate our very rich existing English-language content into a language that will attract other audiences,” observes Chip Cummins, *The Wall Street Journal’s* chief Newswires editor.⁸⁵ The service reaches professionals “who prefer to read their business news in their native language or who are suddenly tapping into the US market as a place where they want to invest, or their clients are investing.”

Similarly, the US television station WCPO 9 uses AI to automatically reformat content for different platforms. “AI helps us take horizontal video storytelling you might see on our broadcast or YouTube and convert it to vertical on Instagram or TikTok,” notes senior manager PJ O’Keefe, “so that, again, we can make as big of an impact as possible, on as many platforms as possible.”⁸⁶

Dow Jones’ approach includes transparency measures: each translated article carries a warning about AI translation and links to the original story. However, as WBEZ reporter Araceli Gómez-Aldana observes, current models still struggle with the “accuracy of the translation for news stories, cultural understanding, translating quotes for tone and understanding – not only a literal word-by-word translation – and adapting to various dialects and literacy levels.”⁸⁷ In other words, there are still risks associated with this kind of functionality.

Despite these powerful capabilities, what AI cannot do is liberate news organizations from platform dependencies. No matter how sophisticated their AI tools become, newsrooms cannot alter the ethical, economic, or political decisions made by platform owners regarding content moderation, algorithmic demotion of news, or sudden rule changes. Platforms can unilaterally reshape their algorithms – as Meta has done repeatedly in prioritizing and then de-prioritizing news content on Facebook – leaving news organizations scrambling to adapt their carefully crafted AI-optimized strategies to new, often opaque, rules. This structural

power imbalance constrains AI's potential to transform news distribution in ways that truly benefit journalism – as we discuss in Chapter 3.

New Codes of Ethics for Journalism

The multifaceted integration of AI into newsrooms has exposed important gaps in journalism's ethical foundations and forced professionals to update ethical frameworks that were designed for a previous era. The result has been a patchwork of revisions to existing institutional codes of ethics and new organizational guidelines.

Consider that journalists can now treat AI-generated content as source material, yet LLMs bear little resemblance to human sources. They cannot be vetted through traditional means, questioned about their reasoning, or held accountable for errors. This becomes especially troubling given AI's tendency to hallucinate.

This raises thorny questions about responsibility.⁸⁸ When an algorithm recommends an inappropriate story, generates misleading content, or amplifies harmful material, who answers for it? The engineer who built the system? The news organization that deployed it? The journalist who relied on its output? Traditional accountability frameworks were not designed with this intersectional paradigm in mind, leaving both journalists and audiences navigating uncertain ethical territory.

Transparency, Harm, and Privacy

Journalism's commitment to transparency collides directly with how the most widespread AI systems operate. Commercial algorithms function as black boxes, revealing little about their training data, decision-making processes, or limitations. How can journalists exercise transparency when they themselves cannot fully explain the tools they use?

This opacity becomes even more problematic when considering bias. AI systems readily absorb and amplify biases embedded in their training data, yet journalists often lack the technical capacity to audit these systems for prejudice. The ethical mandate to serve all communities fairly grows more complex when the tools themselves may be systematically unfair. Moreover,

important ethical considerations emerge when we know AI systems are trained on abusive and non-consensual images, virulent hate speech, and violent extremism.⁸⁹ Such horrors can therefore always bleed into the open – or flood out, given sufficient automation – which is not only an ethical failure but also a significant reputational risk for any news organization.

AI also forces journalism to expand its conception of harm beyond traditional concerns like misidentification or invasion of privacy. Algorithmic data analysis can aggregate publicly available information to expose individuals to new risks – from reidentification in anonymous datasets to the construction of detailed behavioral profiles that reveal intimate details about people’s lives.

Then, there is the risk of inadvertent exposure. When journalists input confidential information into proprietary AI tools, they may unknowingly compromise sources if those data become part of the system’s training corpus. This requires journalists not only to develop better technical literacy but also to make decisions about how trustworthy AI companies’ claims are, since they can’t verify those claims themselves. While enterprise contracts can guard against this, individual use of commercial platforms may not.

An Illusion of Control

Keeping a “human in the loop” has become journalism’s refrain when defending AI adoption, but this principle proves more complicated in practice. When algorithms suggest story ideas, recommend sources, draft sentences, and identify target audiences, they function as invisible collaborators. Agentic systems, discussed in Chapter 6, only add to that complication.

Scale intensifies the challenge. Can spot-checking AI output or deploying rapid-response teams to correct errors constitute sufficient oversight when systems generate content at industrial speed? And how should journalists disclose AI involvement to audiences? Little consensus exists on what level of AI assistance merits disclosure, how such disclosure should be presented, or whether different standards should apply to reporting versus editing versus distribution.

And when news organizations rely on AI tools provided by technology companies, they enter relationships that may com-

promise journalistic autonomy. As we discuss in Chapter 3, platform dependencies operate less visibly and are more fraught than the traditional advertiser influence that journalism has developed strategies to manage. AI companies are poised to exert pressure on even more of journalism's pain points, thereby taking even more control away from journalists and news organizations.

New Principles for a New Age

Recognizing these gaps, the Society of Professional Journalists (SPJ) announced in November 2025 its intention to develop an updated code of ethics – only the sixth revision in more than a century for America's most influential professional ethics code.⁹⁰ “There's a critical need for more guidance on the ethics of using AI to do journalism,” noted SPJ Ethics Committee Chair Dan Axelrod. This planned revision acknowledges that journalism's ethical cornerstones require substantial adaptation for the algorithmic present.

While the SPJ deliberates, individual news organizations have forged ahead with their own AI-specific guidelines. The Associated Press, an early AI adopter, published guidelines in 2023 requiring staff to treat AI-generated material as unvetted source content, protect confidential information, and maintain vigilance against AI-generated misinformation.⁹¹ *The New York Times* followed in 2024 with guidelines that encourage AI use when it advances the organizational mission while requiring human editorial review of all public output and mandating disclosure of significant AI involvement.⁹² Sweden's *Aftonbladet* updated its guidelines in 2025, explicitly encouraging journalists to use AI and disclose its consequential applications while restricting them from entering sensitive information into proprietary platforms like ChatGPT.⁹³

A recent study identified common principles across a number of codes of ethics developed by individual news organizations: human oversight, explainability of AI systems, disclosure of automated content, and protection of user data.⁹⁴ However, the study also revealed significant limitations. Many guidelines lacked important details for how to address specific ethical challenges and rarely included concrete examples. Moreover, these frameworks emerged overwhelmingly from North America and Europe, with Global South media being far less likely to offer

AI-specific guidance. This raises concerns about whether emerging ethical standards adequately represent diverse journalistic contexts and values.⁹⁵

But even thoughtfully crafted ethical guidelines face a stubborn reality: many newsrooms operate in fast-paced environments where journalists are asked to do more with less, often relying on technologies and platforms they have limited control over or insight into. Whether these new codes of ethics will function as effective self-governance tools or merely serve as aspirational signposts remains an open question.⁹⁶

Building and Breaking Knowledge Silos

As AI tools proliferate across newsrooms, another paradox is emerging: a technology that promises to democratize information production may instead widen existing knowledge gaps. Organizational scholars call these “knowledge silos”: structures that emerge when information and expertise become isolated within particular teams or individuals, rendering that knowledge inaccessible to others who might benefit from it.⁹⁷

How Knowledge Silos Form

Knowledge silos can be established deliberately or emerge unintentionally. Journalists who learn to use AI tools gain a deeper understanding than colleagues who rarely use them. This differential accumulation of knowledge isn’t inherently problematic; specialization has long been part of journalism. The trouble begins when these knowledge disparities calcify into rigid boundaries that produce compartmentalization, diminished collaboration, and organizational fragmentation.

Knowledge silos can take multiple forms within organizations.⁹⁸ Individual silos occur when certain employees possess valuable expertise but fail to share it. Departmental silos arise when teams hoard information and fail to communicate effectively with one another. Technological silos develop when information is trapped in incompatible systems, databases, or file formats. Cultural silos emerge when organizational norms and structures actively discourage collaboration and knowledge sharing. Often, these forms overlap and reinforce one another.

Once established, knowledge silos become increasingly difficult to break.⁹⁹ Individuals with extensive AI expertise may feel they have less to gain from exchanges with less technologically fluent colleagues, particularly when such conversations require explaining unfamiliar concepts. Those with advanced capabilities may perceive that sharing knowledge demands more effort than the collaboration yields in return.

To be sure, many journalists take genuine pleasure in sharing discoveries and mentoring colleagues.¹⁰⁰ But contemporary journalism operates under severe resource constraints, where, as noted, staff are consistently asked to do more with less. In such contexts, the time and energy required to bridge knowledge gaps can feel like luxuries that overwhelmed professionals cannot afford. Moreover, journalism's precarity has already eroded the mentoring and professional development culture that relied on news workers' ability to spend large portions of their career in a stable newsroom. This creates a self-reinforcing cycle: solidifying silos widen knowledge gaps, which makes future collaboration feel even more costly and further entrenches the silos.

Incentive structures can also compound the problem. Knowledge gaps mean that workers with AI capabilities are perceived as more valuable, and such perceptions carry material weight: salary increases, promotion opportunities, and – perhaps most importantly in an era of widespread layoffs – enhanced job security. Positioning oneself as possessing irreplaceable AI skills becomes a rational self-preservation strategy. Conversely, journalists who decline to develop basic AI literacy risk being perceived as resistant to innovation, potentially leading to material disadvantages in competitive job markets.

Breaking Down Barriers

As scholars examining digital transformation in media organizations have noted, the fundamental issue isn't the technology itself, but the unequal ability to use it effectively.¹⁰¹ A recent report on AI adoption in Southeast Asian journalism puts it succinctly: "AI is not the problem; the concern should be how AI is being used. Those who can effectively utilize AI will perform better than those who lack AI knowledge, potentially widening the gap in the future."¹⁰²

When only select individuals or teams can leverage AI tools effectively, organizations cannot realize the technology's broader potential. Innovation and transformation thus remain trapped within silos. Addressing this challenge requires rethinking organizational structures and processes to actively promote cross-team collaboration and knowledge sharing.

First, newsrooms can support protected experimentation time that allows employees to explore new AI systems without production deadline pressures.¹⁰³ This hands-on exploration enables organic learning, creates opportunities for informal dialogue about what is and isn't possible with various tools, and allows knowledge to circulate more freely. When experimentation is treated as legitimate work rather than a luxury, it signals organizational commitment to collective learning.

Second, formal tool-testing processes serve both technical and social functions.¹⁰⁴ Systematically evaluating AI technologies before integrating them into standard workflows ensures they operate ethically and efficiently. More importantly, this process creates structured opportunities for cross-functional collaboration, building shared understanding while increasing confidence that thoughtful decisions are being made about technological implementation.

Third, in-house training initiatives offer a directed approach to leveling the playing field.¹⁰⁵ Newsrooms around the world have developed programs ranging from basic literacy workshops to advanced technical courses. These initiatives close capability gaps by providing baseline competencies that facilitate future collaboration and discussion. When everyone shares a common vocabulary and fundamental understanding, the barrier to entry for meaningful conversation drops significantly.

Finally, newsrooms can design for collaboration by revisiting both organizational processes and physical spaces.¹⁰⁶ This might involve reorienting office layouts to increase proximity between those who are comfortable with new technologies and those who aren't. The goal isn't conversion – forcing reluctant journalists to adopt unsuitable technologies – but conversation by creating opportunities for ongoing exchange about what AI can and cannot accomplish, where it adds value, and how to blend the respective advantages offered by humans and machines.

Breaking down knowledge silos is essential for ensuring that AI is adopted, adapted, and even refused in ways that advance

both individual interests and collective missions. When knowledge flows freely, organizations can make better decisions about which technologies to embrace, how to implement them ethically, and where to best direct human effort. More broadly, such organizational and procedural reorganization may be necessary if journalists are to step off their hamster wheels and fundamentally reimagine their roles, orientations, and practices.

Finding New Meaning: From Augmentation to Reimagination

As AI becomes more tightly woven into the fabric of journalism, it forces an existential reckoning: What does it mean to be a journalist in the age of artificial intelligence?

In recent decades, it has meant being on a relentless hamster wheel of content production where reporters and editors are expected to do more with dwindling resources. The drudgery of manufacturing cheap stories, the exhaustion of constant output, the frustration of seldom having time for deeper investigation – all this and more have led some to flee the profession and others to unhappily persist in the hope that things might improve. While the arrival of GenAI no doubt threatens to accelerate the wheel, it may also offer an escape route.

The question is urgent because AI can now demonstrably produce quality content and make sound decisions – not always, but often enough to matter. This capability forces news professionals to abandon the outdated idea of the human journalist alone exerting creativity, shaping how the public understands reality, and conducting professional journalism. Simply put, the traditional boundaries delineating journalistic expertise are being redrawn by technologies capable of performing tasks once thought to require distinctly human capacities.

To be sure, AI's limitations are equally significant, and some appear insurmountable in the foreseeable future. GenAI still hallucinates with alarming regularity and creates reputational risks. Much newsworthy activity occurs outside digital spaces or in forms that are hard to digitize. And the most impactful journalism involves highly complex situations demanding nuanced awareness and great sensitivity to produce responsible reporting.

Beyond the technical considerations, overinvestment in AI – and, by some, *any* investment – risks engulfing organizations in backlashes against AI by those concerned with its many costs, including environmental destruction and labor exploitation. Indeed, as news audiences seek genuine human connections, they might see news organizations that use AI as untrustworthy. Thus, not everything about journalism is poised to change, and journalists won't simply be replaced wholesale by AI technologies.

So, what does this mean for working journalists, editors, and other members of the newsroom?

Augmenting Journalistic Work

At minimum, journalists will need to engage with AI, whether that involves low-level tasks or high-level cognitive work. If AI can handle routine production tasks – summarizing documents, generating social media posts, transcribing interviews, producing derivative content for different platforms – then journalists will be realizing what David Caswell has described as augmentation: using the technology to enhance existing journalistic practices without fundamentally changing them.¹⁰⁷ This is likewise satisfied when AI tools are used not to do journalists' work for them but instead to help them exercise creativity, diversify sourcing, and fact-check information more efficiently. A tool that summarizes documents or suggests story angles is ostensibly just making journalists faster at tasks they were already performing.

But augmentation does not necessarily mean freedom from the hamster wheel. History suggests caution, after all. Previous technological advances promised liberation but often delivered intensification, with journalists simply expected to produce more and faster. Whether AI-driven augmentation leads to genuine relief or merely raises productivity expectations depends on human choices, not technological capabilities alone. (We explore these institutional dynamics in the next two chapters.)

Enabling New Capabilities

Journalists aren't just *adopting* AI; they are also *adapting* it: developing new skills and practices that extend their professional capabilities and extending the technology to unforeseen uses. This connects to Caswell's second category: innovation. Here, AI

creates new capabilities that fit within existing journalistic paradigms while expanding what's possible.

For example, journalists with no programming background can now use GenAI to “vibe code” and mock up interactive web projects. Or, they can use it to help them scrape existing information and reshape it into new datasets. They can even prototype digital products they could never have created independently. The use-cases aren't fundamentally changing what journalism is; journalists are still telling stories, but simply finding new ways to examine documents or engage audiences. However, they're doing so with an expanded toolkit that crosses traditional skill and professional boundaries.

These adaptive approaches go beyond individual uses and reflect broader restructuring of workflows that can lead to true innovation. Journalists increasingly work alongside AI systems in iterative processes, such as by using algorithms to identify patterns in datasets, then investigating the most promising leads; employing AI to generate multiple headline options, then selecting the most appropriate ones; and having systems flag potential factual inconsistencies, then verifying through traditional reporting methods. In these hybrid workflows, journalists are learning to work symbiotically with AI rather than just seeking efficiency gains.

Our hamster wheel metaphor shifts here. Rather than running on the wheel or stepping off it entirely, journalists have the opportunity to reconfigure the apparatus itself – changing the relationship between human expertise and machine capability to create space for work that feels more meaningful than mechanical.

Not all AI integration deserves adoption or adaptation, though. Strategic refusal is an appropriate response when integration undermines professional and ethical values, erodes journalistic quality, or appears likely to harm sources or audiences. This necessitates discussions about the boundaries that need to be drawn around what should (and shouldn't) be delegated to machines, recognizing that some journalistic functions need to remain human because of both technical limitations and social complications. Moreover, if the technology is being used to exploit people or simply make humans more efficient cogs in a machine – to extract more content from fewer workers – then refusal becomes a logical response not only for professional self-preservation but also for the defense of quality journalism.

Reimagining Journalism

And there is the possibility of something more fundamental and exciting: transformation. This represents Caswell's third category, where AI doesn't merely enhance or expand journalism but fundamentally reimagines its essential nature and forms.¹⁰⁸ It's also an opportunity to take this moment to forge something more than just "AI journalism."

Glimmers of this transformation are already emerging. Conversational AI agents enable large-scale dialogic journalism that was simply impossible before. This isn't mere novelty. It may well become a significant part of the future of storytelling, where audiences engage not with the *products* journalists create but with the *knowledge* journalists have generated. Rather than consuming a fixed article or broadcast, they might speak with AI agents trained on an organization's or collective's reporting, be empowered to ask follow-up questions and explore tangential threads, or request information tailored to their specific context and needs. This represents a fundamental shift in the relationship between the journalist and their audience away from broadcasting and ever more toward conversation.

Similarly, pairing increasingly sophisticated content generation with increasingly detailed audience data profiles will likely yield remarkably personalized news environments. While concerns about personalization and filter bubbles aren't new, this represents a leap toward realizing the paradigmatic shift from journalists manufacturing news for mass audiences to working alongside synthetic agents to produce news for individuals. A breaking event might be reported not as a single story but as numerous variations, each calibrated to the recipient's location, interests, literacy level, and prior knowledge. This represents a transformative change in journalism's fundamental economic and creative model – the culmination of a trend previously identified by scholars but only now made fully possible.¹⁰⁹

The broader implication is that being a journalist in the age of AI may mean less emphasis on content creation and more on gathering the ingredients necessary for quality journalism. This need not be seen as demotion or subservience to algorithms – after all, broadcasters have long employed journalists whose job was to monitor sites like the courthouse and call in breaking news information so it could be rewritten and presented on-air

by an anchor or another journalist. Rather, it means reorienting professional identity toward distinctive human advantages.

This includes understanding which powers need to be held accountable and which voices need amplification; focusing on unearthing information that would otherwise remain hidden without a journalist's diligence, tenacity, and source cultivation; employing critical thinking to identify when errors arise or when things aren't as they appear; knowing how to treat information and individuals with empathy and care; and knowing how to ultimately tailor journalistic content to foster the emotional connections that increase interest and impact. These are distinctly human capabilities requiring embodied experience, ethical reasoning, social intelligence, and contextual wisdom. AI does not make the journalist redundant – it makes the humanity of journalism paramount.

This reorientation may ultimately lead to journalism's reinvigoration rather than its obsolescence. In an information environment increasingly saturated with AI-generated content, journalism's value proposition shifts from mere information provision to verification, contextualization, and sense-making. The journalist's role becomes less about being the first to report and more about being the most trustworthy; less about reactivity and more about wisdom; and less about volume and more about meaningful impact.

Choosing the Future

We believe future journalists will remain motivated by their work's profound societal value. We hope they will see AI not as a replacement but as a tool they feel empowered to adopt where beneficial, adapt when it can extend capabilities, and refuse where it undermines their professional and ethical values – all in service of their role as auditors and protectors of truth. The matter of empowerment is key since the most consequential decisions about the use of AI are often being made by news executives and managers – not the journalists themselves.

Whether AI's integration into journalism ultimately results primarily in augmentation, innovation, or transformation will thus depend not on the technology itself but on the choices that journalists and news organizations make about implementation. It will depend on their willingness to revisit and reevaluate

journalism's purpose, practices, and workflows, and their ability to do so in a way that meets the challenges and conditions of a new age. The hamster wheel will remain unless journalists and newsroom leaders collectively insist on using AI not to run faster, but to run differently – and toward work that is meaningful.

3

The New Architecture of Media Power

“The extinction-level event is already here,” warned Helen Havlak, the publisher of *The Verge* – a tech news site and Vox Media’s most-visited brand – in July 2025. “And a bunch of small publishers have already gone out of business.”¹

The catalyst for this crisis was Google’s AI Overviews feature, which has quietly but fundamentally rewritten the relationship between search and journalism. For years, a core element of the digital news economy operated according to a familiar ritual: users seeking to learn more about some news they had heard about would enter some relevant terms into Google’s search box, encounter a curated selection of headlines – typically from major publishers through Google News, supplemented by additional search results – and, crucially, click through to read the full articles. This click-through moment was the economic lifeline connecting audiences to publishers in the modern news environment, enabling news organizations to monetize their work by either showing advertisements to transient users or, ideally, securing new subscribers.

For years, news organizations had adapted their production logic to this ecosystem, crafting headlines and structuring content to match the keywords users typed into search bars – a practice known as search engine optimization, or SEO.² But AI Overviews severed this relationship at its most vulnerable point. By synthesizing answers directly within Google’s interface – and

often drawing from publishers' original reporting – the feature satisfies users' information needs without requiring them to leave the platform.³ While news sources may be credited in footnotes, the economic transaction is largely eliminated. Users get their answers. Google retains their attention and advertising revenue. Publishers receive nothing but attribution, a poor currency for covering the operational costs of a newsroom.

For Havlak, the problem had already hit home: "The Verge's Google traffic has been declining," she told NPR. "I would say a lot of that decline has lined up pretty clearly with the rise of AI Overviews."

Havlak's diagnosis was backed by data from Similarweb, a web analytics company.⁴ Between April 2022 and April 2025, *The Wall Street Journal* saw its organic search traffic share decline by nearly 5 percentage points. *The New York Times* experienced a steeper drop of 7.5 percentage points over the same period. Most dramatically, *Business Insider* experienced a 55 percent collapse in organic search traffic – a shortfall so severe it led to 21 percent of its staff being cut. The company's CEO, Barbara Peng, framed the layoffs as essential for the publication to "endure extreme traffic drops outside of our control," a telling acknowledgment of its diminished agency. Another study similarly found that Google traffic from organic search to over 2,500 sites declined by 33 percent globally between November 2024 and November 2025.⁵

Within media circles, this phenomenon already has a name: "Google Zero," a future in which search engines no longer function as gateways to external content but as self-contained "answer engines."⁶ Google is not the sole architect of this transformation, though. As conversational AI agents proliferate across digital infrastructure – embedded in web browsers like Dia and Perplexity Comet, woven into operating systems from Apple and Microsoft, and integrated into social platforms like Reddit and X – users increasingly complete their information journeys without ever leaving their point of origin.⁷ The platform is no longer an intermediary; it is the genesis and the terminus.

This reconfiguration poses an existential challenge to publishers, who find themselves surrendering ever more control over their products and infrastructure, and ultimately their destinies.⁸ They are increasingly dependent on a constellation of external actors – most operating entirely outside journalism's institutional boundaries – to sustain core functions across the full range of the

journalistic activities described in Chapter 2.⁹ Neil Vogel, CEO of Dotdash Meredith (which oversees prominent brands including *People* and *Southern Living*), offered a stark assessment: “AI was not the thing that was changing everything, but it will be going forward. It’s the last straw.”¹⁰

We argue that the struggle unfolding among publishers, platforms, and AI companies is not merely a skirmish over search traffic or advertising revenue. It is much bigger and more consequential: a battle over the architecture of media power – over who gets to set the terms by which news is produced, distributed, and consumed in an age where the intermediaries between information and audiences are becoming increasingly algorithmic, conversational, and invisible. It is a contest over infrastructure as much as content, over visibility as much as voice, and over the very conditions that make public communication possible.

Media power is no longer a question only of who owns what but also of who controls the conditions for media work, content distribution, and audience engagement. For most of the twentieth century, media power debates revolved around ownership structures: who controlled newspapers, radio stations, and television channels.¹¹ Media concentration was measured in terms of market share, cross-ownership rules to avoid monopolies or duopolies, and the number and type of voices available to audiences. Countering media concentration was a fundamental democratic ideal. As media critic Ben Bagdikian put it, “The safest way to ensure diversity of opinion is diverse ownership.”¹²

In the age of industrialized artificial intelligence, however – where the most powerful AI systems are typically offered as a service by a few providers – power is less about owning news outlets and dictating editorial lines than about controlling the infrastructures and relational conditions that mediate access, distribution, and engagement.¹³ The disruption of traditional power arrangements wrought by the digital revolution in news has produced a shift from ownership-based concentration to infrastructure-based concentration. This requires us not only to rethink core concepts in journalism, such as autonomy and professional authority, but also to ask what becomes of journalism when the infrastructure it depends upon is controlled by others.

Power, Journalism, and Media Concentration

Sociologist Max Weber defines power as “the probability that one actor within a social relationship will be in a position to carry out his own will despite resistance.”¹⁴ Power is not just about possessing resources but also about relationships – specifically, one’s capacity to overcome opposition within those relationships. Power, according to Weber, can manifest both through coercion and through consensual acceptance of authority.¹⁵

News organizations exercise what scholars call “opinion power”: the capacity to influence which issues warrant public attention, how those issues should be understood, and which voices get amplified in the public sphere.¹⁶ This power largely rests on journalists’ socially sanctioned role as credible narrators of current events and arbiters of factual truth.¹⁷ When the public perceives this authority to be legitimate, journalism is able to wield considerable power to shape public discourse, hold institutions accountable, and influence policy agendas.¹⁸

Journalism’s power, though, is neither permanent nor self-sustaining. It requires continuous maintenance through practices that reinforce public trust, such as diligent fact-checking and verification, demonstrations of independence and transparency, and commitments to serving the public rather than partisan or commercial interests.¹⁹ While journalism’s authority rests largely on its historical claim to professional expertise over news narration – think of Walter Cronkite’s mythical status in the United States – that authority is today being actively challenged by declining confidence in institutions and by news organizations’ struggle to remain relevant in a rapidly changing media landscape.²⁰

The Structural Foundations of Journalistic Power

Journalism’s power has historically extended beyond consensual authority to encompass forms of structural control. Throughout the twentieth century, news organizations exerted coercive power through their control over access to mass audiences and media production infrastructure.²¹ Reaching a mass public required passage through media organizations that possessed either a large printing press and a sizable distribution network or a broadcast license. By vertically integrating their operations – controlling

everything from newsgathering to delivery – news organizations secured remarkable autonomy over the news production process and, in many cases, created oligopolistic markets that served their interests.²²

Recognizing this structural positioning, media scholar Timothy E. Cook described journalism as a political institution in its own right.²³ Journalism, in his view, did not just check power; it was itself a site of power. In the mass media era, journalists had greater ability to set the agenda for public discourse and, by extension, the nature of public opinion. As such, journalism's opinion power was inextricably linked to its control over infrastructure.

This capacity to exercise power explains why media scholars devoted decades to examining ownership concentration and the power of proprietors, be they wealthy individuals or corporate conglomerates. Bagdikian's documentation of accelerating media consolidation in the United States revealed how a handful of corporations were gaining control over the majority of news outlets, raising concerns about the diversity of voices and perspectives available to the public.²⁴ Social critics Edward Herman and Noam Chomsky even argued that structural factors, including ownership concentration, led news organizations to filter news content in ways that served elite interests while marginalizing dissent.²⁵

These concerns reflected an understanding of the inverse relationship between media concentration and pluralism: as ownership becomes more concentrated, the range of perspectives in public discourse tends to narrow.²⁶ Media diversity – encompassing diversity of ownership, content, exposure, and sources – became understood as essential for democratic functioning. For this reason, regulators in liberal democracies have at different times broken up monopolies, prevented mergers, and implemented rules intended to promote competition and fairness in news media.

News organizations, and journalism more broadly, do not just exercise power, though. They are themselves subject to multiple influences that shape their values, norms, and practices. For example, media scholars Pamela Shoemaker and Stephen Reese's Hierarchy of Influences Model offers a useful framework for understanding how those forces operate at five levels that move from micro to macro: individual journalists' characteristics and backgrounds; professional media routines; organizational structures and cultures; social institutions (e.g., advertising and

regulatory bodies); and, finally, broader ideological and social system forces.²⁷

Sociologist Pierre Bourdieu's field theory offers complementary insight, conceptualizing journalism as operating within a "field" characterized by specific forms of capital, positions, and struggles for dominance.²⁸ Journalistic autonomy, from this perspective, is not a fixed state but rather an ongoing negotiation against economic, political, and social pressures. These two frameworks underscore that journalism, news organizations, and even individual journalists exist within relational networks that both enable and constrain their power.

The Centrality of Autonomy

Within this context, the concept of autonomy – particularly editorial independence from commercial and political pressures – has served as a cardinal value in professional journalism in many countries.²⁹ Autonomy became institutionalized through formal and informal structures. For example, many newsrooms erected metaphorical (and sometimes even physical) "walls" between editorial and business operations to insulate journalistic decisions from publishers' interests, advertisers, and external political actors.³⁰ Professional codes of ethics were likewise institutionalized within the profession to reinforce norms of independence and the pursuit of truth.

This institutional architecture reflected a recognition that journalism's authority, legitimacy, and opinion power rested fundamentally on its ability to fulfill its claim as a credible narrator of current events and arbiter of truth.³¹ That claim, in turn, required autonomy. Journalism's structural power – its control over infrastructure and audience access – created the conditions enabling such autonomy. News organizations could afford to maintain editorial independence because their position within informational networks provided economic security and negotiating leverage. Overall, many legacy news organizations in the mass media era enjoyed a privileged position of cultural influence combined with financial success that afforded them (and their journalists) a protective buffer against intrusive forces.

What happens, then, when that foundation shifts and new actors begin capturing the essential infrastructure for provisioning and distributing news?

Platforms and the Loss of Autonomy

The digital transformation of journalism fundamentally reorganized its power relations.³² The rise of tech giants and platform companies – search engines, social media networks, and content aggregators – has upended the vertical integration that once characterized the news industry, replacing it with a platform-dependent ecosystem.³³ In this new configuration, news organizations occupy an increasingly subordinate position as their autonomy is constrained by dependencies on technical systems and market structures they neither own nor control.

Erosion and Economic Unraveling

The loss of journalistic autonomy is evident across news production, news distribution, and news consumption. In terms of production, news organizations have been compelled to adapt their content to satisfy platforms' technical requirements and algorithmic preferences.³⁴ Search engine optimization exemplifies this dynamic, as news organizations restructure their content to align with ranking algorithms or include commonly used search keywords in their headlines – thereby ceding some of their editorial authority to fit algorithmic logics.³⁵ The industry-wide “pivot to video” of the mid-2010s likewise represents a notable example of platform-induced strategic shifts.³⁶ News organizations hired videography teams and redirected resources to multimedia production to chase Facebook metrics, which they later learned had been systematically overstated. Overall, substantial investments of increasingly scarce resources were being made based on incomplete and self-serving information provided by those platforms.

The distribution challenge cuts even deeper. News aggregators, search engines, and social media platforms have become the primary gateways through which audiences encounter journalism.³⁷ This dependency has created a structural vulnerability: news organizations must maintain visibility on these platforms to reach audiences, yet they possess no control over the algorithms that determine which content surfaces and to whom. As a result, they must continually adapt to unpredictable and unilateral platform changes, as well as opaque algorithms, to maintain their

relevance in an attention economy characterized by increased fragmentation and competition.³⁸

In the realm of news consumption, digital affordances have also made it dramatically easier for audiences to excerpt, remix, and recontextualize news content in ways that diverge sharply from journalists' intended meanings.³⁹ Quote tweeting, for instance, allowed users to frame news stories through their own commentary, effectively decontextualizing the journalist's work. Moreover, news organizations now rely primarily on platform-provided analytics to understand their audiences – information essential for both editorial strategy and advertiser relations – even though they cannot independently audit the accuracy of those metrics.⁴⁰ In effect, platforms have not only transformed how journalists understand and relate to their audiences but also limited their control over how such audiences interact with journalistic content.⁴¹

This loss of autonomy has severely undermined news organizations' economic sustainability and their ability to be self-sufficient in ways that extend beyond the oft-cited migration of advertising from print to digital. The more fundamental issue lies in who controls the infrastructure of digital advertising.⁴² Google and Meta now capture most digital advertising spending, not just because they offer ad space but also because they control the data, targeting capabilities, and programmatic systems that constitute contemporary advertising infrastructure. News organizations increasingly function as content suppliers within this advertising ecosystem designed and governed by platform companies that can extract substantial shares of the ad revenue.

This creates a destructive logic. To maintain visibility in an intensely competitive attention economy, news organizations must make increasing amounts of content freely available on platforms, thereby undermining their ability to convert audience attention into sustainable revenue through subscriptions, paywalls, and other direct monetization tactics.⁴³ The economics of platform dependency thus incentivizes a vicious cycle in which news organizations sacrifice economic sustainability and self-sufficiency for visibility, while platforms capture the value generated by journalistic labor.

Several organizations, especially at the local level, have rejected this logic, sacrificing visibility to retain their ability to

monetize content among a core audience willing to pay for access to information. However, that strategy is subject to its own risks. Consequently, as media scholar Rasmus Kleis Nielsen argues: “Many news media organizations are not disentangled from platforms, and they will not prioritize disentangling themselves from platforms, because while they might in principle want to, in practice there are other things they want *more*” – namely, the opportunity to reach more people and generate revenue.⁴⁴

Conceptualizing Infrastructural Power

How might we conceptualize this reconfiguration? Platform companies’ growing influence can be understood through the lens of infrastructural power – power exercised not through direct editorial control, as media owners wielded in earlier eras, but through the provision and governance of technical systems, algorithmic architectures, and market structures on which journalism now depends.⁴⁵ This infrastructural position grants platforms considerable power to shape journalism without direct involvement in news production because they effectively control, to a greater degree than ever, the digital spaces in which journalists work, the tools available to them, and the mechanisms through which they relate to audiences.⁴⁶

This represents another profound inversion of twentieth-century power arrangements. While media power has never been evenly distributed, large news organizations were previously the primary beneficiaries of structural asymmetry: they largely controlled the means of production and distribution, and profited handsomely from that position – in terms of both financial success and cultural authority.⁴⁷ The platformization of media infrastructure has reversed these relations. News organizations now need platforms more than platforms need any individual news organization.⁴⁸ While journalism provides value to platforms – namely, free, high-quality content that generates engagement – that content is largely exchangeable (if not expendable) from the platforms’ perspective, as any individual news organization can presumably be easily replaced.⁴⁹ When Australia and Canada tried to force platform companies to share revenue with news organizations, some platforms responded by simply removing news from their products or by pursuing uneven agreements based on organizations’ relative bargaining power.⁵⁰

To be sure, there have been historical parallels to this moment. In the United States, Western Union's near monopoly over telegraphs allowed it to subsidize preferred wire services and box out others by refusing them access to the company's infrastructure at competitive rates.⁵¹ However, what is different in the present moment is scope. A company like Google accounts for the majority of search traffic, holds a dominant position in programmatic advertising, provides the world's most popular web browser and mobile operating system, and hosts the most commonly used email, video, and web analytics services. This gives it power over numerous pressure points throughout news production, distribution, and consumption, making disentanglement risky if not impossible.⁵²

And this isn't just happening in the United States, either. In Chile, for example, the media company Copesa sued Google for "abusing its near-monopoly power by controlling advertising and selecting news content in search results."⁵³ South Africa's Competition Commission similarly recommended Google contribute millions of dollars to a journalism sustainability fund annually to compensate for its negative impacts on the country's news media.

These new dependencies largely constrain journalistic autonomy, as news organizations must accept platforms' terms of service, algorithmic logics, and business models or risk invisibility.⁵⁴ Moreover, these new dependencies compound: they build upon and reinforce previous ones, creating path dependencies that make it harder for news organizations to extract themselves from platform ecosystems or reassert autonomy over their core functions. The result is a media landscape in which the power to determine how journalism functions has substantially shifted from journalists and news organizations to external actors who control the digital infrastructure on which contemporary journalism unavoidably depends.

The Great Houses of AI

The age of AI exacerbates a complication news organizations were already struggling to deal with amid the rise of platforms: What happens to journalistic autonomy when essential news production tools become owned and operated by the same tech

companies that now govern news distribution and the audiences' experiences with news? Or when those tools require scale and capacities that news organizations are simply not in a position to provide?

These questions lie at the heart of what media scholar Fernando van der Vlist and his colleagues call the "industrialisation of artificial intelligence," or the centralization of the resources and capabilities needed to build, train, and deploy large-scale AI models within organizations, including today's newsrooms.⁵⁵ Whereas the platform era enabled tech companies like Google, Facebook/Meta, and Twitter/X to become pivotal gatekeepers by virtue of mediating access between news organizations and their audiences, the age of AI threatens to *directly* restructure much of the work that most people associate with "journalism." At stake is the cognitive work associated with journalism, such as how key activities like sourcing and fact-checking are performed, and who gets to shape that work.

The Rise of Infrastructural Capture

The infrastructure required for today's powerful AI agents represents a qualitative break from previous technological dependencies.⁵⁶ Whereas social media platforms operate primarily through their network effects and programmatic advertising platforms through the user profiles they have captured, large-scale AI systems demand three resources at unprecedented scale and cost: vast datasets, immense computational power, and highly specialized technical expertise. Training a single frontier model can cost hundreds of millions of dollars and require energy consumption equivalent to a small city. As a result, news organizations – many of which already operate on skeletal budgets – face an insurmountable infrastructure gap.

This gap has given rise to what technology researcher Kate Crawford terms the "Great Houses of AI," a concentrated oligopoly of firms – the likes of Amazon, Apple, Google (Alphabet), Meta, and Microsoft – that leverage their existing dominance in cloud computing, processing power, and hardware to control the advancement of AI and its practical integrations within organizations.⁵⁷ Their power comes not just from their economic might but also from their infrastructural control. As Van der Vlist and colleagues observe, "[T]he dominance of these companies

is intrinsically linked to their control over infrastructure, which is crucial for AI development and deployment, raising concerns regarding access and competition.”⁵⁸ Legal scholars Anna Gerbrandy and Pauline Phoa further contend that big tech’s corporate power is so extraordinary that many countries – including the United States – currently lack the legal vocabulary needed to deal with its potential abuses outside the economic domain.⁵⁹

Microsoft, for example, secured its AI foothold through massive investment in OpenAI while providing essential infrastructure through its Azure cloud computing platform.⁶⁰ Google did so by custom-building AI hardware and leveraging the data it has accumulated through dominance over multiple digital services, including email, search, and mobile computing.⁶¹ Amazon’s AWS cloud services likewise provide the computational backbone for a vast portion of the digital economy, including many AI startups that have effectively become tenants on its industrial-scale platform.⁶²

This is not mere market concentration. It represents a form of infrastructural capture, wherein a handful of actors not only dominate market share but also accumulate sufficient power to deeply structure the means of production, set technical standards, and reshape ethical boundaries for entire fields.⁶³ While the release of technologies like DeepSeek, which dramatically reduced the cost and complexity of implementing AI systems, has helped democratize access to AI – including making it more economically viable for smaller news organizations – most commercial offerings remain tied to the Great Houses.

Journalism’s Deepening Dependencies

Unlike the earlier platform era, where dependencies centered on distribution and monetization, the industrialization of AI is making news organizations dependent on the Great Houses for the very tools they need for content production and cognitive labor. It is not about adapting journalism to new platforms but about using these platforms to make journalism.

Media organizations, most of which lack the resources to develop their own AI tools – never mind their own foundational models – are forced to license AI systems from the very firms that already mediate their access to audiences and advertising revenue.⁶⁴ Faced with these constraints, a typical newsroom

might adopt Microsoft's Copilot tools for summarizing internal documents and drafting stories, Google's models for optimizing headlines and generating multimedia content, and Amazon's AWS services to host it all. This creates new, compounding path dependencies.

This entanglement fundamentally alters the nature of editorial autonomy. The "wall" separating the editorial and business sides of newsrooms had already been critiqued for becoming more of a "curtain" amid the proliferation of audience analytics.⁶⁵ Now it faces another challenge: the opaque, algorithmic logics embedded in the AI infrastructures that news organizations increasingly depend on. Those logics embody values – such as maximizing engagement, avoiding controversy, and promoting proprietary services – that sharply diverge from core journalistic principles of public service, accuracy, transparency, and accountability. The concern deepens with the vendor lock-in that tightly integrated systems create: switching providers can become prohibitively expensive, leaving news organizations continually susceptible to changes in pricing, terms of service, and infrastructural logics.

This infrastructural capture revives longstanding concerns about media pluralism, but with a crucial difference. When Herman and Chomsky articulated their propaganda model in 1988, they focused on the concentration of media ownership.⁶⁶ Today's threat is more subtle but potentially more pervasive: the concentration of ownership over the underlying logic of media production itself.

If most news organizations begin relying on the same few providers for core newsgathering, analysis, and writing activities, we collectively risk a dangerous homogenization not of ownership but of editorial perspective and content. Source diversity becomes less meaningful when organizations deploy similar tools with analogous embedded assumptions to direct the production of news – especially as news organizations face pressures to shift away from high-quality, human-driven investigations to maximizing efficiency by leveraging the AI technologies they have licensed to produce content at scale.

Google's deployment of AI Overviews, described at the beginning of this chapter, further illustrates news organizations' uphill battle in reclaiming their power – and the autonomy that comes with it. By exercising its power over the infrastructure of

distribution through its market-leading search engine, Google is able to drive users to its own generative AI (GenAI) technology. That technology effectively becomes news infrastructure – extracting content from news organizations while giving users enough of an answer that they never click through to the source.

Thus, for many users, the news product becomes not the journalist's article but Google's AI-generated synthesis. Journalists become *de facto* suppliers of the ingredients for Google's product, which now reflects the production logic Google chooses to implement. The infrastructural control, in this case, extends beyond the technologies that journalists use for reporting to include the technologies producing the news that audiences ultimately encounter. This displacement threatens news organizations' authority as the primary narrators of events, with cascading consequences for their legitimacy and the cultural and economic power they can exercise.

The industrialization of AI represents an escalation of, not a departure from, the platform dependencies that have progressively constrained journalism's autonomy and economic independence.⁶⁷ The Great Houses of AI are rapidly building a deeper layer of infrastructural control – one that stretches beyond distribution and monetization to include the very processes of journalistic production and its logics.

Journalists' Reckoning with AI Giants

At the time of writing, journalism's institutional response to this challenge has largely been disorganized. Some organizations have attempted to forge alliances with AI companies; others have taken them to court. A few have leapt headfirst into adoption; others have refused even to engage with AI integration. The profession finds itself caught in a tug-of-war between its own values and the fear of obsolescence, between the preservation of editorial independence and the seductive promise of technological salvation.

This disarray is not merely organizational confusion; it reflects a deeper structural predicament. Journalists and news organizations are finding that their range of possible responses – whether cooperative or adversarial – is constrained by power asymmetries built over years. The infrastructure now controlled by the Great Houses of AI and other tech companies is forcing many news

organizations to make uncomfortable compromises and strategic decisions that are often based more on hope than on a solid foundation.⁶⁸

The Licensing Gambit

The most visible manifestation of this reckoning has been the rush toward licensing agreements between media organizations and AI companies. Confronted with the realization that AI models have already eaten up their archives without their permission, many legacy media publishers have opted for pragmatism over principle, signing multi-million-dollar deals to license their content to AI developers.⁶⁹ As media scholar Felix Simon observes: “Disruptive technologies and their developers have a habit of testing and pushing the boundaries of copyright. And so it is in the case of AI.”⁷⁰

A wave of such agreements followed the launch of ChatGPT in 2022.⁷¹ *Le Monde*, Prisa Media, and *The Washington Post* all partnered with OpenAI. CNN and Fox News partnered with Meta. *Der Spiegel*, Gannett, and the *Los Angeles Times* partnered with Perplexity. And others, like Axel Springer, *TIME*, and The Associated Press, have partnered with multiple AI companies.

News and AI executives framed these deals as a way to unlock new revenue streams while ensuring accurate attribution. “As AI technology becomes increasingly integrated into the information ecosystem, we are committed to ensuring that our content is properly attributed and that we are fairly compensated,” declared Gannett Chairman and CEO Michael Reed after the news conglomerate signed an agreement with conversational search engine company Perplexity AI in July 2025.⁷² He continued: “This deal allows us to further accelerate AI opportunities as we share advertising revenue and leverage data to deliver shareholder value while providing credible content for users of the Perplexity platform.”

But scholars have been quick to point out that these agreements function less as genuine partnerships than as retroactive legitimization of data extraction that has already occurred. As communication scholar Nicholas Diakopoulos puts it: “In order to guard against how the courts may decide, maybe [AI companies] want to go out and sign licensing deals so you’re guaranteed legal access to the material you’ll need.”⁷³

The logic behind these deals appears to be survival through preemptive concession – accepting some payment for what has already been taken. As technologist Chris Yoko explains: “The raw material that AI systems are trained on – the news reports, interviews, and investigations – are all the product of human journalists. If these jobs disappear, the flow of new, verified information would dry up, and AI would eventually be left synthesizing and regurgitating old information.”⁷⁴ For newsrooms facing audience migration to chat interfaces, licensing seems not only pragmatic but perhaps even idealistic: If readers will encounter journalism through AI intermediaries like chatbots and search engine overviews, shouldn’t news organizations ensure their content shapes those outputs?

While this may seem like a win-win for everyone involved, at least in the short term, the reasoning both obscures and contributes to the ongoing shift in power relations between journalists and AI companies.⁷⁵ By accepting payment for their content, news organizations are effectively conceding that their role in the information ecosystem is transforming from being the primary destination for news to being providers of raw materials for algorithmic synthesis. The deals themselves, often opaque and shrouded in confidentiality clauses, risk commodifying journalism as AI-ready data.

More critically, these deals do nothing to alter the underlying infrastructural dependency. While news organizations gain short-term capital infusions – welcome lifelines, no doubt – they remain entirely reliant on AI companies for access to those audiences. These licensing deals, then, are more than just business transactions between equals. They formalize the asymmetric relationships.

The Adversarial Alternative

Not all news organizations have chosen accommodation, though. *The New York Times*’ lawsuit against OpenAI, Microsoft, and Perplexity represents one of the most significant adversarial responses to date.⁷⁶ The *Times* has sought to use copyright law to reassert control over its journalistic labor and products, arguing that the unauthorized use of its content constitutes not only intellectual property infringement but also unfair competition. In effect, the *Times* is drawing a line regarding the value

of human intellectual labor in a time of pervasive automated extraction.

This legal battle goes beyond lost revenue, too. It represents the struggle over defining “transformative use” in the age of AI. If courts determine that training AI on high-quality journalism constitutes fair use, the industry’s dominant business model – already precarious after the collapse of advertising revenue – will face an existential threat. Moreover, the litigation highlights a profound institutional anxiety: the fear that the conditions for media work will become usurped by tech companies with no inherent commitment to journalism’s democratic function.

This adversarial approach carries its own limitations, though. First, it assumes that the copyright framework, refined largely around twentieth-century questions, is prepared to handle the realities posed by GenAI. Even if the *Times* (and other major publishers, like Brazil’s *Folha* and Japan’s *Yomiuri Shimbun*) should prevail in their lawsuits,⁷⁷ the victory may prove pyrrhic, as the conditions for distribution and use have perhaps irreversibly shifted toward generative delivery. A legal victory may force AI companies to pay for data, but it will not reverse changing audience habits or restore journalism’s infrastructural control – leaving news organizations in a weak position to negotiate with the Great Houses of AI.

Second, this strategy is available mostly to well-resourced news organizations. Smaller outlets, which lack the *Times*’ scale (and legal capacity), generally must watch from the sidelines as elite institutions negotiate the rules that will govern the entire industry’s future. This adversarial response, while perhaps necessary to establish boundaries, does not fundamentally redistribute power back to news organizations; it merely contests its terms.

The AI Babysitters

Away from courtrooms and licensing negotiations, a different reckoning is unfolding within the daily workflows of newsrooms: the industrialization of AI threatens to further accelerate the hamster wheel of news production.

The promise, articulated by AI companies and eagerly adopted by cost-conscious news publishers, is that AI will “free up” journalists for high-value work by handling routine tasks like document summarization, headline writing, search engine optimization,

and transforming content to fit different platforms.⁷⁸ But journalism's history offers a sobering precedent: efficiency gains are rarely reinvested in high-quality journalism, such as the laborious investigations seen as the profession's crown jewels. Such work also includes computational reporting that leverages AI's massive scale – the kind of work that has recently been recognized with Pulitzer Prizes.⁷⁹ Instead, those efficiency gains typically enable increased output with reduced staff. The hamster wheel spins faster, and journalists must find ways to utilize new technologies just to remain in place.

When newsrooms integrate tools like Google's Gemini or Microsoft's Copilot into their daily workflows, they are not just adopting software. They are importing an epistemological framework – a way of thinking and knowing – that prioritizes certain values over others, such as plausibility over accuracy, volume over depth, and speed over consideration. This shift fundamentally alters the journalist's role and the broader nature of journalistic autonomy. As AI engages in more of the creative work of journalism – ideation, synthesis, storytelling – journalists are forced into the role of being AI babysitters, with their expertise reduced to quality control: checking the machine's output for errors or aberrations while keeping up with the accelerating hamster wheel.

The cruel irony is that this intensification is occurring precisely when journalism can least afford it. Many journalists report an increasing sense of precarity and anxiety over their work's rapidly changing nature, leading nearly half of those surveyed in one 2025 study of journalists in the United States to say they had considered leaving the profession altogether over the past year.⁸⁰ Local news organizations in many countries are struggling to survive, and budgets for several public service broadcasters are being cut, if not eliminated.⁸¹ Enrollments in undergraduate communications programs have sagged in some countries as students worry about low starting salaries and hostile public opinion.⁸² For many, the path off the hamster wheel becomes self-selection out of the profession altogether, while those who stay (or enter) are forced to run faster and adjust to new dependencies, each individually justified as necessary for the industry's survival.

Exporting the Hamster Wheel to the Global South

The industrialization of AI is a global phenomenon, and the emergence of new technological dependencies cannot be understood separately from broader geopolitical dynamics, particularly those relating to what communication scholars Nick Couldry and Ulises A. Mejias call “data colonialism.”⁸³

The Great Houses of AI driving this infrastructural shift are housed almost exclusively in the Global North. When news organizations in Africa, Latin America, and Southeast Asia adopt these tools, they are importing systems trained predominantly on English-language data and Western cultural norms.⁸⁴ This technological adoption creates new forms of epistemic dependency and labor exploitation. Economic value and data extracted from the Global South flow northward to those technology centers, while the North’s values, assumptions, and ways of understanding the world are simultaneously imposed on the South. The infrastructural capture reshaping media architecture thus also produces global structural inequalities, with consequences for how information is processed and knowledge is constructed.⁸⁵

This dynamic becomes visible in the proliferation of AI training programs and funding initiatives launched by tech giants like Google and Meta in Global South countries. While these programs are framed as philanthropic efforts to modernize developing newsrooms – and they no doubt can help those newsrooms with basic infrastructure – they also lock news organizations around the world into proprietary ecosystems, creating dependencies that are beneficial in the short term but difficult to escape from later.⁸⁶ As a result, the journalist in Lagos or Lima faces a difficult proposition: refuse the tools and risk obsolescence relative to competitors who embrace them, or adopt them and surrender even more professional autonomy.

The responses from the Global South reflect this tension. While some journalists are opting to use free or subsidized tools to bridge resource gaps, there is also a growing call among scholars and professionals for “sovereign AI,” such as locally trained foundational models that respect indigenous languages and cultural contexts.⁸⁷ In Chile, for example, engineers and researchers

began building Latam-GPT, a collaborative large language model (LLM) intended to better reflect local languages and cultures than the models imported from the Global North.⁸⁸ However, these initiatives face the same structural constraints as their media counterparts in the United States: they also require vast datasets, immense computational power, and highly specialized technical expertise. Most news organizations around the world thus remain dependent on AI infrastructure they cannot audit, control, or meaningfully contest.

The hamster wheel thus appears likely to be exported to the Global South. Even if it spins at different speeds and under different conditions, the governing logic will likely be influenced by the Great Houses of AI and the commercial pressures they will no doubt exert.

Regulatory Interventions in the United States, China, and the European Union

Despite widespread recognition that the emergence of the Great Houses of AI posed distinctive challenges to news media and the previous architecture of media power, translating this awareness into effective policy has proven difficult. Traditional regulatory frameworks designed around broadcast licenses and cross-media ownership restrictions are inadequate for addressing the seismic shifts in power dynamics over the past two decades.⁸⁹

The problem, after all, is no longer just about media concentration in the conventional sense. It is also about who controls the computational resources necessary to develop and deploy AI, the data pipelines feeding machine learning systems, the algorithmic distribution mechanisms that determine content visibility, and the technical standards that shape news production and consumption.⁹⁰ As communication technology scholar Christoph Busch contends, “[C]ompetition and media law need to be complemented by a new type of ‘platform infrastructure law’ that covers digital platforms as societal infrastructures and increasingly important actors.”⁹¹

The United States, China, and the European Union offer illustrations of distinct legislative and regulatory responses to confront these new challenges.⁹² However, although those governments approach the challenge from divergent traditions and

employ different tools, scholars have identified significant shortcomings in their interventions – particularly when it comes to protecting journalistic autonomy, mitigating infrastructural capture, and meaningfully addressing the emerging power asymmetries affecting the media sector.

The United States: Fragmentation and Uncertainty

The United States has adopted a notably fragmented approach to regulating AI that largely sidesteps questions of infrastructural power. At the time of writing, it has no comprehensive federal legislation governing the development and use of AI, and certainly not in relation to news media.⁹³ The Biden administration's executive order on "Safe, Secure, and Trustworthy Artificial Intelligence" and its earlier "Blueprint for an AI Bill of Rights" policy document arguably represented the most significant federal interventions to date, directing agencies to develop safety standards, conduct impact assessments, and provide guidance on issues such as algorithmic discrimination and data privacy.

However, neither of those interventions engaged with media-specific concerns about infrastructural control. They did not mandate compensation when journalistic content is used to train LLMs, require source attribution in AI-generated summaries, institute interoperability standards, or implement safeguards for media plurality. Moreover, they lacked durability since executive orders and policy blueprints can be revoked when administrations change. Indeed, the succeeding Trump administration dismantled them shortly after taking office in 2025, opting instead for a hands-off approach to AI governance and explicitly pointing to competition with China as justification for "removing barriers to American leadership in Artificial Intelligence."⁹⁴

Congressional efforts have similarly failed to produce meaningful protections for journalistic autonomy. The Content Origin Protection and Integrity from Edited and Deepfaked Media Act would grant content creators rights to control the use of their content and likeness by AI, require platforms to label AI-generated content, and compel preservation of content provenance information. However, it has stalled legislatively since being introduced in July 2024. The few AI bills that have been signed into law, such as the TAKE IT DOWN Act prohibiting nonconsensual online publication of intimate visual depictions,

do not primarily engage with news media or address the structural power asymmetries impacting the news industry.

Individual states have stepped into the regulatory void, creating geographic disparities. By the end of 2025, nearly all states had introduced at least one measure to regulate AI – such as by limiting the collection of certain personal information, protecting individuals’ image and likeness, requiring greater transparency from companies, and imposing liability for some harms. But this patchwork system has been criticized by policy analysts as failing to address the more fundamental questions of infrastructural control. Furthermore, the Trump administration has actively opposed state-level AI regulation and even threatened to direct the Department of Justice to sue states that regulate in this domain.⁹⁵

In short, the United States’ light-touch regulatory approach has provided minimal protection for journalists and news organizations, failed to address emerging power asymmetries, and created regulatory uncertainty. More importantly, it has failed to prevent the capture of the infrastructure on which contemporary journalism increasingly depends.

China: State-Centered Control and Development

In contrast to American fragmentation, China has employed comprehensive regulatory frameworks that signal that AI innovation is encouraged but must occur within state-defined boundaries.⁹⁶ The question is not whether to regulate AI but how to integrate it within (or make it serve) existing structures of state control.

China’s most comprehensive frameworks – the “Interim Measures for the Management of Generative Artificial Intelligence Services,” enacted in August 2023, and the “Measures for Labeling AI-Generated Content,” enacted in September 2025 – require AI-generated content to comply with “core socialist values”; penalize subversion of state power; mandate that training data be accurate and diverse to prevent the generation of false information; compel labeling of AI-generated content; establish mechanisms to prevent the generation of illegal content; and order adherence to intellectual property and personal information protections.⁹⁷ These measures followed earlier regulations establishing the supervisory responsibility of algorithm service providers and for combating deepfakes.⁹⁸ Collectively,

these frameworks reflect China's dual objectives of becoming the world leader in AI by 2030 while maintaining state control over information flows.⁹⁹

However, analysts have argued that the Chinese approach to regulation includes some intentional ambiguity and has cumulatively promoted AI development at the expense of its news industry.¹⁰⁰ Communication scholars Jingyan Yuan and Lin Zhang characterize this as an “infrastructuralization” of platforms, wherein commercial platforms are co-opted into essential state infrastructure for information control.¹⁰¹ Additionally, while regulatory efforts to impose transparency requirements and interoperability provisions might at first glance appear designed to reduce platform dominance, they have paradoxically created compliance barriers that favor large, well-resourced firms like Alibaba, Baidu, and Tencent – China's own Great Houses of AI and the entities most capable of implementing extensive content filtering systems that satisfy state requirements.¹⁰² Infrastructural power thus concentrates around state-approved providers.

To partially counterbalance this concentration, news aggregators and tech companies in China have been compelled to establish formal partnerships with news organizations and share the revenue derived from news content.¹⁰³ However, these partnerships operate within a broader context where journalistic autonomy is itself constrained by direct state oversight that pressures news organizations to promote pro-government narratives and uncritical AI discourse.¹⁰⁴ This dynamic further deepens the power imbalances that weaken news organizations' bargaining positions.

In short, China has been proactive in regulating AI, but arguably in ways that perpetuate market concentration, deepen platform dependencies around state-approved AI systems, and restrict news organizations' ability to use AI in ways that might challenge official narratives. These changes add to the restrictive ways in which China has long controlled the press.

The European Union: Comprehensive Yet Insufficient

The European Union has pursued the most ambitious regulation agenda, establishing multiple comprehensive frameworks through three landmark pieces of legislation: the AI Act, the European Media Freedom Act (EMFA), and the Digital Services

Act (DSA). These represent genuine attempts to govern AI's societal implications while protecting media pluralism and journalistic autonomy.

The AI Act classifies non-exempt AI applications by their risk of causing harm. Applications with unacceptable risks (e.g., social scoring by authorities) are banned. High-risk applications (e.g., automated exam scoring) must comply with security, transparency, and quality obligations. Low-risk applications (e.g., chatbots) have only transparency obligations, and minimal-risk systems (e.g., recommendation engines) are generally unregulated. The EMFA aims more narrowly to protect media pluralism and editorial independence, including provisions that protect journalistic sources against spyware and safeguard against political and economic interference. The DSA regulates platforms' content moderation and the algorithmic systems that affect news distribution.

Although the EU's regulatory frameworks are less ambiguous than China's and more comprehensive than the United States', scholars have identified notable limitations in their capacity to address the fundamental issues arising from the new architecture of media power. The frameworks' focus on direct harms to individuals – discrimination, manipulation, privacy violations, and the like – fails to address systemic power imbalances between large AI companies and news organizations. For example, most AI tools that scrape and repackage journalistic content fall into lower-risk categories with minimal obligations, and the frameworks lack provisions addressing bargaining power differentials when media outlets negotiate content licensing deals with tech giants.¹⁰⁵

Media policy scholar Theresa Seipp argues that the EU frameworks exhibit “noticeable blind spots” in tackling platform power and that “neither the infrastructural power of platforms . . . nor their relational power . . . are effectively addressed in the current EU regulatory frameworks, despite both forms of power driving digital media concentration.”¹⁰⁶ By focusing on content moderation, transparency requirements, and risk assessments, these policies fail to address more problematic structural dependencies and, in turn, “neglect the growing dependence in journalism production and the essential infrastructures for journalism, particularly related to AI.”¹⁰⁷

Paradoxically, the EU's frameworks may inadvertently accelerate journalism's dependence on external technology providers.

Media policy scholar Natali Helberger observes that under the AI Act, “deployers that *do* more and modify or adapt the system to their own organization-specific context and requirements risk being considered providers in their own right and having to comply with the extensive legal obligations that apply to providers.” The additional legal risks and substantial compliance costs can become “one additional factor in the decision not to invest in in-house technology and instead procure from external technology providers.”¹⁰⁸ In attempting to ensure AI safety through comprehensive regulation, the EU may be unintentionally discouraging the very efforts to establish infrastructural autonomy that might counterbalance the growing power asymmetries between AI companies and news organizations.

In short, under the EU’s regulatory frameworks, as Seipp notes, “journalism is inadequately protected as a public good and is instead treated more like a consumer good, overlooking journalism’s fundamental public value.”¹⁰⁹ It remains unclear whether a transparency- and risk-based approach can adequately counterbalance the structural power asymmetries arising from the emergence of the Great Houses of AI – especially when the dominant platforms and digital infrastructure operators are headquartered outside the EU.

Structural Solutions for Structural Problems

The limitations of these regulatory approaches – whether fragmented US non-regulation, state control-focused Chinese frameworks, or comprehensive but insufficient EU legislation – reveal a broader failure to address AI’s structural transformation of journalism. They also do little to address journalism’s hamster wheel woes: regulations that focus on labeling AI-generated content, preventing discriminatory algorithms, or ensuring transparency in content moderation hardly slow the wheel or allow journalists to step off it. They may even unintentionally accelerate it by imposing compliance burdens that make external dependence more attractive than building the internal capacity that gives news organizations the necessary autonomy to chart new paths.

If states want to confront the more fundamental reconfiguration of the news ecosystem, what would be required? Certainly, there would need to be a clearer focus on creating effective counterbalances to tech companies’ infrastructural power. One promising

approach involves directing public funds toward creating public infrastructure designed to advance democratic objectives and diffuse power across more actors, as we further discuss in Chapter 4.¹¹⁰ This is especially important given the significant barriers to entry in large-scale computing and foundational AI modeling, which limit the effectiveness of competition-oriented interventions that aim to reduce the power of dominant technology companies. There have been growing calls in Africa, Europe, and Latin America for greater infrastructural sovereignty and counterbalances to the Great Houses of AI – highlighting that infrastructural control is as much a political matter as a technical one.¹¹¹

Regulatory responses must also reconceptualize journalism as a public good, rather than as just another consumer good. Seipp contends that “to effectively address media concentration and protect pluralism, (local) journalism needs to be assigned a higher level of protection, to avoid vulnerabilities of market-driven and consumer logics,” which requires acknowledging that journalism “operates differently from typical markets, holding ‘value even if not consumed.’”¹¹² Digital society scholar José van Dijck and her colleagues further propose “expand[ing] the notion of consumer welfare to citizen wellbeing” and “shift[ing] attention from markets as level playing fields towards societal platform infrastructures where hierarchies and dependencies are built into their architecture.” Doing so recognizes that “users are also *citizens* who for their democratic and civic duties have come to depend on services offered by platform companies.”¹¹³

These are not just academic distinctions. Treating journalism as a consumer good further subjects it to market logic, where efficiency, scale, and consumer preference predominate – precisely the logic that drives infrastructural dependence. Treating journalism as a public good justifies more fundamental regulatory interventions, including public investment in alternative infrastructure, interoperability requirements that democratize infrastructure, and even limits on commercial entities’ dominance over the requisite infrastructure for modern journalism. At minimum, regulatory provisions need to either center the public good in market-oriented and individual-risk assessments or forgo such approaches altogether in favor of alternatives that better protect core democratic actors and activities – including smaller, independent, and local news outlets least able to circumvent infrastructural capture.

The challenge is therefore not simply about regulating new technology but also about ensuring that the new architecture of media power does not undermine the vital democratic functions journalism serves. Regulatory approaches that focus on media concentration or on managing AI's outputs rather than governing its structural role in the production, distribution, and consumption of news leave fundamental power disparities largely unaddressed. Only by confronting the infrastructural dimension of AI power, creating public alternatives to commercial dependencies, and reconceptualizing journalism's status can regulation effectively safeguard the pluralistic, independent media ecosystem democratic governance requires.

Journalism's Position in the New Media Architecture

The path we have traced in this chapter is a story of demotion – and, with it, pressing questions around what comes next for news. Throughout the twentieth century, large news organizations stood atop the hierarchy of media power through their control of key communication infrastructure. The platform era knocked media organizations off their pedestal, increasingly turning them into tenants dependent on digital landlords who controlled access to audiences and advertising revenue. The industrialization of AI now threatens further descent: from tenant to raw material supplier, with journalism's products scraped, synthesized, and repackaged by systems that satisfy users' information needs without requiring them to encounter the original journalism at all.

This shift has tangible consequences for journalism as an industry and the news it produces. A growing number of online users today will see a Google AI-generated summary of a news story but not the original work that was necessary to produce it, severing crucial opportunities for connection between news organizations and the audiences they seek. Journalism's grip on the reach and relevance it needs to survive therefore grows increasingly tenuous.

At the heart of this new architecture lies a paradox: journalism remains functionally indispensable even as individual news organizations become increasingly disposable. The Great Houses

of AI need verified information to train their models and make their outputs useful. Platforms need credible content to attract engagement. Users need someone to perform the hard work of reporting – bearing witness to events, gaining access to key source material, and offering narrative meaning to public life. Journalism as a set of practices – as an epistemic instrument with commitments to truth and ethics – thus remains essential in the information ecosystem, even if it is only obliquely acknowledged as such.

But under this new management, one news organization can be readily replaced by another – its content exchanged for that of a competitor, its journalists laid off without disrupting the broader system. This might explain why licensing deals can be framed as partnerships even as they formalize asymmetric relationships, why platforms can profess to value journalism while building systems designed to circumvent it, and why news organizations find themselves with so little leverage despite doing part of the work the new media architecture depends on. It also might explain why the specter of a faster-spinning hamster wheel feels so real to journalists anxious to keep up.

What does this repositioning mean for journalism's traditional claims – its autonomy, authority, and legitimacy? The answers are troubling. Autonomy has been progressively ceded, first to platforms and now to AI providers. Authority depends on direct engagement with audiences – engagement that weakens in a world of AI-crafted snapshots that substitute for news itself. And legitimacy becomes more fraught in an environment where it's harder for people to "see through" the opacity of algorithmic intermediation to recognize and appreciate the journalist working on the other side. The infrastructural capture described in this chapter is not merely an economic problem; it is a challenge to the very foundations on which journalism's democratic role depends.

The consequences of these changes in media power extend beyond journalism as a profession. The same architecture that repositions news organizations as raw material suppliers also reshapes how citizens encounter information, engage with one another, and participate in public life. If the Great Houses of AI increasingly mediate what people know, how they deliberate, and whether their voices are heard, the stakes are not only journalistic but also democratic. And the logics currently driving AI

development – capitalist imperatives in the United States, authoritarian impulses in China – are not oriented toward democratic flourishing. Whether this trajectory is inevitable, and what it would take to bend AI in favor of democracy rather than against it, is the question we take up next.

4

Designing AI for Democracy

“Investment in AI is helping to make the US Economy the ‘HOTTEST’ in the World,” President Donald J. Trump boasted in a Truth Social post on November 18, 2025.¹

There was some truth to Trump’s post. Nvidia – the chipmaker powering the AI boom – had just become the world’s first \$5 trillion company.² OpenAI was valued at \$500 billion, making it the most valuable private company globally.³ Five tech giants accounted for nearly one-third of the S&P 500’s market capitalization and were driving its 2025 gains.⁴

“[B]ut overregulation by the States is threatening to undermine this Major Growth ‘Engine,’” Trump warned in the same social media post. “Some States are even trying to embed DEI ideology into AI models, producing ‘Woke AI’ (Remember Black George Washington?). We MUST have one Federal Standard instead of a patchwork of 50 State Regulatory Regimes. If we don’t, then China will easily catch us in the AI race.”

Trump’s warning was a direct response to increasing oversight by several US states, the most recent of which was California’s “Transparency in Frontier Artificial Intelligence Act,” which requires large AI developers to disclose safety protocols, report incidents, and create whistleblower protections. It also came as his administration was cozying up to tech leaders and even considering directing the Department of Justice to sue states over their AI regulations.⁵

Trump's statement was also an explicit acknowledgment that the United States (much like China) is among the countries viewing AI as key to establishing or maintaining economic and military supremacy in the twenty-first century. This positioning highlights a significant tension: How can societies ensure that companies and researchers are free to develop AI while protecting the industries, organizations, and individuals increasingly governed by it?

The risks of the status quo were evident. Ahead of California's landmark regulation, an AI chatbot was accused of contributing to a teenager's suicide.⁶ The world's largest music group opted to partner with an AI music generator start-up after settling a lawsuit against it.⁷ A large language model (LLM) provider faced copyright infringement accusations from the largest US news media trade group.⁸ And news organizations were bracing for "Google Zero" – the point when referrals from search engines and platforms would dry up due to AI-generated overviews (discussed in Chapter 3) – raising fears that the new architecture of media power would make it even harder for news organizations to perform their democratic functions.⁹

We argue in this chapter that AI will play a role in making individuals' interactions with news, information, and one another more interactive, individualized, and centralized as the technology proliferates within and beyond journalism. This will profoundly impact how people understand the social world, deliberate over news and knowledge, and participate in democratic processes. While these impacts are not inherently good or bad, the competing capitalist and authoritarian logics at the heart of current AI development are more likely than not to challenge democratic underpinnings by rewarding designs that further reorient information networks toward isolation, control, and appeal to interest and predispositions rather than meaningful free debate or the pursuit of truth.

We are not claiming that AI will only produce bad outcomes for democracy. Some individuals will benefit from AI-enhanced news products and improved experiences – although these benefits will likely be unevenly distributed and reinforce existing informational divides. Rather, we are calling attention to a fundamental problem: communicatively capable AIs are currently shaped largely by extractive, dominating systems that are antithetical to democratic progress.¹⁰ The infrastructures being built,

the incentive structures being reinforced, and the power concentrations being enabled carry profound implications not only for news but also for democracy writ large.

These anti-democratic possibilities are not inevitabilities. It is still possible to promote democracy-oriented AI development. Indeed, this could be an opportune moment to break from the hamster wheel many democracies find themselves on, and use AI to help revitalize democratic virtues. Doing so will require moving away from the status quo, though. It will take intentionality and action to adjust incentive systems and realign AI development with democratic virtues. But while news organizations have a role to play, the challenge – and its potential solutions – lies well beyond them.

Furthermore, the relationship between journalism and democracy is not merely symbolic – though it is an important part of the professional ideology of journalism in many democracies. Journalism also benefits culturally and financially when citizens perceive it as indispensable to their democratic participation.¹¹ When news consumers believe that quality reporting can genuinely inform their civic choices and that their engagement matters, they invest time, attention, and money in journalism. Conversely, when democracy feels unresponsive – or, worse, like a farce – citizens see less value in journalism, or in the work of staying informed and engaged. In other words, when AI undermines democracy, it simultaneously erodes the financial viability and cultural capital of journalism itself.

The Public Sphere and Public Arenas

At its core, democracy simply means rule (*kratos*) by the people (*demos*). It is often operationalized as a system where people express their will through predetermined aggregation mechanisms like voting. Modern democratic nation-states, especially in Europe and the Americas, use representative democracy as their system of government. Under this system, citizens periodically elect representatives who then make decisions about specific laws and policies on their behalf, with the public holding them accountable for their governance. While democratic practice is far messier than this proceduralist ideal suggests, the basic premise offers a useful foundational anchor.

Democracy can also be examined through a communication lens. Media scholar Christian Fuchs conceptualizes democracy as an organized set of “processes of communication where humans inform themselves, debate, and take collectively binding decisions.”¹² Media philosopher Mark Coeckelbergh similarly contends that “communication is not only necessary (or must be presupposed) for democracy to work; communication *is what democracy is all about*.”¹³

These perspectives are valuable because they highlight that communication is the primary means for developing, disseminating, and collectively refining ideas at scale. Communication is central to how people understand phenomena beyond their personal experience, express their preferences, and engage with one another. It is also one of the most important vehicles for meaningful democratic participation.

Social theorist Jürgen Habermas’ notion of the *public sphere* offers a useful conceptual starting point.¹⁴ Habermas describes the public sphere as a metaphysical space where critical-rational discourse informs public opinion, which should then translate into more sound democratic decision-making. This space must be autonomous from the state and economy so citizens can engage freely on their own terms, even as it is inevitably structured by those external forces.

Free, high-quality deliberation among diverse participants is critical to the public sphere’s success. All affected parties must have access to relevant, quality information and to the debate stage itself, with equal opportunity to argue. The desired outcome is consensus – or, more practically, sufficiently broad agreement – resulting from reasoned evaluation of competing ideas.

Habermas’ conceptualization has been critiqued for being insufficiently attuned to power asymmetries, assuming a single arena rather than multiple competing ones, and presuming that participants engage in good faith. Nevertheless, as media scholar Zizi Papacharissi suggests, the public sphere remains a useful normative metaphor. It captures idealized civic interaction within a communication-oriented view of democracy.¹⁵

Communication scholars Andreas Jungherr and Ralph Schroeder address these concerns by relying instead on the concept of *public arenas*, which focuses on how media structures – including news media, digital platforms, and other discursive

institutions – enable and constrain the publication of information and opinion and the contestation that follows.¹⁶ Unlike the public sphere, public arenas do not require normative ideals like rationality or civility. While such ideals remain desirable, they often fail to describe modern mediated exchanges.

What matters most to Jungherr and Schroeder is that the structures undergirding public arenas help society become visible to itself and create spaces – both mainstream and alternative – where citizens can work toward the common good, express ideas and concerns, and build shared political identities. Visibility is crucial: actors and claims can only become salient if they are first seen. If a journalist's revelation of corruption reaches few people, citizens have diminished ability to check elite power. If some elites' arguments are more visible than their competitors', even the elites cannot check one another. Visibility is therefore central to the dynamics of public debate.

But arenas are also sites where people discover one another – through agreement or shared disagreement – and develop collective identities. Spaces must exist for people to formulate and express views and subsequently find community. Broad common spaces, like the editorial pages and comment sections of large mainstream news outlets, have historically served this function. But so too have fragmented and fringe spaces: niche internet forums and subreddits where marginalized or minority communities can gather, discuss issues, form their own identities, and develop alternative viewpoints apart from mainstream debate.

While majorities naturally enjoy greater visibility and more identity-building spaces, media structures should not tilt further in their favor. Key arenas, especially common spaces where masses congregate, must remain open to public scrutiny, and control over them must be sufficiently diffused to allow that scrutiny to yield corrective action – or, at minimum, open movement toward alternatives.

These perspectives offer structural considerations for assessing the health of common and fragmented spaces critical to public deliberation as well as normative ideals to aspire toward when thinking about communication within democracy.

Technology and the Public Arena

Many scholars – ourselves included – initially viewed the proliferation of digital, networked technologies as positive for both the public sphere’s normative ideals and public arenas more broadly.¹⁷ These technologies, many hoped, would advance more transparent deliberative models that promised to reduce prejudice, lower barriers to equal participation, and encourage informed dialogue through easier, cheaper access to information.¹⁸

Legal scholar Yochai Benkler embodied this early optimism, arguing that technological advances around the turn of the century would radically transform how information is produced and shared.¹⁹ This shift, he believed, would move us away from a public conversation dominated by traditional media owners toward a more democratic online environment where citizens have greater opportunities to participate. In recent years, however, a growing chorus of scholars – including Habermas himself – have voiced concerns about the state of the public sphere and the health of public arenas as technology becomes ever-present in daily life and increasingly shapes how we communicate.²⁰

Social media aptly illustrates the challenge. It has been credited with giving previously marginalized voices access to more prominent debate stages, playing significant connective roles during movements like the Arab Spring, #MeToo, and Black Lives Matter, and creating spaces for counter-publics to form.²¹ But it has also been charged with advancing numerous harms: making harassment ubiquitous, promoting cultures of outrage, enabling epistemic silos, and easing the spread of misinformation and disinformation.²²

Why has technology like social media failed to live up to this early optimism? Here, it is important to reiterate a point central to science and technology studies: technologies and platforms all have politics, both in their design and in their application.²³ The dominant Western social media platforms are commercial enterprises driven not by any intention to foster a healthy public sphere or advance the common good – despite occasional claims to the contrary – but by exploiting the resources they entangle: user attention, user-generated content, and the network effects that lock users in.²⁴

Consequently, the algorithms powering platforms like Facebook and X have never been neutral. They have always been designed and refined with intent – chiefly to monetize online engagement. Even when individuals within these organizations have pursued public interest goals, the overriding logic has been to commodify the elements of deliberation and commercialize the spaces where it occurs. Admittedly, some platforms like Mastodon break from this commercial logic entirely, and meaningful deliberation still emerges even on for-profit social media. But these remain exceptions to the norm.

Scholars like Fuchs argue that social media thus acts as an anti-democratic force.²⁵ This isn't because social media itself is a bad idea or lacks democratic potential. Rather, it's because the companies controlling these platforms and the way people typically use them discourage important democratic virtues – like mutual respect, rational and inclusive deliberation, truth-seeking, and accountability. When a technology undermines these virtues, it deserves to be criticized as anti-democratic.

Technologies do not develop in a vacuum, nor are they directed by any single controlling interest. It is more useful to situate technological development within broader systems of incentives – examining how design choices intersect with social, political, economic, and cultural conditions. Low trust in traditional institutions like government and news media, for example, may push designers toward user-led moderation systems, while high trust may favor algorithmic regimes that privilege information from those same institutions.²⁶

By the same token, technologies are not static. They can be steered – in both design and implementation – toward new ideals, including those at the heart of Habermas' public sphere, and toward media structures that support productive public arenas.

Logics Driving How AI Is Developed and Deployed

The United States and China are both at the forefront of AI development and deployment, with each nation vying for supremacy. However, the logics driving these efforts differ substantially, reflecting distinct political and economic arrange-

ments.²⁷ Critically, neither framework positions AI to primarily serve democratic goals.

In the United States, the most visible and pervasive AI applications are entangled with the blend of capitalist economic logic and libertarian political logic associated with Silicon Valley.²⁸ While considerable basic AI research originated in educational institutions, recent leaps in generative AI (GenAI) and public-facing applications have emerged predominantly from tech start-ups like OpenAI and established giants like Google.

Media scholar Nick Dyer-Witheford and his colleagues argue that AI has largely been developed and deployed in the United States as an instrument of capital, seeking to exploit human labor and concentrate power and wealth.²⁹ This, they contend, represents a deliberate reorganization of power toward tech companies and their shareholders – directly contributing to societal inequalities and positioning AI as an inherently anti-democratic tool.

China presents a contrasting model, where AI development is deeply entangled with the nation-state.³⁰ The Chinese government has also positioned artificial intelligence as a central pillar of national development and international supremacy, investing substantial resources in AI companies – accepting that many will fail so long as some excel – while actively embedding AI technologies throughout society.

Unlike the United States' emphasis on self-regulation, China has promoted direct state intervention in AI development. Jungherr argues that autocracies possess competitive advantages in AI development, as they can readily align centrally located resources, experts and their firms, and regulatory instruments toward unified objectives.³¹ This creates perverse incentives even in democratic environments: AI business leaders may favor solutions that tilt toward autocracy, given the efficiencies that centralized control brings to systems reliant on large-scale data and technical infrastructure – particularly when facing competitors operating within such systems.

Across these divergent contexts, a troubling pattern emerges: advancing a common good takes a backseat to other interests. While each environment nudges development in different directions through investments, regulatory efforts, and social expectations, neither actively works toward positioning AI as a fundamentally democratic force.³²

Technological design is not deterministic, though. As demonstrated in earlier chapters, while developers can initiate trajectories and structure product use in line with their objectives and environmental pressures, users exercise agency to deploy technologies in unintended ways. Additionally, the release of open-source AI technologies like DeepSeek, which serves as a viable and less computationally intensive alternative to proprietary LLMs, opens up even more user-led (and democratic) possibilities. We can therefore engage only in educated speculation about AI's future; no one can speak about it with certainty.

With this in mind, we turn now to how AI development and use might impact three key aspects of news-related communication in democracy: public understanding, deliberation, and participation.

How AI Influences Public Knowledge

Artificial intelligence is poised to fundamentally reshape how people understand news and the social world by transforming their information environments.³³ AI already plays a major role in structuring how individuals access information and how that information is presented. It also alters the conditions for creating and provisioning news, including how news professionals monetize their work. These parallel developments carry profound implications for democratic public understanding.

Individualization and Personalization

The most apparent manifestation of AI's influence is the increasing sophistication of information personalization. Personalization itself is not new; news recommendation systems have created personalized collections for years based on expressed and inferred interests. However, GenAI represents a qualitative and quantitative shift, further scaling the automated production of personalized derivative information products.

Organizations like Particle exemplify this evolution, gathering articles from multiple news organizations and using GenAI to create distinct summaries based on user preferences.³⁴ Users can request simplified explanations through an "Explain Like I'm 5" option, focus on core facts via "The 5Ws [Who? What?

Where? Why? When?],” examine diverging perspectives through “Opposite Sides,” or access direct quotes from newsmakers. If users have questions, Particle’s chatbot can provide answers drawn from additional sources.

Such personalization can benefit public understanding. When news becomes more relevant through personalized selection, more accessible through multiple formats and improved machine translation, and more diverse through deliberate introduction of distinct viewpoints, people may be more likely to engage with it, make sense of it, and become better informed. Evidence suggests AI applications that make news easier to understand are especially appealing to news avoiders, potentially drawing previously disengaged citizens back into contact with journalism.³⁵

However, these possibilities are tempered by evidence that while people are comfortable with personalization for things like music recommendations, they grow more skeptical when it comes to news. Media researcher Nic Newman and his colleagues found that news consumers’ primary concerns center on bias and manipulation, fear of missing important information, exposure to irrelevant content due to algorithmic shortcomings, and loss of privacy.³⁶ Yet those same individuals were attracted by the potential for less biased content and decision-making, as well as the novel possibilities for discovering more relevant and viewpoint-diverse news. One interpretation is that significant portions of the public maintain faith in algorithms’ technical neutrality while remaining deeply skeptical of the motives of those implementing them.

Despite audiences expressing limited desire for personalization – in their words, if perhaps less so in their actions – nearly three-fourths of newsroom leaders surveyed by Newman indicated they are actively exploring format personalization, and numerous organizations are experimenting with content personalization.³⁷

Epistemic Siloing and Disinformation

The anti-democratic implications of personalization become evident when considering how AI enables users to construct increasingly insular information environments. At the time of writing, Particle had beta-tested functionality allowing users to create highly customized news feeds based on natural language

expressions – potentially enabling requests to curate stories about specific topics, “but without any stories from liberal media” or similar filtering.

It is not far-fetched to anticipate that sizable segments of users would exercise this agency to seek psychological safety through familiarity and confirmation bias. This raises the prospect of further epistemic siloing, as algorithmic sorting and self-selection combine to reinforce one another.³⁸ Because this presents a market opportunity, even if Particle or professional news organizations decline to offer such content filtering out of ethical reservations, more market-oriented actors will likely step in to fill the gap despite the negative implications of that feature for public understanding.

Moreover, the same AI models enabling intelligent content creation also facilitate the generation of “AI slop” and sophisticated disinformation at an unprecedented scale. NewsGuard had identified over 2,000 undisclosed AI-generated news and information websites at the time of writing, some offering traditional-sounding brands like the *Dixie Sun News* and the *Carroll County Observer* – both formerly legitimate outlets that closed, with the new proprietors claiming their names.³⁹ Such websites collectively produce thousands of articles, some of which feature fabricated events or recycled information primarily designed to game programmatic advertising systems – but that nevertheless confuse real people.

GenAI image models are likewise rapidly becoming more convincing. While images generated in 2024 could often be identified through their uncanny qualities or anatomical impossibilities like extra fingers, models released just one year later had already addressed several of those shortcomings. The viral image of Pope Francis in an expensive white puffer jacket illustrated this challenge, prompting the Pope himself to declare that “[t]o navigate the complexities of AI, governments and businesses must exercise due diligence and vigilance.”⁴⁰

Distinguishing AI-generated content from human-created content is thus becoming increasingly difficult, presenting novel opportunities for manipulation. Political scientists Sarah Kreps and Douglas Kriner provide some evidence that AI-generated propaganda may be just as believable as human-written propaganda, and they contend that propaganda will only become more powerful as it is further micro-targeted through datafication and

AI capabilities.⁴¹ This view isn't unanimous, to be sure. Media scholar Felix Simon contends that such fears may be overblown and underestimate how savvy citizens actually are.⁴²

While concerns about whether publics can distinguish deep-fakes from reality might capture the headlines, the more pernicious threat may well be that anything – genuine or fabricated – can now be more easily dismissed as fake.⁴³ Though proactive individuals can turn to fact-checkers to resolve ambiguity, motivated reasoning ensures many will remain vulnerable to both new and old efforts to misinform and disinform – and disinclined to question attempts to remake reality.

Cognitive Overload and Intermediation

More broadly, the proliferation of AI-generated information products is likely to confuse or cognitively overwhelm individuals, at least initially and with negative consequences for public understanding. Faced with vast quantities of conflicting information, some citizens are already choosing to disengage altogether.

As Kreps and Kriner aptly observe: “A reasonable cognitive prophylactic measure in such a media environment would be to believe nothing, a nihilism that is at odds with vibrant democracy and corrosive to social trust.”⁴⁴ Those who do not withdraw will nevertheless likely rely even more heavily on shortcuts such as partisanship and what friends, family, and other influencers say, potentially divorcing them further from empirical reality.⁴⁵

Compounding this problem is the fact that AI appears likely to continue transforming the economic conditions for producing professional journalism, primarily by further positioning tech companies between news producers and news users through their ownership of key infrastructure and often superior technology.⁴⁶ While these economic dimensions were explored in detail in Chapter 3, it bears repeating that the prospect of original, professionalized news reporting becoming less economically viable and more consolidated poses clear risks for public understanding in democratic societies.

There is reason for some economic optimism, however. As Jungherr argues, concerns about information quality or technological mediators could drive some citizens to purposefully seek news from trusted news brands, thereby offering “a reversal of fortune compared to their economic and ideational challenges

of the last 20 years.”⁴⁷ In other words, news organizations can capitalize on concerns around synthetic media and algorithmic mediation to build new economic models that support public service journalism.

AI and Deliberation

When considering AI’s impact on democratic deliberation, it is essential to recall that the Enlightenment principle of reason is fundamentally rooted in social activity. As Coeckelbergh notes, the very notion of thinking for oneself requires subjecting one’s thoughts to scrutiny by others.⁴⁸ This demands stepping outside bubbles and spaces curated for similarity, and is most effective when there is at least some directed engagement toward those who think differently.

AI is increasingly poised to alter how people interact with one another and deliberate over news and knowledge, primarily through the more sophisticated forms of human–machine communication it enables, its regulation of the news ecosystem, and its nudging of individuals toward other individuals, groups, and information based on profiled estimates and algorithmic logics.

Chatbots as Deliberative Tools

While human–machine communication predates AI – humans have been interacting with computers as social actors for decades – GenAI has significantly expanded what agents like chatbots can accomplish.⁴⁹

When they work properly, such as by correctly understanding a user’s query and responding with relevant and accurate information, chatbots can not only produce better-informed individuals but also help them practice deliberation.⁵⁰ Chatbots can potentially help citizens assess the accuracy and validity of claims made during deliberation and even help them construct persuasive arguments that draw on widely accepted knowledge.⁵¹ They can also help citizens practice arguments through role-playing and contradiction detection, allowing them to rehearse deliberative skills in a face-saving space.

However, this assumes a level of veracity that the most common chatbots – built on LLMs developed primarily through

identification of statistical language patterns rather than genuine reasoning – do not currently possess. Chatbot “hallucinations” – instances when they unintentionally produce incorrect or fabricated information and present it as fact – can just as easily produce flawed arguments, especially when given prompts that include keywords likely to elicit false information or draw on questionable sources. The presumption of accuracy and machine neutrality can make such hallucinations all the more damaging to a citizen’s ability to deliberate with peers, as users may unknowingly incorporate false or misleading information into their own arguments and understanding.

Chatbots also have the capacity to create very different, if not truly novel, experiences for news users. As discussed in Chapter 2, they can transform previously linear and largely static products into dynamic conversations where the AI agent responds to specific information needs expressed by the user, including ones journalists may not have anticipated. This potential is on newsroom leaders’ minds; according to a recent survey, more than half said they are actively exploring AI chatbots and search interfaces.⁵² Some of these tools are already widely available to users of major news publications, including the *Financial Times*, *TIME*, and *The Washington Post*.

Chatbots created by tech companies are already engaging in that task as well.⁵³ Some are becoming embedded in web browsers and operating systems, thereby empowering agents that are even more sensitive to information a user has already been exposed to or their preferences.⁵⁴ The fact that users are far more likely to already be interacting with those tools raises the likelihood that this novel form of news engagement will ultimately be led by non-journalistic entities – such as Google and OpenAI – that have different standards for evaluating information provenance and veracity.

Weaponization and the Marketplace of Ideas

GenAI tools can also be weaponized by malicious actors to make deliberation more challenging, such as by sowing confusion among participants, diluting the salience of certain viewpoints, and harassing specific individuals. Online deliberation has long been characterized as a negative experience, with acts like “flaming” and “trolling” being fixtures for decades.⁵⁵ Social media

companies have thus been repeatedly faulted for not doing more to effectively moderate deliberative spaces.⁵⁶

The challenge is complicated by the ease with which GenAI can flood deliberative spaces with both plausible-sounding arguments intended to disinform and utter junk contributions designed to drown out certain arguments. While news organizations have previously employed human moderators or ceded responsibility to social media platforms – such as by embedding Facebook discussions into news websites – those solutions are currently struggling to keep up with GenAI. And AI doesn't just churn out more content. It can spawn larger networks of sock puppet accounts (fake identities designed to deceive) that create the impression of consensus or controversy, making fringe ideas appear mainstream or popular opinions seem divisive.

Although AI is now being actively used to perform more context-sensitive moderation and has been shown to improve the quality of large-scale moderation over previous systems, those who wish to maliciously use GenAI have a built-in advantage.⁵⁷ Moderation systems need to be careful not to exclude too much human-created communication, since that risks alienating users. They must also be careful not to filter out valid, if uncommon, arguments, since that runs against deliberative ideals. Malicious actors can capitalize on those liberties and gray areas, and have done so, resulting in more polluted information environments.⁵⁸

The broader consequence is that the already naive notion of the marketplace of ideas – the belief that the best arguments and ideas naturally win out – is becoming increasingly impractical because of the structural advantages held by the entities that have the expertise and resources to use GenAI in anti-democratic ways.⁵⁹ Currently, the lack of transparency around the use of AI in moderation makes it difficult to measure or assess the full scope of its impacts on speech.

Echo Chambers and Epistemic Bubbles

AI is also being increasingly used to recommend a range of things to users, from news content to relevant social groups. As Coeckelbergh argues, this interferes with individuals' freedom and autonomy to form their own beliefs, since those recommendation systems – now ubiquitous and baked into news apps, social platforms, and operating systems – are generally designed

to quietly nudge users, subverting their attempts to think for themselves.⁶⁰

Of course, individuals did not have full freedom and autonomy before the advent of algorithmic recommendation systems. Social influence over news habits has long existed. Moreover, there are compelling arguments that algorithmic systems can be leveraged to improve democratic deliberation.⁶¹ One approach is to design AI-infused recommendation systems that identify distinct, valid arguments and intentionally recommend opposing views to users – as Particle already offers. Alternatively, AI could generate contrasting views when users hover over specific arguments or when a summary of an argument in a specific arena is generated.

However, such designs must be intentional, and current incentive structures in places like the United States and China do not favor such intent. Within capitalist logic, systems are incentivized to cater to users' pre-existing beliefs based on profiled information to stimulate positive psychological responses and increase engagement. Within autocratic logic, systems are incentivized to marginalize arguments that counter those made by the ruling regime.

In both cases, there is an incentive to lead users toward what scholars call echo chambers (social structures in which dissenting opinions are not necessarily absent but actively undermined) or epistemic bubbles (knowledge structures that simply lack information or sources that would be relevant or important to the user).⁶² While evidence of the extent and effects of these phenomena remains mixed, there are still valid reasons for concern that at least sizable segments of the public end up in such informational enclosures, and that the ways in which AI is currently being designed and implemented amplify that risk.⁶³ Ultimately, such isolation can promote entrenchment of specific ideas – including conspiracy theories and beliefs lacking empirical support – and a broader unwillingness to engage with others, striking at the heart of democratic deliberation.

AI and Participation

Journalism has long served as democracy's connective tissue, linking citizens to pathways for political action. News organizations provide the information and context necessary

for civic engagement – from explaining candidates’ positions to identifying relevant advocacy groups to highlighting avenues for involvement. They also advance participatory goals by creating mediated spaces for mobilization, whether through publishing community opinions or offering platforms where audiences can coordinate action.⁶⁴

The Authenticity Problem

AI’s potential to enhance democratic participation seems obvious at first glance.⁶⁵ LLMs can help non-native speakers craft persuasive opinion pieces or letters to the editor. AI-powered conversational or embedded agents can instantly bring up information about the political history of an organization mentioned in a news story, enabling more informed engagement. These tools promise to democratize civic participation by making sophisticated communication accessible to everyone.

Participation is fundamentally relational, though. It requires not just the ability to speak but also the capacity to be heard and recognized. For participatory acts to matter, their targets – elected officials, candidates, government agencies – must be aware of them and interpret them accurately. Here lies a particularly insidious participatory threat: by flooding democratic channels with synthetic content, AI makes it nearly impossible for public officials to discern their constituents’ genuine priorities.⁶⁶

If ordinary citizens, journalists, and platform moderators struggle to distinguish machine-generated content from human-authored content, why should we expect politicians and their staffers to fare any better? This isn’t hypothetical. Kreps and Kriner sent randomized AI-generated and human-written emails to 7,200 state legislators and found equivalent response rates for both.⁶⁷ This evidence underscores that the threat posed by AI is not confined to fake social media posts or confused news consumers. AI can be weaponized to crowd out individuals’ genuine attempts to convey preferences within a representative democracy.

Perhaps more troubling still, the mere possibility of AI manipulation provides politicians with a convenient excuse to dismiss dissent.⁶⁸ If they don’t like what their constituents are saying, they can simply claim their expressed viewpoints are manufactured by bots. This represents yet another instance where AI does

not create problems on its own but, rather, intersects with existing anti-democratic tendencies to amplify and accelerate their effects.

Old Problems, New Scale

Using technology to simulate grassroots support is hardly new. “Astroturfing,” or manufacturing the appearance of organic public support, has plagued democratic processes for decades. What’s changed is the scale and sophistication of such efforts, made possible by rapidly evolving GenAI.

Consider the 2017 debate over network neutrality in the United States. During a four-month public comment period, the Federal Communications Commission received 21.7 million submissions. However, a study revealed that only 6 percent of them were unique.⁶⁹ Thousands of identical comments appeared at precisely the same moment, suggesting the use of rudimentary bots to manipulate a legitimate participatory process. While such simple patterns were detectable then, today’s GenAI makes similar campaigns far harder to identify and easier for bad actors to deny when exposed.

The problem compounds when we recognize, as some scholars have observed, that public opinion is being increasingly reconceived as quantified aggregations of private opinions and behaviors captured or inferred through AI-powered data mining.⁷⁰ This shift offers potential benefits, to be sure. Automated aggregation might register perspectives from people reluctant to speak in traditional forums or unable to engage with official processes, such as town halls. It could therefore yield more representative signals by incorporating marginalized voices.⁷¹

But it also poses a democratic threat by reducing public opinion to what can be easily quantified and counted. This emphasis on metrics creates perverse incentives to game such systems through sophisticated AI agents that simulate easily registered behaviors and phrases. Likewise, quantified data can be more easily instrumentalized to rationalize predetermined political ends rather than genuinely inform democratic decision-making.

Algorithmic Governance and Accountability

These participation challenges intersect with a broader transformation: the rise of “smart” societies where AI systems administer essential social services and governmental functions.⁷² While a growing body of evidence shows how these technologies can enable more responsive and effective governance, there are also well-documented cases where algorithmic systems have inflicted considerable harm on individuals and specific communities.

For example, dozens of people have been wrongly detained since police in the Brazilian city of Rio de Janeiro implemented an AI-powered facial recognition pilot in 2019 – with one report discovering that 90 percent of them were Black.⁷³ Six years later, *The Washington Post* documented similar failures across multiple US police departments.⁷⁴ These incidents reveal how algorithmic systems, deployed without adequate safeguards, may perpetuate and amplify existing inequalities.

Robust oversight of these fallible systems is therefore critical for democracy. Yet mechanisms for meaningful oversight often remain opaque to ordinary citizens or are ceded to actors with inherent conflicts of interest.⁷⁵ This gap in democratic accountability is precisely where journalism matters most.

Journalists must do more, however, than investigate algorithmic failures and biases – crucial as that work is. They must help citizens understand how these systems are governed and how people can exercise participatory agency to ensure AI advances rather than undermines democratic ideals. This responsibility grows more urgent as AI becomes embedded in the infrastructure of public life. When algorithms determine who receives welfare benefits, whose children are flagged for visits by child protective services, or who gets stopped by police, the stakes for democratic oversight become existential. Citizens cannot meaningfully participate in governing systems they don’t understand.

The technical complexity of AI systems necessitates expert involvement, certainly. But democratic legitimacy also requires the meaningful participation of the governed.⁷⁶ Journalists serve as essential translators and intermediaries in this process, making the technical comprehensible and the opaque visible. More broadly, they have a key role to play to ensure that technology serves rather than subverts the participatory foundations that democracy requires.

Toward Democracy-Oriented AI

To understand AI's relationship with democracy, we must look beyond what the technology can do and critically engage with how and why it is being built. Technical design choices, institutional incentives, and the social systems embedding these technologies determine whether AI promotes democratic or anti-democratic outcomes. Context, in short, matters as much as capability.

Coeckelbergh's argument that "AI as it is currently developed and used *undermines the fundamental principles and knowledge basis on which our democracies are built and does not contribute to the common good*" therefore rings true to us.⁷⁷ His critique cuts deeper than concerns about malicious actors weaponizing AI to manipulate public opinion or disrupt civic participation. The fundamental problem is that the forces driving AI development – primarily private corporations, investors, and autocratic leaders – are not pursuing democratic objectives.⁷⁸

Instead, AI systems are designed to secure competitive advantages, achieve operational efficiencies, and consolidate control.⁷⁹ The infrastructure powering these systems belongs to private entities that increasingly operate outside democratic accountability. The public information flowing through their proprietary algorithms is not subject to much public oversight. This results in societies becoming locked into particular tools and frameworks, ceding democratic governance over public communication to unaccountable private actors.

These harms are not inevitable consequences of disruptive technology. They result from deliberate choices to prioritize profit and power over democratic values. Such design decisions have profound implications for democracy: they determine what information becomes visible, how citizens access communicative spaces, whether these spaces serve diverse audiences, and ultimately whether they support or undermine democratic engagement.

This does not mean AI will only harm democracy. Some news organizations are genuinely attempting to implement AI in democracy-enhancing ways, and many citizens are exercising new forms of agency to produce positive outcomes. But who benefits most from the integration of AI in journalism? Who stands to lose?

Citizens already motivated to seek quality news and diverse viewpoints – those with existing strong media literacy skills – will likely benefit most. They stand to gain from both the novel news experiences and the higher-quality journalism made possible through AI-powered augmentation, innovation, and transformation. Thoughtfully designed personalization algorithms could help these engaged citizens navigate complex information landscapes while maintaining exposure to diverse perspectives. Better information access can lead to deeper understanding, more informed deliberation, and more meaningful democratic participation.

Conversely, AI's risks fall heaviest on those already less engaged with news and public affairs, especially in media systems dominated by commercial outlets with weak public service alternatives. These citizens already face mounting cognitive overload from the volume (and pollution) of news today, algorithmic nudges toward preferred content, and social pressures to remain in comfortable ideological spaces.⁸⁰ When personalization optimizes for engagement rather than public understanding, it risks further fragmenting the public sphere into isolated echo chambers and epistemic bubbles. The result is the erosion of shared knowledge, trust, and democratic involvement – undermining the very foundation upon which democratic engagement depends.⁸¹

These benefits and harms will not be randomly distributed. Better-educated, economically secure citizens living in stronger democracies will likely reap the most of AI's democratic benefits. Meanwhile, AI will likely deepen informational divides that scholars have long documented, further stratifying citizens' capacity for meaningful democratic participation.⁸² In other words, technical design choices embedded in AI systems – from content recommendation algorithms to the interfaces that mediate access to public arenas – are likely to interact with existing social inequalities to produce systematically unequal democratic outcomes.

These developments cannot be separated from anti-democratic waves currently threatening multiple democracies worldwide. AI is not *the* cause of rising authoritarianism, but it appears increasingly likely to amplify existing problems and serve anti-democratic agendas – especially if, as Jungherr argues, authoritarian contexts offer systemic advantages to AI development.⁸³ Many democratic and institutional leaders recognize these risks,

as detailed in Chapter 3. Their responses, however insufficient, represent attempts to reclaim democratic governance over technologies shaping both journalism and public life.

Reimagining Democracy with AI

Many citizens in democratic societies today feel trapped on a hamster wheel. They learn, deliberate, and participate – but find themselves no closer to meaningful change. Others feel democracy is itself locked in a hamster wheel as public leaders are rarely held to account for broken promises or failed policies, individual voices are drowned out by organized lobbying and political donations, and the very democratic processes meant to enable change are weaponized to slow progress and protect the well connected.

This widespread disillusionment is neither accidental nor inevitable. It stems from significant flaws in how contemporary democracies function – or fail to function.⁸⁴ While we embrace democracy as a normative good in this book, we need not accept Western democracy in its current form as the pinnacle of what is possible. Indeed, the status quo has evidently failed large groups of people in numerous ways, creating fertile ground for distrust and disengagement.

Here lies both a challenge and an opportunity for AI to help us step off the hamster wheel entirely and fundamentally rethink how we organize our democratic systems and hold them accountable. This requires moving beyond passive responses to AI's societal and professional threats, and actively harnessing its capabilities to address the very problems that have fostered such deep frustration.⁸⁵

Escaping the hamster wheel demands intentionality. We need both positive incentives for developing AI in democracy-oriented directions and meaningful penalties for perpetuating anti-democratic patterns. This work must happen at multiple levels simultaneously, each addressing different aspects of how AI is shaping journalism and public life.

Individual and Community Action

At the most basic level, individuals can act through their interactions with AI and their participation in democratic governance.

Hacker communities prove particularly instructive here, illustrating possibilities for repurposing anti-democratic features of technologies and directing them toward pro-democratic ends through creative tinkering.

Consider Hacks/Hackers, a nonprofit bridging journalists and technologists. In June 2025, they launched a Newsroom AI Lab aimed at building lasting technical capacity through hands-on collaboration and development of AI tools specifically designed for journalism. Part of the goal, according to Hacks/Hackers founder Burt Herman, is to “support AI developing in ways to augment human abilities in alignment with democratic ideals.” This includes, as Lab advisor Jake Kara notes, “learning to identify when AI tools are or are not the right tool for the job.”⁸⁶

Citizens can also exercise responsibility through deliberate non-participation, thereby economically punishing actors whose AI implementations undermine the common good. When *Vogue* ran a two-page spread in July 2025 featuring an AI-generated model, the backlash centered not just on displaced fashion photographers but also on how such choices reinforce unrealistic beauty standards and accelerate creative deskilling.⁸⁷ The economic and reputational punishments send broader signals that matter.

Civic responsibility also involves becoming educated about how AI structures social life. This means understanding not just individual harms but also systemic impacts on democratic engagement itself. Organizations like The Markup – founded in 2018 and operating with the aim of “challenging technology to serve the public good” – offer citizens such windows by pairing industrious journalism with engineering expertise to reveal hidden technical design choices and their consequences.⁸⁸ Such understanding is essential for citizens to meaningfully signal governance preferences to representatives and hold power to account.

Oversight by News and Civic Organizations

News and civic organizations occupy a crucial middle ground between individual action and government intervention. Several news organizations, mostly in the Global North, have created AI implementation guidelines emphasizing values like human oversight, algorithmic transparency, and data protection.⁸⁹ These

guidelines, detailed in Chapter 2, represent welcome starts, but questions remain about whether they become actual standards or constitute merely aspirational documents.

As journalism scholar Mathias Felipe de Lima Santos and his colleagues observe, “In their current form, these guidelines may primarily serve as a public statement of the company’s beliefs, rather than a practical guide for media practitioners.”⁹⁰ The challenge intensifies when organizations depend on external developers and platforms, leaving many ethical choices beyond their control.⁹¹ Well-intentioned ethics codes risk becoming symbolic gestures if the underlying incentive structures remain unchanged and journalists are stripped of their autonomy – offering, at best, another turn of the hamster wheel.

A genuine democracy-oriented approach requires intention and transparency about how organizations factor in democratic values when selecting AI tools, performing internal audits, and establishing vendor contracts. More fundamentally, it demands structural insulation from market pressures through cooperative ownership models, collective bargaining with tech platforms, and participation in shared oversight bodies that include journalists, civic groups, and technologists. These arrangements directly impact what technical design choices become possible and whether those choices serve public interests.

Newsroom employees themselves play crucial accountability roles. When *POLITICO* deployed its “Live Summaries” feature during the 2024 Democratic National Convention – which produced content riddled with factual errors, missing context, and stylebook violations – unionized journalists took the organization to arbitration.⁹² They argued that “the protection of newsroom ethical standards [is] an integral part of their labor rights.”

News organizations also shape public understanding of AI itself. While journalists’ agenda-setting power has diminished in many contexts, they remain elite tastemakers in narratives about disruptive technologies.⁹³ Democracy-oriented journalism not only challenges AI hype but also critically examines impacts on citizens’ autonomy and participation – including how personalization algorithms affect viewpoint visibility and the capacity for meaningful deliberation.

Unfortunately, research finds that journalists often fail to interrogate fundamental assumptions about AI beyond displacement concerns. Communication researcher Nirit Weiss-Blatt urges

journalists to provide context around AI prognostications, recognize nuance beyond binary extremes, avoid excessive anthropomorphization, and highlight diverse voices with evidence-based claims.⁹⁴ Critical journalism is essential for fostering the public understanding necessary for meaningful democratic deliberation about AI's role in society.

Systemic and Governmental Interventions

At the highest level, escaping the hamster wheel requires industry-wide or governmental interventions that address structural problems rather than symptoms. We need infrastructure oriented toward the common good that can compete with private offerings from Amazon, Google, OpenAI, and X. This addresses more than infrastructural sovereignty – it maintains genuinely accessible and empowering public arenas that support the ideals of the public sphere.

Such interventions might include creating data commons that serve as resources for training algorithms or sourcing information from reputable sources; public LLMs purpose-built to advance democratic objectives; and other foundational elements that can enable the creation of AI technologies that improve public understanding, deliberation, and participation. Open-source solutions that all members of society can benefit from are particularly important here, as they enable broader participation in shaping technical design choices that structure public communication. These infrastructures could ensure that design choices like personalization serve democratic ends, such as helping citizens access relevant information and diverse perspectives, rather than merely optimizing for engagement or commercial imperatives.

These endeavors should pair with regulations promoting healthy competition between private and public actors. As law professor Paul Nemitz argues, self-regulation alone is anti-democratic because it cedes society-wide decisions to unelected private individuals.⁹⁵ At the same time, we must remain vigilant of excessive lobbying and regulatory capture – themselves profoundly anti-democratic forces that keep the hamster wheel spinning.

Philosopher Daniel Innerarity contends that comprehensive strategies for steering AI development toward democratic virtues require multiple stakeholders and complementary efforts across

four areas: education and awareness; regulation and legislation; public participation and inclusivity; and transparency, explainability, and contestability.⁹⁶ These efforts should ideally coalesce into global frameworks, though this is no doubt complicated by nation-states viewing AI as key to gaining competitive advantages.

Coeckelbergh adds that democratic development of any serious intervention must involve active participation by both experts and citizens and remain forward-looking to anticipate likely harms.⁹⁷ He advocates for “democratic AI development” through democratically elected bodies of citizens and experts with both informal (e.g., social pressure) and formal (e.g., rule-making authority) power over AI-related technologies. Such participatory governance enables meaningful citizen input into both the technical design choices and the implementations of technology.

Experimenting to Step Off the Hamster Wheel

The hamster wheel of contemporary democracy – where effort produces civic exhaustion but not democratic progress – was built through countless design choices, incentive structures, and power arrangements. AI, as a powerful design tool, can very well accelerate that wheel, but it can also help dismantle it. The difference lies in intentionality, accountability, and genuine democratic participation in shaping how these systems are developed and implemented.

Defining “democracy-oriented AI” is admittedly complex and politically fraught work. Reasonable people will disagree about what constitutes the common good, what democracy should look like, and what role AI should play. Moreover, challenging the status quo is especially difficult when systemic forces work against those pursuing the common good, as we have shown here and in Chapter 3. But this difficulty cannot justify paralysis or ceding responsibility to entities not designed to advance the public interest. Instead, we need an experimental mindset – one that accepts we’ll get some things wrong before getting them right but insists we must try.

We can begin with broad principles that command sufficient agreement and then refine them through practice and reflection.

Perhaps we might find general agreement that public arenas should accommodate both mass communication reaching broad audiences and segmented communication serving particular communities, while maintaining sufficient interconnection between the two to enable collective deliberation about shared concerns. Perhaps the notion of treating the digital infrastructure supporting AI and modern platforms as a public utility rather than private property might also find sufficiently broad support.

These need not be final answers or principles that must be set in stone before we can proceed. Rather, they are productive starting points for democratic experimentation. Under a deliberative model, such principles must remain fluid, continually redefined through communication and participation as technologies and societies evolve. We also recognize that they reflect a particular vision of democracy – one that stands apart from the very different values and governance systems shaping significant AI development and investment in countries such as China and Saudi Arabia.

Although this experimental mindset offers no guarantees, it provides a path toward addressing the frustrations with the hamster wheel. And the possibilities before us remain exactly that: possibilities, not inevitabilities. If we act deliberately across multiple levels – as individuals, through professional and civic organizations, and via systemic reforms – we can collaboratively orient AI's development toward democratic ends in both journalism and public life. But discovering what democratic AI looks like demands moving beyond status quo thinking and hoping things work themselves out. It requires intentionality and an actionable commitment to using AI to improve, if not outright reimagine, democracy – to foster public understanding, enable meaningful deliberation, expand possibilities for participation, and maintain public arenas that serve the ideals of a democratic public sphere.

5

Learning to Think in the Age of AI

The question seemed straightforward enough: Had any of the students used AI?

History professor D. Graham Burnett posed the question to his class of 30 undergraduates at Princeton, spanning a dozen majors, in early 2025.¹ No hands went up. He tried his graduate students. Nothing.

To put students more at ease sharing, Burnett offered gentle prodding and encouragement: “Hey! I use these tools! They’re incredible! Let’s talk about this!” Still nothing.

The silence wasn’t deception, he later reflected – it was fear. As one student confided afterward, their syllabi in other courses had made the stakes plain: dare to touch ChatGPT or another AI tool like it, and you could be reported to the academic deans. Another student had heard a rumor that a major AI platform was blocked on the campus network, but was too afraid to even check whether it was true.

Burnett was stumped. Of all the things that should have been top of mind for students and faculty alike in 2025, he reasoned, it should be AI. More precisely, how to use AI. As a historian of science and technology who had spent years studying universities, from their medieval origins in clerical training to their modern incarnation as multidisciplinary knowledge factories, he recognized the magnitude of the moment. While universities in countries like the United States were facing a range of

threats – significant funding cuts, federal investigations, encroachment on faculty autonomy, and their forced embedding within culture wars, to name a few² – Burnett argued there was an even more fundamental disruption at play. “The juggernaut actually barrelling down the quad,” he wrote in *The New Yorker*, “is AI, coming at us with shocking speed.”³

Yet the institutional response at Princeton and elsewhere, Burnett contended, resembled something closer to collective denial infused by technological panic. One university department drafted an anti-AI policy that initially was so restrictive it would have prohibited faculty from assigning work that engaged with AI at all. When external reviewers urged Burnett’s department to confront AI’s implications for teaching and scholarship, the recommendation “got a notably cool reception.” As Burnett summed it up: “Everyone seems intent on pretending that the most significant revolution in the world of thought in the past century *isn’t happening*. The approach appears to be: ‘We’ll just tell the kids they can’t use these tools and carry on as before.’ This is, simply, madness. And it won’t hold for long.”⁴

The paralysis and pretending that Burnett describes are not unique to higher education. It exemplifies a broader societal failure to grapple with the full scope of disruption unfolding in how we create, validate, and transmit knowledge.

To be sure, universities face genuine dilemmas around how to respond to the rapid proliferation of AI. Some critics make compelling arguments that caution is warranted. AI tools no doubt facilitate academic dishonesty, compromise learning outcomes, and raise thorny questions about intellectual ownership. Such concerns deserve serious consideration.

But framing the challenge primarily as a policing problem – how to prevent academic dishonesty, detect violations, and enforce compliance – misconstrues what is at stake. The more consequential question AI poses for education is not “how do we prevent cheating?” but rather “how do we develop human capacity?” That is, the capability to act effectively, the competence to think independently, and the creativity to produce novel things.

In this chapter, we argue that the central imperative – whether in the classroom, the newsroom, or society at large – is not to prohibit AI tools or embrace them uncritically, but to cultivate the human capacity necessary for people to use them wisely. This

demands more than new policies; it requires rethinking how we prepare both journalists and the citizens they serve.

Journalism education is a particularly revealing site for this reckoning. The work of reporting sits at the intersection of craft and cognition, technical skill and ethical judgment, and individual creativity and institutional constraint. Journalists must gather original information, evaluate sources, and construct narratives – capacities AI can increasingly approximate. Meanwhile, journalism educators prepare citizens to navigate an information ecosystem increasingly rife with synthetic content both through their work in the classroom and through the news literacy materials they produce. Ultimately, the age of AI forces the question of whether education can become a site of genuine empowerment – helping people develop the literacy and judgment to work with AI rather than become diminished by it.

Journalism Education's Hamster Wheel

For the better part of a century, journalism schools have taken as their mission the preservation and replication of professional wisdom. Professors of journalism, many drawing on years of experience in the news industry, have sought to imbue the next generation of journalists with the skills, values, ethics, and habits of mind (“thinking like a journalist”) considered essential for the proper practice of the craft. Their pedagogical models benefited from some stability: even as the news industry struggled economically and journalists were asked to take on additional duties, the core competencies – writing effective summary ledes, organizing information using the inverted pyramid structure, adhering to AP style, to name a few – remained reasonably steady. They required some adaptation as the internet and social media became disruptive forces, but content was still being produced largely by humans, for humans, and within recognizable institutional frameworks.

The rapid rise of AI has unsettled that stability. Many journalism educators now wonder, with a mixture of exhaustion and exasperation: *Now what?* How are we to teach journalism when the future feels even more uncertain? Which skills, practices, and forms of knowledge should be prioritized when the very definition of news work seems so unstable?

These questions cannot be understood apart from the hamster wheel metaphor we have described throughout this book. Educators now face their own version of the dilemma: Are they ready to run faster on the wheel, cramming a whole new set of skills, competencies, and values into rigid and already packed journalism curricula? Or might this moment offer an opportunity to step back and fundamentally rethink what journalism education should accomplish – and perhaps even reckon with the anxieties and frustrations that professionals and audiences alike express about modern journalism? These questions carry particular urgency given declining enrollments in journalism programs across higher education institutions in countries like the United States.⁵

Why Education Matters as a Site of Investigation

The questions being debated in journalism classrooms – Should AI tools be permitted in the classroom? Which journalistic activities should remain human-centric? How do we assess learning when assignments can be algorithmically generated? – matter because they reveal important human dimensions sometimes obscured in techno-deterministic narratives. They bring to the fore how people negotiate what it means to live, work, and learn when machines are increasingly able to help us think, and even do some of our thinking for us.

Examining education also forces us to confront folk theories about how aspiring professionals are using AI. The prevailing narrative in sensational media accounts of cheating on college campuses is that students across disciplines are happy to outsource their homework to AI. But a growing number of students appear to be rejecting the encroachment of AI on ethical or environmental grounds.⁶ Some even see instructors' use of AI or incorporation of AI tools as effectively shortchanging their learning.

Most consequentially, examining education allows us to pursue a question at the heart of many visible debates about AI: What happens to human thinking when we delegate cognition to systems designed to do more and more of our thinking for us? Journalism education, positioned at the intersection of professional training, humanistic inquiry, and democratic responsibility, offers no easy answers. But it does offer a revealing lens

through which to observe how individuals and institutions struggle to preserve what matters while adapting to what's coming.

The Confidence–Competence Gap

When AI can generate fluent explanations of complex material, engage in seemingly thoughtful dialogue, and produce journalistic work indistinguishable from human output, we face a question that cuts deeper than concerns about technologies like calculators ever did: Can students distinguish between *feeling* knowledgeable and *being* knowledgeable?

A 2024 classroom experiment by Cindy Royal, a digital media professor, helps illustrate this distinction.⁷ Royal introduced her students to Google's NotebookLM shortly after its release. The AI-powered tool allows users to upload documents and generate podcast-style discussions about the material led by human-sounding "hosts" who banter about the subject matter with the easy familiarity of NPR personalities. Royal uploaded nearly a dozen of her own essays on the future of journalism and technology to see how the system would synthesize complex ideas. "I was blown away by what it produced and more than a bit impressed by how it engaged with my work," she reflected. "I couldn't wait to give students a chance to try it."

When students uploaded their own writing and generated AI podcasts discussing it, the results seemed promising. Hearing the "hosts" engage with their work gave students more confidence in their ideas. They then critiqued where AI succeeded and where it faltered. But did the students actually understand their material better after hearing AI discuss it, or did they simply *feel* more confident?

Research from cognitive psychology suggests this distinction matters profoundly. A phenomenon known as the "illusion of explanatory depth" shows that people routinely overestimate how well they understand complex systems – from indoor plumbing to government policies – until they are asked to explain them in detail, at which point the gaps in their knowledge become apparent.⁸ AI tools may amplify this illusion by providing fluent, confident explanations that create a sense of understanding without requiring the friction and cognitive work that genuine comprehension demands. Students may feel they grasp a concept because

they've heard it explained clearly, but they've been spared the confusion, false starts, and hard-won resolution that helps embed that understanding into their own cognitive architecture.

Methods, Epistemology, and Ethics in Learning

A 2024 dialogue between scholars published in *Higher Education Research & Development* frames this tension explicitly. While some view AI as “an indispensable tool for academic research,” others worry it may “undermine the very essence of academic inquiry.”⁹ The debate highlighted how AI potentially implicates three critical elements of academic study.

First, AI complicates *methods* – the “how-to” of research inquiry – because it can now capably perform literature reviews, analyze data, and even generate hypotheses. Second, AI challenges *epistemology* – the philosophy of how knowledge is defined and organized – because its outputs blur the line between human understanding and machine pattern-matching. Third, AI raises questions about *ethics* – the shared normative beliefs about what counts as appropriate practice – as traditional standards of authorship, originality, and intellectual labor become harder to apply. The optimistic position sees AI as a tool that reduces some biases and augments human capabilities. The skeptical view worries that AI changes not just *how* we find answers but *what counts* as an answer worth finding in the first place.

This question of epistemic authority – of what counts as valid and factual – holds special significance for journalism. Journalists treat truth-seeking as their ultimate North Star. That idealist notion might generate eye rolls among those who find the news to be frustrating, even infuriating, because of the biases and blind spots they perceive in journalism, but it doesn't change the fact that journalists remain professionally obsessive about accuracy, even if they invariably get some things wrong. This commitment leads journalists to prioritize factors such as presence (e.g., being on the scene or gathering eyewitness testimony), originality (e.g., offering a “fresh” angle on a news story), and accountability (e.g., having the confidence to stand behind one's reporting).

AI potentially complicates each of these factors. It possesses no sense of presence, produces inherently derivative material, and cannot be held accountable for errors. Moreover, each of these

factors presupposes practice – as in the case of student journalists practicing through multiple assignments the craft of interviewing sources, checking documents, and accepting liability when their work is published. The introduction of AI as a collaborator thus complicates, for better or worse, how we think about the methods, epistemology, and ethics of journalism education, much like it does for academic inquiry as a whole.

Evidence of Cognitive Erosion

Many of the worries raised about artificial intelligence in education rest on what happens when students become too reliant on AI to do the work for them. This is no hypothetical concern.

A systematic review found that excessive dependence on AI dialogue systems leads to measurable deterioration in decision-making, critical thinking, and analytical reasoning.¹⁰ There are multiple mechanisms involved: AI systems confidently spread misinformation, embed algorithmic biases that users internalize, facilitate plagiarism that shortcuts learning, and operate in ways that are so opaque that users never understand why and how they receive particular answers. Each of these mechanisms can widen the gap between what students believe they know, what they ought to know, and what they actually understand.

Research on AI-driven personalized learning identifies a related problem. Despite AI's ability to customize learning paths to individual students, computer scientists Kristjan-Julius Laak and Jaan Aru argue that existing systems fail to foster “learner agency, cognitive engagement and self-regulated learning skills” or the broader competencies needed for intellectual flourishing. The emphasis on personalization, they find, “struggles to align with the broader goals of modern education.”¹¹ In other words, AI excels at delivering content efficiently but stumbles when it comes to developing the very capabilities that distinguish genuine expertise from surface-level competence, such as knowing when to question an answer or how to persist through confusion. AI can tell students *what to think* about a topic more effectively than it can teach them *how to think* about it.

This raises an uncomfortable question that echoes through the research literature: If AI optimizes for speed and efficiency in knowledge acquisition but undermines the development of judgment and autonomy – qualities essential across disciplines

but especially in journalism – are we cultivating a generation of knowledge workers who know many things but struggle to think clearly about them?

As higher education increasingly integrates AI into learning environments, we must distinguish between students who exhibit cognitive *confidence* and those who demonstrate cognitive *competence*. This distinction becomes central to evaluating what kinds of knowledge students are truly building – and what traditional learning pathways, designed to develop depth and resilience through productive friction, are being bypassed in favor of tools designed for speed and delivery.

Why Cognitive Struggle Matters

Journalist Derek Thompson uses a fitness metaphor to illustrate the value of productive struggle: “time under tension” in weightlifting shows that muscles grow through sustained resistance, not quick movements.¹² A 10-second squat builds more strength than a two-second one. The same applies to thinking: sitting patiently with disconnected ideas allows us to braid them into something combinatorially new.

When AI tools promise to offload mental labor, the work feels easier. But we’re bypassing cognitive labor and the intellectual skill-building that comes with it. Writer Nicholas Carr warns that this lets B students produce A-level work while slowly turning them into C students. “To automate learning,” he argues, “is to subvert learning.”¹³

This is why Thompson believes we have our fears about AI backwards. Instead of worrying that technology will steal jobs from capable humans, we should worry about how we’ll fritter away our own capabilities chasing ease and efficiency. “We are so fixated on how technology will out-skill us,” he writes, “that we miss the many ways that technology can de-skill us.”¹⁴

Medical student Benjamin Popokh embodied these concerns.¹⁵ His professors at the University of Texas Southwestern Medical Center encouraged students to use large language models (LLMs) for diagnosing patients, despite the concerns it might erode future doctors’ independent judgment.

“I went to medical school to become a real, capital-‘D’ doctor,” he reflected. “If all you do is plug symptoms into an AI, are you

still a doctor, or are you just slightly better at prompting AI than your patients?”

Popokh’s concerns were reinforced when he found both himself and his classmates becoming increasingly dependent on AI for virtually every patient interaction. “I started to feel dirty presenting my thoughts to attending physicians, knowing they were actually the AI’s thoughts,” he confessed. One day brought a disquieting realization: leaving the hospital, he realized he “hadn’t thought about a single patient independently that day.”

We can just as easily imagine a similar scenario for journalism education. The struggle to come up with an enterprising story idea or find narrative coherence in scattered interview notes helps develop editorial ingenuity. The struggle to cultivate sources and uncover little-known details helps develop verification abilities. The struggle to craft an arresting lede or identify a striking kicker helps develop a writing voice. AI can increasingly handle these frustrating parts: making the blinking cursor on a blank page go away, generating a contact list of potential sources, and instantaneously producing “good enough” grabbers. What essential capacities are being forfeited when we outsource these struggles?

Troubling Parallels, or Just Another Moral Panic?

Such concerns may be overstated, though. From novels to bicycles to radio and television, new tools and technologies tend to arrive accompanied by moral panic: dire warnings about their effects on minds and morals – warnings that often prove overblown.¹⁶ Humans adapt, usually far more productively than pessimists predict.

Learning through imitation has always been central to creative development, too. How many painters began by replicating the classics? One of us spent many teenage afternoons using the newfangled tools of the 1990s – Netscape Navigator, dial-up modems, newspapers suddenly available globally through the wonder of the World Wide Web – to study the work of great sportswriters like Frank Deford, Rick Reilly, and Mitch Albom, reverse-engineering their articles sentence by sentence, studying word choice and turns of phrase like the brush strokes of the masters. That process of becoming a good writer by first becoming an attentive reader, then reconstructing what you’ve observed

through your own effort, remains crucial for developing voice and style.

And we should be honest: LLMs can produce prose that is grammatically pristine, logically structured, and often more polished than most humans' first drafts. We have used these tools ourselves while writing this book to refine text, improve flow, and strengthen our argumentation. We are well past the questions of whether AI can write well or be genuinely helpful in different domains. It demonstrably can, even if many in newsrooms and academia are loath to admit this.

Perhaps, then, the tradeoffs of incorporating AI into education, journalism, and other zones of knowledge work will differ from what we imagine as new capabilities emerge where others wane. Even Thompson, the journalist bemoaning "the end of thinking," acknowledges the real professional value in using AI – from drawing on ChatGPT's Deep Research function to generate graduate-level reports on niche topics to using its basic features to summarize how average readers might interpret passages in a story.¹⁷ This practical reality may explain why even though UK journalists are broadly pessimistic about the impact of AI on journalism, more than half report using it weekly at work.¹⁸

Acknowledging the reality that moral panics are common doesn't require dismissing the concern entirely. The question isn't whether AI will destroy human capability wholesale. It's which capabilities require productive struggle to develop, and what is lost when that struggle is bypassed.

Distinguishing Necessary from Unnecessary Friction

Not every frustration serves formative learning, though. Some *might*, but it's often hard to tell productive struggle from unnecessary friction. Two examples illustrate how this distinction might play out in practice.

In journalism, students traditionally learn headline writing through practice: studying conventions, internalizing constraints (like headline length and the use of active verbs), and trying to catch the reader's attention until it becomes almost instinctual. But several content management systems – including those used in newsrooms – can now suggest AI-generated headline options instantly. Does this change what we should teach?

Perhaps journalism educators should invert the process: rather than teaching students how to write headlines from scratch, teach them to select the best ones from AI-generated options. Through selection, students might learn to identify what makes a headline work: the verb choice, rhythm, and value proposition to a reader. They will likely develop discernment and judgment along the way, which will eventually enable them to write strong headlines independently. The entry point is inverted, but the destination might be the same.

GPS navigation offers a parallel case from everyday life. Many of us now struggle to navigate without the likes of Google Maps, and research suggests heavy GPS dependence reduces spatial awareness.¹⁹ But consider what else is true: less time and cognitive energy spent memorizing street names and efficient routes means more capacity for other tasks. The app signals real-time traffic conditions that would otherwise be unknowable and suggests alternate routes we may have never discovered on our own. Regular commuters eventually learn familiar routes and don't need to rely on the app. App users also develop judgment about when to override questionable GPS directions, such as suspicious detours.

This is admittedly far from the expertise required by London's "Knowledge" test, considered one of the world's hardest driving exams, which dates to the era of horse-drawn carriages.²⁰ No GPS-dependent driver could pass that test, but they don't need to. The question is what level of navigational competence is necessary for the task at hand, and whether the tool helps or hinders the development of that competence.

Some frustrations may therefore prove less essential than imagined – such as agonizing over headlines when AI routinely offers multiple good alternatives to start from. Others remain precisely where productive struggle matters: recognizing when tools fail, cultivating judgment to override them, and developing a distinctive voice that machine-generated averages cannot replicate. The hard work of education in the age of AI perhaps lies in distinguishing between the two – identifying which tensions, held long enough, build the muscles that matter.

Journalism Education at a Crossroads

Against this backdrop, what should journalism educators do? If the research suggests that AI poses risks to human cognitive development, how are institutions charged with training future journalists actually responding? The answer, according to a spate of recent studies, reveals an educational landscape characterized by cautious experimentation, structural constraints, and considerable anxiety about who – if anyone – is getting it right.

Structural Challenges for Implementation

A survey of the heads of accredited journalism programs in the United States illuminates the scope of the challenge. It found that only 20 percent of the programs had incorporated AI into their curricula by 2022–3, even as 51 percent indicated plans to do so within the next three to five years.²¹ We can presume that figure has increased substantially amid the surge in AI adoption following ChatGPT's public release, but the point remains: there is a gap between intention and implementation that reflects deeper structural realities than mere institutional inertia so familiar to academia.

The program heads themselves identified structural issues as key: large majorities of those leaders surveyed identified the need for long-term investment, funding for equipment, and the difficulty of hiring faculty with necessary technological expertise as some of their biggest concerns for staying current with digital transformations. Notably, 84 percent of them rated managing the continual process of change as very or somewhat challenging. By their own recognition, the task facing program heads involves much more than just adding a new course to their catalog.

These patterns extend beyond the United States. Research on journalism education in Egypt and the surrounding region found that programs still largely “equip students with traditional skills” despite the need for “a solid understanding of technology, algorithms, and programming.”²² A study of journalism faculty in Spain uncovered similar concerns. This led the researcher, Marcos Mayo-Cubero, to propose a “techno-humanist” educational model that would blend foundational journalistic competencies with emerging digital ones, all while upholding ethical

standards and social responsibility. As one Spanish professor told him: “We are going through the most complex moment in the history of journalism education, and this readaptation requires a balance between understanding, empathy, and acceptance of change.”²³

Even this aspiration confronts practical obstacles. A separate curricular analysis of 1,775 courses across 41 Spanish journalism programs found that only seven explicitly addressed data journalism or big data as standalone courses, while another 19 only included AI topics as components of broader modules. The insufficient inclusion of AI in curricula, the researchers concluded, “constitutes a deficiency that must be addressed to ensure that future journalists are adequately prepared for a rapidly evolving media environment.”²⁴

A Lack of Expertise

For many journalism programs, there’s also a personnel challenge: Who would teach such courses even if they were offered?

In a qualitative study involving interviews with 14 journalism program administrators, Debora Wenger and her colleagues found that faculty expertise – more than AI policies, industry partnerships, curriculum design, or ethical considerations – appears to be most decisive in shaping institutional responses to AI. “The biggest thing I worry about,” one director of a small journalism program told them, “is whether our faculty can adapt fast enough – figure out the right balance between responding to the technology but not overreacting to the technology.” Another administrator, at a large journalism program, captured the knowledge gap more starkly: “Most of our faculty, if they’ve had time, they’ve worked with one or two of the [AI] tools. But there are very few people who’ve had the intellectual time, let alone the expertise to say, ‘I’ve worked with four or five different ones, here’s what we want to be able to teach.’”²⁵

This has cascading effects. One program chair, acknowledging that no one on their faculty had the requisite expertise, contemplated whether an adjunct hire might address the gap – likely a temporary fix given the market demand for such capabilities. Small programs face particular challenges here since dedicating funds for professional development is perpetually difficult, and some faculty seem reluctant to experiment with AI at all.²⁶

Within specialized areas, the expertise deficit becomes more pronounced. A survey of video journalism educators found that while 93 percent had spent at least some time working with ChatGPT, only 10 percent considered themselves very or extremely familiar with AI video production capabilities.²⁷ More than half reported feeling extremely or somewhat unprepared to teach students about AI video tools – even as local TV news companies already use AI tools to index and retrieve video content, automate translations, and recommend specific clips to video editors.

The “Digital Native” Fallacy

We often assume that these younger, “digitally native” students arrive on campus equipped to navigate these technologies. Might they just intuitively grasp how AI systems function, rendering such deficits moot?

The evidence suggests this assumption is profoundly mistaken. Interviews with students at one large US university found that while they view LLMs as inevitable elements of their future careers, “they often misunderstand how these tools generate content, in part because they lack more fundamental knowledge about digital tools.”²⁸ Many conflated ChatGPT with Google, believing it simply “scans everything” on the internet – a misconception that misrepresents how LLMs function.

This pattern resonates with the confidence–competence gap discussed earlier. Students operated with self-assurance despite lacking actual understanding of how these systems work. The study also found another concerning dynamic: “[S]tudents seem unaware of what they do not know about LLMs.”²⁹ This meta-ignorance – not knowing what one doesn’t know – poses particular risks in journalism, where confidence without competence can lead to major reporting errors.

The Distributed AI Curriculum

How, then, are programs actually integrating AI instruction? Most appear to favor distributing AI content across existing courses rather than developing separate courses dedicated to AI.³⁰

The logic behind this approach is reasonable: treating artificial intelligence as a tool relevant across multiple domains –

newswriting, multimedia production, ethics, and specific beats, to name a few – prevents it from becoming siloed as a technical niche. It shows that AI is not just a specialization (“AI journalism”) but an infrastructural reality permeating the profession.

However, this approach also risks incoherence. Students may receive conflicting messages from different faculty members, with some prohibiting AI use entirely while others embrace it as standard practice. One student recalled this pedagogical whiplash to researchers, noting they had “one professor who directly equated LLM use with plagiarism,” but then “another professor, a working journalist, [who] described how she uses LLMs in her work.”³¹ Such contradictions aren’t just confusing; they leave students without clear professional norms to internalize. Establishing policies around appropriate AI use in journalism education remains a moving target, even years after ChatGPT’s launch. As one administrator noted, “[m]ost of our degree programs have not yet agreed on a specific language” regarding acceptable AI use.³²

Trust in AI “Vibe Checks”

The challenges facing educators mirror those confronting working journalists, for whom the need for professional development proves just as urgent. Many newsrooms have established AI policies of their own (see Chapter 2), but according to New York University journalism professor Hilke Schellmann, the guidelines often prove too abstract or broad to be clearly applicable to a journalist’s daily work. In the absence of clear standards, she observes, “journalists have largely been left to figure things out for themselves.”³³

The risk of that self-guided approach is that many reporters and editors consequently default to what data journalist Cynthia Tu calls “vibe checks”: just messing around with tools to get a sense for whether they are useful or not.³⁴ But as *Washington Post* journalist Jeremy Merrill contends: “Vibes are not enough. You’re not taking a good enough look at your real data. Is it 60 percent accurate? Seventy? Ninety-five? You just don’t know.”³⁵

To move beyond intuition, Schellmann led a research team that systematically tested AI tools across two categories that appear immediately useful for reporting: chatbots for summarizing meetings and documents, and scholarly research tools for

assessing whether new studies are as newsworthy as their press releases claim. The results are insightful for both practitioners and educators.

The summarization tests compared four major chatbots against human-generated summaries of local government meeting transcripts. For short summaries of roughly 200 words, the AI tools actually outperformed humans, highlighting more facts with few hallucinations. ChatGPT-4o's error rate, for example, remained consistently below 1 percent. But for longer summaries of around 500 words, the results proved "surprisingly poor." The AI-generated versions captured only about half the facts included in the human summaries, and the hallucinations increased. True, the tools could produce a summary much faster than their human counterparts, but what value does speed hold when quality degrades so substantially?

The research tool results were far worse. Schellmann's team tested five AI tools marketed for quickly creating literature reviews, comparing their output against human-authored reviews from four award-winning academic papers across diverse fields. The results were, in Schellmann's blunt assessment, "a disaster." Most tools identified less than 6 percent – often zero – of the papers cited in human-authored literature reviews. The tools also disagreed wildly with one another, suggesting hundreds of seemingly related papers with minimal overlap across tools. The results also proved inconsistent when identical queries were executed days later – a surprising finding given that scientific consensus does not shift overnight. The implications extend beyond individual reporting errors. "If a journalist relies on these tools to understand the context surrounding new research," Schellmann warns, "they risk misunderstanding and misrepresenting scientific breakthroughs, omitting published critiques, and overlooking prior work that challenges the findings."³⁶

Now, it's true that the leading frontier AI models as we write this in 2026 are considerably more accurate and less error-prone than the ones that Schellmann's team analyzed in 2025.³⁷ But the principle remains: for educators and professionals alike, the lesson is that relying on gut feelings or "vibe checks" when evaluating AI tools is insufficient, especially given these tools' confident presentation of their results. At minimum, prompts should be tested across multiple competing tools to see if they yield complementary results – a tactic that assumes, of course, that they

are powered by distinct infrastructure. More broadly, this underscores that students need frameworks for systematic assessment, not just exposure to technologies.

Toward an AI-Ready Journalism Curriculum

What, then, should an AI-ready journalism curriculum include? A synthesis of emerging scholarship and evolving program approaches suggests a three-part framework – one that balances preservation, adaptation, and critical engagement.

First, certain competencies remain foundational and cannot be automated. These include source cultivation and interviewing skills – the interpersonal capacities required for building trust and eliciting information that sources would not otherwise share. They include verification methods and fact-checking protocols, encompassing the critical thinking and investigative know-how required to determine the veracity of documents and claims. They also include ethical reasoning in ambiguous situations and an understanding of journalism's democratic function. These are not artifacts of a bygone era; they constitute the deeply human capacities that give journalists their enduring value as purveyors of truth and sense-makers of current events. As Mayo-Cubero's Spanish study observed, professors continue to insist that such "cornerstones of journalism . . . are nonnegotiable." Mastery of the craft still means "telling stories, contextualizing, and critically expressing current events. Innovation should enrich education, but not compromise essentials."³⁸

Second, certain competencies must be added or redefined. Prompt engineering now functions as a form of editorial judgment – not just the ability to craft queries that elicit useful responses from AI systems, but also the capacity to work iteratively with different types of prompts to address different needs as they arise. Computational thinking represents another baseline literacy. This does not require coding expertise; instead, it means cultivating the capacity to think systematically about data, patterns, and processes. It means grasping the logics of algorithms well enough to recognize when computational approaches lend themselves to particular problems and where they are more likely to fall short. Among these emergent competencies, too, is algorithmic accountability as a distinct reporting beat. This means

teaching students to investigate automated systems with the same rigor they would apply to any institution shaping public life, such as by digging into whether AI systems are appropriately adjudicating bail decisions, screening insurance claims, or determining eligibility for government assistance.

Third, certain knowledge domains require new emphasis. Understanding how AI systems are trained and where they fail becomes foundational to using them responsibly. Students and journalists need not become machine learning engineers, but they should at least grasp the fundamentals of how the leading generative AI tools and agentic systems work. This includes comprehending where and how those tools gather their training data; the biases and errors those data are likely to have, and that models are consequently likely to reproduce; the limits to those models' apparent reasoning abilities, and that their plausible-sounding outputs aren't anchored in any human-analogous understanding of truth or falsity; and that confident-sounding output bears no relation to accuracy. This knowledge is essential for how journalists – present and future – approach the AI-generated content they encounter and may yet help produce.

In all, what becomes apparent is that simply teaching students to *use* AI is not only inadequate but also misguided. It is liable to give them the illusion of confidence rather than the assurance of actual competence. Human–AI workflow management should be taught as an emergent professional competency, and one with cross-cutting relevance for knowledge work across disciplines. This means teaching students and professionals how to stress-test AI to understand both how algorithmic systems work generally and how such tools might facilitate (or hinder) the particular kind of work that reporters and editors do in the day-to-day. Some AI tools are genuinely useful and could save time for higher-order cognitive tasks, while others, like those Schellmann's team examined, may just project usefulness while being riddled with errors and inconsistencies. Only with a deep appreciation for *interrogating* tools – and not just being familiar with them – can students be equipped to learn the difference.

Scholars have warned that “universities risk losing their purpose if they focus solely on software training and neglect critical thinking.”³⁹ This caution echoes throughout the emerging consensus on AI-ready journalism education. The goal is not technological fluency for its own sake, but rather a form of “AI literacy.”

AI Literacy for Journalism – and Beyond

What does it mean to be genuinely literate about AI? Not merely familiar with the tools, not merely capable of using them, but able to recognize their limits, understand their tendencies, and exercise judgment about when they help and when they harm?

This question echoes past debates about media literacy that arose during major technological shifts – from print, to TV and radio, and to the internet. However, it feels especially urgent today. AI-generated content is becoming harder to distinguish from professional work, and while AI's influence on journalism and society continues to grow, much of it remains invisible. The question matters not just for journalists but also for students learning to navigate an AI-saturated information environment and for a public whose capacity for democratic engagement depends on distinguishing reliable information from synthetic noise.

A Framework for AI Literacy

Twentieth-century media literacy focused on decoding persuasion techniques in advertising and propaganda. Twenty-first-century news literacy has emphasized cross-referencing and “reading laterally” across multiple reputable sources for verification. Today's AI literacy demands something different, according to media scholars Mark Deuze and Charlie Beckett: recognizing, negotiating, and applying AI tools through a constant process of reflection – about the tools and about oneself in relation to them.⁴⁰ It's an ongoing effort in calibration, not a box-checking test of one-time competency.

Their framework identifies three interconnected components. The first is *knowledge about AI*: about its genealogy, its actual capabilities, and its limitations, moving beyond both the utopian promises and dystopian fears that dominate popular discourse. The second is *recognition*: the ability to identify instances where AI might be usefully and creatively applied, and, equally important, when it should be avoided. The third is *skill-based*: the capacity to coach or teach others in understanding, imagining, developing, and implementing AI. Together, these components suggest that AI literacy is less about mastering a technology than

about cultivating judgment – a capacity as relevant for citizens evaluating AI-generated content as it is for the journalists who may produce it.

Deuze and Beckett draw a crucial distinction between instrumental and imaginative approaches to AI. The instrumental approach – which dominates most industry discourse – treats AI as a tool for efficiency. It asks: How can AI help us do what we already do, but cheaper and faster? The imaginative approach asks a different question: What new forms of journalism become possible – ones we had never even imagined before?

The instrumental approach isn't necessarily wrong. Efficiency matters, particularly for resource-strapped news organizations still struggling to adapt to new realities. But when it becomes the only approach, something valuable is lost. The result is a kind of defensive crouch: AI is something to react to, to manage, to survive – rather than something to shape, to direct, to advance journalism's public mission.

This ability to develop a healthy skepticism combined with an imaginatively creative and critically self-reflexive engagement with AI extends beyond newsrooms. AI literacy matters as much for society as for media professionals. Journalists, then, face a dual challenge: cultivating their own AI literacy while also helping the public do the same through editorial choices about what to cover and how to cover it.

There are multiple ways to address that challenge. Journalists might disclose not only *when* they use AI in their work but also *how and why* – helping themselves and their audiences better understand AI's strengths and weaknesses, as well as how their uses align with professional ethical values. They might talk more freely about what it means to investigate AI as an institution – to hold platforms, applications, and designers accountable for their work in situating algorithms to operate in ways that may help or harm society. They might counter the utopian hype of industry marketing and the dystopian fears in popular culture, articulating instead a more realistic and responsible take on AI's actual workings as well as its impacts on people's everyday lives – on jobs, education, and the environment.

As AI-generated content increasingly fills our screens, this responsibility to train the public in AI literacy becomes more urgent. It adds new components to media and news literacy, from recognizing synthetic material to appreciating how bad actors can

weaponize tools for disinformation and trolling at a scale made possible only by AI.⁴¹ This process of educating the public about AI begins with journalists first gaining a fuller understanding of it for themselves – a task complicated by AI’s often counterintuitive capabilities, as we show next.

The Jagged Technological Frontier

Researchers have used the term “jagged technological frontier” to describe the irregular boundary between tasks where AI excels and those where it fails unpredictably.⁴² The metaphor captures something important: AI’s capabilities form an uneven, hard-to-map terrain where high performance in one area doesn’t guarantee similar results in related areas.

Management professor Ethan Mollick likens this jagged technological frontier to an irregular fortress wall, with towers jutting out (AI strengths) and sections folding back (AI weaknesses).⁴³ In some areas, AI is unreliable; in others, even closely related ones, it seems superhuman. The trick is that it’s hard to decipher which type of knowledge tasks fall into which category. The challenge is compounded as improvements in LLMs and other AI systems continually reshape this invisible boundary over time.⁴⁴

A 2023 study Mollick helped lead demonstrated how this jagged technological frontier operates in practice.⁴⁵ The researchers asked 758 management consultants to complete a series of realistic, knowledge-intensive tasks. For tasks within AI’s capabilities – creativity, ideation, and content creation – consultants using AI completed more tasks, worked faster, and produced higher-quality results. But for tasks outside the frontier – say, a business problem requiring deeper analysis of financial data and interview notes – consultants using AI were 19 percentage points less likely to arrive at optimal solutions. The AI didn’t just fail to help; it actively degraded performance.

The underlying problem is a mechanism that the researchers called “miscalibrated trust”: participants overrelied on AI precisely where it was weakest. AI’s mistakes come wrapped in the same confidence as its successes, without as much of the hedging language that human collaborators would likely include to convey their uncertainty. This creates an epistemic hazard extending well beyond professional contexts. Individuals researching health questions, students seeking homework help, and citizens hoping

AI can help them make sense of political information all face the same challenge: the machine's confident tone does not correspond to its actual capacity. This is, fundamentally, a matter of AI literacy for journalists and everyone else alike.

The Pull toward the Average

Beyond the jagged technological frontier lies another consequence that is significant for creative work: AI tends to homogenize output. The study that Mollick helped lead found that while AI-assisted work achieved higher quality on average, outputs became markedly more similar to one another. When everyone accesses the same AI trained on the same data, work converges toward a statistical mean.

Music composer Brian Eno captures this problem in reflecting on his decades of experimenting with generative systems: “[W]hat you spend nearly all your time doing is trying to stop the system becoming mind-numbingly mediocre. You really feel the pull of the averaging effect of AI, given that what you are receiving is a kind of averaged out distillation of stuff from a lot of different sources.” Designer Frank Chimero extends this observation: “Take what the AI gives without question and you’re not producing, you’re consuming.”⁴⁶ Without critical engagement, the model’s limits become our own.

This homogenization effect intersects with existing anxieties about the commodification of news. Competitive pressures, social media metrics, and algorithmic distribution were already pushing at least some news coverage toward a lowest-common-denominator sameness as journalists were forced to run faster on their hamster wheel. AI threatens to accelerate this flattening, undercutting the capacities needed for originality. In other words, it undermines the capacity to produce journalism that actually stands out and is impactful.

AI literacy must therefore include awareness of this averaging tendency and the discipline to resist it. An AI-literate individual knows not only where AI fails outright but also where it succeeds in ways that flatten rather than elevate. Both forms of awareness require what Deuze and Beckett call “positional reflexivity,” or ways of knowing ourselves in relation to machines.⁴⁷ This reflexivity matters for journalists producing content and consumers receiving it, helping both recognize when the diversity of voices

essential for public life gets averaged away into algorithmic sameness.

Personal Empowerment in the Age of AI

If AI literacy is principally about understanding, our attention now turns to empowerment – the crucial next step in acting on that understanding. What does it mean to use AI in ways that develop rather than diminish human capacity?

Chimero again proves insightful here in offering a useful conceptual framework for distinguishing different modes of technological engagement. He proposes thinking about AI not as a tool but as an instrument – something that requires practice and rewards skill. “Instruments can surprise you with what they offer, but they are not automatic,” he writes. “In the end, they require a touch. You use a tool, but you *play* an instrument.”⁴⁸

This distinction maps directly onto the core debates in journalism education. Tools are optimized for efficiency; instruments are designed for capability. And the capabilities that journalism demands – both old and new – are forged through productive friction. “There is no purpose to better machines,” Chimero argues, “if they do not also produce better humans.”⁴⁹

Working under the Machine

To understand what empowerment *isn't*, consider what Chimero calls being “under the machine.” This means accepting outputs without question or letting algorithmic capabilities define the frontiers of one’s work. Chimero associates this with “vibe coding”: the notion that, with the right sensibility and prompt engineering, a person can build applications without understanding underlying systems. Vibe coding suggests taste alone is enough; you can let the machine handle the hard work.

This isn’t inherently bad, of course; vibe coding, as we show in the next chapter, can speed up experimentation and innovation. But this approach of working “under the machine” brings its own hazards, particularly for novices. Management professor Nan Jia and her colleagues found that using AI to partially automate customer sales calls actually widened the gap between top and junior salespeople.⁵⁰ Because AI took over the routine calls that junior

workers had traditionally learned from, they never developed the expertise needed to close sales – the one task that still required a human touch. As a result, top performers pulled even further ahead. Mollick aptly captures the main point in arguing that “AI is often most useful where we’re already expert enough to spot its mistakes, yet least helpful in the deep work that made us experts in the first place.”⁵¹

If AI handles routine journalistic activities – summarizing press releases, generating draft copy from structured data, conducting preliminary research – where do students encounter the low-stakes practice opportunities essential to developing judgment? The on-ramp to expertise cannot be automated away without consequence. “Time saved is not strength gained,” as Chimero observes.⁵²

Working beside the Machine

To work beside the machine, according to Chimero, is to treat it as a collaborator: shaping outputs rather than simply accepting them. He references the aforementioned composer to make the point: “Eno often says while making music he feels like a gardener: planting loops and textures, then watching them sprout into something unexpected with the potential to become incredibly beautiful with a little bit of care and pruning.”⁵³ The machine makes the material, yes, but the creativity comes in how that raw material is reshaped into a more finished form.

For journalism pedagogy, this suggests approaches that maintain human agency and judgment while strategically leveraging AI capabilities. Students might use AI to generate draft material, then actively reshape it – not merely editing for accuracy but also pushing against averaged-out prose, adding editorial voice, fighting algorithmic predictability. They might employ AI to reveal patterns in data, then apply journalistic judgment about which patterns matter and why. As Deuze and Beckett suggest, time freed from routine tasks could redirect attention toward “investing in more real-world and grassroots reporting, more human to human interaction, more creative experimentation with storytelling formats and topics, more investigation and structured reflection.”⁵⁴

The jagged technological frontier study identified two successful patterns of human–AI collaboration.⁵⁵ “Centaurs”

strategically divide labor between humans and machines, handling certain tasks independently while delegating others to the machine. “Cyborgs” integrate AI continuously into their workflow with intensive real-time oversight. Both approaches outperformed passive AI users, but only because they featured substantial domain expertise and judgment. Successful journalistic interactions with AI thus require an understanding of the profession’s core competencies.

This points toward a specific pedagogical stance: meaningful integration means teaching AI as an instrument requiring mastery, not a tool that offers shortcuts. It means designing curricula that preserve friction where friction serves learning, while introducing AI fluency as an additional layer of expertise rather than a substitute for foundational skills.

The Questions Journalism Education Must Answer

Chimero concludes with a meditation on limits: “The question isn’t just ‘what can this machine do?’ but ‘what should it serve?’ and, most importantly, ‘when should we stop?’”⁵⁶ This final question – when to stop – carries particular urgency for journalism education.

Journalists and educators alike find themselves on a relentless hamster wheel, tasked with doing more with less even as their social relevance comes under question. If the status quo has proven inadequate, why optimize it further? For what orienting goal are we using AI in the first place?⁵⁷

Educators must determine where to stop in teaching AI itself – equipping students for contemporary newsrooms without eliminating the productive friction essential to genuine learning. They must also reckon with curricular limits: finite classroom hours, faculty expertise gaps, and governance structures that resist rapid transformation. Adding something in often means having to take something out. How do we choose?

There are no sure answers to those questions, nor a one-size-fits-all model that successfully addresses the challenges posed by the proliferation of AI in the academy, profession, and society. Engaging with them will require experimentation and the humility to accept missteps along the way. However, that means breaking our paralysis and breaking from inertia. It means deliberately choosing and acting to step off our own hamster wheel.

Empowerment in the age of AI means preparing students, practicing journalists, and the public not merely to use these tools, but also to interrogate them. It requires a clear-eyed understanding of how AI systems function and what it means to use them imaginatively to serve a higher purpose. The goal is not simply teaching people how to engage with AI – it is cultivating the judgment to know when to set it aside altogether. That means learning when to stop.

6

Reimagining Journalism for the Age of AI

“The way that I look at the applications of AI today are very much focused on very small, very practical problems,” Chase Davis told a room full of data journalists in October 2025.¹

But Davis, an early adopter of artificial intelligence who formerly led *The New York Times*’ Interactive News Desk, was less interested in using AI to increase the efficiency of day-to-day newsroom tasks. Instead, he was more excited about its prospects to help journalists do something new – something *better*.

One emerging phenomenon had especially piqued his interest: vibe coding, a new way to build software that became a viral sensation in 2025. “To vibe code,” as technologist Paul Ford puts it, “is to make software with prompts sent to a specialized chatbot – not coding, but telling – and letting the bot work out the bugs.”² The user expresses their goal in plain English (or another natural language) within a simple interface, and *voilà*: an AI-powered system writes the code for them. If the user isn’t happy with the output or wants to keep adding to it, they prompt the system accordingly and it makes the changes. Davis found that idea empowering.

“It allows us to spend our brain calories thinking more about the how and the why and the methodological questions [of data analysis],” he said, “and less about the ‘Wait a minute, what’s the second argument in the VLOOKUP function again?’”

At the E.W. Scripps Company, one of the largest operators of local television stations in the United States, Kerry Oslund was also watching journalists with no programming experience create working software together using natural language. “This is the revenge of the English major,” the vice president of AI strategy told an industry conference in 2025.³

If previously the barrier to doing data journalism was technical skills, Davis argued, “now it’s a little bit more about higher-order thinking and creative problem-solving.”⁴ More broadly, as creative technologist Kawandee Virdee forecasted in his 2026 prediction for Nieman Lab, it meant “the most significant newsroom innovation won’t come from dev teams – it’ll come from journalists who can now create their own tools.”⁵

Vibe coding is about more than just making coding accessible to non-programmers, though. It represents a glimpse of something bigger: a step toward new forms of agentic AI, or advanced systems designed to operate autonomously to achieve goals. AI agents, in effect, do “all sorts of work a human might do on a computer,” such as online shopping or email, but without human involvement.⁶ The agents in such systems are effectively AI-powered programs that break complex tasks down into sub-tasks, employ a variety of tools and techniques to complete them, and “learn” how to navigate increasingly complex, multi-step processes through rapid iteration – like how to build a software application from plain English descriptions.⁷

This vision of an agentic future was already making the rounds in the news industry by late 2025. At Scripps, Oslund talked about deploying more than 300 AI agents to handle adaptive tasks, imagining “agent swarms” where multiple AI systems pass work back and forth, compiling weekly reports, summarizing changes, and assembling dashboards with minimal human intervention.⁸ Other news organizations have also explored how agents might someday coordinate across the entire news-gathering process, with one agent performing internet searches, another editing copy, and others handling distribution – all while chaining tasks and communicating with each other to stay in sync.⁹

We are, by mid-2026, already seeing a growing number of newsroom leaders gravitating toward the idea of a future where different AI agents perform narrow journalism tasks but, crucially, coordinate with one another intelligently, much like a func-

tional team of human professionals. It is a vision that not only rearranges the labor of journalism – and, perhaps, its values and purposes – but also further challenges the longstanding insistence on having a “human in the loop” to ensure quality control.

Accelerating or Interrupting the Hamster Wheel?

Several years after the release of ChatGPT, the excitement around AI – and the imagined futures for generative and agentic AI – remains palpable.

There are many reasons for this, but we argue that an important motivator is the opportunity to step off the hamster wheel that Dean Starkman wrote about back in 2010.¹⁰ As we have argued throughout this book, that metaphor captures not only what it has meant to do journalism in the twenty-first century but, in some similar and dissimilar ways, the status quo of media industries, democratic life, and education. In this context, AI doesn’t just represent a new technology; it represents disruption at a time when change feels *necessary*.

But why does change feel so urgent?

Consider the quality crisis in journalism, which at minimum reflects a widespread belief that journalists are often chasing clicks and producing minimally viable news products – stories with fewer sources and even less fact-checking – just to meet growing productivity demands.¹¹ Or the growing rates of burnout as journalists and audiences alike express intense emotional distress and exhaustion in the face of that pursuit.¹² Or the economic precarity in journalism that increasingly makes it feel financially unsustainable, especially as the competition for the finite amount of attention becomes overwhelming against competitors who have more entertaining media products to sell.¹³ Or the widespread belief that democracy isn’t just backsliding but becoming captured and non-responsive to the needs of citizens and the common good.¹⁴ Or even the growing perception that higher education no longer offers the same value proposition it once did, especially given the eye-watering costs of attendance at many institutions and the perceived misalignment with industry needs.¹⁵

Against this backdrop, AI’s promise of relief sounds appealing. But a central aspect of the critique that Starkman offered – and that other scholars have substantiated – is that new technologies

that arrive bearing promises of relief more often than not end up delivering new forms of acceleration.

What's to say AI won't do the same?

An eight-month ethnographic study of a technology company in 2025 offers one potential preview of the answer to that question.¹⁶ Researchers Aruna Ranganathan and Xingqi Maggie Ye found that AI tools did not reduce work – they consistently intensified it. Employees worked at a faster pace, took on broader tasks, and extended work into more hours of the day. AI simply made “doing more” feel more possible and even rewarding. However, what looked like a productivity surge carried with it quiet workload creep.

In other words, the hamster wheel began spinning faster. While it provided a *sensation* of progress, those running on it seemingly remained in place. As one engineer told the researchers: “You had thought that maybe, oh, because you could be more productive with AI, then you save some time, you can work less. But then really, you don't work less. You just work the same amount or even more.”

From Deprofessionalization to Epistemic Thinness

But the evidence points to more than just the likely acceleration of the existing wheel for journalism. The arrival of AI also accelerates a set of interlocking dynamics that, in concert, propose to reshape the institution from several directions at once. We have traced four such dynamics in this book.

The first is *deprofessionalization* (Chapter 2). AI continues to erode the boundaries around what constitutes distinctly professional human labor and challenges core ideas about what it means to be a journalist. When vibe coding allows anyone to build journalistic data visualizations, when AI can draft articles in seconds, and when automated systems can fact-check and edit copy, it becomes harder to delineate the specialized expertise or jurisdiction that defines a journalist.

The second is *infrastructural capture* (Chapter 3). AI continues to strip journalists of their economic power and control over how content is created, distributed, and used, transferring it instead to a small, powerful group of tech companies. When

newsrooms rely on the same handful of large language models, when platform companies increasingly mediate between journalists and audiences through AI-powered intermediaries, and when the costs of developing competitive alternatives are measured in hundreds of millions of dollars, the asymmetry of power becomes an existential threat to several news organizations.

The third is *democratic corrosion* (Chapter 4). AI-generated disinformation, cognitive overload from the exponential growth of synthetic content, and the structural logics of AI development – whether capitalist or authoritarian – have degraded the epistemic and communicative conditions that democratic deliberation requires. When synthetic content becomes indistinguishable from authentic material, when there are multiple incentives to steer citizens away from genuine and thoughtful deliberation, and when the common good becomes an afterthought in the systems that increasingly govern public arenas and civic life, the fragile environment on which democracy depends is further corroded.

The fourth is *cognitive atrophy* (Chapter 5). AI-powered tools that shortcut productive struggle in education and learning may lead to journalists and citizens becoming fluent in using the technology but deficient in the contextual awareness, critical thinking, and experiential wisdom that AI cannot replicate. When students use AI to complete assignments without wrestling with difficult concepts, when journalists rely on AI summaries rather than reading source documents, and when “time under tension” – that essential period of cognitive struggle that builds expertise – gets optimized away, we risk producing individuals skilled at prompting but weak at thinking.

Running through each of these four dynamics is a common thread: *epistemic thinness*, or the growing gap between *feeling* informed and actually *being* informed. It represents the potential apotheosis of hamster wheel acceleration: volume without knowledge, information without wisdom, and surface-level awareness of many things without deep contextual understanding of anything.

When Google’s AI Overviews synthesize an answer from multiple news sources, users may *feel* their information needs have been met, but they have ultimately just encountered a simplified version of the reporting that typically offers little of the contextual richness, sourcing detail, or qualifying hedges found in the

original work. When journalists use AI to quickly summarize documents, they may feel they've grasped the key points, but they are often missing the nuances, contradictions, and context that emerge from careful reading – sometimes with disastrous results for their own stories. When students use AI to draft an essay, they may feel they've learned something, but they are more likely to have outsourced the cognitive work that actually builds understanding – the wrestle through which knowledge becomes one's own.

This isn't some entirely new phenomenon, of course. We've always relied on summaries, headlines, and shortcuts to navigate the overwhelming flow of information. But the risk of feeling you know something solid while skating on epistemically thin ice will likely grow in the age of AI. And this thinness reinforces those four aforementioned dynamics. Deprofessionalization weakens the expertise behind information production. Infrastructural capture undermines content diversity and the notion of a marketplace of ideas. Democratic corrosion undercuts the deliberative environment as systems optimize for engagement rather than understanding. And cognitive atrophy limits the capacity to recognize the thinness itself – when people lose the habit of productive intellectual struggle, they may lose the very faculty that helps them recognize when they have confidence without equivalent competence.

Reimagining Journalism for the Age of AI

It doesn't have to be this way, though. Technology is not deterministic, nor have journalists lost their ability to implement technologies in ways that align with their professional values.¹⁷ While AI is more than a neutral tool, it's far from almighty. Its trajectory will continue to be shaped by human choices – by the decisions that developers, journalists, policy-makers, educators, and citizens make about how to build, implement, and regulate these systems.

This constitutive moment for journalism and society may very well go in a different and better direction. For one thing, the arrival of AI could yet “break” the hamster wheel in a good way. The rise of AI already challenges the hamster wheel's fundamental premise – that humans must power the wheel – and thereby

creates space for a long-overdue reckoning about what it's all for, anyway: What, exactly, do we want from journalism?

But if we are to guide AI away from mere acceleration, we need to give it a new direction. That requires us to reimagine what journalism is for, what it means to do journalism, and how the ecosystem supporting journalism ought to be restructured.

One framework for this reimagination may come from a largely forgotten source. Two decades ago, the Harvard management theorist Clayton Christensen worked with the American Press Institute on a \$2 million initiative called Newspaper Next that placed his “jobs to be done” (JTBD) framework at the center of a blueprint for industry transformation.¹⁸ The core insight was deceptively simple: people don't consume news because they want “articles” or “broadcasts” as such. They turn to journalism because they have needs they want fulfilled – understanding what's happening in their community, making sense of complex issues, deciding how to act on problems that affect their families. They have, in short, jobs to be done. We all “hire” and “fire” products and services constantly as we seek the right fit for the job at hand.

The JTBD perspective means starting not from what newsrooms produce but from what communities actually need – a radically bottom-up reorientation. It forces journalists to reckon with a reality: as much as they might know how to craft compelling news products, they often know much less about exactly *why* people might hire or fire those products – principally because news media have never been especially good at “identifying jobs that are poorly performed in customers' lives and then designing products, experiences, and processes around those jobs.”¹⁹

When Newspaper Next launched in 2006, it generated genuine excitement: more than 5,000 people attended workshops, publications piloted the approach, and editors spoke of a future “more in our control than the popular wisdom would have us believe.”²⁰ But that was also the year when daily newspaper circulation fell below 50 million for the first time since 1945. Then the Great Financial Crisis hit in 2008–9, and survival imperatives overwhelmed experimentation. The JTBD approach became a casualty of the very acceleration that made reimagination so necessary.

Given AI's accelerationist tendencies, it would be foolish to assume it might somehow *slow down* journalistic work enough

to provide space for innovative reorientation. And yet, it would be equally foolish to overlook how AI might enable new and potentially better ways of thinking and doing journalism – or how it might allow news audiences to accomplish “jobs” in their lives that had scarcely been imagined before. Ultimately, news organizations will succeed to the extent that they take more seriously the need to create and sustain value in people’s lives. AI could undermine such efforts, no doubt, but it also could unlock fresh approaches and newfound sources of value creation.

Objectives for a Reimagination of Journalism

We believe that a reimagination of journalism, given the present and more realistic futures for AI, can center many different objectives, but we offer four as a starting point.

The first is making journalism more *accessible*. This means increasing availability across languages, formats, and modalities, shaped not by what the newsroom finds convenient to produce but by how communities actually consume and use information. AI-powered translation is already enabling newsrooms like *Le Monde* to publish dozens of stories daily in English alongside their French editions.²¹ In India, *Scroll* built Factivo, a tool that converts text articles into video in local languages, reaching audiences who might never encounter a traditional news story.²² These are not just distribution efficiencies; they represent a fundamental reorientation from “read our article” to “understand what is happening in the format, language, and style that is most useful to you.”

The second is reshaping journalism to become more *disinformation-resistant*. This means treating verification as a core organizational capability – as many newsrooms already do – while embracing new workflows that can better capture the forms of disinformation that news users are being subjected to, more expediently address that disinformation, and better reach those users with the truthful or corrective information in a more helpful way. Disinformation resistance is expensive, technically complex, and relentless – so, developing this capability will be no easy feat and will require coordination with partners across technology and civil society. One promising direction is the rise of open-source investigations (OSINV), sometimes called open-

source intelligence (OSINT), which involves drawing on publicly available digital traces, geolocation techniques, and collaborative verification communities (such as Bellingcat) to complement traditional newsroom reporting, altogether enabling faster, distributed, and more transparent fact-finding.

The third is rebuilding journalism around *micro-relevance*. This means enhancing the capacity to connect societal information with individual concerns. A community member dealing with contaminated water may not need a 1,200-word feature on municipal infrastructure policy; they may just need to know whether their water is safe and whom to call, now. AI makes it possible to ask questions about that story, or to render it in a multitude of versions, each highlighting the information most relevant to a particular individual or their community. This isn't just about hyper-personalizing content – and any personalization runs the risk of siloing users or ceding editorial authority – but rather about recognizing that different individuals have different resources (e.g., attention, knowledge) and needs. The line between serving individual needs and fragmenting shared understanding is fraught and requires constant navigation, but micro-relevance can reinvigorate the perceived value of journalism and is impossible to enact in a human-centric production environment.

The fourth is refashioning journalism to become more *conversationally flexible*. This means freeing it from conventions that have long served the profession's internal logic more than its audiences' needs. Consider the stylistic conventions of news, such as the inverted pyramid narrative structure or the "voice of God" delivery of broadcasters. People are familiar with these conventions, but they are not well suited for the current informational environment, which rewards informality, authenticity, and unstuffy storytelling. While experiments in chatbots serving up news have had a mixed track record of success, the widespread popularity of ChatGPT, Claude, Gemini, and more suggests that news organizations have much to learn about how people seek information through multiple forms of back-and-forth. The notion of a more "conversational journalism" has been around for years, but it has never been realistic: human journalists simply cannot develop relationships with every member of their audience, much less do so adaptively and in real time.²³ AI at least makes a more dialogical journalism possible – a human-machine form of communication that might allow news to reach audi-

ences in tones, styles, and contexts suited to them, even if it simulates rather than replicates genuine human engagement.

These four objectives don't begin by focusing on what journalists already do – they start from addressing jobs people need done. A fuller reimagination will require news professionals to find renewed purpose, including a more bottom-up understanding of how to create lasting value and help people live fuller, better-informed lives. In a profession that has spent several decades feeling adrift – necessary for society but hardly appreciated, and on increasingly shaky footing socially and economically – this renewed sense of purpose may be as important as any specific technological capability that develops in the years to come.

Crucially, this process of reimagination doesn't end with a single vision. Its outcome, whatever it looks like, will no doubt require ongoing renegotiation as conditions change: as new technologies emerge, new barriers are erected, and new opportunities become attainable.

For example, if the visions for agentic AI articulated at the beginning of this chapter come to pass, we can begin reimagining journalism's objectives in even more fundamental ways. This will become doubly necessary if hamster wheels are further accelerated and the conditions for journalism, democracy, education, and social life are further disrupted.

Indeed, such a future might necessitate having journalists produce news for machines and in machine-readable forms. As media researcher Daniel Trielli speculated to Nieman Lab about an agentic AI future: "AI systems do not need ledes, nut-graphs, or narrative flows; they need user-relevant, novel, and machine-readable content."²⁴ That format might be bulleted lists or JSON – whatever helps machines "read" and reformat content for AI outputs. Perhaps even more fundamentally, rather than prompting AI agents to help them complete a task, human journalists may find themselves being prompted to help the agent complete whatever task it has determined itself to be incapable of performing.

These are incredibly uncomfortable notions for many journalists. But they're worth exploring if what matters isn't how journalism gets done but whether it leads to outcomes that better serve the common good in light of the new possibilities before us.

The Choices before Us

This call to reimagine journalism in ways that interrupt or break the hamster wheel and guide it toward desirable ends may seem naive to many – especially given our own arguments about the power imbalances that increasingly take agency away from journalists and news organizations. As we’ve argued, the dominant logics driving AI development are built around efficiency, extraction, commercialization, and control. Karen Hao’s *Empire of AI* documents what those logics look like in practice: Kenyan content moderators sifting through psychologically devastating material; Chilean communities resisting the resource extraction required by AI development; and artists whose work is scraped without consent or compensation.²⁵ As Hao told one interviewer: “Every single person [who] first encountered the ‘empire of AI’ felt the same feeling, which is, I have no more agency to self-determine my future.”²⁶

However, the story of inevitability is one we often tell to avoid the responsibility of making difficult choices. It’s a form of giving in to nihilism and giving up whatever agency we still have. This is the deeper danger of what the sociologist Mona Sloane calls the “prediction paradigm”: the way AI systems, by extrapolating forward from patterns in the past, present one future as already determined and thereby crowd out the collective deliberation through which societies might choose a different one. By design, she observes, such systems can offer only “old solutions to old problems”; the work of imagining genuinely new ones remains stubbornly, necessarily human.²⁷ With AI’s predictive tendencies increasingly embedded into more and more of everyday social life, we need journalism, of all institutions, to be a counterweight, keeping the door open to an expansive set of potential futures.

So, it is fine to acknowledge that redirecting what appears to be a runaway train being driven by conductors who have a very different destination in mind is a hard – no, *really hard* – task. But it is necessary, and it will only get harder with each passing day of inaction. Indeed, even as we write in 2026, several safety researchers are either departing AI labs or seeing their safety teams eliminated, with some explicitly warning that their employers are moving too fast and downplaying the technology’s

shortcomings.²⁸ In other words, the extractive dynamics Hao documents are not self-correcting; if anything, commercial and competitive pressures are only accelerating them.

We likewise recognize that our call for reimagination is quite aspirational, and would add that it only stands a chance to succeed with outside help. It cannot just be an effort where journalists come together to chart a new vision. It must involve citizens, regulators, educators, and the many others who intersect with journalism and thereby help shape the conditions for its operation. The asymmetry of power between individual newsrooms and platform giants makes collective action essential. It requires coalition-building across stakeholders who may have different interests but share a concern about journalism's future role in democratic life.

Furthermore, we accept that our call may end up being wasteful, if not misguided. AI may turn out far less transformational than even tempered voices believe. But as we have demonstrated in this book, even if AI doesn't fulfill all its hype, it has already left a lasting impact on journalism that will be felt for years to come. Nevertheless, it is important to acknowledge that every choice comes at a cost, and the AI environment may look very different by the time this book reaches your hands. (One of the authors was already working on a book about journalism and AI when ChatGPT was released in 2022. That draft looks rather quaint now.)

Even with these uncertainties, we argue that fear – whether of the unknown or of being wrong – must be cast aside, for to do nothing is more than likely to be the gravest mistake. Journalism has learned hard lessons from inaction in the face of technological change. Best not to repeat that error.

While the specific systems and examples we describe in these pages may look dated within a year or two – superseded, abandoned, or absorbed into something we cannot yet name or imagine – we believe the broader questions we've raised will remain enduring and vital regardless of the specific technologies involved. What we've offered is a conceptual vocabulary for thinking about the choices journalism faces: acceleration versus reimagination, working under versus beside the machine, replacing versus complementing human labor, and infrastructural autonomy versus capture. These are structural tensions that will remain pivot points for the future of news and journalism.

The subtitle of this book – *From Acceleration to Reimagination* – therefore captures the choice before us. We’ve lived through decades of acceleration. We’ve watched journalism speed up, hollow out, and lose its footing. We’ve seen technologies arrive bearing promises of salvation and deliver new dependencies and pressures instead. The hamster wheel has spun faster and faster, and all of us – journalists and news consumers alike – are exhausted.

The invitation of this constitutive moment is to step off the wheel – not to abandon journalism but to reclaim and reinvent it. To ask what it’s for, whom it serves, and how it might be rebuilt around those purposes. To use the extraordinary capabilities of artificial intelligence not to run faster in place but to push for something truly better. The path forward is not predetermined. The choice – between acceleration and reimagination – is at least partly ours to make.

Acknowledgments

This project began with an email. On December 12, 2023 – 377 days after OpenAI launched ChatGPT and sparked the GenAI frenzy – Mary Savigar, senior editor for media and technology at Polity Press, reached out to Seth Lewis.

“I am developing an ambitious and cutting-edge new book series on AI and Society,” she wrote. The invitation was to write what she hoped would be the definitive book on what AI means for journalism, news, and the public arena. It would bring together the emerging research and synthesize it into a readable narrative. “Because AI is changing so rapidly,” she added, “it will be important to develop strong analytical frameworks and a persuasive argument.”

Mary was prescient in recognizing the significance – and the speed – of what was coming. Since 2023, a cascading series of developments in AI has touched, and in some cases transformed, a great many sectors of digital media and information. We owe a great deal to her vision, and to her invitation. We hope this book has lived up to her ambition.

At that time in late 2023, Seth was researching how journalists were reacting to the rise of LLMs and chatbots. Over the past decade, he had studied automation in news and novel forms of human-machine communication, so extending that work to include GenAI felt like a natural next step. Even so, he was skeptical of the hype surrounding AI and questioned whether the

timing was right for a book: Was there too much on this topic already, or, paradoxically, too little to draw from? Would the AI bubble burst in 2024, making the whole thing moot? (It seems quaint to ask that now, in 2026, as the AI hype train barrels ahead even faster. This year, more than half a trillion dollars is being invested into AI data center infrastructure by tech giants including Alphabet, Amazon, Microsoft, and Meta – and the adoption of AI tools has outpaced almost anything previously seen in consumer technology, in both speed and scale.)

Above all, Seth questioned whether he could take on such a complex, fast-moving subject alone. Artificial intelligence is multifaceted and multipolar. It defies easy definition, in part because “intelligence” itself remains so hard to pin down. And its manifestations across industries, domains, countries, and political systems have a Hydra-like quality: With so many moving parts, which aspects of AI deserve attention, at the inevitable expense of others? Add to that the fact that public fascination with AI has followed a familiar pattern for nearly three-quarters of a century – dazzling hype, bewildering disappointment, repeat – and the challenge felt even more daunting.

And yet, in 2023 (as in 2026), something felt different about this moment in AI development. There had been a real step-change in computing capabilities, to be sure, but also a demonstrated shift in social, economic, political, and cultural impact that exceeded anything seen in previous hype cycles for synthetic intelligence – both in general and in the context of journalism in particular.

Thus, our project began as a collective effort to make sense of this moment.

Seth knew he couldn’t tackle this project alone – and, thankfully, he didn’t have to. He first turned to Rodrigo Zamith, his former doctoral advisee at the University of Minnesota and more recently the chair of the Journalism Department at the University of Massachusetts Amherst. Rodrigo had actually started writing his own book on AI and journalism back in 2022, but mounting administrative duties eventually pushed it to the back burner. Rather than letting that groundwork go to waste, he channeled it into this collaboration. He brought not only a strong foundation in the topic but also deep expertise in how technology intersects with journalism, a sharp editorial eye, and an almost unrivaled attention to detail. Those qualities proved essential – both for

keeping pace with a rapidly expanding body of research and for weaving it into the book's core arguments.

But even two people weren't enough for such a wide-ranging project. In spring 2024, Seth and Rodrigo recruited Tomás Dodds, who was the most active among us in studying, in a deeply grounded way, what was happening in newsrooms: how AI was being adopted, adapted, or refused. Tomás also brought a critical perspective that the project sorely needed, one attentive to power imbalances, both between tech platforms and media organizations, and between the Global North and the Global South.

With our trio in place in 2024, we spent the next 18 months debating, outlining, organizing, and writing this book. A lot happened during that time. GenAI advanced rapidly: capabilities that were promising but crude in 2023 – automatically creating realistic and convincing photos, videos, text, and audio, among them – had become remarkably good by the time we finished writing in 2026. Further development of AI “agents,” which allow computers to direct other computers in accomplishing complex tasks, often with limited human involvement, took steps toward what may be the next leap. It is already crystallizing fears about widespread job displacement in fields like law, finance, and insurance, triggering even more fear and anxiety about the future of white-collar work.

Our personal lives were equally eventful during this period. Two of us accepted new faculty positions. Seth went from the West Coast to the East Coast, leaving the University of Oregon after a decade to become the inaugural Elcan Jefferson Scholars Foundation Distinguished Professor of Artificial Intelligence and Media Studies at the University of Virginia. Tomás made an even bigger move, departing Leiden University in the Netherlands and trekking across the Atlantic to join the University of Wisconsin-Madison as an Assistant Professor and founding director of his newly created Public Tech Media Lab. And Rodrigo began winding down his service as department chair.

Reflecting on the project now, the three of us are deeply grateful to the many collaborators, colleagues, students, research assistants, and friends and family. It's to them that we turn next.

We begin with our research assistants. Daniel Barth, at Columbia University, merits special praise for his contributions

to examples throughout the book and particularly to Chapters 1 and 5. University of Wisconsin students Kate Warren, Dale Edwardson Jr., and Soyeong Cho assisted with Chapter 2. And several research assistants at the University of Oregon – James Phillips, Bismark Owusu-Yeboah, Ramtin Ranjpour, Ethan Donahue, and Will Arangelov – helped us organize our material and carry out the final fact-checks on the manuscript.

We each benefited from working through ideas for the book in conversation with many people – on panels, in hallway chats at conferences, in interviews with journalists and podcasters, and especially with colleagues both at our home universities and at places where we spent fellowships while working on the book. In 2025, Seth spent a sabbatical month in Germany, where he was a visiting fellow with the Center for Digitized Public Spheres Research (ZEDOE) in the Department of Communication at the University of Münster. There, he benefited greatly from conversations with Thorsten Quandt and his research team studying online communication, with Nina Springer and her group focused on journalism studies, and with Jakob Jünger, who taught him how to sail on the lovely Lake Aasee. Meanwhile, Rodrigo worked on drafts that would find their way into this book while spending the fall of 2022 as a visiting researcher at the Ander Centre for Research on News and Opinion in the Digital Era at Karlstad University – exchanging ideas with Cornelia Brantner, Michael Karlsson, Joanne Kuai, Henrik Örnebring, and others while discovering the Swedish tradition of *fika*. In 2025, Tomás was a Research Fellow at the Institute for Advanced Study at the University of Amsterdam and a DigiPol Visiting Fellow at the University of Liverpool, and both experiences were key to developing his thinking about breaking knowledge silos in the newsroom.

We are grateful for the tremendous support we enjoyed from our institutions. At the University of Oregon, Seth was privileged to hold the Shirley Papé Chair in Emerging Media for a decade – an endowed faculty position that funded his research assistants and made possible the conference travels and research collaborations that were crucial for this project. He also received a small grant from the UO's Center for Science Communication Research to support this research. Rodrigo was awarded a year-long sabbatical leave from UMass in 2022 to begin working on a book about AI. And Tomás moved to the University of Wisconsin-Madison

thanks to its newly created RISE AI Initiative, a campus-wide effort created to tackle major societal challenges.

We are grateful, too, to all of our co-authors on related projects who helped us develop the ideas that shaped this one: Rana Arafat, Jon Benedik A. Bunquin, Nicholas Diakopoulos, Maxwell Ely, Lucas Graves, Andrea L. Guzman, Michael Koliska, Bibo Lin, David M. Markowitz, Gregory P. Perreault, Valeria Resendez, Thomas R. Schmidt, Guillén Torres, and Wang Ngai Yeung.

We are especially indebted to our editor, Mary, her editorial assistant, Maddie Hazelden, and the thoughtful reviewers they recruited to assess the book. The team at Polity has been wonderful to work with. Justin Dyer was a brilliant copy editor at the end. We are equally grateful to the many generous friends and colleagues who read and commented on an earlier draft of this manuscript. Their varied expertise and cross-national perspectives led to feedback that proved invaluable in the final stages of production. To each of them, we convey a heartfelt thanks: Joshua Braun, Matt Carlson, David Caswell, Bryan Dunford, Natali Helberger, Heesoo Jang, Bronwyn Jones, Brett Oppegaard, Sena Ozkurt Azimli, John Russial, Sue Robinson, Nadja Schaetz, Theresa J. Seipp, Felix Simon, Mona Sloane, Debora Wenger, and Oscar Westlund.

Our loved ones – family and dear friends – deserve special recognition for listening patiently as we talked through our plans, ideas, and frustrations, and for providing steady encouragement throughout the past 18 months as we worked diligently on this book. Thank you, to each of you.

Finally, it seems only fitting, in a book about AI, to acknowledge our own use of it. As AI models improved over the course of this project, we found them increasingly useful for our authorial needs. We leaned on them – particularly Anthropic’s Claude (Sonnet 4.5, 4.6, and Opus 4.6 models) and, to a lesser extent, Alphabet’s Gemini and OpenAI’s ChatGPT – to help us structure our arguments and identify gaps in our thinking. We used them to suggest phrasing when we struggled to express an idea, and more often as an editorial aid that helped us say what we meant more clearly and concisely. We used them to help crystallize conceptual connections across the work that we describe in Chapter 6. We even used AI to quickly prototype cover designs for this book – something none of us felt otherwise especially equipped

to do – which we then passed on to guide a professional designer. We were, in short, working “beside the machine,” as we describe in Chapter 5. The experience confirmed for us both the power of human–AI partnerships and the indispensability of the human judgment that should guide them.

Notes

Chapter 1 Journalism's Constitutive Moment

- 1 Max Tani, "New York Times Goes All-in on Internal AI Tools," *Semafor*, February 17, 2025.
- 2 Joshua Benton, "The New York Times Will Let Reporters Use AI Tools While Its Lawyers Litigate AI Tools," Nieman Lab, February 18, 2025.
- 3 Imran Rahman-Jones, "AI Chatbots Unable to Accurately Summarise News, BBC Finds," BBC, February 11, 2025.
- 4 Ibid.
- 5 Dade Hayes, "Washington Post Teams With OpenAI To Make 'High-Quality News More Accessible' Via ChatGPT," *Deadline*, April 22, 2025.
- 6 Nic Newman, *Journalism and Technology Trends and Predictions 2026* (Reuters Institute for the Study of Journalism, 2026).
- 7 Isabella Simonetti and Katherine Blunt, "News Sites Are Getting Crushed by Google's New AI Tools," *The Wall Street Journal*, June 10, 2025.
- 8 Charlotte Tobitt, "Wired and Business Insider Remove 'AI-Written' Freelance Articles," *Press Gazette*, August 21, 2025.
- 9 Ibid.
- 10 Felix Simon et al., *Generative AI and News Report 2025: How People Think about AI's Role in Journalism and Society* (Reuters Institute for the Study of Journalism, 2025).
- 11 Ethan Mollick, *Co-Intelligence: Living and Working with AI* (Portfolio, 2024).

- 12 Seth C. Lewis et al., “Generative AI and Its Disruptive Challenge to Journalism: An Institutional Analysis,” *Communication and Change* 1 (2025): 9.
- 13 Andrea L. Guzman and Seth C. Lewis, “Artificial Intelligence and Communication: A Human–Machine Communication Research Agenda,” *New Media & Society* 22, no. 1 (2020): 70–86.
- 14 Felix M. Simon, *Artificial Intelligence in the News: How AI Retools, Rationalizes, and Reshapes Journalism and the Public Arena* (Tow Center for Digital Journalism, 2024).
- 15 Dean Starkman, “The Hamster Wheel,” *Columbia Journalism Review*, September 14, 2010.
- 16 Rena Kim Bivens, “The Internet, Mobile Phones and Blogging,” *Journalism Practice* 2, no. 1 (2008): 113–29.
- 17 Seth C. Lewis and Logan Molyneux, “A Decade of Research on Social Media and Journalism: Assumptions, Blind Spots, and a Way Forward,” *Media and Communication* 6, no. 4 (2018): 11–23; Jorge Vázquez-Herrero et al., “Let’s Dance the News! How the News Media Are Adapting to the Logic of TikTok,” *Journalism* 23, no. 8 (2022): 1717–35.
- 18 Bivens, “The Internet, Mobile Phones and Blogging.”
- 19 Cal Newport, *A World without Email: Reimagining Work in an Age of Communication Overload* (Penguin, 2021).
- 20 Caitlin Petre, *All the News That’s Fit to Click: How Metrics Are Transforming the Work of Journalists* (Princeton University Press, 2021).
- 21 Valérie Bélair-Gagnon et al., eds., *Happiness in Journalism* (Routledge, 2023).
- 22 Karen Hao, *Empire of AI: Dreams and Nightmares in Sam Altman’s OpenAI* (Penguin Press, 2025).
- 23 Mary L. Gray and Siddharth Suri, *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass* (Harper Business, 2019).
- 24 Clark Merrefield, “AI and the News: What Researchers Learned from the AP + the BBC,” *The Journalist’s Resource*, March 4, 2025.
- 25 Abrar Al-Heeti, “AI Slop Has Turned Social Media into an Antisocial Wasteland,” CNET, November 19, 2025.
- 26 “And The Economist’s Word of the Year for 2025 Is . . .,” *The Economist*, December 3, 2025.
- 27 Alan Mathison Turing, “Computing Machinery and Intelligence,” *Mind* LIX, no. 236 (1950): 433–60.
- 28 Norbert Wiener, “Cybernetics,” *Scientific American* 179, no. 5 (1948): 14–19.
- 29 Claude Elwood Shannon, “A Mathematical Theory of Communication,” *Bell System Technical Journal* 27, no. 3 (1948): 379–423.

- 30 John Von Neumann, "First Draft of a Report on the EDVAC," *IEEE Annals of the History of Computing* 15, no. 4 (1993): 27–75.
- 31 Hao, *Empire of AI*.
- 32 Marvin Minsky, *Semantic Information Processing* (MIT Press, 1968).
- 33 Andrew Butterfield and Gerard Ekembe Ngondi, eds., *A Dictionary of Computer Science*, 7th ed. (Oxford University Press, 2016).
- 34 Pei Wang, "On Defining Artificial Intelligence," *Journal of Artificial General Intelligence* 10, no. 2 (2019): 1–37.
- 35 Nicolas Sabouret, *Understanding Artificial Intelligence* (CRC Press, 2021).
- 36 Butterfield and Ngondi, eds., *A Dictionary of Computer Science*.
- 37 JournalismAI, *AI Journalism Starter Pack* (2023), <https://docs.google.com/document/d/1pWwbqPERg0bUbMHMbYYDWmFQmWJYvK8N2Dmbenp4Qu0/edit>.
- 38 Aimee Rinehart and Ernest Kung, *Artificial Intelligence in Local News* (The Associated Press, 2022).
- 39 Rodrigo Zamith, *The American Journalism Handbook: Concepts, Issues, and Skills* (UMass Amherst Libraries, 2022).
- 40 Seth C. Lewis, "Journalism," in *The International Encyclopedia of Journalism Studies*, ed. Tim P. Vos et al. (Wiley, 2019).
- 41 Stephen D. Reese, "The Institution of Journalism: Conceptualizing the Press in a Hybrid Media System," *Digital Journalism* 10, no. 2 (2022): 253–66.
- 42 Ira Chinoy, *Predicting the Winner: The Untold Story of Election Night 1952 and the Dawn of Computer Forecasting* (Potomac Books, 2024).
- 43 Vanessa Richter et al., "Imaginaries of Artificial Intelligence," in *Handbook of Critical Studies of Artificial Intelligence*, ed. Simon Lindgren (Edward Elgar Publishing, 2023).
- 44 Anthony Smith, *Goodbye, Gutenberg: The Newspaper Revolution of the 1980s* (Oxford University Press, 1980).
- 45 Will Mari, "The Pre-History of News-Industry Discourse around Artificial Intelligence," *Emerging Media* 2, no. 3 (2024): 499–522.
- 46 Philip Meyer, *Precision Journalism: A Reporter's Guide to Social Science Methods* (Indiana University Press, 1973).
- 47 Mark Coddington, "Clarifying Journalism's Quantitative Turn," *Digital Journalism* 3, no. 3 (2015): 331–48.
- 48 Wang Ngai Yeung and Tomás Dodds, "Automated Journalism," in *Elgar Encyclopedia of Political Communication*, ed. Alessandro Nai et al. (Edward Elgar Publishing, 2025).
- 49 Matt Carlson, "The Robotic Reporter," *Digital Journalism* 3, no. 3 (2015): 416–31.

- 50 Dylan Byers, “AP to Automate Earnings Stories,” *POLITICO*, June 30, 2014.
- 51 Will Oremus, “The First News Report on the LA Earthquake Was Written by a Robot,” *Slate*, March 17, 2014.
- 52 Tim Adams, “And the Pulitzer Goes To... a Computer,” *The Guardian*, June 28, 2015.
- 53 Carlson, “The Robotic Reporter”; Seth C. Lewis, “Journalism in an Era of Big Data: Cases, Concepts, and Critiques,” *Digital Journalism* 3, no. 3 (2015): 321–30.
- 54 Christer Clerwall, “Enter the Robot Journalist,” *Journalism Practice* 8, no. 5 (2014): 519–31 (525).
- 55 Lucia Moses, “The Washington Post’s Robot Reporter Has Published 850 Articles in the Past Year,” *Digiday*, September 14, 2017.
- 56 Chris Fox, “General Election 2019: How Computers Wrote BBC Election Result Stories,” BBC, December 13, 2019.
- 57 “Uh-Oh, Reuters Has a Robot Writer Who Can Churn Out Earnings Reports,” *Digiday*, March 13, 2018.
- 58 Rasmus Kleis Nielsen and Sarah Anne Ganter, *The Power of Platforms: Shaping Media and Society* (Oxford University Press, 2022).
- 59 Charlie Beckett, *New Powers, New Responsibilities: A Global Survey of Journalism and Artificial Intelligence* (London School of Economics, 2019).
- 60 Jeffrey T. Hancock and Jeremy N. Bailenson, “The Social Impact of Deepfakes,” *Cyberpsychology, Behavior, and Social Networking* 24, no. 3 (2021): 149–52.
- 61 Rachel E. Moran and Sonia Jawaid Shaikh, “Robots in the News and Newsrooms: Unpacking Meta-Journalistic Discourse on the Use of Artificial Intelligence in Journalism,” *Digital Journalism*, June 14, 2022, 1–19.
- 62 Dan Milmo, “Sam Altman Issues ‘Code Red’ at OpenAI as ChatGPT Contends with Rivals,” *The Guardian*, December 2, 2025.
- 63 Damian Radcliffe, *Journalism in the AI Era: Opportunities and Challenges in the Global South and Emerging Economies* (Thomson Reuters Foundation, 2025).
- 64 Frank Landymore, “CNET Is Quietly Publishing Entire Articles Generated by AI,” *Futurism*, January 11, 2023.
- 65 Maggie Harrison Dupré, “USA Today Owner Pauses AI Articles after Butchering Sports Coverage,” *Futurism*, August 29, 2023.
- 66 Maggie Harrison Dupré, “Sports Illustrated Union Says It’s ‘Horrorified’ by Publication of AI-Generated Writers,” *Futurism*, November 28, 2023.

- 67 Kim Björn Becker et al., “Policies in Parallel? A Comparative Study of Journalistic AI Policies in 52 Global News Organisations,” *Digital Journalism* 13, no. 9 (2025): 1578–98.
- 68 Benton, “The New York Times Will Let Reporters Use AI Tools While Its Lawyers Litigate AI Tools.”
- 69 Charlotte Tobitt, “Who’s Suing AI and Who’s Signing: CNN Latest to Sue Perplexity,” *Press Gazette*, May 29, 2026.
- 70 David Caswell, “On the Cusp of Abundance? How AI May Redefine Our Relationship with News,” Reuters Institute for the Study of Journalism, October 13, 2025.
- 71 David Caswell, “Audiences, Automation, and AI: From Structured News to Language Models,” *AI Magazine* 45, no. 2 (2024): 174–86.
- 72 Sabouret, *Understanding Artificial Intelligence*.
- 73 Rodrigo Zamith, “Algorithms and Journalism,” in *Oxford Encyclopedia of Journalism Studies*, ed. Henrik Örnebring et al. (Oxford University Press, 2019).
- 74 Butterfield and Ngondi, eds., *A Dictionary of Computer Science*.
- 75 Emily Öhman, *Introduction to Text Analytics: A Guide for Digital Humanities & Social Sciences* (SAGE, 2024).
- 76 Zamith, “Algorithms and Journalism.”
- 77 Emily M. Bender et al., “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? 🦜,” *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, March 3, 2021, 610–23.
- 78 Öhman, *Introduction to Text Analytics*.
- 79 Sourav Banerjee et al., “LLMs Will Always Hallucinate, and We Need to Live with This,” in *Intelligent Systems and Applications*, ed. Kohei Arai (Springer, 2025).
- 80 Josh Tyrangiel, “America Isn’t Ready for What AI Will Do to Jobs,” *Technology, The Atlantic*, February 10, 2026.
- 81 Taina Bucher, “The Algorithmic Imaginary: Exploring the Ordinary Affects of Facebook Algorithms,” *Information, Communication & Society* 20, no. 1 (2017): 30–44.
- 82 Tarleton Gillespie et al., *Media Technologies: Essays on Communication, Materiality, and Society* (MIT Press, 2014).
- 83 Kevin Roose, “AI Poses ‘Risk of Extinction,’ Industry Leaders Warn,” *The New York Times*, May 30, 2023.
- 84 Tomás Dodds et al., “The AI Turn in Journalism: Disruption, Adaptation, and Democratic Futures,” *Journalism* 27 (June 2026): 530–44.
- 85 Seth C. Lewis et al., “Technological Hype and AI in Journalism: Five Functions and Why They Matter,” *Digital Journalism*, September 15, 2025, 1–12.

- 86 Beckett, *New Powers, New Responsibilities*.
- 87 Neil Thurman et al., *AI Adoption by UK Journalists and Their Newsrooms: Surveying Applications, Approaches, and Attitudes* (Reuters Institute for the Study of Journalism, 2025).
- 88 Ben Steward, “EBU News Report 2025: Leading Newsrooms in the Age of Generative AI,” European Broadcasting Union, April 23, 2025.
- 89 Gartner, “Gartner Hype Cycle Research Methodology,” <https://www.gartner.com/en/research/methodologies/gartner-hype-cycle>.
- 90 Bryan Pfaffenberger, “Technological Dramas,” *Science, Technology, & Human Values* 17, no. 3 (1992): 282–312.
- 91 Zamith, “Algorithms and Journalism.”
- 92 Alfred Hermida, “Twittering the News: The Emergence of Ambient Journalism,” *Journalism Practice* 4, no. 3 (2010): 297–308.
- 93 Rich Ling, *The Mobile Connection: The Cell Phone’s Impact on Society* (Morgan Kaufmann, 2004).
- 94 Trevor J. Pinch and Wiebe E. Bijker, “The Social Construction of Facts and Artefacts: Or How the Sociology of Science and the Sociology of Technology Might Benefit Each Other,” *Social Studies of Science* 14, no. 3 (1984): 399–441.
- 95 Simon, *Artificial Intelligence in the News*.
- 96 Jack Myers, “Technological Determinism Theory: Media as a Driver of Social Change,” in *Media Ecology for the 21st Century: Theories of Culture, Communications, and Consciousness* (Palgrave Macmillan, 2025).

Chapter 2 The Newsroom Accelerates

- 1 Jonathan Gerhardson, “Reporter Bots Come to Western Mass,” *The Shoestring*, November 19, 2025.
- 2 MassLive.com Staff, “‘We’re Coming for You’: Officials Claim Mass. Woman Threatened to Kill ICE Agents,” News, MassLive, October 16, 2025.
- 3 Gerhardson, “Reporter Bots Come to Western Mass.”
- 4 Ibid.
- 5 Andrew Deck, “How This Year’s Pulitzer Awardees Used AI in Their Reporting,” Nieman Lab, May 22, 2025.
- 6 Felix M. Simon, *Artificial Intelligence in the News: How AI Retools, Rationalizes, and Reshapes Journalism and the Public Arena* (Tow Center for Digital Journalism, 2024).
- 7 Rowan Philip, “How Newsrooms Are Using AI Chatbots to Leverage Their Own Reporting – and Build Trust,” Global Investigative Journalism Network, August 7, 2025.

- 8 Seth C. Lewis et al., “Generative AI and Its Disruptive Challenge to Journalism: An Institutional Analysis,” *Communication and Change* 1 (2025): 9.
- 9 David Caswell, “Producing News in the Age of Artificial Intelligence,” *MedienWirtschaft* 2022, no. 3 (2022): 12–15.
- 10 Jane B. Singer, “The Political J-Blogger: ‘Normalizing’ a New Media Form to Fit Old Norms and Practices,” *Journalism* 6, no. 2 (2005): 173–98.
- 11 Fabia Ioscote et al., “Artificial Intelligence in Journalism: A Ten-Year Retrospective of Scientific Articles (2014–2023),” *Journalism and Media* 5, no. 3 (2024): 873–91.
- 12 J. E. (Hans) Korteling et al., “Human- versus Artificial Intelligence,” *Frontiers in Artificial Intelligence* 4 (March 2021): 622364.
- 13 Mark Deuze and Charlie Beckett, “Imagination, Algorithms and News: Developing AI Literacy for Journalism,” *Digital Journalism* 10, no. 10 (2022): 1913–18.
- 14 Wiebke Loosen et al., “‘X Journalism’: Exploring Journalism’s Diverse Meanings through the Names We Give It,” *Journalism* 23, no. 1 (2022): 39–58.
- 15 Dilruba Showkat and Eric P. S. Baumer, “Where Do Stories Come From? Examining the Exploration Process in Investigative Data Journalism,” *Proceedings of the ACM on Human–Computer Interaction* 5, no. CSCW2 (2021): 390.
- 16 Tony Harcup and Deirdre O’Neill, “What Is News?,” *Journalism Studies* 18, no. 12 (2017): 1470–88.
- 17 David M. Ryfe, “Structure, Agency, and Change in an American Newsroom,” *Journalism* 10, no. 5 (2009): 665–83.
- 18 Mark Deuze, “On Creativity,” *Journalism* 20, no. 1 (2019): 130–4.
- 19 Mike Ananny and Matt Pearce, “How We’re Using AI,” *Columbia Journalism Review*, May 12, 2025.
- 20 Suzanne Franks et al., “Using Computational Tools to Support Journalists’ Creativity,” *Journalism* 23, no. 9 (2022): 1881–99.
- 21 Emma Gometz, “Key Questions for Journalists to Consider before Using Generative AI,” *The Open Notebook*, February 18, 2025.
- 22 Paul Bradshaw, “Investigative Journalism and ChatGPT: Using Generative AI for Story Ideas,” *Online Journalism Blog*, July 10, 2024.
- 23 Aynur Sarisakaloğlu, “Artificial Intelligence as a Game-Changer or Passing Tech Fad in Journalism? A Qualitative Study on Journalists’ Perceptions and Technology Adoption in a Global South Country,” *Journalism*, October 21, 2025: 1–21.
- 24 Gilles Fauconnier and Mark Turner, *The Way We Think: Conceptual Blending and the Mind’s Hidden Complexities* (Basic Books, 2003).

- 25 Matthew Bird-Meyer et al., “The Role of Serendipity in the Story Ideation Process of Print Media Journalists,” *Journal of Documentation* 75, no. 5 (2019): 995–1012.
- 26 Franks et al., “Using Computational Tools to Support Journalists’ Creativity.”
- 27 Lynge Asbjørn Møller, “Designing Algorithmic Editors: How Newspapers Embed and Encode Journalistic Values into News Recommender Systems,” *Digital Journalism* 12, no. 7 (2024): 926–44.
- 28 Sarisakaloğlu, “Artificial Intelligence as a Game-Changer or Passing Tech Fad in Journalism?”
- 29 Dawn Wheatley, “A Typology of News Sourcing: Routine and Non-Routine Channels of Production,” *Journalism Practice* 14, no. 3 (2020): 277–98.
- 30 Morten Hertzum, “How Do Journalists Seek Information from Sources? A Systematic Review,” *Information Processing & Management* 59, no. 6 (2022): 103087; Ville J. E. Manninen, “If Only They Knew: Audience Expectations and Actual Sourcing Practices in Online Journalism,” *Journalism Practice* 14, no. 4 (2020): 412–28.
- 31 Hertzum, “How Do Journalists Seek Information from Sources?”
- 32 Dhrumil Mehta et al., “UNHEARD: Using Large Language Models to Conduct Source Audits of the News,” *Columbia Journalism Review*, February 18, 2025.
- 33 Alexander Spangher et al., “Identifying Informational Sources in News Articles,” in *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, ed. Houda Bouamor et al. (Association for Computational Linguistics, 2023).
- 34 Ananny and Pearce, “How We’re Using AI.”
- 35 United Robots, “High End Automated Content in Any Team Sport,” <https://www.unitedrobots.ai/content-services/sports>.
- 36 Ruby Voge, “Journalism Will Always Be about People – AI Can’t Change That,” Editorial, *The Daily Free Press*, February 28, 2025.
- 37 Rachel Treisman, “A Magazine Touted Michael Schumacher’s First Interview in Years. It Was Actually AI,” NPR, April 28, 2023.
- 38 Jonathan Noble, “Schumacher Family Wins Legal Action over Fake AI Interview,” *Autosport*, May 23, 2024.
- 39 Drew Harwell, “Jim Acosta Sparks Fury with ‘Interview’ of Dead Parkland Teen’s AI Avatar,” *The Washington Post*, August 5, 2025.
- 40 Tomasz Hollanek and Katarzyna Nowaczyk-Basińska, “Griefbots, Deadbots, Postmortem Avatars: On Responsible Applications of Generative AI in the Digital Afterlife Industry,” *Philosophy & Technology* 37, no. 2 (2024): 63.

- 41 Neil Thurman et al., *AI Adoption by UK Journalists and Their Newsrooms: Surveying Applications, Approaches, and Attitudes* (Reuters Institute for the Study of Journalism, 2025).
- 42 Garance Burke and Hilke Schellmann, “Researchers Say an AI-Powered Transcription Tool Used in Hospitals Invents Things No One Ever Said,” *Technology, AP News*, October 26, 2024.
- 43 Ananny and Pearce, “How We’re Using AI.”
- 44 The Associated Press, “Artificial Intelligence in Local News Case Study: Michigan Radio,” <https://www.ap.org/wp-content/uploads/2024/02/ap-local-ai-michigan-radio-oct-2023.pdf>.
- 45 Hannes Cools and Nicholas Diakopoulos, “Uses of Generative AI in the Newsroom: Mapping Journalists’ Perceptions of Perils and Possibilities,” *Journalism Practice* 20, no. 3 (2026): 878–96.
- 46 Sue Robinson, *How Journalists Engage: A Theory of Trust Building, Identities, & Care* (Oxford University Press, 2023).
- 47 Lucas Graves and Michelle A. Amazeen, “Fact-Checking as Idea and Practice in Journalism,” in *Oxford Research Encyclopedia of Communication*, ed. Jon F. Nussbaum (Oxford University Press, 2019).
- 48 Rachel E. Moran, “The So-Called ‘Crisis’ of Trust in Journalism,” in *The Routledge Companion to News and Journalism*, 2nd ed., ed. Stuart Allan (Routledge, 2022).
- 49 Matt Carlson, “The Information Politics of Journalism in a Post-Truth Age,” *Journalism Studies* 19, no. 13 (2018): 1879–88.
- 50 Sanha Lim, “A Better ClaimReview to Grow a Global Fact-Check Database,” *Reporters’ Lab*, April 18, 2019.
- 51 Andrew Dudfield, “The Web Just Got a Little Harder to Trust,” *Full Fact*, June 26, 2025.
- 52 Ethan Porter and Thomas J. Wood, “The Global Effectiveness of Fact-Checking: Evidence from Simultaneous Experiments in Argentina, Nigeria, South Africa, and the United Kingdom,” *Proceedings of the National Academy of Sciences* 118, no. 37 (2021): e2104235118.
- 53 Graves and Amazeen, “Fact-Checking as Idea and Practice in Journalism.”
- 54 Kevin K. King and Bin Wang, “Diffusion of Real versus Misinformation during a Crisis Event: A Big Data-Driven Approach,” *International Journal of Information Management* 71 (August 2023): 102390.
- 55 Henk van Ess, “Reporter’s Guide to Detecting AI-Generated Content,” *Global Investigative Journalism Network*, September 1, 2025.

- 56 Lucas Graves, “Anatomy of a Fact Check: Objective Practice and the Contested Epistemology of Fact Checking,” *Communication, Culture and Critique* 10, no. 3 (2017): 518–37.
- 57 Manish Kumar Singh et al., “A Comprehensive Review on Automatic Detection of Fake News on Social Media,” *Multimedia Tools and Applications* 83, no. 16 (2024): 47319–52.
- 58 Full Fact and Fundación Maldita.es, “Prebunking at Scale: The Next Chapter for European Fact Checking,” November 5, 2025.
- 59 Kate Wilkinson, “Keeping Pace with Misinformation: Why We Built an AI Team,” Full Fact, September 22, 2025.
- 60 Daniela Mahl et al., “‘We Follow the Disinformation’: Conceptualizing and Analyzing Fact-Checking Cultures across Countries,” *The International Journal of Press/Politics* 31, no. 2 (2026): 264–90.
- 61 Nestor Maslej et al., “Artificial Intelligence Index Report 2025,” arXiv:2504.07139, preprint, *arXiv*, November 20, 2025.
- 62 Isabelle Augenstein et al., “Factuality Challenges in the Era of Large Language Models and Opportunities for Fact-Checking,” *Nature Machine Intelligence* 6, no. 8 (2024): 852–63.
- 63 Houjiang Liu et al., “Human-Centered NLP Fact-Checking: Co-Designing with Fact-Checkers Using Matchmaking for AI,” *Proceedings of the ACM on Human–Computer Interaction* 8, no. CSCW2 (2024): 423.
- 64 Graham Fraser, “BBC Complains to Apple over Misleading Shooting Headline,” BBC, December 13, 2024.
- 65 Tripp Mickle, “Apple Plans to Disable AI Features Summarizing News Notifications,” *The New York Times*, January 16, 2025.
- 66 Rodrigo Zamith, *The American Journalism Handbook: Concepts, Issues, and Skills* (UMass Amherst Libraries, 2022).
- 67 Luba Kassova, “21st Century Journalism Needs Emotions to Survive,” *Journalism UK*, October 1, 2025.
- 68 Nicholas Diakopoulos, *Automating the News: How Algorithms Are Rewriting the Media* (Harvard University Press, 2019).
- 69 Christer Clerwall, “Enter the Robot Journalist,” *Journalism Practice* 8, no. 5 (2014): 519–31.
- 70 Kerry Flynn and Sara Fischer, “Time Launches New AI Agent,” *Axios*, November 10, 2025.
- 71 Ibid.
- 72 Jodie Hopperton, “What Is Vibe Coding and Why Should News Organisations Care?,” International News Media Association, November 30, 2025.
- 73 Joe Amditis, “Vibe Coding for Newsrooms,” *Medium*, December 1, 2025.

- 74 Maggie Harrison Dupré, “Sports Illustrated Published Articles by Fake, AI-Generated Writers,” *Futurism*, November 27, 2023.
- 75 David Bauder, “Sports Illustrated Is the Latest Media Company Damaged by an AI Experiment Gone Wrong,” The Associated Press, November 28, 2023.
- 76 Andrew Deck, “Florida Nonprofit News Reporters Ask Board to Investigate Their Editor’s AI Use,” Nieman Lab, November 19, 2025.
- 77 Joshua A. Braun, *This Program Is Brought to You By . . .* (Yale University Press, 2015).
- 78 T. Dodds Rojas, “Newsroom Dissonance: How New Digital Technologies Are Changing Professional Roles in Contemporary Newsrooms” (doctoral thesis, Leiden University, 2022), <https://hdl.handle.net/1887/3270873>.
- 79 Tomás Dodds et al., “Popularity-Driven Metrics: Audience Analytics and Shifting Opinion Power to Digital Platforms,” *Journalism Studies* 24, no. 3 (2023): 403–21.
- 80 Logan Molyneux et al., “Media Work, Identity, and the Motivations That Shape Branding Practices among Journalists: An Explanatory Framework,” *New Media & Society* 21, no. 4 (2019): 836–55.
- 81 Nic Newman and Federica Cherubini, *Journalism and Technology Trends and Predictions 2025* (Reuters Institute, 2025).
- 82 Michael Savage, “BBC News to Create AI Department to Offer More Personalised Content,” *The Guardian*, March 6, 2025.
- 83 Emily Roseman and Ta Ha, “How Nonprofit News Outlets Are Using AI to Save Time and Money,” *INNights*, August 15, 2024.
- 84 Katie Robertson, “Bloomberg Has a Rocky Start with AI Summaries,” *The New York Times*, March 29, 2025.
- 85 Andrew Deck, “A Pressure Test for AI: Dow Jones Makes a Translation Push for Real-Time Financial News,” Nieman Lab, July 17, 2025, <https://www.niemanlab.org/2025/07/a-pressure-test-for-ai-dow-jones-makes-a-translation-push-for-real-time-financial-news/>.
- 86 PJ O’Keefe, “Here’s How WCPO 9 Is Using Artificial Intelligence in Our Journalism,” WCPO 9 Cincinnati, June 19, 2025.
- 87 Ananny and Pearce, “How We’re Using AI.”
- 88 Seth C. Lewis et al., “Libel by Algorithm? Automated Journalism and the Threat of Legal Liability,” *Journalism & Mass Communication Quarterly* 96, no. 1 (2019): 60–81.
- 89 Karen Hao, *Empire of AI: Dreams and Nightmares in Sam Altman’s OpenAI* (Penguin Press, 2025).
- 90 Joe Mandese, “An AI Journalistic Code of Ethics (Proposed by an AI),” *MediaPost*, November 17, 2025.

- 91 Nicole Meir, “Standards around Generative AI,” The Associated Press, August 15, 2023.
- 92 The New York Times Company, “Principles for Using Generative AI in The Times’s Newsroom,” Press Room, May 9, 2024.
- 93 Martin Schori, “Aftonbladet’s AI Policy,” *Aftonbladet*, August 19, 2025.
- 94 Mathias-Felipe de-Lima-Santos et al., “Guiding the Way: A Comprehensive Examination of AI Guidelines in Global Media,” *AI & Society* 40 (2025): 2585–603.
- 95 Damian Radcliffe, *Journalism in the AI Era: Opportunities and Challenges in the Global South and Emerging Economies* (Thomson Reuters Foundation, 2025).
- 96 Colin Porlezza and Aljosha Karim Schapals, “AI Ethics in Journalism (Studies): An Evolving Field between Research and Practice,” *Emerging Media* 2, no. 3 (2024): 356–70.
- 97 Tomás Dodds et al., “Knowledge Silos as a Barrier to Responsible AI Practices in Journalism? Exploratory Evidence from Four Dutch News Organisations,” *Journalism Studies* 26, no. 6 (2025): 740–58.
- 98 Ibid.
- 99 Marco Valeri, “AI and Knowledge Management in Modern Organizations,” in *Organizations and Knowledge Management*, ed. Marco Valeri (Springer, 2025).
- 100 Álvaro Villagrán Sánchez and Fernando López Pan, “Journalism, the Sociology of Knowledge, and Science Studies: A Narrative Literature Review,” *Journalism*, July 3, 2025: 1–18.
- 101 Dodds et al., “Knowledge Silos as a Barrier to Responsible AI Practices in Journalism?”
- 102 Raphael Lachkar and Chatrine Siswoyo, *AI and Journalism in Southeast Asia: A Survey of Opportunities and Challenges* (Vero, 2024).
- 103 Dodds et al., “Knowledge Silos as a Barrier to Responsible AI Practices in Journalism?”
- 104 Ibid.
- 105 Simangele Mkhize and Melanie Lourens, “Elevating Knowledge Sharing and Communication through Artificial Intelligence: An Organizational Perspective,” *International Journal of Research in Business and Social Science* 14, no. 4 (2025): 103–14.
- 106 Ana Cecilia Bisso Nunes et al., “Sandbox Journalism: The Role of Media Labs as Innovation Drivers in European News Organisations,” *Journalism Studies* 27, no. 4 (2026): 460–81.
- 107 Caswell, “Producing News in the Age of Artificial Intelligence.”
- 108 Ibid.

- 109 Nick Couldry and Joseph Turow, “Advertising, Big Data and the Clearance of the Public Realm: Marketers’ New Approaches to the Content Subsidy,” *International Journal of Communication* 8 (2014): 1710–26.

Chapter 3 The New Architecture of Media Power

- 1 Bobby Allyn and John Ruwitch, “Online News Publishers Face ‘Extinction-Level Event’ from Google’s AI-Powered Search,” NPR, July 31, 2025.
- 2 Joshua A. Braun, “News Distribution,” in *Oxford Research Encyclopedia of Communication*, ed. Jon F. Nussbaum (Oxford University Press, 2019).
- 3 Athena Chapekis and Anna Lieb, “Google Users Are Less Likely to Click on Links When an AI Summary Appears in the Results,” Pew Research Center, July 22, 2025.
- 4 Isabella Simonetti and Katherine Blunt, “News Sites Are Getting Crushed by Google’s New AI Tools,” *The Wall Street Journal*, June 10, 2025.
- 5 Nic Newman, *Journalism and Technology Trends and Predictions 2026* (Reuters Institute for the Study of Journalism, 2026).
- 6 Nilay Patel, “Google Zero Is Here – Now What?,” *The Verge*, May 30, 2024.
- 7 Chapekis and Lieb, “Google Users Are Less Likely to Click on Links When an AI Summary Appears in the Results.”
- 8 Felix M. Simon, “Escape Me If You Can: How AI Reshapes News Organisations’ Dependency on Platform Companies,” *Digital Journalism* 12, no. 2 (2024): 149–70.
- 9 Felix M. Simon, “Uneasy Bedfellows: AI in the News, Platform Companies and the Issue of Journalistic Autonomy,” *Digital Journalism* 10, no. 10 (2022): 1832–54.
- 10 Simonetti and Blunt, “News Sites Are Getting Crushed by Google’s New AI Tools.”
- 11 Theresa Josephine Seipp et al., “Between the Cracks: Blind Spots in Regulating Media Concentration and Platform Dependence in the EU,” *Internet Policy Review* 13, no. 4 (2024): 1–25.
- 12 Ben Bagdikian, “The 50, 26, 20. . . Corporations That Own Our Media,” *Extra!*, June 1987.
- 13 Theresa Josephine Seipp et al., “Dealing with Opinion Power in the Platform World: Why We Really Have to Rethink Media Concentration Law,” *Digital Journalism* 11, no. 8 (2023): 1542–67; Fernando van der Vlist et al., “Big AI: Cloud Infrastructure Dependence and the Industrialisation of Artificial Intelligence,” *Big Data & Society* 11, no. 1 (2024): 1–16.

- 14 Max Weber, *Economy and Society: An Outline of Interpretive Sociology*, ed. Guenther Roth and Claus Wittich (University of California Press, 1978), 53.
- 15 Ibid.
- 16 Natali Helberger, “The Political Power of Platforms: How Current Attempts to Regulate Misinformation Amplify Opinion Power,” *Digital Journalism* 8, no. 6 (2020): 842–54.
- 17 Matt Carlson, *Journalistic Authority: Legitimizing News in the Digital Era* (Columbia University Press, 2017).
- 18 Timothy E. Cook, *Governing with the News: The News Media as a Political Institution* (University of Chicago Press, 1998).
- 19 Jacob L. Nelson et al., “‘Money Is the Root of All Evil.’ How the Business of Journalism Shapes Trust in News,” *Journalism* 26, no. 2 (2025): 251–68.
- 20 Matt Carlson et al., *News after Trump: Journalism’s Crisis of Relevance in a Changed Media Culture* (Oxford University Press, 2022).
- 21 Ben H. Bagdikian, *The Media Monopoly* (Beacon Press, 1983).
- 22 Robert Waterman McChesney, *The Problem of the Media: US Communication Politics in the Twenty-First Century* (Monthly Review Press, 2004).
- 23 Cook, *Governing with the News*.
- 24 Bagdikian, *The Media Monopoly*.
- 25 Edward S. Herman and Noam Chomsky, *Manufacturing Consent: The Political Economy of the Mass Media* (Pantheon Books, 1988).
- 26 Philip M. Napoli, “Deconstructing the Diversity Principle,” *Journal of Communication* 49, no. 4 (1999): 7–34.
- 27 Pamela J. Shoemaker and Stephen D. Reese, *Mediating the Message: Theories of Influences on Mass Media Content*, 2nd ed. (Longman, 1996).
- 28 Pierre Bourdieu, “The Political Field, the Social Science Field and the Journalistic Field,” in *Bourdieu and the Journalistic Field*, ed. Rodney Benson and Erik Neveu (Polity, 2005).
- 29 Henrik Örnebring and Michael Karlsson, *Journalistic Autonomy: The Genealogy of a Concept* (University of Missouri Press, 2022).
- 30 Mark Coddington, “The Wall Becomes a Curtain,” in *Boundaries of Journalism*, ed. Matt Carlson and Seth C. Lewis (Routledge, 2015).
- 31 Carlson, *Journalistic Authority*.
- 32 John V. Pavlik, *Disruption and Digital Journalism: Assessing News Media Innovation in a Time of Dramatic Change* (Routledge, 2021).
- 33 Emily J. Bell et al., *The Platform Press: How Silicon Valley Reengineered Journalism* (Tow Center for Journalism, 2017).

- 34 Efrat Nechushtai, “Could Digital Platforms Capture the Media through Infrastructure?,” *Journalism* 19, no. 8 (2018): 1043–58.
- 35 Murray Dick, “Search Engine Optimisation in UK News Production,” *Journalism Practice* 5, no. 4 (2011): 462–77.
- 36 Caitlin Petre, *All the News That’s Fit to Click: How Metrics Are Transforming the Work of Journalists* (Princeton University Press, 2021).
- 37 Rasmus Kleis Nielsen and Sarah Anne Ganter, *The Power of Platforms: Shaping Media and Society* (Oxford University Press, 2022).
- 38 Nechushtai, “Could Digital Platforms Capture the Media through Infrastructure?”
- 39 Alfred Hermida, “Social Media and the News,” in *The SAGE Handbook of Digital Journalism*, ed. Tamara Witschge et al. (SAGE Publications Ltd., 2016).
- 40 Tomás Dodds et al., “Popularity-Driven Metrics: Audience Analytics and Shifting Opinion Power to Digital Platforms,” *Journalism Studies* 24, no. 3 (2023): 403–21.
- 41 Rodrigo Zamith, “Quantified Audiences in News Production: A Synthesis and Research Agenda,” *Digital Journalism* 6, no. 4 (2018): 418–35.
- 42 Rodrigo Zamith and Joshua A. Braun, “Technology and Journalism,” in *The International Encyclopedia of Journalism Studies*, ed. Tim P. Vos et al. (Wiley, 2019).
- 43 Merja Myllylahti, “Paying Attention to Attention: A Conceptual Framework for Studying News Reader Revenue Models Related to Platforms,” *Digital Journalism* 8, no. 5 (2020): 567–75.
- 44 Rasmus Kleis Nielsen, “Do News Media Want to Disentangle from Platforms?,” *Digital Journalism*, February 1, 2026: 1–9 (1, original emphasis).
- 45 Nechushtai, “Could Digital Platforms Capture the Media through Infrastructure?”
- 46 Nielsen and Ganter, *The Power of Platforms*.
- 47 Cook, *Governing with the News*.
- 48 Helberger, “The Political Power of Platforms.”
- 49 Nechushtai, “Could Digital Platforms Capture the Media through Infrastructure?”
- 50 Tim Dwyer et al., “Where to Next with Australia’s News Media and Digital Platforms Mandatory Bargaining Code?,” *Communications* 48, no. 3 (2023): 440–56.
- 51 Richard R. John, *Network Nation: Inventing American Telecommunications* (Harvard University Press, 2010).
- 52 Nielsen, “Do News Media Want to Disentangle from Platforms?”

- 53 Stefan ten Nijmegen, “Google’s AI Overview Is Not Yet Harming Traffic, but Publishers Should Remain Alert,” International News Media Association, August 11, 2025.
- 54 Bell et al., *The Platform Press*.
- 55 Van der Vlist et al., “Big AI.”
- 56 Dieuwertje Luitse, “Platform Power in AI: The Evolution of Cloud Infrastructures in the Political Economy of Artificial Intelligence,” *Internet Policy Review* 13, no. 2 (2024): 1–44.
- 57 Kate Crawford, *The Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence* (Yale University Press, 2021).
- 58 Van der Vlist et al., “Big AI.”
- 59 Anna Gerbrandy and Pauline Phoa, “The Power of Big Tech Corporations as Modern Bigness and a Vocabulary for Shaping Competition Law as Counter-Power,” in *Wealth and Power*, ed. Michael Bennett et al. (Routledge, 2022).
- 60 Deepa Seetharaman et al., “Microsoft, OpenAI Reach Deal Removing Fundraising Constraints for ChatGPT Maker,” Reuters, October 28, 2025.
- 61 Jennifer Elias, “How Google Put Together the Pieces for Its AI Comeback,” Technology, CNBC, November 27, 2025.
- 62 Marcus Law, “AWS Remains \$330bn Cloud Market Leader, Driven by AI Growth,” *Technology Magazine*, February 17, 2025.
- 63 Helle Sjøvaag et al., “Capture beyond the Platforms: The Material and Infrastructural Conditions for Digital Journalism,” *Digital Journalism* 13, no. 7 (2025): 1331–50.
- 64 Katie Duffy, “A Complete List of Publishers and Their AI Licensing Deals,” *Futureweek*, September 4, 2025.
- 65 Coddington, “The Wall Becomes a Curtain.”
- 66 Herman and Chomsky, *Manufacturing Consent*.
- 67 Dodds et al., “Popularity-Driven Metrics.”
- 68 Nielsen, “Do News Media Want to Disentangle from Platforms?”
- 69 Emma Roth, “Microsoft and OpenAI Are Giving News Outlets \$10 Million to Use AI Tools,” *The Verge*, October 22, 2024.
- 70 Felix Simon, “We Need Clarity about the Deals between AI Companies and News Publishers. Here’s Why,” Reuters Institute for the Study of Journalism, October 10, 2024.
- 71 Charlotte Tobitt, “Who’s Suing AI and Who’s Signing: CNN Latest to Sue Perplexity,” *Press Gazette*, May 29, 2026.
- 72 Gannett Co., Inc., “USA TODAY Network and Perplexity Announce Strategic AI Content Licensing Agreement,” *Business Wire*, July 30, 2025.
- 73 Matt O’Brien, “ChatGPT-Maker OpenAI Signs Deal with AP to License News Stories,” The Associated Press, July 13, 2023.

- 74 Chris Yoko, “AI’s ‘Thoughts’ on AI Stealing Journalism,” Chris Yoko, September 17, 2025.
- 75 Simon, “We Need Clarity about the Deals between AI Companies and News Publishers.”
- 76 Ashley Capoot, “The New York Times Sues Perplexity, Alleging Copyright Infringement,” CNBC, December 5, 2025.
- 77 Tobitt, “Who’s Suing AI and Who’s Signing.”
- 78 Neil Thurman et al., “Journalists May See AI as a Threat to the Industry, but They’re Using It Anyway,” Nieman Lab, December 2, 2025.
- 79 Andrew Deck, “How This Year’s Pulitzer Awardees Used AI in Their Reporting,” Nieman Lab, May 22, 2025.
- 80 Muck Rack, *The State of Work–Life Balance in Journalism: A Study on Stress, Burnout and Support in 2025* (Muck Rack, 2025).
- 81 Reporters Without Borders, “Trump’s Foreign Aid Freeze Throws Journalism around the World into Chaos,” February 3, 2025.
- 82 Suzanne Blake, “One of the Most Popular College Majors Is Seeing Salaries Decline,” *Newsweek*, April 14, 2025.
- 83 Nick Couldry and Ulises A. Mejias, “Data Colonialism: Rethinking Big Data’s Relation to the Contemporary Subject,” *Television & New Media* 20, no. 4 (2019): 336–49.
- 84 Karen Hao, *Empire of AI: Dreams and Nightmares in Sam Altman’s OpenAI* (Penguin Press, 2025).
- 85 Gregory Gondwe, “Artificial Intelligence, Journalism, and the Ubuntu Robot in Sub-Saharan Africa: Towards a Normative Framework,” *Digital Journalism* 13, no. 4 (2025): 826–44.
- 86 Allen Munoriyarwa et al., “The Philanthrocapitalism of Google News Initiative in Africa, Latin America, and the Middle East – Empirical Reflections,” *International Journal of Cultural Studies* 28, no. 1 (2025): 56–79.
- 87 Freddie Hiney, “Developing AI for Africa: Why Digital Sovereignty Matters and Who Is Leading the Charge,” *Forbes Africa*, August 31, 2025.
- 88 Eduardo Levy Yeyati, “Latam-GPT and the Search for AI Sovereignty,” The Brookings Institution, November 25, 2025.
- 89 Jo Pierson et al., “Governing Artificial Intelligence in the Media and Communications Sector,” *Internet Policy Review* 12, no. 1 (2023): 1–28.
- 90 Ugur Aytac, “Big Tech, Algorithmic Power, and Democratic Control,” *The Journal of Politics* 86, no. 4 (2024): 1431–45.
- 91 Christoph Busch, *Regulation of Digital Platforms as Infrastructures for Services of General Interest*, no. 09/2021, WISO Diskurs (Friedrich-Ebert-Stiftung, 2021), 20.

- 92 Kena Zheng, “Antitrust in Artificial Intelligence Infrastructure – between Regulation and Innovation in the EU, the US, and China,” *Computer Law & Security Review* 59 (November 2025): 106211.
- 93 Ibid.
- 94 Donald J. Trump, “Removing Barriers to American Leadership in Artificial Intelligence,” The White House, January 23, 2025.
- 95 Jesse Bedayn, “What to Know about Trump’s Draft Proposal to Curtail State AI Regulations,” The Associated Press, November 21, 2025.
- 96 Zheng, “Antitrust in Artificial Intelligence Infrastructure.”
- 97 Sara Migliorini, “China’s Interim Measures on Generative AI: Origin, Content and Significance,” *Computer Law & Security Review* 53 (July 2024): 105985.
- 98 Ibid.
- 99 Huw Roberts et al., “The Chinese Approach to Artificial Intelligence: An Analysis of Policy, Ethics, and Regulation,” *AI & Society* 36 (2021): 59–77.
- 100 Joanne Kuai et al., “AI ≥ Journalism: How the Chinese Copyright Law Protects Tech Giants’ AI Innovations and Disrupts the Journalistic Institution,” *Digital Journalism* 10, no. 10 (2022): 1893–912; Zheng, “Antitrust in Artificial Intelligence Infrastructure.”
- 101 Jingyan Elaine Yuan and Lin Zhang, “From Platform Capitalism to Digital China: The Path, Governance, and Geopolitics,” *Social Media + Society* 11, no. 1 (2025): 1–9.
- 102 Ibid.; Zheng, “Antitrust in Artificial Intelligence Infrastructure.”
- 103 Joanne Kuai, “Unravelling Copyright Dilemma of AI-Generated News and Its Implications for the Institution of Journalism: The Cases of US, EU, and China,” *New Media & Society* 26, no. 9 (2024): 5150–68.
- 104 Xuenan Cao and Lucas Power, “Controlled Carefully: How Consumer Care Legitimizes China’s AI Regulations,” *Information, Communication & Society* 28, no. 15 (2025): 2740–56.
- 105 Natali Helberger, “FutureNewsCorp, or How the AI Act Changed the Future of News,” *Computer Law & Security Review* 52 (April 2024): 105915; Seipp et al., “Between the Cracks.”
- 106 Seipp et al., “Between the Cracks”: 1.
- 107 Ibid.: 10.
- 108 Helberger, “FutureNewsCorp, or How the AI Act Changed the Future of News”: 9, original emphasis.
- 109 Seipp et al., “Between the Cracks”: 13.
- 110 Aytac, “Big Tech, Algorithmic Power, and Democratic Control”; Seipp et al., “Between the Cracks.”

- 111 Fernando Filgueiras, “Artificial Intelligence and Governance Challenges in Latin America – the Game between Decolonization and Dependence,” in *Handbook on Governance and Data Science*, ed. Sarah Giest et al. (Edward Elgar Publishing, 2025).
- 112 Seipp et al., “Between the Cracks”: 16.
- 113 José van Dijck et al., “Reframing Platform Power,” *Internet Policy Review* 8, no. 2 (2019): 1–18 (4–5, original emphasis).

Chapter 4 Designing AI for Democracy

- 1 Donald J. Trump, Truth Social, November 18, 2025, <https://truthsocial.com/@realDonaldTrump/115573090200730933>.
- 2 Blake Montgomery and Nick Robins-Early, “Nvidia Becomes World’s First \$5tn Company amid Stock Market and AI Boom,” *The Guardian*, October 29, 2025.
- 3 Krystal Hu, “OpenAI Hits \$500 Billion Valuation after Share Sale to SoftBank, Others, Source Says,” Technology, Reuters, October 2, 2025.
- 4 George Steer, “Big Tech’s Market Dominance Is Becoming Ever More Extreme,” *Financial Times*, October 31, 2025.
- 5 Gerrit De Vynck and Nitasha Tiku, “White House Drafts Order Directing Justice Department to Sue States That Pass AI Regulations,” *The Washington Post*, November 19, 2025.
- 6 Jesse Barron, “A Teen in Love with a Chatbot Killed Himself. Can the Chatbot Be Held Responsible?,” *The New York Times Magazine*, October 24, 2025.
- 7 Wendy Lee, “Universal Music Group Settles with AI Music Startup Udio,” Hollywood Inc., *Los Angeles Times*, October 30, 2025.
- 8 Sara Fischer, “News Publishers Sue Cohere for Copyright and Trademark Infringement,” *Axios*, February 13, 2025.
- 9 Nilay Patel, “Google Zero Is Here – Now What?,” *The Verge*, May 30, 2024.
- 10 Karen Hao, *Empire of AI: Dreams and Nightmares in Sam Altman’s OpenAI* (Penguin Press, 2025).
- 11 Thomas E. Patterson, *Doing Well and Doing Good* (Shorenstein Center on Media, Politics and Public Policy, 2000).
- 12 Christian Fuchs, *Digital Democracy and the Digital Public Sphere* (Routledge, 2023), 12.
- 13 Mark Coeckelbergh, “Democracy as Communication: Towards a Normative Framework for Evaluating Digital Technologies,” *Contemporary Pragmatism* 21, no. 2 (2024): 217–35 (219, original emphasis).

- 14 Jürgen Habermas, *The Structural Transformation of the Public Sphere: An Inquiry into a Category of Bourgeois Society*, trans. Thomas Burger (MIT Press, 1989).
- 15 Zizi Papacharissi, “The Virtual Sphere 2.0: The Internet, the Public Sphere and Beyond,” in *Routledge Handbook of Internet Politics*, ed. Andrew Chadwick and Philip N. Howard (Routledge, 2008).
- 16 Andreas Jungherr and Ralph Schroeder, “Artificial Intelligence and the Public Arena,” *Communication Theory* 33, nos. 2–3 (2023): 164–73.
- 17 Daniel Innerarity, *Artificial Intelligence and Democracy* (UNESCO, 2024).
- 18 Thorsten Quandt, “Dark Participation,” *Media and Communication* 6, no. 4 (2018): 36–48.
- 19 Yochai Benkler, *The Wealth of Networks: How Social Production Transforms Markets and Freedom* (Yale University Press, 2006).
- 20 Jürgen Habermas, “Reflections and Hypotheses on a Further Structural Transformation of the Political Public Sphere,” *Theory, Culture & Society* 39, no. 4 (2022): 145–71.
- 21 Sarah J. Jackson et al., *#HashtagActivism: Networks of Race and Gender Justice* (MIT Press, 2020).
- 22 Seth C. Lewis et al., “Online Harassment and Its Implications for the Journalist–Audience Relationship,” *Digital Journalism* 8, no. 8 (2020): 1047–67.
- 23 Tarleton Gillespie, “The Politics of ‘Platforms,’” *New Media & Society* 12, no. 3 (2010): 347–64.
- 24 Fuchs, *Digital Democracy and the Digital Public Sphere*.
- 25 Ibid.
- 26 Maria D. Molina and S. Shyam Sundar, “Does Distrust in Humans Predict Greater Trust in AI? Role of Individual Differences in User Responses to Content Moderation,” *New Media & Society* 26, no. 6 (2024): 3638–56.
- 27 Kena Zheng, “Antitrust in Artificial Intelligence Infrastructure – between Regulation and Innovation in the EU, the US, and China,” *Computer Law & Security Review* 59 (November 2025): 106211.
- 28 Joseph Jones, “Don’t Fear Artificial Intelligence, Question the Business Model: How Surveillance Capitalists Use Media to Invade Privacy, Disrupt Moral Autonomy, and Harm Democracy,” *Journal of Communication Inquiry* 49, no. 1 (2025): 6–26.
- 29 Nick Dyer-Witheford et al., *Inhuman Power: Artificial Intelligence and the Future of Capitalism* (Pluto Press, 2019).
- 30 Joanne Kuai et al., “AI ≥ Journalism: How the Chinese Copyright Law Protects Tech Giants’ AI Innovations and Disrupts the Journalistic Institution,” *Digital Journalism* 10, no. 10 (2022): 1893–912.

- 31 Andreas Jungherr, “Artificial Intelligence and Democracy: A Conceptual Framework,” *Social Media + Society* 9, no. 3 (2023): 1–14.
- 32 Zheng, “Antitrust in Artificial Intelligence Infrastructure.”
- 33 Jungherr, “Artificial Intelligence and Democracy.”
- 34 Particle, “Introducing Particle: The News, Organized,” *Medium*, November 12, 2024.
- 35 Nic Newman et al., *Digital News Report 2025* (Reuters Institute, 2025).
- 36 Ibid.
- 37 Nic Newman and Federica Cherubini, *Journalism and Technology Trends and Predictions 2025* (Reuters Institute, 2025).
- 38 Jungherr, “Artificial Intelligence and Democracy.”
- 39 Ben Paviour, “AI-Generated News Sites Spout Viral Slop from Forgotten URLs,” Nieman Lab, October 16, 2025; NewsGuard, “Tracking AI-Enabled Misinformation,” NewsGuard, October 24, 2025.
- 40 Joshua McElwee, “Pope Decries ‘Crisis of Truth’ in AI after His Own Deepfake Image,” *The Independent*, January 24, 2025.
- 41 Sarah Kreps and Douglas L. Kriner, “The Potential Impact of Emerging Technologies on Democratic Representation: Evidence from a Field Experiment,” *New Media & Society* 26, no. 12 (2024): 6918–37.
- 42 Felix M. Simon et al., “Misinformation Reloaded? Fears about the Impact of Generative AI on Misinformation Are Overblown,” *Harvard Kennedy School Misinformation Review* 4, no. 5 (2023): 1–11.
- 43 Christian von Sikorski and Michael Hameleers, “Disinformation in the Age of Artificial Intelligence (AI): Implications for Journalism and Mass Communication,” *Journalism & Mass Communication Quarterly* 102, no. 4 (2025): 941–57.
- 44 Sarah Kreps and Doug Kriner, “How AI Threatens Democracy,” *Journal of Democracy* 34, no. 4 (2023): 122–31 (123).
- 45 Jacob L. Nelson and Seth C. Lewis, “Only ‘Sheep’ Trust Journalists? How Citizens’ Self-Perceptions Shape Their Approach to News,” *New Media & Society* 25, no. 7 (2023): 1522–41.
- 46 Helle Sjøvaag, “The Business of News in the AI Economy,” *AI Magazine* 45, no. 2 (2024): 246–55.
- 47 Jungherr, “Artificial Intelligence and Democracy,” 6.
- 48 Mark Coeckelbergh, *Why AI Undermines Democracy and What to Do about It* (Polity, 2024).
- 49 Seth C. Lewis et al., “Generative AI and Its Disruptive Challenge to Journalism: An Institutional Analysis,” *Communication and Change* 1 (2025): 9.

- 50 Alessandro Volpe, “Toward an Artificial Deliberation? On Google DeepMind’s Habermas Machine,” *Ethics and Information Technology* 27, no. 4 (2025): 45.
- 51 Nikos Karacapilidis et al., “Generative AI and Public Deliberation: A Framework for LLM-Augmented Digital Democracy,” *Proceedings EGOV-CeDEM-ePart Conference* (Ghent, Belgium), September 1, 2024.
- 52 Newman and Cherubini, *Journalism and Technology Trends and Predictions* 2025.
- 53 Felix M. Simon, “Rationalisation of the News: How AI Reshapes and Retools the Gatekeeping Processes of News Organisations in the United Kingdom, United States and Germany,” *New Media & Society*, May 8, 2025: 1–21.
- 54 John Ruwitch, “OpenAI’s New Web Browser Has ChatGPT Baked in. That’s Raising Some Privacy Questions,” NPR, November 7, 2025.
- 55 Rodrigo Zamith and Seth C. Lewis, “From Public Spaces to Public Sphere: Rethinking Systems for Reader Comments on Online News Sites,” *Digital Journalism* 2, no. 4 (2014): 558–74.
- 56 Quandt, “Dark Participation.”
- 57 Tao Huang, “Content Moderation by LLM: From Accuracy to Legitimacy,” *Artificial Intelligence Review* 58, no. 10 (2025): 320.
- 58 Mia Hoffman, *The Mechanisms of AI Harm* (Center for Security and Emerging Technology, 2025).
- 59 Kreps and Kriner, “How AI Threatens Democracy.”
- 60 Coeckelbergh, *Why AI Undermines Democracy and What to Do about It*.
- 61 Karacapilidis et al., “Generative AI and Public Deliberation.”
- 62 Cody Turner, “Online Echo Chambers, Online Epistemic Bubbles, and Open-Mindedness,” *Episteme* 22, no. 1 (2025): 232–57.
- 63 David Hartmann et al., “A Systematic Review of Echo Chamber Research: Comparative Analysis of Conceptualizations, Operationalizations, and Varying Outcomes,” *Journal of Computational Social Science* 8, no. 2 (2025): 52.
- 64 Zamith and Lewis, “From Public Spaces to Public Sphere.”
- 65 Karacapilidis et al., “Generative AI and Public Deliberation.”
- 66 Kreps and Kriner, “How AI Threatens Democracy.”
- 67 Kreps and Kriner, “The Potential Impact of Emerging Technologies on Democratic Representation.”
- 68 Ibid.
- 69 Paul Hitlin et al., “Public Comments to the Federal Communications Commission About Net Neutrality Contain Many Inaccuracies and Duplicates” (Pew Research Center, 2017).

- 70 Slavko Splichal, *Datafication of Public Opinion and the Public Sphere: How Extraction Replaced Expression of Opinion* (Anthem Press, 2022).
- 71 Renée Sieber et al., “What Is Civic Participation in Artificial Intelligence?,” *Environment and Planning B: Urban Analytics and City Science* 52, no. 6 (2025): 1388–406.
- 72 Ibid.
- 73 Leonardo Coelho, “Brazil Turns Facial Recognition on Rioters despite Racism Fears,” *World, Reuters*, January 12, 2023.
- 74 Douglas MacMillan et al., “Arrested by AI: Police Ignore Standards after Facial Recognition Matches,” *The Washington Post*, January 13, 2025.
- 75 Sasha Costanza-Chock et al., “Who Audits the Auditors? Recommendations from a Field Scan of the Algorithmic Auditing Ecosystem,” *2022 ACM Conference on Fairness, Accountability and Transparency*, June 21, 2022, 1571–83.
- 76 Ugur Aytac, “Big Tech, Algorithmic Power, and Democratic Control,” *The Journal of Politics* 86, no. 4 (2024): 1431–45.
- 77 Coeckelbergh, *Why AI Undermines Democracy and What to Do about It*, 5, original emphasis.
- 78 Jones, “Don’t Fear Artificial Intelligence, Question the Business Model.”
- 79 Kreps and Kriner, “How AI Threatens Democracy.”
- 80 Turner, “Online Echo Chambers, Online Epistemic Bubbles, and Open-Mindedness.”
- 81 Coeckelbergh, *Why AI Undermines Democracy and What to Do about It*.
- 82 Nikki Usher, *News for the Rich, White, and Blue: How Place and Power Distort American Journalism* (Columbia University Press, 2021).
- 83 Jungherr, “Artificial Intelligence and Democracy.”
- 84 Matthew Wood, “The Political Ideas Underpinning Political Distrust: Analysing Four Types of Anti-Politics,” *Representation* 58, no. 1 (2022): 27–48.
- 85 Paul Nemitz, “Constitutional Democracy and Technology in the Age of Artificial Intelligence,” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 376, no. 2133 (2018): 20180089.
- 86 Burt Herman, “Hacks/Hackers Launches New Lab to Empower Newsrooms to Build AI Tools,” *Hacks/Hackers*, June 12, 2025.
- 87 Taijuan Moorman, “‘AI Models in Vogue?’ Guess Ad Sparks Uproar,” *USA TODAY*, July 28, 2025.
- 88 Nabiha Syed and Sisi Wei, “The Markup Is Joining Forces with CalMatters,” *The Markup*, April 18, 2024.

- 89 Nicholas Diakopoulos et al., *Generative AI in Journalism: The Evolution of Newswork and Ethics in a Generative Information Ecosystem* (The Associated Press, 2024).
- 90 Mathias-Felipe de Lima Santos et al., “Guiding the Way: A Comprehensive Examination of AI Guidelines in Global Media,” *AI & Society* 40 (2025): 2585–603.
- 91 Simon, “Rationalisation of the News.”
- 92 Dominic Ponsford, “Politico Journalists Fight Back over AI Tools That Published ‘Glaring Errors,’” *Press Gazette*, December 2, 2025.
- 93 Seth C. Lewis et al., “Journalists, Emotions, and the Introduction of Generative AI Chatbots: A Large-Scale Analysis of Tweets before and after the Launch of ChatGPT,” *Social Media + Society* 11, no. 1 (2025): 1–13.
- 94 Nirit Weiss-Blatt, “Media Has a Blind Spot When Covering the AI Panic,” *Freethink*, May 15, 2025.
- 95 Nemitz, “Constitutional Democracy and Technology in the Age of Artificial Intelligence.”
- 96 Innerarity, *Artificial Intelligence and Democracy*.
- 97 Coeckelbergh, *Why AI Undermines Democracy and What to Do about It*.

Chapter 5 Learning to Think in the Age of AI

- 1 D. Graham Burnett, “Will the Humanities Survive Artificial Intelligence?,” *The Weekend Essay*, *The New Yorker*, April 26, 2025.
- 2 Oscar Westlund et al., “Public Knowledge and Expertise under Authoritarian Siege: A Defense of Academic Freedom from Digital Journalism Studies,” *Digital Journalism* 13, no. 5 (2025): 869–92.
- 3 Burnett, “Will the Humanities Survive Artificial Intelligence?”
- 4 Ibid., original emphasis.
- 5 R. Glenn Cummins et al., “2021 Survey of Journalism & Mass Communication Enrollments,” *Journalism & Mass Communication Educator* 78, no. 1 (March 2023): 69–83.
- 6 Julie Jargon, “7 Reasons Teens Say No to AI,” *The Wall Street Journal*, January 31, 2026.
- 7 Cindy Royal, “AI Inspires Innovation in Journalism Education,” Nieman Lab, December 2024.
- 8 Leonid Rozenblit and Frank Keil, “The Misunderstood Limits of Folk Science: An Illusion of Explanatory Depth,” *Cognitive Science* 26, no. 5 (2002): 521–62.
- 9 Russell Butson and Rachel Spronken-Smith, “AI and Its Implications for Research in Higher Education: A Critical Dialogue,” *Higher*

- Education Research & Development* 43, no. 3 (2024): 563–77 (563).
- 10 Chunpeng Zhai et al., “The Effects of Over-Reliance on AI Dialogue Systems on Students’ Cognitive Abilities: A Systematic Review,” *Smart Learning Environments* 11, no. 1 (2024): 28.
 - 11 Kristjan-Julius Laak and Jaan Aru, “AI and Personalized Learning: Bridging the Gap with Modern Educational Goals,” arXiv:2404.02798, preprint, *arXiv*, March 21, 2025.
 - 12 Derek Thompson, “The End of Thinking,” October 2, 2025, <https://www.derekthompson.org/p/the-end-of-thinking>.
 - 13 Nicholas Carr, “The Myth of Automated Learning,” *New Cartographies*, May 27, 2025.
 - 14 Thompson, “The End of Thinking.”
 - 15 Dhruv Khullar, “If AI Can Diagnose Patients, What Are Doctors For?,” Brave New World Dept., *The New Yorker*, September 22, 2025.
 - 16 James P. Walsh, “Social Media and Moral Panics: Assessing the Effects of Technological Change on Societal Reaction,” *International Journal of Cultural Studies* 23, no. 6 (2020): 840–59.
 - 17 Derek Thompson, “What Happens When AI Learns to Do Our Jobs,” *Plain English with Derek Thompson* (podcast), October 24, 2025.
 - 18 Neil Thurman et al., *AI Adoption by UK Journalists and Their Newsrooms: Surveying Applications, Approaches, and Attitudes* (Reuters Institute for the Study of Journalism, 2025).
 - 19 Louisa Dahmani and Véronique D. Bohbot, “Habitual Use of GPS Negatively Impacts Spatial Memory during Self-Guided Navigation,” *Scientific Reports* 10, no. 1 (2020): 6310.
 - 20 Isabella Kwai, “He Wants a New Start. So He Is Taking the Hardest Driving Test in the World,” *The New York Times*, November 24, 2025.
 - 21 Md Sazzad Hossain and Debora Wenger, “Has Journalism Education Kept Up with Digital Transformation?,” *Journalism & Mass Communication Educator* 79, no. 2 (June 2024): 129–46.
 - 22 Abdelmohsen Hamed Okela, “Artificial Intelligence and the Future of Journalism Education: Opportunities and Challenges in Egypt,” *Journalism & Mass Communication Educator* 80, no. 1 (March 2025): 54–69 (57).
 - 23 Marcos Mayo-Cubero, “Exploring Journalism Education in the AI Era: A Mixed Analysis of University Training in Spain,” *Journalism & Mass Communication Educator* 81, no. 1 (March 2026): 19–36 (31).

- 24 Santiago Tejedor et al., “Integrating Artificial Intelligence and Big Data in Spanish Journalism Education: A Curricular Analysis,” *Journalism and Media* 5, no. 4 (2024): 1607–23 (1621).
- 25 Debora Wenger et al., “AI and the Impact on Journalism Education,” *Journalism & Mass Communication Educator* 80, no. 1 (March 2025): 97–114 (104).
- 26 Ibid.
- 27 Debora Wenger and Iveta Imre, “Integrating AI in Video Journalism Education: Current Trends and Challenges,” *Journalism & Mass Communication Educator* 81, no. 1 (March 2026): 3–18.
- 28 Aaron S. Veenstra et al., “Journalism Students’ Use, Expectations, and Understanding of Generative AI Tools,” *Journalism & Mass Communication Educator* 80, no. 1 (March 2025): 115–28 (115).
- 29 Ibid.: 124.
- 30 Wenger et al., “AI and the Impact on Journalism Education.”
- 31 Veenstra et al., “Journalism Students’ Use, Expectations, and Understanding of Generative AI Tools”: 122.
- 32 Wenger et al., “AI and the Impact on Journalism Education”: 107.
- 33 Hilke Schellmann, “I Tested How Well AI Tools Work for Journalism,” *Columbia Journalism Review*, August 19, 2025.
- 34 Ibid.
- 35 Ibid.
- 36 Ibid.
- 37 Anita Kirkovska, “GPT-5 Benchmarks,” *Vellum*, December 3, 2025.
- 38 Mayo-Cubero, “Exploring Journalism Education in the AI Era”: 31.
- 39 Ibid.: 31.
- 40 Mark Deuze and Charlie Beckett, “Imagination, Algorithms and News: Developing AI Literacy for Journalism,” *Digital Journalism* 10, no. 10 (2022): 1913–18.
- 41 Justin Hutchens, *The Language of Deception: Weaponizing Next Generation AI* (Wiley, 2024).
- 42 Fabrizio Dell’Acqua et al., “Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality,” SSRN Scholarly Paper no. 4573321 (SSRN, September 15, 2023).
- 43 Ethan Mollick, “Centaurs and Cyborgs on the Jagged Frontier,” *One Useful Thing*, September 16, 2023.
- 44 Ethan Mollick, *Co-Intelligence: Living and Working with AI* (Portfolio, 2024).
- 45 Dell’Acqua et al., “Navigating the Jagged Technological Frontier.”
- 46 Frank Chimero, “Beyond the Machine,” October 20, 2025, <https://frankchimero.com/blog/2025/beyond-the-machine/>.

- 47 Deuze and Beckett, “Imagination, Algorithms and News”: 1917.
- 48 Chimero, “Beyond the Machine,” original emphasis.
- 49 Ibid.
- 50 Nan Jia et al., “When and How Artificial Intelligence Augments Employee Creativity,” *Academy of Management Journal* 67, no. 1 (2024): 5–32.
- 51 Ethan Mollick, “15 Times to Use AI, and 5 Not To,” *One Useful Thing*, December 9, 2024.
- 52 Chimero, “Beyond the Machine.”
- 53 Ibid.
- 54 Deuze and Beckett, “Imagination, Algorithms and News.”
- 55 Mollick, “Centaurs and Cyborgs on the Jagged Frontier.”
- 56 Chimero, “Beyond the Machine.”
- 57 Bibo Lin and Seth C. Lewis, “The One Thing Journalistic AI Just Might Do for Democracy,” *Digital Journalism* 10, no. 10 (2022): 1627–49.

Chapter 6 Reimagining Journalism for the Age of AI

- 1 Clara Harter, “An Early AI Pioneer Shares How the ‘Vibe Coding’ Revolution Could Reshape Data Journalism,” USC Center for Health Journalism, October 9, 2025.
- 2 Paul Ford, “The AI Disruption Has Arrived, and It Sure Is Fun,” *The New York Times*, February 18, 2026.
- 3 Jon Accarrino, “Agent Swarms and Vibe Coding: Inside the New Operational Reality of the Newsroom,” *TV News Check*, December 19, 2025.
- 4 Harter, “An Early AI Pioneer Shares How the ‘Vibe Coding’ Revolution Could Reshape Data Journalism.”
- 5 Kawandee Virdee, “Rise of the Vibecoding Journalists,” Nieman Lab, December 2025.
- 6 Lila Shroff, “AI Agents Are Taking America by Storm,” *Technology, The Atlantic*, February 17, 2026.
- 7 Ford, “The AI Disruption Has Arrived, and It Sure Is Fun.”
- 8 Accarrino, “Agent Swarms and Vibe Coding.”
- 9 Dawn McMullan, “The Rise of Agentic AI Signals the Biggest Shift in News Companies Yet,” International News Media Association (INMA), October 2, 2025.
- 10 Dean Starkman, “The Hamster Wheel,” *Columbia Journalism Review*, September 14, 2010.
- 11 Tomás Dodds et al., “Popularity-Driven Metrics: Audience Analytics and Shifting Opinion Power to Digital Platforms,” *Journalism Studies* 24, no. 3 (2023): 403–21.

- 12 Bumsoo Kim and Nicholas R. Buzzelli, “‘I Feel Burnout as a Journalist’: Testing Two Interaction Models on Burnout in a Hyper-Connected Society,” *Journalism Practice* 20, no. 3 (2026): 823–42.
- 13 Matt Carlson et al., *News after Trump: Journalism’s Crisis of Relevance in a Changed Media Culture* (Oxford University Press, 2022).
- 14 Liz David-Barrett, “State Capture: How Democracy Can Be Systematically Corrupted,” *Political Insight* 15, no. 1 (2024): 16–19.
- 15 Chris Quinn, “Journalism Schools Are Teaching Fear of the Future,” *Cleveland.com*, February 14, 2026.
- 16 Aruna Ranganathan and Xingqi Maggie Ye, “AI Doesn’t Reduce Work – It Intensifies It,” *Harvard Business Review*, February 9, 2026.
- 17 Tomás Dodds et al., “On Controlled Change: Generative AI’s Impact on Professional Authority in Journalism,” *Journalism Studies* 27, no. 7 (2026): 1051–68.
- 18 Seth C. Lewis et al., “Jobs-to-Be-Done and Journalism Innovation: Making News More Responsive to Community Needs,” *Media and Communication* 12 (2024): 7578.
- 19 Clayton M. Christensen et al., “Know Your Customers’ ‘Jobs to Be Done,’” *Harvard Business Review*, September 2016.
- 20 Lewis et al., “Jobs-to-Be-Done and Journalism Innovation”: 5.
- 21 Charlotte Tobitt, “Le Monde’s Olympian Effort to Attract More English-Language Subscribers,” *Press Gazette*, May 9, 2024.
- 22 Gretel Kahn, “These Pioneers Are Working to Keep Their Countries’ Languages Alive in the Age of AI News,” Reuters Institute for the Study of Journalism, May 27, 2025.
- 23 Donghee Shin, “The Perception of Humanness in Conversational Journalism: An Algorithmic Information-Processing Perspective,” *New Media & Society* 24, no. 12 (2022): 2680–704.
- 24 Daniel Trielli, “The Rise of Agentic Journalism,” Nieman Lab, December 2025.
- 25 Karen Hao, *Empire of AI: Dreams and Nightmares in Sam Altman’s OpenAI* (Penguin Press, 2025).
- 26 Nicholas Demba, “Journalist Karen Hao Warns against the ‘Empires of AI’ and Their Impact,” *Student Life*, November 3, 2025.
- 27 Mona Sloane, *Predicted: How AI Is Restructuring Social Life* (University of California Press, 2026), 133.
- 28 Zoë Hitzig, “OpenAI Is Making the Mistakes Facebook Made. I Quit,” *The New York Times*, February 11, 2026.

Index

- acceleration, 5, 7–8, 10, 30, 148–9, 151, 156–7
 - see also* hamster wheel metaphor; reimagination
- accessibility, 48–50, 152
- accuracy, 1–3, 38, 41–3, 133–6
 - hallucinations, 23, 40, 43, 47, 51, 105, 134
 - limitations of, 23, 42–3, 57
- advertising, online, 63–5, 70–2, 77
- Aftonbladet*, 53
- agentic AI, 5, 24
 - agent swarms, 146
 - autonomous task completion, 146–7
 - implications for journalism, 154
- AI (artificial intelligence)
 - algorithmic foundations, 21–3
 - authenticity, 108–9
 - capabilities and limitations, 23–4, 36–51, 104–7, 133–5, 137–41
 - definitions and history of, 11–13
 - hype surrounding, 25–7, 115, 156
 - public attitudes toward, 4, 14, 19, 25–8
 - public opinion, 67, 95, 109
 - see also* agentic AI; automation, journalistic; generative AI; large language models (LLMs)
- AI literacy
 - need for, 137–41
 - positional reflexivity, 140–1
- AI Overviews, *see* Google
- AI policies, newsroom, 19–20, 51–4
 - limitations of, 53–4, 115, 133
- AI slop, 10, 24, 32, 102
- algorithmic imaginary (Bucher), 25
- algorithms
 - accountability, 110
 - bias, 18–19, 23, 51–2, 92–4, 97–8, 101, 110, 125, 136
 - distribution via, 48–51, 75–6, 85–7
- Amazon, 73–5, 116

- Amditis, Joe, 46–7
 American Press Institute, 151–2
 Anthropic: Claude, 5–6, 9, 19, 153
 Apple, 43–4, 73
 Associated Press (AP), 16–17, 20, 45, 53, 77
 see also Wordsmith
 audience engagement, 29, 33, 36–7, 45–7, 153–4
 authorship, 5–6, 19, 31–2, 47–8, 108–9, 124
 automation, journalistic, 14–20, 31–51
 augmentation vs. transformation, 33–4, 57–62
 automated content generation, 16–20, 31–3, 35–7, 45–50
 autonomy
 journalistic, 29, 52–3, 66–8, 74–6, 80, 83–8
 loss of, 69–72, 90, 115, 125–6
 Axel Springer, 20, 77

 Bagdikian, Ben, 65–7
 BBC, 1–2, 18, 43–4, 49
 Beckett, Charlie, 137–8, 140, 142
 see also AI literacy
 Bourdieu, Pierre, 68
 Bucher, Taina, 25
 Burnett, D. Graham, 119–20
 burnout, journalistic, 8, 10, 147
 bylines, *see* authorship

 Caswell, David, 33, 58–60
 see also automation, journalistic
 ChatGPT
 accuracy of, 1, 134
 impact on education, 119–21, 132
 impact on journalism, 2, 19–20, 26, 39, 77, 128
 see also OpenAI
 Chimero, Frank, 140–3
 instrument vs. tool distinction, 141
 “working under/beside the machine,” 141–3
 see also personal empowerment
 China
 AI development in, 92–3, 98–9
 regulatory environment, 84–5
 Christensen, Clayton, 151
 civic organizations, AI oversight by, 114–16
 ClaimReview, 41
 Claude, *see* Anthropic
 CNET, 19
 CNN, 77
 Coeckelbergh, Mark, 95, 104, 106, 111, 117
 cognitive function: atrophy and erosion of, 125–6, 149–50
 confidence–competence gap, 123–6, 132
 overload of, 103, 112, 149
 value of struggle in, 126–9
 constitutive moment, journalism’s, 1, 5–6, 28, 30, 150, 157
 see also reimagination
 content moderation, 50, 86–7, 106
 conversational journalism, 33, 153–4
 Copilot, *see* Microsoft
 copyright
 lawsuits, 20, 78–9, 93
 licensing agreements, 77–8
 corporate consolidation, media, 67
 Crawford, Kate, 73–4
 critical thinking, 125–6, 135–6, 149

- data commons, 116
- data journalism, 15–16, 131, 145–7
- deepfakes, 19, 42, 83–4, 103
see also disinformation
- deliberation, democratic, 94–6, 98, 104–7, 116–18
- democracy
 - AI development, 92–4, 107–17
 - corrosion of, 89–91, 149
 - disillusionment with, 113
 - participatory governance, 113–14
 - reimagining of, 113–18
- deprofessionalization, 148–50
 - boundaries of professional journalism, 17, 57, 148
 - deskilling, 19, 114, 126
- Deuze, Mark, 137–8, 140, 142
see also AI literacy
- Diakopoulos, Nicholas, 77
- disinformation, 19, 42, 44, 97, 139
 - AI-generated, 19, 43, 101–3, 105–6, 149
 - disinformation-resistant journalism, 152–3
 - see also* fact-checking; misinformation
- echo chambers, 106–7
- editorial judgment, 31–4, 44–8, 135–6
- education
 - AI in, 119–44
 - higher education, 119–21, 128–36
 - see also* journalism education
- epistemic authority, 124
- epistemic silos, 101–3
- epistemic thinness, 148–50
- ethics, 41–4, 51–4, 57–62, 68, 89–91, 115, 124–5
- European Union: regulatory environment, 85–7
- Facebook, *see* Meta
- fact-checking, 41–2, 66, 135
 - AI tools for, 42–3, 73, 152–3
 - credibility, 43–4
 - see also* ClaimReview
- Factive (Scroll tool), 152
- faculty expertise, *see* journalism education
- Fox News, 77
- friction, productive vs. unnecessary, 128–9
- Gannett, 19, 77
see also LedeAI
- Gemini, *see* Google
- generative AI
 - definition of, 21–4
 - impact on journalism, 1–6, 19–20, 31–62
 - public opinion on, 3, 25–6
 - see also* training data
- Global South, journalism in, 8, 53–4, 81–2
- Google
 - AI Overviews, 2, 63–4, 75–6, 89, 149–50
 - DeepMind, 12
 - evolution of, 2, 18, 41–2, 74, 81
 - Gemini, 1, 5, 8, 19, 80, 153
 - Google Zero, 64, 92
 - lawsuits, 72
 - NotebookLM, 1, 123
 - “Great Houses of AI,” 72–82, 87–90
 - see also* Amazon; Apple; Google; Meta; Microsoft
- Hacks/Hackers, 114
- hallucinations, *see* accuracy
- hamster wheel metaphor

- democracy, 94, 113, 117–18
- education, 121–2
- escaping the wheel, 140, 143, 147–50, 154–7
- newsroom acceleration, 32, 34, 57–62
- origin (Starkman), 6–10
- platform power, 79–82, 87, 90
- Hao, Karen, 155–6
- headline writing, 33, 45, 59, 79, 128–9
- Helberger, Natali, 87
- homogenization, creative, 75, 140–1
- human–AI collaboration, 35–50
see also personal empowerment
- ideation, journalistic, 35–7
- infrastructural capture, 73–6, 148–9
 - dependency on Big Tech, 63–72, 148–9
 - resistance to, 76–9
- intellectual property, *see* copyright
- interviewing assistance, 39
- iTromsø (Norway), 32
- jagged technological frontier, 139–40, 142–3
- “jobs to be done” (JTBD) framework, 151–2
- Jones, Bronwyn, 9
- journalism
 - definition of, 13–14
 - democratic function of, 89–91, 114–16
 - economic pressures on, 7–8, 74–6, 147–8
 - evolving practice of, 14–20, 34–5, 58–62, 89–91, 147, 152–7
 - history of automation in, 14–20
 - open-source investigations (OSINV), 152–3
 - power of, 66–8
 - public trust in, 4, 10, 19, 41–4, 57–8, 61, 66, 98, 103–4
 - reimagining of, 152–7
 - see also* automation, journalistic; news industry; newsrooms
- journalism education
 - AI curriculum integration, 130–6, 143–4
 - AI expertise of faculty, 131–2
 - “digital native” fallacy, 132
 - emphasized competencies in, 135–6
 - enrollment declines, 80, 122
 - international perspectives, 130–1
 - structural challenges, 130–4
- Jungherr, Andreas, 95–6, 99, 103–4, 112
- knowledge production, 100–4, 119–26
- knowledge silos, 54–7
- large language models (LLMs)
 - chatbots, 104–5
 - definition of, 21–3
 - errors and limitations, 23–4, 133–4
 - see also* AI (artificial intelligence); ChatGPT; generative AI
- LedeAI, 19
- licensing, news content, 20, 77–8, 86, 90
- literacy, *see* AI literacy
- local news, 31–2, 40–1, 46–7, 49, 132, 152
- Los Angeles Times*, 17, 77
see also Quakebot

- machine learning, 22–4
 - see also* AI (artificial intelligence); generative AI; large language models (LLMs)
- Markup, The, 114
- MassLive, 31–2
- Mayo-Cubero, Marcos, 130–1, 135
- media ownership concentration, 65–76
- Meta
 - AI development, 73, 81
 - news distribution, 18, 42, 48, 50, 69–70, 73, 98
 - partnerships with, 77
- micro-relevance, 153
- Microsoft
 - AI development, 73–4
 - Copilot, 1, 75, 80
 - NYT lawsuit against, 1, 20, 78–9
- miscalibrated trust, *see* jagged technological frontier
- misinformation, 41–2, 44, 53, 97, 125
 - see also* disinformation
- Mollick, Ethan, 139–40, 142
- Monde, Le*, 77, 152
- moral panic, 127–8
- natural language processing, *see* large language models (LLMs)
- news industry
 - economic crisis in, 7–8, 74–6, 147–8
 - platform dependency, 74–80
 - see also* journalism
- Newspaper Next, *see* American Press Institute; “jobs to be done” framework
- newsrooms
 - AI adoption in, 1–5, 18–20, 31–62
 - AI policies, 19–20, 41–4
 - workforce changes, 51–7
- New York Times, The*
 - AI tools adoption, 1, 20, 53
 - lawsuit against OpenAI/Microsoft, 1, 20, 78–9
- Nielsen, Rasmus Kleis, 71
- OpenAI, 92, 99
 - lawsuits against, 1, 20, 78–9
 - partnerships with, 1–2, 74, 77
 - see also* ChatGPT
- open-source AI, 100, 152
 - see also* data commons
- participation, democratic, 93–5, 108–9, 112–17
- Perplexity, 2, 77–8
- personal empowerment, 141–4
- personalization, algorithmic, 18, 49, 100–1, 112, 115–16, 125, 153
- POLITICO, 115
- Prisa Media, 77
- professionalism, journalistic, *see* deprofessionalization
- propaganda, 102–3
- ProPublica, 32
- public sphere
 - digital infrastructure as, 97–8, 116–17
 - Habermasian ideals, 94–6
 - impact of AI on, 94–8, 117–18
- Pulitzer Prizes, 15, 32, 80
- Quakebot (*Los Angeles Times*), 17
- Reese, Stephen, 67
- regulation of AI: governmental interventions, 82–9, 92–3, 116–17
- reimagination, 57–62, 150–7
- Reuters, 3, 18

- robot journalism, *see* automation, journalistic
- Royal, Cindy, 123
- Rural News Network, 49
- Schellmann, Hilke, 133–4
- search engines, 2, 48, 63–5, 69, 75–6, 93
 - traffic decline, 2, 63–4
 - see also* Google
- Seipp, Theresa J., 86–8
- Shoemaker, Pamela, 67
- Simon, Felix, 77, 103
- Sloane, Mona, 155
- social media
 - AI content on, 108
 - democracy, 97–8, 105–6
 - journalism, 14, 48–51, 121
- Society of Professional Journalists, 53
- source cultivation, 37–41, 135
- Spiegel, Der*, 77
- Sports Illustrated*, 19, 47
- Starkman, Dean, 7, 147
 - see also* hamster wheel metaphor
- storytelling, 44–8, 153
- summarization, AI tools for, 1, 40, 79–80, 133–4
- Suncoast Searchlight, 47
- TIME*, 46, 77, 105
- training data, 21–3, 40, 43, 51–2, 136
- consent, 39, 155
- copyright disputes, 20, 78–9, 93
- transcription tools, 40
- translation, AI-powered, 50, 152
- Trielli, Daniel, 154
- Trump, Donald, 41, 83–4, 92–3
- United States
 - AI development, 92–3, 98–9
 - journalism programs, 130–5
 - regulatory environment, 82–4, 92–3
- universities, *see* education; journalism education
- van der Vlist, Fernando, 73–4
- van Dijck, José, 88
- vibe checks (AI evaluation), 133–4
- vibe coding, 46–7, 59, 141, 145–6, 148
- Vogue*: AI-generated model controversy, 114
- Washington Post, The*
 - AI adoption, 2, 77
 - Ask the Post AI chatbot, 46
 - Heliograf (robot reporter), 18
- Wordsmith (Automated Insights), 16–17
- X (formerly Twitter), 26, 73, 116

POLITY END USER LICENSE AGREEMENT

Go to <http://politybooks.com/eula/> to access Polity's ebook EULA.