

AN ANTHOLOGY BY NORDIC AI JOURNALISM & FOJO MEDIA INSTITUTE

Futures of Journalism

Eleven invited essays on what may lie ahead for media

Nordic AI Journalism

Fojo. Medieinstitutet

NOVEMBER 2026

CONTENTS

Eleven essays, three themes

EDITOR'S INTRODUCTION

00 The responsibility to imagine what’s next
Agnes Stenbom Swedling

I · RETHINKING JOURNALISM

01 The information revolution will not be televised
Shuwei Fang

02 Four different responses to public skepticism around AI in news
Rasmus Kleis Nielsen

03 Journalism as Science
David Caswell

II · REINVENTING VALUE

04 Towards sustainable publisher models for the AI-driven information ecology
Kasper Lindskow

05 The next era of journalism will be defined by service, not publication
Nikita Roy

06 Journalism with room to move
Ezra Eeman

III · REBUILDING NEWSROOMS

07 The noble art of flogging dying horses
Gard Steiro

08 Look beyond the AI tools
Martin Schori

09 From low adoption to real enablement – a journey
Olle Zachrison

10 Futures friction: a new approach to how media organisations remember, learn and reimagine
Sofie Hvitved

11 Rewiring organisations starts in our heads
Sara Inkeri Vardar



00

The responsibility to imagine what’s next

Agnes Stenbom Swedling

Strategic Advisor, Schibsted and Co-founder, Nordic AI Journalism

Journalism has never been short of people willing to explain what is wrong with it. This anthology seeks to go beyond diagnosing problems to offer visions of what may come next. In this introduction, project initiator and lead Agnes Stenbom Swedling (Strategic Advisor, Schibsted and Co-founder, Nordic AI Journalism) argues that editorial media has a responsibility to imagine desirable futures – and act towards them.

Journalism has never been short of people willing to explain what is wrong with it.

Declining revenues. Disappearing audiences. Eroding trust. Powerful technology companies capture an ever-greater share of the economic value created on the information market. And now, AI threatens to disrupt not only how journalism is distributed, but how it is produced, experienced and valued.

The challenges are real, and many are profound. Yet sometimes I wonder whether we have become a little too comfortable talking about them. We seem to have developed an industry identity around being under siege: our work is indispensable, yet insufficiently appreciated; our competitors are ruthless, our audiences distracted, and the future perpetually threatening.

Perhaps that is understandable. Journalism is, after all, a profession largely built around scrutinising issues. But I worry that our ability to diagnose what is wrong has outpaced our willingness to imagine what could be right.

The responsibility to imagine what's next

I believe those of us who care about journalism have a responsibility to do more than identify challenges, defend established practices and assign blame. To me, that includes undertaking the intellectually demanding work of imagining alternatives.

What might journalism become? What could it offer people that it does not today? What new forms of value might emerge as technology changes the possibilities for producing, sharing and using information? And what would we need to change about ourselves and our organisations to realise them? These are questions requiring us to challenge assumptions that have shaped our professional identities, organisations and business models for decades, if not centuries.

As David Caswell writes in his essay for this anthology, the situation facing journalism "is deeply uncomfortable and is made more so by the lack of any clearly defined path forward". But the absence of a clear route ahead does not absolve us of the responsibility to find one.

Technology companies routinely invest enormous resources in imagining radically different futures and building towards them. Meanwhile, editorial media too often finds itself responding to developments initiated elsewhere. We debate the implications of technologies, products and business models that others have already decided to pursue. I would like us to spend more time setting the agenda ourselves.

Journalism needs more competing ideas, more willingness to be proven wrong, and perhaps most of all, more courage to act in new ways. That requires leadership in editorial media organisations to create the conditions for people to update their skills and behaviours to meet the demands of new realities. And on an individual level, journalism practitioners across roles and contexts must seize the opportunity to

expand their portfolios of journalistic skills.

Why this anthology?

In 2020, Olle Zachrison and I founded Nordic AI Journalism to enable collaborative exploration of AI and journalism. Our ambition has always been to bring together people from different disciplines, organisations and perspectives to exchange knowledge, challenge assumptions and explore shared opportunities in pursuit of responsible AI developments in our region and beyond. Since 2023, our annual Nordic AI in Media Summit in Copenhagen has served as one of the main arenas for those conversations.

At this year's summit, taking place in May 2026, we had invited a number of keynote presenters and workshop facilitators to challenge our community to think more radically different about what might lie ahead for journalism. Together with Fojo Media Institute, we have now invited them to develop their ideas into essays.

The brief was the same for everyone: articulate a compelling vision for the future of journalism in an AI-mediated world. The writers’ responses were anything but homogenous. Some contributors offer ambitious theoretical frameworks for understanding an entirely new information economy. Others imagine future newsrooms, products and professional roles building on personal newsroom experience. Some propose concrete updates in business models or organisational practices. Others question the assumptions that underpin our ambitions in the first place.

As the project lead and editor, I have deliberately chosen not to impose a uniform style, level of abstraction or degree of optimism. The contributors bring different experiences, intellectual traditions and professional perspectives, and I want those differences to remain visible. That is, after all, part of the beauty of human creativity. Give eleven people the same question, and you are unlikely to receive eleven versions of the same answer.

I have not asked the invited writers to disclose the details of their use of AI in the process building towards these essays. As their editor, I trust that the writers have used the best biological and technological tools they have available to ideate, iterate and deliver their essays.

Different visions, shared responsibilities

The essays do not present a unified roadmap for journalism's future, nor should they. In fact, I find that some of the most interesting insights emerge from the tensions between them.

David Caswell argues for a more scientific form of journalism to exist alongside the humanistic traditions of the profession. Martin Schori, meanwhile, imagines a future in which distinctly human, subjective qualities such as voice and relationships

become journalism's most important competitive advantages.

Shuwei Fang envisions an information economy in which machines increasingly consume and act upon information on behalf of people. Ezra Eeman asks us to consider what might be lost if we optimise information experiences so thoroughly that we leave little room for surprise, reflection and independent judgement.

Nikita Roy imagines newsrooms that maintain ongoing relationships with the questions people need answered, while Sofie Hvitved explores how organisations might become better at remembering, learning and questioning their own assumptions.

These are not necessarily opposing visions, but they emphasise different possibilities, priorities and risks. And they remind us that technological capability does not determine what journalism ought to become.

What strikes me most, however, is how many of the contributors return to the importance of deliberate choices.

Rasmus Kleis Nielsen asks news organisations to consider not only how they use AI, but how they choose to explain those choices to the public. Kasper Lindskow outlines distinct strategic positions that publishers might pursue in a changing information economy. Gard Steiro calls for constructive paranoia: an honest recognition of the threats ahead, coupled with the determination to act despite uncertainty.

Their arguments differ, but they share an important premise: The future depends on the choices we make here and now.

Three ways of looking ahead

I have organised this anthology into three broad themes:

Rethinking Journalism begins with the most fundamental questions. As AI transforms the production, circulation and use of information, what becomes of journalism's role in society? What should remain, what might change, and how should journalism relate to the public it exists to serve?

Reinventing Value explores what journalism could offer in an information economy no longer necessarily organised around articles, websites or even direct human consumption. The essays consider new economic models, journalistic services and audience experiences, and challenge us to reconsider what makes journalism valuable.

Rebuilding Newsrooms turns the attention inward. Imagining a different future is one thing; creating organisations capable of realising it is another. Here, contributors explore the leadership, capabilities, cultures and ways of working that meaningful transformation might require.

These categories are neither exhaustive nor mutually exclusive. Most essays could comfortably belong to more than one theme. But taken together, the essays offer perspectives to help us understand what is changing, imagine what could emerge, and create the conditions to make it happen.

This collection is inevitably incomplete. For one, journalism encompasses far more than the written article, and the futures of audio, video, visual storytelling and other forms of media deserve more attention than they receive in this anthology. Further, eleven essays, however thoughtful, cannot meaningful represent the full diversity of perspectives needed to shape journalism's future. With that in mind, consider this anthology a contribution to an ongoing conversation, rather than an attempt to settle it.

An invitation to act

I neither expect nor wish for readers to agree with every argument made in this anthology.

Instead, my hope is that you encounter at least one idea that makes you reconsider an assumption, question a familiar practice or imagine a possibility you had not previously considered. And more importantly, that you do something with it.

As media professionals, we can continue explaining why the world is changing in ways that threaten our industry or try to defend the practices that have served us well in the past. But we could also choose to imagine something better, and begin the difficult work of making it real.



CHAPTER I

Rethinking Journalism

01 Shuwei Fang
02 Rasmus Kleis Nielsen
03 David Caswell



01

The information revolution will not be televised

Shuwei Fang

Co-founder, Signals at Scale; ’25-’26 Shorenstein Fellow, Harvard Kennedy School; Senior Fellow, CUNY-Newmark

What happens to journalism when information is increasingly consumed by machines? In this essay, Shuwei Fang (Co-founder, Signals at Scale; 25-’26 Shorenstein Fellow, Harvard Kennedy School; Senior Fellow, CUNY-Newmark) explores the emergence of an agentic information economy, where journalistic knowledge becomes input to decisions and actions rather than a finished product for human consumption. She argues that journalism has an opportunity to reinvent itself as a vital part of the infrastructure that makes this new economy trustworthy and accountable.

Within ten years, AI could leave society far better informed. Almost nobody will notice. That is the point, and the challenge.

There are no altars to the gods of showers, card payments or traffic lights.

Nobody wakes up in grateful wonder that the water runs hot, that money moves at the tap of a card, or that a busy junction does not descend into carnage. These are among the great achievements of modern life, and they are almost entirely invisible. Wake up to find any one of them broken, though, and your day is ruined. Possibly worse.

The information environment of the near future should work the same way. Most people do not think about it, and they should not have to.

Too tired to be informed

Consider a new mother today, six months in and running on broken sleep. This week's headlines told her the baby-food pouches she has been relying on are loaded with sugar; she feels guilty rather than helped. The government has just changed its childcare entitlements, and she stands to save thousands. She has seen the news. But her circumstances are slightly complicated, and none of the coverage quite answers whether, or how, the change applies to her. The answer exists. It is public, accurate and free, sitting in a policy document she never has the time or energy to read. She is not misinformed. She is too tired to be informed. The money goes unclaimed.

Ten years from now, another new mother is just as tired. But an agent acting on her behalf knows the rules have changed, works out what she is entitled to given her particular circumstances, and drafts the application for her to check. When a local investigation exposes failings at the maternity unit where she is booked for a follow-up, she hears about it that morning, with her options laid out. The decisions remain hers. If she asks how it knows, it can show her: where each claim came from, who checked it, and how other parents nearby have tested it, questioned it or disagreed. Then she goes back to sleep.

By 2036, there could be a step change in how well informed people are, and organisations, communities and governments too, delivered by an agentic information economy that works. Most of society's decisions will rest on it, and much of society’s critical functions will depend on it. Yet to the average person, it will barely seem to exist. Nobody will build an altar to it. Ingratitude will be the sign that it works; we might notice only if it breaks.

Emergence of an agentic information economy

Two linked shifts make this possible. The first is that machines are becoming the primary audience for information. A growing share of what is published, recorded or measured will be read first, and often only, by AI systems acting for people, and increasingly by agents acting for other agents. That leads to the second shift: information changes category. For most of its history, news, information and knowledge have been a final good, often produced antecedent to consumption and read by a person who may or may not act on it. In an agentic economy, they become an intermediate input into a decision, a transaction or an automated process. In most cases, they may never be rendered for a human at all.

Together, these shifts could give rise to an agentic information economy: the epistemic bedrock on which the wider agentic economy rests. It could be far bigger than the information economy of today, because machines consume on people's behalf without the limits of human attention. It could be far more complex: glance at a map of a language model's embeddings, or at an event-centric knowledge graph, and imagine most of our knowledge and context interlinked like that.

The price of being right

Unlike today's information economy, an agentic one could run in favour of accuracy, truth-seeking and self-correction.

To see why, start with what an error is. When a person misreads something, the error becomes an impression. When an agent acts on it, the error becomes an action: a wrong purchase, a missed entitlement, a recalled product back on the shelf. In agentic systems, outcomes loop back as inputs, so errors compound at machine speed and scale. And because so many systems draw on the same few models, they can fail together. This is more than "garbage in, garbage out": correlated, automated error is a different class of risk from misinformation.

It is also a different class of opportunity. For as long as there has been a digital information market, being right and sounding right have fetched roughly the same price, except in finance and a few other high-stakes niches. Once the cost of being wrong becomes measurable, and at potentially far greater cost, that need no longer hold, and the exception could soon extend to almost all information. Provenance, accuracy and the capacity to self-correct could become things that are priced. Across much more of the information economy, being right could start to pay.

But here’s the catch. The best version of an agentic information economy would span almost everything we know: personal data, niche expertise, the knowledge inside organisations, government records, real-time events and the findings of investigations. For it to be truth-seeking and self-correcting, it needs certain conditions, such as ownership and control that are spread out rather than concentrated, and no moral hazard in how information is gathered and verified. This

will need a myriad of native protocols, institutions and rules to govern it. That is no easy feat, and the best case is unlikely to emerge by default.

Complex information systems are not new

Finance offers a useful comparison. An agentic information economy, in which news, information and knowledge are sourced, verified, exchanged, priced and recombined by agents on behalf of people and of each other, could become as complex as, if not more complex than, the current global financial system, and far less legible. Most of us have only a folk understanding of finance, yet we hand it our salaries and pensions. We don't fully understand it, because we don't really need to. We can rely on it because of protocols, institutions and regulation to which we delegate trust, legibility and risk: shared identifiers and messaging standards, custodians, rating agencies, insurers, clearing houses, and the regulators that oversee them.

The agentic information economy will need analogous systems, and probably more. Finance has SWIFT and shared identifiers, which let banks message each other and know exactly which party and which asset a transaction refers to; information needs a shared protocol and a standard unit, so that a claim can be identified and traced. In the agentic information economy, agent protocols such as MCP and content credentials such as C2PA are early versions of comparable systems. Finance has custodians, who hold assets on their owners' behalf; information needs trusted holders of dispersed, niche knowledge. Data cooperatives such as Switzerland's MIDATA, whose members decide who may use their health data, are an early version. Finance has rating agencies; information needs calibrated confidence in its sources. Probabilistic news, which attaches explicit likelihoods to claims, is an early version. Finance has clearing houses; information needs corrections that reach every system that used the original. Crossref's Crossmark, which flags corrected and retracted research, is an early version.

These systems also decide where value flows and where it is captured. In an agentic information economy, value is likely to pool at the interface, with whoever knows the most about the user. With the right rails, market mechanisms and institutions, some share of it, and of the demand signal (knowing what the user wants), could flow back along the chain to those who originate, verify and maintain information. A crude example from finance shows it can be done: card interchange fees send a slice of every card payment back to the bank that issued the card.

The right blueprints matter

Finance also shows how such institutions fail. The rating agencies that blessed mortgage-backed securities before 2008 were paid by the issuers whose products they rated. And the same infrastructure that makes finance work has also enabled extraordinary concentrations of wealth and power. Journalism and science encode the opposite instincts: independence from those you cover, disclosure of sources and methods, scrutiny by peers, public correction. Those are the right blueprints. Whether they are used is being decided now, while the standards are still soft. The agentic information economy is so early that it does not yet have a standard way to represent a piece of information, let alone the institutions to build on one. In October 2026, Meta and Sierra announced a protocol for how personal agents deal with businesses, with Walmart, Shopify and Stripe on board. The foundations of the agentic economy are being laid, but the information systems it runs on are lagging behind.

This is a space journalism should occupy, though not in the form it takes today. Its institutions already hold much of the DNA: verification, editorial standards, archives, corrections policies and a working ethic of accountability. Within ten years they could be taking on functions that look little like journalism today: originating machine-ready, verified events with their provenance attached; acting as custodians of a community's local and niche knowledge; rating sources and attaching calibrated confidence to claims; underwriting the information agents act on, with liability attached; and running the clearing systems that carry a correction to every system that used the original.

Journalists will also need to watch this new system. A system this consequential and this opaque will need watchdogs of its own. Covering the agentic information economy will become a beat, much as finance and technology are today. Someone will still need to hold information power to account; the power will simply be harder to see.

The invisible fourth estate

None of this displaces storytelling as we know it. Stories are how people come to care, and they will flourish alongside. But they will not be the only way people are informed. There is room for a new discipline of information and knowledge to emerge, one that no existing field, from writing and journalism to economics to computer science, fully addresses today.

So what should newsroom leaders do now? Start from first principles. That means asking what an agentic information economy will actually need, and looking well beyond the news industry for answers, rather than searching for a best-in-class version of something that already exists. It means rethinking what the unit of information, and therefore of value, actually is, building the economics around that, and developing the skills to work with information in ways the newsroom was never

designed for. And it means a level of coordination and getting into the rooms where the standards for this economy are being written, while they are still soft.

If the future agentic information economy works for a better-informed society, with its value and power more evenly distributed, most of what it does will be invisible, and most people will never think about it. They should be able to trust it, without fully understanding it or inspecting it for themselves. But invisibility cuts both ways. A system in the background that can quietly leave you better informed can just as quietly manipulate you, and you would probably never even know it was happening. That is all the more reason to build it to serve citizens rather than to steer them. It also sits uneasily with journalism's cultural identity. As the fourth estate, journalism is used to occupying a visible seat in society, and wearing its trust on its masthead. That will have to be reckoned with. But our future selves deserve what journalism has always offered at its best: accountability, legibility, a willingness to be corrected, and the nerve to bark at power. There is a role for journalism in that future. It just might be unrecognisable.



02

Four different responses to public skepticism around AI in news

Rasmus Kleis Nielsen

Professor of Communication, University of Copenhagen and Senior Research Associate, Reuters Institute, University of Oxford

How should journalism respond when the public is more sceptical of AI in news than journalists might like to believe? Drawing on international audience research, Rasmus Kleis Nielsen (Professor of Communication, University of Copenhagen and Senior Research Associate, Reuters Institute, University of Oxford) outlines four possible responses to AI, ranging from reluctant non-adoption to openly embracing technological transformation. He argues that the most important choice is not necessarily how much AI newsrooms use, but whether they can explain their decisions and maintain the public's trust.

The value of news is based on its relation with the public

The value of any kind of news is in large part in the eye of the beholder. It's a relationship between what journalists and news publishers do and the people who engage with it, members of the public.

That relationship is generally important, and specifically crucial for journalists and news publishers trying to navigate widespread public skepticism around the use of AI in news.

This relationship has always been important, but will, I expect, grow more important in the context of an ever-more competitive marketplace for attention.

This context is one in which generative AI represents an additional disruption, layered on top of an already radically disruptive move from a relatively low-choice offline media environment dominated by competing legacy media organizations controlling the main channels of communication, to a high-choice environment where many of the commanding heights are occupied by competing technology companies who own and operate the platforms most of us rely on to find, access, and engage with all sorts of content, including news.

News publishers who have lost much of their control over distribution, discovery, data, and advertising already now face the prospect of also losing ground when it comes to content. True, some of what AI can generate at scale – cheaply, near-instantaneously and in personalized fashions – today falls short of the best journalism has to offer. It is entangled in problematic political economics, and some of it is generic low-information generalities, misleading, or outright false and fabricated. But it is also demonstrably already the case that many generative AI users find a fair share of synthetic content useful and valuable. Generative AI use is currently growing roughly twice as fast as internet use did in the 1990s and early 2000s.

In this evolving environment, many kinds of content are likely to grow ever more commoditized. But relations are not necessarily so.

That is why I expect that the relationship between news and journalism on the one hand and the public it aims and claims to serve on the other hand will only grow more important going forward.

That relationship is in many ways challenged by external forces – by AI, by already-dominant platforms, in many countries by political attacks and much more.

But it is also challenged by an inability or unwillingness by the industry to be mindful of and care for relations with the public.

At worst, there is a tendency to take the relationship for granted, assume people will accept journalism as valuable because journalists think it is, and that they will defer

to the authority it aspires to. Even absent such outright arrogance, the profession and the industry does not always put in the work to maintain the relationship with the public, understand what trust does or does not look like from different people's point of view, let alone ensure that journalists and news publishers earn and retain it on the public's terms.

In the rest of this piece, I will, first, briefly summarize key points from recent research I have been involved in documenting how the rapid development of AI and public perception of AI when it comes to news provides one such external challenge. Then, I will outline four different responses journalists and news media publishers might embrace when it comes to public skepticism around AI in news as part of efforts to be more active stewards of their relationship with the public, before ending with the question 'what do you do next?'

The public is skeptical of AI in news

Together with colleagues, I have for several years been conducting research on how people use and think about generative AI and news. A few key findings provide important evidence of how skeptical the public is. The results I summarize below are based on online surveys of nationally representative samples of adults conducted in 2025 across a range of mostly high-income countries in Europe, North America, and Asia. They will not be representative of every country or every community, but still capture important aspects of the context many journalists and news publishers navigate. (Those interested in details on the research can read more here.)

First, when asked whether news produced mostly by artificial intelligence with some human oversight is likely to be more or less of various qualities we might associate with news, compared to news produced entirely by human journalists, people tend to be pretty pessimistic. While there are more respondents who say they think such news will be more up to date, easier to understand, and informative than say the opposite, there are significantly more pessimists than optimists when it comes to how transparent and, crucially, trustworthy people think such news will be. The quality most people expect to characterize news mostly produced by artificial intelligence? Cheaper to make. So, in short, less costly to churn out, but less transparent and trustworthy, all in a context where journalists and news publishers already struggle with suspicious minds and low trust or even active distrust.

Second, some journalists and news publishers might protest "we do not produce news mostly by artificial intelligence!" and assume that this public skepticism might provide a dark background against which their human journalism can shine. Perhaps that is so for some of them, and for some parts of the public, but not for most. Because a majority of respondents, when asked how frequently, if at all, they think generative AI is being used by news media, answer "always" or "often", at a rate behind only search engine and social media companies. So, irrespective of whether or not a given journalist or news publisher is in fact using AI, much of the public is likely

to assume they are, and their (fairly skeptical) beliefs about what that will mean for them (less transparent, less trustworthy, cheaper) is likely to be "priced in" whether or not AI is actually used.

Third, maybe a "sinking tide" of trust driven by public skepticism might lower all boats in the marketplace of attention? Declining trust may not be great, but, after all, trust is not only relational, but also relative, so if citizens are growing more and more skeptical about everything in what is sometimes assumed to be a general crisis of trust, perhaps journalism and news media is losing ground "only" in an absolute sense and not in comparison to other organizations? That is unlikely for two reasons. First, when it comes to trust in general, while trust in news media is in overall long-term decline in much of the world, and the same applies for some of the institutions most frequently featured in news coverage, including governments and political parties, this is not a general crisis of trust. Trust in other institutions, including the civil service, the legal system, and for that matter science, has generally held up or even increased in the same period in many countries. (As shown here.) Second, when asked how much better or worse the use of generative AI will make their experience of interacting with various sectors of society, the optimists currently actually outnumber the pessimists for most sectors – except news media, governments, and political parties. The contrast between the preponderance of optimists when it comes to search engine companies and even (!) social media companies versus the preponderance of pessimists when it comes to news media is particularly striking, and particularly worrying for journalists and publishers who care about their relation with the public.

So people are skeptical of what the use of artificial intelligence in news will mean for them, often assume news media are already using generative AI often or even all the time, and are more pessimistic about the implications of AI use in news media for them than they are when it comes to AI use by technology companies. Journalists and news publishers will not like this, but we have to navigate the world as it is, not as we wish it was.

Four different responses to public skepticism around AI in news

So what are they to do? One way to think of the options available here is to imagine two stylized axes, one of which concerns whether a given journalist or news publisher is a less or more intense user of generative AI, and the other of which concerns whether they are leaving their choice of what to do with AI implicit or explicitly telling the public a story about what they are doing and why. These two axes can then help us identify four ideal type responses. No real situation will ever be so clear and neat, but hopefully the typology is useful to think with.

The first is what we might call "resigned journalism". This involves (a) limited use of generative AI and (b) abstaining from explicit communication about how it is used, if

at all, and why. Some news publishers will end up here by default because they are not changing their editorial practices, and communicating about that probably seems superfluous. Others may decide not to change (perhaps hoping for a fast-follower option based on learning from others' mistakes, and successes) and, perhaps understandably, not interested in drawing attention to their holding pattern. In both cases, the position is complicated by the tendency for people to assume that journalists and news media are already intense users of AI.

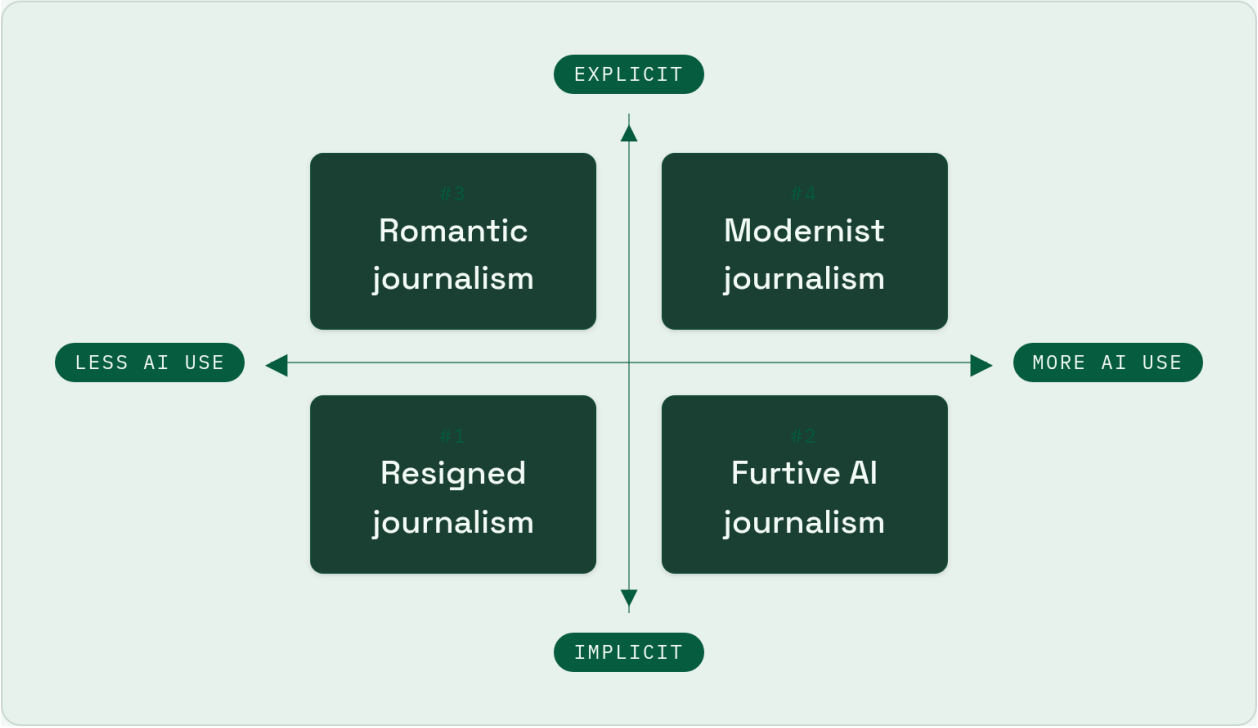
The second is what we might think of as "furtive AI journalism". This involves (a) increasingly intense use of generative AI and (b) abstaining from explicit communication about how it is used and why. Some news publishers will end up here because they may feel that the work should speak for itself, and while they are keen to use generative AI, they are not more keen to commit scarce resources to explicitly communicating how and why they use it than they are to commit scarce resources to explicitly communicating any other part of their editorial practices, workflow, or internal operations. Click deep enough, and you might find something, but it may well be about as detailed and publicized as the organization's policy on, say, the use of anonymous sources or press releases and other handouts, or the inner operations of their dynamic paywall. Other news publishers may end up here because they want to realize the "cheaper to make"-part of what much of the public thinks generative AI has to offer when it comes to news, but may privately think that the public might be right about the "less transparent"- and "less trustworthy"-bits too, and are therefore reluctant to draw attention to the issue. Given our research findings, in the absence of explicit communication, much of the public will assume news publishers are engaged in furtive AI journalism.

The third then represents what some have called "romantic journalism". This involves (a) limiting the use of generative AI, at least on the editorial side and (b) explicitly communicating about this decision and what it is based on. An example of this, and explicitly named that way, would be the "romantic journalism manifesto" launched at the 2026 International Journalism Festival in Perugia by Lea Korsgaard from Zetland in Denmark and Joshi Hermann from Mill Media in the UK. It is a nuanced and thoughtful view that I encourage everyone to consider. It is also a clear example of explicit communication to differentiate one's position from a competitive set many of whom are at least experimenting with generative AI use, also in editorial, and all of whom much of the public, as said, may assume is already using generative AI quite intensely. It is easy to understand the attraction here, and it is effectively the mirror twin of the response above – if much of the public thinks using generative AI will lead to news that is cheaper, less transparent, and less trustworthy, and you fear that they may in many cases be right, it is important to really clearly differentiate what you do from the rest of the profession and industry, to stand apart from being judged by the lowest common denominator.

The fourth and final is what I would call "modernist journalism", which then involves investing scarce resources in (a) increasingly intense use of generative AI and (b)

explicitly communicating about this decision and what it is based on. This would be a modernist view in the sense that the term “modernist” has in some forms of architecture, art, and literature, a broad ethos of “make it new”, a belief that our best selves might be ahead of us and can be enabled by new approaches, new forms, new technologies. Leaving aside some AI-enthusiastic start-ups, it is hard to identify examples of this among established news publishers, though there are clearly journalists and other news media professionals who are effectively pursuing this kind of response. A record number of Pulitzer Prize awardees disclosed AI use in 2026, the New York Times has published a human-computer interaction paper on the development and use of the Epstein Files Engine as an inhouse A.I. agent using Google BigQuery SQL queries across various corpora, Bloomberg’s AI researchers contributed 8 papers to the 2026 Association for Computational Linguistics conference on issues ranging from agentic systems to responsible AI to financial NLP, and journalists at Die Zeit and Der Spiegel have taken time to explain how they used AI to build a searchable database of Nazi Party members. These are not examples only of using AI. They are also steps of explicitly communicating how it was used, what for, and why. While in large part oriented at fellow professionals, they contain at least the kernel of a more explicit communication about what generative AI in news can also mean – beyond making it cheaper to produce.

Resigned journalism, furtive AI journalism, romantic journalism, and modernist journalism, as presented in the figure below, thus represent four different stylized responses journalists and news publishers might opt for in a context where the public is generally skeptical of generative AI in news.



What do you do next?

So what do you do next? I strive to be impartial as a scientist, and stay clear of professorial preaching. But I also care about journalism and news publishing, and about helping journalists and news publishers think through the context they are in, some of the options they have, and what the implications might be.

Perhaps for some, this piece reads like a call for modernist journalism. It is not meant to be. Each of the positions outlined above could be named differently, resigned journalism might be called “stoic”, furtive AI journalism could be called “stealthy”, romantic journalism “retro”, and modernist journalism “techno-solutionist”. Each comes with risks and possible rewards.

My concrete proposal here is not that journalists and news publishers should or ought to commit to one specific response out of these four. Instead, my concrete proposal is that they decide which one comes closest to what they want to do, and consider whether and how they communicate with members of the public about what they do, how, and why.

Some paths require little change. Others, especially those committed to greater use of generative artificial intelligence and/or more explicit communication require more.

Here it is useful to step away from AI and a changing competitive situation as the proximate cause for caring about the relationship between journalism and news publishers on the one hand and on the other hand the public that their work is supposed to serve, and who provide the basis for everything they are trying to achieve – whether political impact, social significance, or for that matter financial sustainability.

Because anything that involves more explicit communication probably requires a culture shift in how much of journalism is often practiced and how many newsrooms often work. To be blunt, research suggests that, journalists and news publishers still, even after all the profession and the industry has been through over the past few decades, often have a pretty partial picture of how the public sees their work and do not prioritize investing in their relationship with the public as much as they prioritize investing in many other things.

Perhaps this is the cultural heritage of a profession and an industry that came of age in a context where it could be assumed that the work would speak for itself, and that if it did not, ultimately it did not matter very much as members of the public did not have that many alternatives to journalism and news media to choose from. But that was then, this is now. As I see it, empirical research suggests that journalism today does not speak for itself, and reminds us that people have an abundance of alternatives.

However one chooses to approach this situation, and however one chooses to communicate about it, generative AI is yet another radical disruption of a profession

and an industry that has taken more than its fair share of knocks in recent years. That disruption will no doubt be painful for many journalists and news publishers, even as some also see various opportunities in it.

What you do with AI in this disruption is the immediate question. The wider question is one that transcends AI and centers journalism and news publishers' relation with the parts of the public they want to engage with and asks who you are doing it for, how they see it, and whether you can explain that it is valuable not just for you, but also, ultimately more importantly, for them.

This is not (only) about whether or not you use AI, for what, or how much. It is about much more basic questions of ensuring you are clear about what you do, who you do it for, why, and that what you do and what you say you do is aligned, explainable, and intelligible to those parts of the public you want to engage with your work. Communicating your choices about AI is but one component of this; a small part of trying to invest more care and effort into the relationship between news and the public. A simple starting point would be to ask yourself three basic questions: What does the part of the public we serve assume we do with AI? What do we do to explain what we do with AI? Would people see the principles we preach and what we actually practice as aligned - or not - if they knew precisely what we do with AI?

Truly AI re-wiring journalism and news publishers, for those who wish to, will take years, if not decades.

But the wider work on journalism and news publishers' relation with the public is not only or even primarily about AI. It is cultural, in a way more foundational. It could take longer. It could also begin tomorrow.



03

Journalism as Science

David Caswell

Co-founder, Signals at Scale

What if the future of journalism requires us to fundamentally rethink what journalism is? In this essay, David Caswell (Co-founder, Signals at Scale) challenges the profession's humanistic traditions and argues for the emergence of a parallel, scientific form of journalism. Drawing on the distinction between humanistic and scientific cultures, he makes the case for embracing AI-driven, systematic approaches to producing and verifying information, while preserving journalism's democratic mission.

When an idea appears again and again, in different forms, then you can be confident that there is something meaningful behind it. This was the case at NAMS 2026 regarding ideas about the restructuring of journalism due to the inexorably improving capabilities of AI.

In his closing keynote, and in the essay included in this anthology, Rasmus Kleis Nielsen described a 4-part framework for understanding journalism's reaction to AI, with the two primary responses termed 'Romantic Journalism' and 'Modernist Journalism'. The former, articulated in a recent manifesto by Lea Korsgaard and Joshi Herrmann, is centered on the author's "bond of affinity with the people who are reading or listening or watching" and emphatically decenters AI from journalism. The latter, articulated by Nielsen as a form of journalism in which "Everything that can be automated, will be automated", anticipates an information ecosystem that is intensely mediated by AI.

At the same event, Shuwei Fang used different terms in her comprehensive talk on the economics of AI-mediated information. Fang (who I work with) described the 'Barbell Effect' – a generalized and well-studied phenomena in which the appearance of a scaling technology causes an industry to separate into two extremes: a 'differentiated' extreme, based on intangible value; and a 'scale' extreme, based on applying the new technology to its full extent. Businesses in the middle range of the industry decline, in the manner of department stores in an era of boutique brands and e-commerce.

This same essential division shows up repeatedly, in different places, and with different terms. It is, in some ways, built into the organizational structure of every news publisher using the terms 'editorial' and 'technical'. This essay is about that division, about how AI is changing the relationship between these two interpretations of journalism, and about why the journalism industry and journalism's research institutions are extremely ill-equipped to handle that change. It is also about the potential to both clarify and expand journalism by accepting a scientific alternative, and the potential of science to supercharge journalism in the same way it has supercharged so many other fields over centuries. I will try to make my arguments as fundamental as I can, and so I will use a different pair of terms that I believe are closer to the heart of what is going on. Those terms are 'humanistic' and 'scientific'.

The two cultures

In his famous 1959 lecture The Two Cultures, C.P. Snow argued that modern intellectual life had become divided into two increasingly disconnected cultures: the humanistic culture, represented by writers, artists and literary intellectuals, and the scientific culture, represented by scientists, engineers and technologists. Snow argued that these cultures differ not merely in their subject matter, but in their fundamental ways of understanding the world. The humanistic culture emphasizes subjective experience, interpretation, judgment and the 'qualia' of human experience. The scientific culture emphasizes empirical observation, experimentation, quantification and the accumulation of objectively verifiable knowledge. Snow saw these two cultures as equally valid and necessary and bemoaned the division between them as a barrier to human progress.

In his 2004 book Hackers & Painters, Y-Combinator co-founder Paul Graham makes a similar argument about the intellectual differences between the humanities and the sciences, but was less neutral in his interpretation. Graham suggests that while many physicists could probably complete a PhD in French literature, relatively few professors of French literature could complete a PhD in physics. His argument is not that scientists are more intelligent, but that scientific disciplines are more directly constrained by objective tests of truth and falsity, whereas humanistic disciplines allow greater scope for social influence, intellectual fashion and established authority. Scientific disciplines, essentially, have more discipline.

This structural difference between humanistic and scientific cultures can be interpreted as more than just an asymmetry. In Homo Deus, his 2015 book anticipating the AI era, Yuval Noah Harari argues that human civilization has changed its fundamental assessment of the source of ultimate authority in the past, and is currently changing it again. Harari identifies three distinct eras – Theism, Humanism and Dataism. He posits that the printing press caused a change in civilization's source of ultimate authority from God (Theism) to human opinion (Humanism), and that computation, the internet and AI are now causing a further change from human opinion to, essentially, data (Dataism). Before the printing press it was difficult to argue with God, and during the humanist era it was difficult to argue with personal opinions and collective human decisions as expressed in elections and markets. Harari's point is that this authority of human opinion, which humanism treats as foundational, may instead be historically contingent and in decline. If it is data that now settles arguments, then perhaps the era of the dominance of human opinion is behind us.

It is as difficult to understate the significance of these changes as it is to ignore their reality in our direct experience. At some level these ideas are, to use a common idiom, "playing for all the marbles", and for many journalists it will be easy to reflexively identify with the humanistic perspective and even to retreat into an 'us and them' defensiveness. But this need not be the only response. The changing

dynamics between humanistic and scientific cultures in the AI era does not mean that humanistic journalism is ending – on the contrary it may become more valuable than ever, as the 'differentiated' side of the Barbell Effect or the argument for Romantic Journalism make clear. Instead, it can mean that it is time to expand journalism in ways that allow for the meaningful inclusion of scientific culture. If the objective of journalism in the AI era is truly to serve society with useful, accurate information that empowers citizens, then perhaps it is time to build a 'scientific journalism' around AI – a parallel journalism entirely within the scientific culture.

AI's tectonic impact on humanistic culture

AI has deeply unsettled the balance between humanistic and scientific cultures. This is, obviously, because AI significantly encroaches into what was previously the sole preserve of humanism. AI can research, create new ideas and write, and therefore AI is a researcher, a creator and an author. AI has opinions and judgement. It can interpret and contextualize. It can make moral assessments. The usefulness of AI's output can be tested and measured. That output can be compared with the output of humans and is found, in many cases and by many measures, to be superior to that of humans. All of this is true now, and the reach of AI into humanistic domains is likely to continue as AI models improve, as more agentic systems are built and as we learn how to use them.

For journalism, this disruption has been rapid and practical. AI authorship of varying kinds, whether overt or furtive, is now thoroughly embedded in most newsrooms. AI Overviews and similar experiences are providing more and more of the answers that consumers previously sought from journalists. Single individuals now operate news publications that would previously have required a newsroom. Consumers are poised to gain access to personal agentic news editors, and AI platforms are poised to bypass news articles as sources and gather relevant information directly.

More generally, journalism has experienced a 'vibe shift'. The long, slow managed decline of journalism caused by the internet and social media has become a new, urgent realisation that we may be approaching some kind of endgame. The entire industry seems to now operate on a year-to-year basis and is pervaded by anticipation of sudden irrelevance. Journalism, as an industry, is waiting for an AI hammer to fall.

This situation is deeply uncomfortable and is made more so by the lack of any clearly defined path forward. For journalists this means anxiety, career uncertainty, questions about identity and even a kind of grief. For those who study journalism it means dilemmas about what is meaningful scholarship and what will be trivial or irrelevant work. For investors in journalism it means an almost incapacitating inability to allocate capital effectively beyond a few years.

There is no hiding from this situation. It is real, it is happening and the pre-AI status quo is gone forever – regardless of any plausible regulatory or legal development. Even if the observed impact is relatively light now, the ability to automate large

portions of humanistic work is an enormous Sword of Damocles hanging over humanistic journalism's head. The possibility of enormous impact by AI on large parts of humanistic culture must be accepted, but there is a coherent response despite the uncertainty: Journalism must develop a new conception of itself that firmly includes the scientific culture, operating alongside humanistic journalism. To move forward with confidence, despite the uncertainty, we need to clearly delineate between Romantic and Modernist journalism and between the differentiated and scale ends of the barbell.

Humanistic Journalism

Most journalism in our current information ecosystem is humanistic journalism. It is defined by the centrality of the opinions of specific humans called journalists, editors and publishers. It is an artisanal craft in an economy that has been industrializing for two hundred years – a handmade product that is still shaped almost entirely by the intuitive choices of individuals. Its practitioners define themselves through socialisation, identity and values, and not through the acquisition of an accumulated body of knowledge that defines the field. As knowledge work, humanistic journalism has previously been impervious to automation.

Opinion is everywhere in humanistic journalism, from decisions about commissioning and newsworthiness, to editing and word choices. The best humanistic journalists have a 'voice' – an undefinable but unmissable point-of-view or attitude that in many cases approaches literary art and is sometimes of great value. The very best are adept at storytelling – shaping narratives in ways so powerful that they produce rich and vivid experiences in the minds of their consumers. More commonly and mundanely, humanistic journalism often follows 'party line' editorial policies set and enforced by Editors-in-Chief, Publishers and shareholders. As a product largely defined by human opinion, and valued partly for the opinionated judgement of its producers, much of humanistic journalism may not benefit from automation.

As a humanistic endeavour, humanistic journalism is intensely social. It is a person 'speaking' personally to other people, and it carries all of the social considerations and signals present in such a relationship – trust, reputation, status, norms and conventions. Within humanistic journalism success is often defined through relationships – opportunities and status are decided through the favour of an editor or the receipt of an award from peers. In its relationship with audiences, humanistic journalism is often adept at triggering 'limbic brain' responses, like emotion, empathy or outrage, or even, for particularly skilled journalists, curiosity. In its relationship with its own elites, humanistic journalism is often conceived of as a tool for the exercise of influence within human societies. If we define humanistic journalism as a social activity, then automating it may be pointless, or even meaningless.

Humanistic journalism is 'romantic journalism'. It is differentiated by the myriad indefinable things that differentiate an individual or a community, and it is positioned

in opposition to scale on the barbell. It is literally defined by its human origins, and it is possible that those human origins may become far more valuable in an information ecosystem mediated by AI – not because of vaguely defined 'quality', but precisely because it is human. We have seen this before in the popularity of chess following Deep Blue's defeat of Garry Kasparov, and of Go following AlphaGo's defeat of Lee Sedol. Human beings are deeply social and often assign value based on social criteria, regardless of 'quality'. Humanistic journalism has value now, and it will probably have value a hundred years from now – without automation.

Scientific journalism

Scientific journalism exists today only on the margins of journalistic culture. Aspects of it are present in data journalism, in the more numerical parts of sports or financial reporting, in the details of elections results, in investigative projects centred on databases and document repositories, and in fringe innovation projects. It is not a prestigious activity in most news organizations, and some journalists barely consider it to be journalism at all.

Scientific journalism is defined by abstraction, and thus by its independence from human opinion. It is a form of journalism that values precision and objectivity. It seeks to be systematic rather than selective, and to build systems rather than to handcraft individual artifacts one by one. It seeks to be clinically neutral and rational about the facts that it is reporting. Its judgements are technical and methodological – explicitly discussed and written down rather than opaquely socialized. The units of information that it gathers and processes are all treated the same, and it is assessed by its completeness and its accuracy. Its measure of success is its fidelity to an objective and measured reality – and its value to its consumers is judged primarily by the volume and intensity of their engagement with it. The introduction of the opinion of the practitioner of scientific journalism into its output is considered an error, not an attribute. Scientific journalism is not seeking to be socially influential – it is seeking to be useful.

Scientific journalism is eminently automatable. Automation of this kind of journalism began well before the arrival of general-purpose LLMs, especially using template-based natural language generation (NLG) systems. Although early scientific journalism was possible only on highly structured or numerical forms of information, this was a purely technical constraint rather than a fundamental requirement. Like humanistic journalism, scientific journalism is essentially a philosophical perspective about what journalism is and can be. Just as humanistic journalism is an expression of the larger culture of humanism, so scientific journalism is an expression of the larger culture of science, technology and engineering. Scientific journalism is very different from humanistic journalism but it is also, in the eyes of its producers and consumers, real journalism.

The application of a scientific perspective to journalism also applies to journalism's understanding of its audience and their needs and wants. For much of its recent history journalism did not need to care what its customers thought of its product, because near-monopolistic control of pre-digital distribution channels offered few alternatives. Even now, journalism's understanding of its audience is remarkably thin – at worst paternalistic folk theories, and at best simplistic analytics and coarse surveys based on declared preferences. Similarly the academic study of journalism, as a soft social science, is overwhelmingly oriented towards its producers rather than its consumers, and towards its existing form rather than its fundamental function. In 100 years of existence, journalism studies has failed to produce a consensus theory of journalism that could be used to orient the field. Science offers a far richer understanding of why and how human beings value information – cognitive research techniques, functional MRI studies, controlled multivariate testing and theory-building based on neuroscience and evolutionary biology rather than on identity and social conditions. Beyond audiences, science offers many avenues of productive research into the flows of information through societies, including the application of network theory, Agent-Based Modeling, information topology and similar techniques.

As the arrival of powerful AI deeply unsettles the balance between humanistic and scientific cultures, it also dramatically expands the potential scope of scientific journalism, including into what was previously the sole preserve of humanistic journalism. The technical constraints of scientific journalism have been removed, and so a new opportunity to apply a scientific approach to much more of journalism has appeared. This empowerment of a scientific approach to journalism is a positive thing, even for humanistic journalists. Journalism as it stands is clearly stuck, clearly out of ideas, and clearly faced with what appear to be insurmountable challenges ahead. By explicitly and operationally separating journalism into the romantic and the modernist, the differentiated and the scaled, the humanist and the scientific, we can perhaps bring clarity, precision and conviction about the place of each in our AI future.

Society urgently needs scientific journalism

Science works. It works so well that our entire modern civilization exists because of it. The scientific revolution of the 17th century unlocked something immensely powerful in human affairs. The basis of field after field transitioned from socialised folk theories into rigorous evidence, and the results were near miraculous – from medicine to metallurgy. Science built the intellectual platform which made the industrial and agricultural revolutions possible, bringing a way of life that now supports more than 8 billion humans at a level of comfort and security that would have been unimaginable a few centuries ago. We need to seriously consider what that power might do for the empowerment of individuals and collective interests if it were brought to bear on the flows of information that course through our societies.

A technical interpretation of societal information flows is inevitable. The primary reason is that humanistic interpretations cannot plausibly adjust to the scale, scope and velocity of AI-mediated information. The seeds of this inevitability have been sprouting all around us, in start-ups, platforms and academic fields that have essentially zero overlap with journalism as a field. Probably the most significant example so far is X and X.ai, and Elon Musk's explicit ambition to replace journalism – all of it – with an entirely new model of societal information flows, and to spend billions of dollars doing it. This has not worked well (so far) and has not meaningfully changed our information ecosystem, but it is perhaps a harbinger of a new kind of fundamental innovation rooted entirely in the scientific culture. The problem with this inevitable technical alternative to humanistic journalism is not its existence, or its value, or its effectiveness or even its 'quality' – all of that will be decided by consumers. The problem is its values. Whatever the reality of our legacy journalistic institutions, their declared values of truthfulness and independence, and their watchdog and 'checks and balances' incentives have served society very well. The potential dangers of misused scientific power applied to societal information flows is very real, and even urgent. We will need, therefore, a scientific journalism that does not merely operate in a scientific way, but which also pursues objectives and values that are similar to the objectives and values of science itself. As journalism becomes more technical it would probably be best if the values of humanistic journalism – essentially enlightenment values – became the values of an emerging scientific journalism too.

Science is a moral force, as well as a pragmatic one. It is, in fact, the most credible and legitimate source of authority in our societies today. Many people still trust the words of religious leaders, intellectuals, politicians and even of journalists, but not as many and not as much as those who trust science. The best of technology, too, is trusted daily, by billions of people in trillions of practical acts seeking to fulfil their most important needs. Science and technology are trusted because they work – they deliver value – but science is also trusted for the same reasons as journalism wants to be trusted. Like journalism, science seeks to provide truthfulness, objectivity and consistency, but unlike journalism, science is not perceived to be filtered through the gate of human opinion. This neutrality comes from science's reliance on specificity, evidence, data, logic, peer-review, citations and specific, directed criticism. It comes from systematic rather than selective treatment of its subject matter. Of course scientists have opinions too, but science is a system, a perspective and a culture designed to diminish rather than to accentuate the influence of those opinions.

Humanistic journalists sometimes consider scientific approaches to journalism to be somewhat naïve or 'unsavvy', claiming that reality and 'truth' are just too complex and nuanced to be understood, represented and communicated by technology – even technology as capable as advanced AI. This objection has not aged well, but it does have a point. Humanistic journalism excels at reporting what epistemologists call 'inter-subjective' truth or 'social truth' – the truth that is 'true' because it is agreed

upon by a community of humans. There is clearly a significant demand for inter-subjective truth by information consumers, and this may well be a form of coverage in which humanistic journalists can specialise. But it is objective truth – scientific truth – that dominates human affairs. Scientific journalism, with its more systematic coverage and more neutral reporting, may perhaps offer society a way to coalesce around a shared perception of an objective reality in a way that journalism centered on human opinion cannot.

AI now gives us a new possibility for a genuine scientific journalism – for journalism engineering and for technical journalism – and thus for a new deal for journalism as a whole. Humanistic journalism has delivered immense value for societies over the centuries, and it will most likely deliver value for centuries to come. It can't, however, become an end in itself – a sacred cow that must be privileged over the greater goal of open, democratic societies of empowered citizens with access to truthful, trustworthy information. Journalism can't be limited by the desire to preserve familiar jobs and institutions or even cultures. It must expand or it will inevitably shrink.

Accepting and facilitating scientific journalism

Journalism must make space for the scientific culture within its purpose, but it is probably unreasonable to expect some kind of cultural merger with humanists, especially on the terms that legacy journalism institutions have required from technologists during the digital transition. The differences between humanistic and scientific cultures are, in my opinion, too great to be bridged – especially as we begin the exploration of journalistic possibilities from our new AI toolkit and as we automate many of the functions previously reserved for humans.

Instead I think that we need to explicitly recognize and separate humanistic journalism from scientific journalism – the differentiated side of the barbell from the scale side. We need to be clear about which is which – in leadership, in skills, in activities, in products and even in organizations – just as we also need to be clear about the need to transfer the values and mission of humanistic journalism into scientific journalism. This clarity and this separation can relieve both humanistic and scientific journalism of much of their uncertainty and unease, by unambiguously refocusing humanistic journalism on what it does best, while unblocking scientific journalism to stride ahead into the AI future. AI-driven and highly technical forms of journalism are coming, and it would probably be best for society if they developed with the support and encouragement of humanistic journalism rather than with its antagonism.

We also need a research and development ecosystem for scientific journalism. We need first principles applied research on information representation, automated verification, network coordination, incentive structures, socio-technical dynamics and more. We need audience research that applies the full gamut of cognitive science and

neuroscience research tools to human information consumption and sensemaking. We need AI and data science oriented unapologetically towards radically expanded societal understanding, instead of as support functions for humanistic storytelling. We need an information ecosystem in which research institutes for scientific journalism are developing concepts, methodologies, frameworks and communities around the application of AI to societal information flows.

Humanistic culture and humanistic journalism will remain as long as humans remain, just as religion has remained long after the scientific revolution. Romantic journalism has enormous value, and the differentiated side of the barbell is real – claims that are supported by a market for romantic and differentiated physical goods that exceeds half a trillion dollars. This essay is itself a kind of humanistic journalism. We really do, however, stand on the cusp of a fundamentally new era, dominated by machines that can fulfil editorial functions. A more technical journalism is arriving – and again, ideally with the values and mission of humanistic journalism – and it can deliver radical new benefits to all of us, as individuals and as societies. Humanistic journalism and its institutions should help facilitate that arrival.



CHAPTER II

Reinventing Value

04 Kasper Lindskow
05 Nikita Roy
06 Ezra Eeman



04

Towards sustainable publisher models for the AI-driven information ecology

Kasper Lindskow

Head of AI at JP/Politikens Media Group and Co-lead, Nordic AI Journalism

Where will news publishers create and capture value when AI fundamentally reorganises the information economy? In this essay, Kasper Lindskow (Head of AI, JP/Politikens Media Group and Co-lead, Nordic AI Journalism) explores how journalism might evolve from producing finished media products to originating, maintaining and commercialising trustworthy knowledge. He proposes three emerging strategic positions for publishers, and a model for combining them to build sustainable businesses in an AI-driven information ecosystem.

The creation of an AI-driven information ecology

The fundamental impulse that sets and keeps the capitalist engine in motion comes from the new consumers' goods, the new methods of production or transportation, the new markets, the new forms of industrial organization that capitalist enterprise creates … that incessantly revolutionizes the economic structure from within, incessantly destroying the old one, incessantly creating a new one. This process of Creative Destruction is the essential fact about capitalism

Joseph A. Schumpeter (1942), Capitalism, Socialism, and Democracy: 82–84.

For most of the modern era, news publishers have had a fairly stable and easy to define role in society: Journalists gathered information and produced stories, editors bundled them together, publishers distributed those stories in newspapers, broadcasts, websites and apps.

Since the turn of the millennium, digitisation has chipped away at that model – and in the coming decade artificial intelligence (AI) is poised to transform the model in fundamental ways.

The change is often perceived as a threat to news publishers. Declining direct traffic to news publishers, misinformation and fake news, job losses from automation, and the erosion of trust of news. All these developments matter and must be taken seriously, but they describe only the destruction of the existing model. The even more important question is: what is AI creating?

The transition to an AI-driven information ecology is not simply another technological disruption of news publishing. It is a fundamental reorganisation of how information is produced, maintained, discovered, interpreted, and used to create value. In the next five to ten years, AI will not only change how journalism is produced but how information itself flows through society. This will create a new information ecology of agents and flows that is likely to gradually displace the information ecology of articles, websites and apps that we know today.

When this happens, the central strategic challenge for news publishers is not how to use AI to produce and distribute news more efficiently. It is how to redefine their editorial and economic role when static articles, videos or podcasts cease to be the main unit of information exchange, and AI systems increasingly mediate the relationship between information and its users.

In this essay, I outline a positive vision for the strategic position of news publishers in an AI-driven information ecology. I argue that the central opportunity for news publishers is to transform themselves from producers of media products into companies that originate, maintain and commercialise trustworthy information across multiple markets. This vision does not transform the core mission of news publishers. Rather, it resituates the editorial mission in an AI-driven information ecology and

points at ways to realise the mission in an economically sustainable way.

The vision is not a prediction. It is an ambition or strategic goal that publishers can strive for. Further, it is based on my best guess of how the AI-driven information ecology will develop. This guess is inspired by the thinking of Shuwei Fang, David Caswell, Ezra Eeman, Florent Daudens and others, but represents my own synthesis of how publishers might position themselves in an AI-driven future.

The AI-driven information ecology emerges

Several times in modern history, technology has changed how information is produced, distributed and used. In the analogue information ecology, information was packed into finished media products such as newspapers and broadcasts and distributed to mass audiences. In the digital information ecology that emerged in the 2000s, information became digital objects such as articles and videos that were accessible through networks, and increasingly controlled via social platforms and search engines.

AI is now making a more fundamental transformation possible. AI systems are increasingly producing, retrieving, and synthesising information for specific purposes. Individuals are progressively experiencing information through AI systems that select, combine and interpret information on their behalf. Companies are now using AI agents to discover and analyse information that supports their operations and decisions. Communities are developing new information environments to knit them together and set them apart.

The deeper transformation is that the basic unit of information may shift from finished media products and static digital objects such as articles and videos, towards more fundamental units of knowledge that can be recombined to serve different purposes on demand. This makes it possible to separate the production of information from its uses. The same information can flow into numerous services and applications, creating value in different ways that are adapted to the intentions and contexts of individual active or passive users. In the coming years, this transformation is likely to deepen and widen.

As of now, this emerging AI-driven information ecology exists in parallel to (and to some extent, builds on) the traditional digital ecology of websites and apps that rose in the 2000s and matured in the 2010s, gradually displacing the analogue ecology. In the coming years, the emerging information ecology will grow relative to the traditional ecology and eventually become the primary mode for information to flow. Table 1 describes core dimensions of the three successive information ecologies.

Dimension	Analogue	Digital	AI-driven
Approximate period	Before 2000	2000–2020	2020s onward

Dimension	Analogue	Digital	AI-driven
Dominant technologies	Print, radio, television	Internet, search engines, social platforms, smartphones	Foundation models, AI assistants, agents
Principal transformation	Industrialised production and distribution of information	Networked access, algorithmic discovery, and distribution	Machine-mediated interpretation, synthesis, reasoning, and action
Primary human activity	Reading, watching, listening	Searching, browsing, navigating, sharing	Questioning, conversing, collaborating, delegating
Powerful intermediaries	Publishers, broadcasters, distributors	Search engines, social platforms, app ecosystems	Model providers, AI assistants, agent platforms
Core economic resources	Content, distribution capacity, audience reach	Attention, behavioral data, network effects, audience relationships	Trustworthy knowledge, proprietary data, cognitive capabilities, task completion
News publishers' primary role	Gatekeepers and distributors of information: select, verify, interpret, package, and distribute news to mass audiences	Producers and curators of news in a platform-mediated environment: create journalism, build audience relationships, and compete for attention and direct engagement	Providers of trusted journalistic knowledge and interpretation: produce, verify, maintain, and supply knowledge to both human audiences and AI-mediated systems

Table 1. The dimensions of three information ecologies.

How AI reorganises the information value chain

The shift to an AI-driven information ecology involves a reorganisation of the production, distribution, verification, and economic value of information. Leaders such as Shuwei Fang and David Caswell have offered suggestions and scenarios concerning the value chain enabled by AI-driven transformation. Loosely inspired by these authors, I describe the emerging value chain as four steps in a bidirectional chain, with value added via activities related to the production and processing of information.

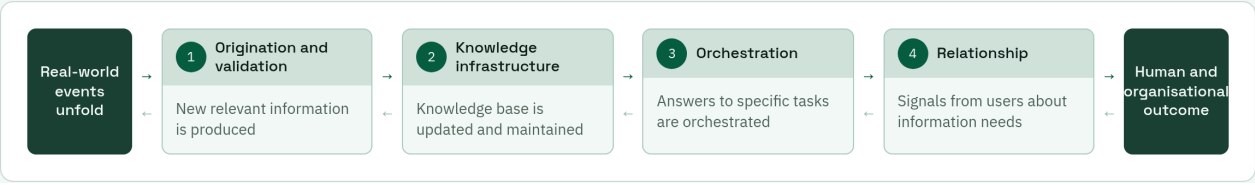


Figure 1. Two-way information value chain in the AI information ecology

In the first step of this transformed value chain, new information is added to the information flows of the information ecology as real-world events unfold. This is done by machines (e.g. via sensors) as well as by humans (empowered by AI), as some information production continues to require physical presence, as well as access and emotional connections to sources and situations.

In the second step, the produced original information is organised in a knowledge infrastructure containing the accumulated information that has been produced over time.

INFO BOX 1

What is a (journalistic) knowledge infrastructure?

In a news publishing context, a knowledge infrastructure can be a database that organises and maintains the facts about the world contained in the publishers' news flow. Such a database would contain information about relationships between entities and events (facts) that is tied to its source (provenance) and its validity (e.g. confidence and temporality). The database would have to be continuously maintained, as the world evolves and facts change. Processes of information addition, updating and extraction are likely to be highly automated, with human supervision utilised as needed to ensure the factuality and trustworthiness of the content of the database.

In the third step, information from the knowledge infrastructure is orchestrated to address specific information needs. The orchestration is carried out by AI systems that operate within guidelines and guardrails set out by the companies or institutions that control them.

In the fourth step, the orchestrated information is presented to users. This happens in AI services that provide answers, understanding, advice, decisions, completed tasks or actions – whatever is actively or passively requested by users.

This proposed new information value chain describes economic activities that can function separately as well as upstream and downstream information flows. Downstream, information is added, organised and transformed to fill specific needs. Upstream, signals about the information needs of users are channelled, informing the orchestration of information synthesis and providing data (new knowledge) that is in demand.

INFO BOX 2

Knowledge flowing downstream, demand flowing upstream in the value chain

The Danish Parliament passes a new law introducing tax deductions for energy-efficient home renovations. A journalist attends the parliamentary debate, interviews politicians and experts, examines the legislation and publishes a verified account of what has changed and why it matters. The reporting is incorporated into a shared knowledge infrastructure, where the law's provisions, eligibility criteria, implementation dates and sources become structured, verifiable knowledge. From here, the knowledge can flow downstream into potentially thousands of AI-powered applications and services.

Meanwhile, a homeowner asks their personal AI assistant: "Should we replace our gas boiler with a heat pump next year?" This individual need generates a demand for information that travels upstream through the value chain. The AI service translates the question into information requirements, and an orchestration system identifies and retrieves relevant knowledge concerning energy prices, installation costs, government subsidies and applicable legislation. Among the sources it draws upon is the journalist's reporting on the new tax deductions, which has already entered the shared knowledge infrastructure.

Note: Example generated with ChatGPT (Astra, high) based on the description of the value chain in the section.

The economics of the new information value chain

The example above illustrates how the same piece of information can create value in different contexts as it flows through the new information value chain. But it does not answer the fundamental economic question: Who creates value and who captures it?

Three developments are likely to shape the economic dynamics of the new information value chain:

AI expands information markets: AI is poised to expand the market for information by reducing the cost of producing and transforming information and increasing the number of purposes it can serve. This creates new information markets along the value chain that publishers and other information producing actors can participate in and (potentially) derive value from.

AI creates new scarcities: The growing abundance of information changes where scarcity – and thus, economic value – is situated. While the types of information that can be produced by everyone with AI are likely to become abundant and commoditised, two types of information are likely to remain scarce. First, original, trustworthy knowledge that AI systems cannot independently produce is likely to become more valuable as AI increases its uses. Second, access to the intentions and

context of users is likely to become an increasingly valuable resource, as it determines what needs to be met.

AI shifts bargaining power: The expansion of markets and creation of new scarcities shifts bargaining power towards organisations that control the new scarcities: Original knowledge and information about user intentions and contexts. Accordingly, organisations that control such resources are likely to be able to capture a large portion of the value created in the AI-driven information ecology.

While there are strong grounds for expecting that AI will expand the information economy and increase the value of journalism, it is not a given that publishers will be able to appropriate more economic value than today. Processes of creative destruction can create temporary or permanent imbalances between the ones who create value and the ones who capture it. Imbalances created by monopolies of the leaders of transformation (from Standard Oil, Google or potentially OpenAI and Anthropic) can become permanent because of scale economies, network effects or switching costs, as observed as early as Schumpeter's analysis of techno-economic transformations. This might create a need for intervention in the market to ensure socially beneficial outcomes.

Whether or not publishers will be aided by interventions in the market, their economic and journalistic future is likely to depend not only on whether journalism remains valuable. It will depend on whether they can translate their existing capabilities into positions in the value chain where they can capture a portion of the value that is created.

Three emerging positions for news publishers

The transition to the AI-driven information ecology challenges the already strained position that news publishers have clung to in the digital information ecology. It does so by further eroding their historical control of news production and news versioning, as well as moving the locus of information use from articles read on websites to multiple uses in AI services.

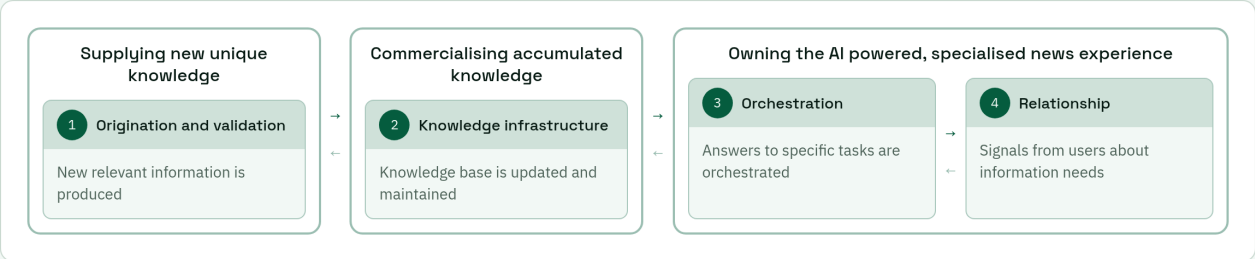


Figure 2. Publishers positions in the new information value chain

However, new opportunities exist for publishers that can convert their existing capabilities in ways that allow them to control scarce resources in the new value chain. While doing so will not be without challenges, three potential positions are

emerging:

News publishers as originators of scarce information (B2A): In this model, publishers continue to provide journalism as real-world events unfold - focusing on the news that does not already exist in the information ecology, and as such, requires human journalistic action and supervision to produce. This journalistic work will likely be increasingly AI-assisted, but focus on information production that requires humans to get physical or emotional access to sources and situations.

News publishers as knowledge providers (B2B): In this model, publishers depart from their traditional focus on producing and publishing news to providing access to their accumulated journalistic knowledge in ways that allow companies and institutions to get answers to the questions they need to perform work. This requires the publisher to build and maintain a knowledge base on top of their existing news production that is structured to make information maintainable (i.e. it must be updated when facts change), verifiable, and retrievable. Direct customers of such knowledge bases are likely to be companies and institutions that need factual knowledge about the domains that the news publisher continuously covers.

News publishers as news experience providers (B2C): In this model, the publisher retains its direct relationship with the broader audience. Instead of exclusively disseminating news stories in articles, videos or podcasts, the publisher provides an intelligent interface through which users can experience and explore news and apply journalistic knowledge to their interests and needs. This makes the news content more personally valuable through interaction, contextualisation, and assistance.

Each position builds on a different journalistic advantage, but each also faces distinct challenges in establishing a defensible market position.

The information originator position is vulnerable to keeping original news content exclusive to ensure bargaining power downstream in the value chain, in competition with other types of information originators and an expanding commodification of news content as AI capabilities grow. To address this competition, publishers must leverage their journalistic capabilities to produce original knowledge that is trustworthy, fresh and relevant to the intentions of individuals, companies and communities while at the same time defending scarcity through control over information access, rights enforcement, uniqueness of reporting, etc. vis-à-vis other information originators. This might also call for publishers to form alliances or consolidate in the originator step of the value chain.

The knowledge provider position will likely face competition from other specialised information services. This must be addressed by protecting the core scarcity – the exclusivity of significant parts of the information in the journalistic knowledge base – through a combination of uniqueness, continuity, and trustworthiness. Also, collecting a broad enough knowledge base might require extending information

collection from journalistic reporting to include other complementary types of information relevant to domain specific companies and institutions.

Finally, the news experience provider position is likely to be vulnerable to competition from general AI services (such as ChatGPT or Claude). This competition must be met with a differentiated experience focused on news, building on the publishers' existing direct relationship with users, which is still strong (and especially so in the Nordic countries). The most powerful differentiation will likely take factors such as editorial judgement, exclusive trustworthy knowledge, and specialised AI-mediated experiences of news content as the point of departure. The specialised AI-mediated news experience can emphasise experiences that prioritises what is deemed editorially important and allows users to discover and explore content in ways that match their intentions and their individual and social contexts. Indeed, a source of valuable differentiation can be tying the personalised news experiences together in an overarching shared experience that lets users know what other users in the community care about and that enables community deliberation.

	Publishers as information originators	Publishers as knowledge providers	Publishers as news experience providers
Position	Publishers originate information that is sold downstream (B2A)	Publishers maintain knowledge base that is sold to professional and institutional users (B2B)	Publishers offer specialised AI powered news experience (B2C)
Value proposition	Provides unique news needed to get relevant knowledge of the world as it evolves	Provides all information relevant to your professional intention and context	Allows you discover and explore news aligned with your intention and context as an individual or member of a community
Source of strategic strength	Journalistic point of departure (trustworthiness to individual, company, community)		
	Unique and trustworthy information	Accumulated proprietary knowledge	Strength of existing relationship with user
Source of strategic challenge	Can scarcity be controlled? (vis- à-vis other AI-powered information originators)	Ability to organise a competitive knowledge base (vis-à-vis specialised information companies)	Unique attractive differentiation (vis-à-vis general AI services)

Table 2. Publisher positions in the AI-driven information ecology.

The vertically integrated knowledge publisher

The three positions are distinct, but they are not mutually exclusive. While some publishers may choose to specialise in one position, larger publishers could benefit from combining all three. Such a combined strategy would enable news publishers to (re)organise to leverage all their relative strengths while creating synergies from economies of scope by serving multiple markets across the new value chain. This model could be called the ABC-model, as it allows the publisher to address Agents, Businesses and Consumers based on the same journalistic production and knowledge infrastructure.

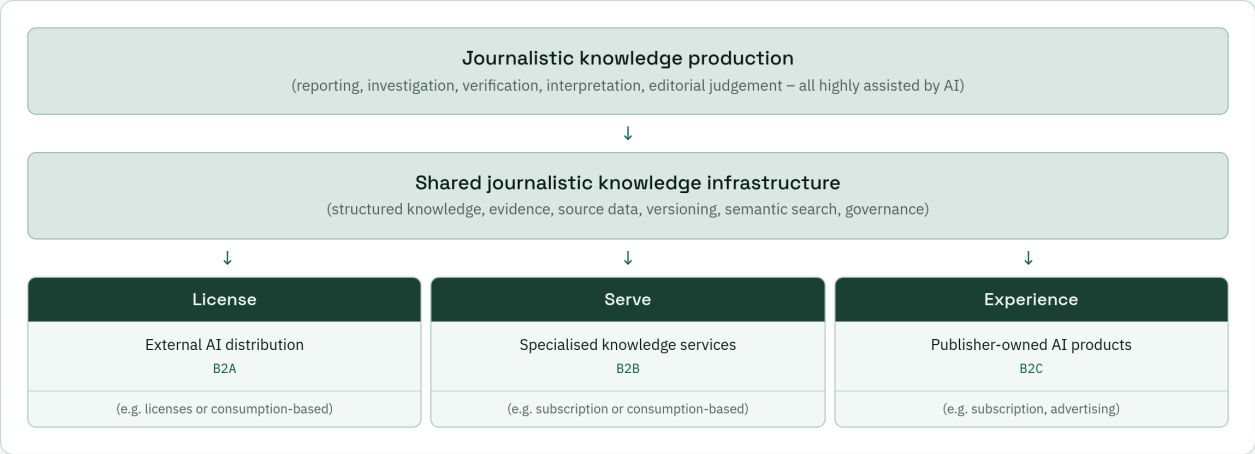


Figure 3. A vertically integrated news publisher in the AI-driven information ecology.

In this integrated model, the publisher continues to build on traditional journalistic values, methods, and capabilities, although increasingly powered by AI. Its core competitive advantage would lie in its ability to consistently and continuously originate information that AI systems cannot easily (re)produce: uncovering previously unknown facts, first-hand access to sources and events through established networks, independently verifying contested claims. Combined with editorial judgement and accountability, these capabilities will allow publishers to produce original, trustworthy and timely knowledge that is scarce and grows in value in an AI-driven ecology where it can be (re)used for multiple purposes.

Rather than being confined to individual pieces of published content, journalistic knowledge is accumulated in a shared infrastructure, continuously maintained, updated, and potentially enriched with other information relevant to audiences and customers. From there, journalistic knowledge can flow downstream to multiple markets: the news flow licensed to external AI services for grounding and verification of their responses, domain-specific knowledge sold to companies for integration in enterprise AI systems, AI-powered news experiences enabling users to access and explore journalism in ways adapted to individual intentions and contexts.

INFO BOX 3

How the vertically integrated news publisher originates knowledge about a local flood and channels it to three markets for journalistic knowledge

Imagine that a severe storm surge floods hundreds of homes along the Danish coast. Journalists from a national news publisher travel to the affected areas, interview residents, collect eyewitness accounts, examine official flood data and investigate how local authorities prepared for the disaster. Their reporting reveals that a municipality had repeatedly been warned about weaknesses in its coastal defences but failed to act. Through traditional investigative methods, assisted by AI tools for analysing documents and data, the journalists produce original, verified information that did not previously exist in the information ecology.

Rather than simply publishing a series of articles, the publisher incorporates this reporting into its shared journalistic knowledge infrastructure. The infrastructure contains structured knowledge about the flooding, affected geographical areas, infrastructure damage, government responses, previous warnings and the responsible institutions. Each piece of knowledge is linked to its underlying evidence, sources and publication history. As new information emerges, the knowledge base is continuously updated, allowing the publisher to maintain a verified and evolving representation of the event.

This shared knowledge can now create value in three different markets:

External AI distribution (B2A). A general-purpose AI assistant serving a homeowner considering buying property in the affected area accesses the publisher's licensed knowledge to explain the flood risks and the municipality's history of managing them. Another AI service uses the same knowledge to answer questions about the political consequences of the disaster. The publisher receives licensing or consumption-based revenue when external AI systems access its verified knowledge.

Specialised knowledge services (B2B). The publisher offers a specialised climate and infrastructure intelligence service to insurance companies, property investors and infrastructure operators. Combining its original reporting with public records, historical flood data and other relevant knowledge, the service continuously monitors risks, regulatory developments and institutional responses. An insurance company might use the service to identify previously undocumented risks to coastal properties, while an infrastructure investor might use it to assess the quality of local flood protection. Customers pay for access to continuously updated, domain-specific intelligence.

Publisher-owned AI news experiences (B2C). In the publisher's own AI-powered news service, a resident of the affected area asks: "Why did the flood barriers fail, and what is being done to prevent this from happening again?" Drawing on the same knowledge infrastructure, the service provides an interactive, personalised explanation, combining investigative findings, maps, timelines, expert assessments and the latest political developments. The resident can explore the evidence, ask follow-up questions and subscribe to updates concerning their neighbourhood. Other users might explore the political accountability, economic costs or broader climate implications of the same event.

Note: Example generated with ChatGPT (Astra, high) based on the description of the vertically integrated news publisher in the section.

If adopted, the integrated ABC-model would realise synergies enhancing the publisher’s ability to confront the strategic challenges associated with each of the three positions it integrates. Indeed, the synergies that arise from the economies of scope that the exploitation of information reusability enables might be essential for recovering the still high costs of original reporting (even when empowered by AI), enabling investments in journalism that distributing to one channel alone could not sustain. However, naturally, by entering three markets, publishers will need to carefully manage channel conflicts by traditional means such as differentiating the products, resources and rights offered through each channel.

The future of the journalistic mission

The vertically integrated news publisher is more than a possible economic answer to the transition to the AI-driven information ecology. It is also an answer to how news publishers could reinvent how they realise their editorial mission. By making original and trustworthy knowledge available across a wider range of services and contexts, publishers potentially strengthen their mission in an economically sustainable way.

However, addressing these opportunities does raise questions of how journalistic methods can be reinvented to address these opportunities in ways that remain true to the mission. When information is continuously updated, transformed and reused for different purposes, publishers must develop new methods that ensure the accuracy and provenance of knowledge both at the time of publication and over time. When knowledge is transformed to address individuals’ intentions and contexts, AI-powered news experiences must balance what the user wants and needs, and what is deemed journalistically important. And when journalistic use of AI at scale is likely to become part of news origination itself, developing new ways of ensuring editorial judgement and emotional connection when humans steer processes at higher levels of abstraction will become necessary. Such inventions are likely required to maintain journalism in the AI era, and could also prove to be essential sources of differentiation and competitive advantage in the future information ecology.

The transformation ahead will not be easy. But it offers the opportunity that the destruction of our familiar ways of being a news publisher can be followed by the creation of something new and stronger: A reinvention of the news publisher that strengthens, rather than merely defends, the journalistic mission.



05

The next era of journalism will be defined by service, not publication

Nikita Roy
Founder of Newsroom Robots Lab

What if journalism's greatest AI opportunity lies not in producing more content, but in becoming a more useful public service? In this essay, Nikita Roy (Founder, Newsroom Robots Lab) challenges the idea that journalism's responsibility ends at publication. She envisions newsrooms that organise their work around the questions people need answered, using AI to build lasting, accountable services that help the public understand what matters, what has changed and what happens next.

We have spent decades making journalism easier to publish, faster to distribute and simpler to consume. Yet we still ask an extraordinary amount of the people we serve: notice the story, recognise why it matters, remember what happened last time, find the update, and figure out what has actually changed.

We do the difficult work of reporting. Then we leave much of the work of staying informed to the reader. That is not just an audience-engagement problem. It is a limitation in how we have designed journalism itself.

An investigation can expose a failing institution. A report can reveal that a city has promised a new bus route, or that a school closure is being considered. But the question people live with rarely ends when the story does. Did the promise hold? Was the closure approved? Who benefited, who lost out, and what happens next?

A finished article is not always a finished public service. The most consequential possibility AI offers journalism is not that we can make more content. It is that we can finally begin to build journalism that stays with the questions people need answered.

The article is a form. Journalism is a function.

I often make this distinction in my work with newsrooms: verification, scrutiny, explanation and helping people navigate uncertainty are journalism's functions. Articles, broadcasts, newsletters and podcasts are forms we use to deliver them.

Those forms matter. A narrative investigation can reveal a pattern no collection of personalised answers could capture. A shared story can create public attention and political pressure. The point is not to replace the article. It is to stop treating the article as the boundary of our responsibility.

Imagine a newsroom organised not only around stories to publish, but around public questions it commits to follow. Will the promised housing be built? What is happening to local water quality? How is a new policy actually affecting families? Each question has a history, sources, unresolved claims, competing accounts and future developments worth watching.

Reporting still creates the knowledge. But in a newsroom organised around questions, that knowledge becomes something more durable than a sequence of links: a living, accountable record that can be revisited, interrogated and updated as the world changes.

The invitation to the audience changes from "read this" to "we will keep watching this." Not because every reader wants a personalized news feed, but because most people do not have the time to become expert navigators of every institution shaping their lives. This is a more demanding promise than an alert whenever a new article appears. It means noticing consequential change even when no new article was

planned. It means sometimes saying that nothing has been resolved. And it means acknowledging when the newsroom can no longer reliably keep watch.

Making continuity easier to sustain

None of this began with AI. Community-centred newsrooms have long answered individual questions, followed public promises and built practical information services. To take one example: Outlier Media in Detroit has run a text-message service since 2016 that delivers housing and utilities information on demand and lets residents reach a reporter directly. It now uses AI for property lookups and to surface what residents are asking about, while journalists still write every message. It is journalism designed around what people need to know, not merely about what gets published.

What changes with AI is the range of connecting work that may become feasible. AI can help sift sprawling agendas, compare revisions to public documents, connect a new development to years of reporting, translate complex language, and adapt verified explanations for different levels of familiarity. It can help someone understand an issue without first knowing which agency, committee or legal term to search for.

The shift is subtle but profound as we move from publishing information and hoping people find their way through it, toward maintaining a service that helps them understand what matters as it unfolds. The newsroom still needs reporters who know which questions to ask, editors who judge significance, and clear records of where each claim came from. AI may lower the cost of connecting information but it does not remove the cost of being responsible for it.

In fact, the better these services become, the more important that responsibility becomes. If an error travels into a text message, translation or audio briefing, a correction cannot live only on the original article. This is where our editorial responsibility becomes product design.

The future of news is bigger than a chatbot

Much of the current conversation about AI and news begins with a familiar interface: a box where someone types a question and receives an answer. That can be useful. But it is an incomplete vision of journalism’s future.

A service built only around answering users’ questions leaves a crucial gap which is that people may not know what they need to ask. Journalism matters precisely because people often do not know what is happening, what information powerful institutions would rather they overlook, or what question they should be asking in the first place. An interface that answers perfectly formulated questions can still leave the biggest blind spots untouched.

This is why I am more interested in systems that combine responsiveness with editorial initiative. A person might want updates about a particular school. The newsroom can explain that school's decision while also revealing how the policy affects other neighbourhoods, what evidence officials have omitted and why the wider choice deserves attention. Personal relevance should be a doorway into public understanding, not an escape from it. People may receive different formats, languages or levels of explanation, but not receive different underlying facts. Nor should an algorithm decide that an inconvenient truth is irrelevant because someone did not ask for it. And the service need not look like AI at all. It might arrive as a text message, a short audio explanation or a carefully timed update. Its form is irrelevant, as long as it makes journalism more useful.

What becomes of the newsroom?

There is a strategic question here as well as a civic one. If more people encounter reporting through assistants and platforms they do not associate with the publisher, newsrooms risk becoming invisible suppliers of raw material for somebody else's product and business model.

Publishing more articles will not, by itself, solve that problem. A newsroom's enduring value may lie increasingly in the knowledge it can stand behind: what it has verified, what remains disputed, what has changed, and which public questions it continues to investigate.

This also changes the work. The reporter remains indispensable, not as a human stamp on machine-produced prose but as the person who understands context, challenges official accounts, notices what is missing and earns trust. Editors become stewards of commitments that extend across time. Product and engineering teams help turn journalistic judgement into systems that can be tested, corrected and sustained. The newsroom becomes more than a place that produces coverage. It becomes an institution that remembers, checks and returns.

Sustaining this work still requires journalists who care whether a public promise was broken, earn trust in their communities and judge which changes matter. The most interesting use of AI is to make those human capabilities reach further.

The journalism we could not build until now

For much of journalism's history, we have had to make difficult choices between depth and reach, between reporting a story and helping each person understand its implications, between following an issue and keeping thousands of people meaningfully informed as it develops. Technology does not eliminate those trade-offs. But it may change them enough to let us imagine services that once seemed impossible to maintain.

A newsroom that remembers every public commitment it has investigated. That connects a new decision to the reporting that preceded it. That can explain what changed in the language and format someone needs, without sacrificing a shared factual record. That notices the consequential development a reader never knew to search for. That knows when it is uncertain, when it has failed, and when it owes someone a correction.

This is not journalism made frictionless or automated into perfection. It is journalism whose value continues between publications. And perhaps the most important change is in how we understand success. It would mean asking more than how many stories we produced, how much attention they attracted or how efficiently we distributed them. Instead, the value of journalism would be measured by whether people were better able to understand the forces shaping their lives, recognise what had changed and exercise their own judgement.

The future worth building is not one in which AI does journalism for us, but one in which AI allows journalism to fulfil more of its purpose.



06

Journalism with room to move

Ezra Eeman

AI & Innovation Director, WAN-IFRA and Founder, Wayfinder

Can journalism become more personalised and useful without taking away our freedom to think, question and discover? In this essay, Ezra Eeman (AI & Innovation Director, WAN-IFRA and Founder, Wayfinder) explores the possibilities of AI-powered journalism that helps people act on information while preserving space for curiosity, surprise and independent judgement. He introduces the idea of the story that acts, arguing that the best future news experiences may be those that know when to help, when to pause and when to leave people alone.

There is a Japanese word that has been on my mind this summer: ***Ma***. It is hard to translate. I would describe it as an interval. A pause. The space between things that gives them their shape. The rest between notes that turns sound into music. The empty room in which something can still happen. It is space that is held open on purpose.

What if AI could do exactly that for journalism in the future? Give back the space to think rather than crowding out our thoughts?

I admit it feels counterintuitive. Everything points to media moving in the opposite direction. More efficient. More scaled across devices and surfaces. More tailored to our context and behaviour.

From attention to intention

For a long time, digital media has organised itself around attention. We have learned to measure it, predict it and sell access to it. Feeds, alerts and recommendations compete for every free second. Even when a newsroom's ambitions are much broader, the practical conversation often returns to the same things: how to get someone to open the story, how to keep them there, how to persuade them to come back.

Agents are now moving that conversation towards intention. What is someone trying to achieve? What can we help them do? Answer engines give you the answer, not the page where you yourself may find it on. Personal agents promise to brief you in your preferred tone, at your preferred length, about your preferred topics. There is an obvious appeal in organising media around what people are trying to understand or do. Helping a person understand a bill or prepare for a public decision feels more useful than finding another reason for them to scroll. But a system that serves our intentions can still enclose us. Especially when it starts anticipating them, deciding what we probably want before we have had the chance to think about it ourselves. Media that fits like a second skin.

It’s all very exciting. It’s also slightly suffocating.

Because if we carry the logic through, we could end up with a media environment that fits us so closely there is very little room left to move. Every interest recognised. Every hesitation interpreted. Every gap filled with something the system thinks we will find useful. It might be a very comfortable environment. It might also become quite hard to encounter anything that changes the shape of it.

There is a word for that in machine learning: overfitting. A model overfits when it becomes so tuned to its training data that it struggles with new examples. It can perform very good on yesterday’s tasks and fail on tomorrow’s.

The neuroscientist Erik Hoel has a beautiful hypothesis about this. He suggests that dreaming may help prevent our brains from overfitting. In his account, the strangeness of dreams, their wrong turns and impossible combinations, is the point. It injects noise into a mind that would otherwise become too tuned to the narrow routine of its days. The idea is that the brain breaks its own fit so it can keep generalising to a world it hasn't seen yet.

An overfitted media environment is one that knows you so well it can no longer surprise you. I wonder what that does to us over time. How do we learn to cope with a changing world if so little of what we encounter is unfamiliar?

Making room for judgment

But it doesn’t need to be this way. The same technology that games our attention and nudges us towards a familiar path could be used to give us some of that room back. Protecting the space the attention economy wants to fill, and the agentic delegation wants to close.

The future I would like journalism to help build, over the next three to five years, leaves more space than that. By 2030, I would like us to have products that make reporting much more useful in people's lives, while allowing for surprise, unfinished thoughts and the freedom put it down without being drawn into another interaction. Journalism that helps us navigate a changing world and gives us a chance to notice something we had never thought to ask about.

That means being more ambitious about what stories can do. It also means being more careful about what we optimise for.

Sometimes we have a poorly formed question. Sometimes we change our mind while trying to answer it. And sometimes something matters precisely because we had no idea it existed. The unknown unknowns. You cannot ask for what you do not know you are missing. Journalism has always had a role in bringing those things into view. An investigation can introduce a problem nobody had recognised. A report from another country can shift how we see our own. A good interview can leave us with a question that stays around for days.

This is where the idea of Ma can guide us. It asks what might happen in the interval between what a system offers and what a person makes of it. Where there is room for interpretation, disagreement, curiosity. We can design that interval out of the experience. Or we can choose to protect it.

Make no mistake. I am not pleading to make information unnecessarily difficult to use. A person who needs a clear explanation should get one. A person struggling with language, time or access should not have to perform extra intellectual work to qualify for journalism. There is plenty of friction worth removing. The question is where removing it helps someone, and where it starts to take over the work of forming a judgment.

A story might help me compare two proposals without selecting the one it thinks I should prefer. It might explain what is known while leaving a disputed question visibly unresolved. It might give me a place to begin without immediately arranging the rest of my journey. There is a difference between leaving someone lost and giving them room to find their own way.

The story that acts

I call this idea the story that acts. A story that both has and gives agency. It doesn’t only tell you what happened but can also offer a service grounded in the reporting: help with a calculation, a comparison or a practical next step. Rather than a generic chat box under every article, it begins with what this particular piece of journalism has established and what someone might be able to do with it.

As the story becomes an active participant, it gets better at anticipating questions and helping with next steps. It should never make the reader less involved in deciding where to go. There should be room to refuse its suggestion, follow a different thread, or arrive at a conclusion the newsroom did not expect. If there isn’t, we have built a more capable product but perhaps a less capable relationship.

Imagine a local newsroom in 2030 reporting on plans to redevelop an old industrial site. It has examined the council's documents, spoken to residents, questioned the developer's claims and looked at the promised housing and transport connections. It publishes an article with a clear editorial account of what the proposal would change, who stands to benefit and which questions remain unanswered.

You read it because the site is near your home. Alongside the story is a small offer: if you choose a street, the reporting can help you explore the proposed changes around it. You might look at the new walking route, compare the current plan with the previous version, or ask what evidence supports a claim about traffic. The answers draw on the newsroom's work and show where that work runs out.

Perhaps your first concern is whether the road to your child's school will become busier. While exploring it, you encounter reporting about residents who may be displaced by the development. You had not gone looking for that. The story makes the connection visible because it matters to the public decision, even though it sits outside your original question. You can stay with it. You can also return to your original query. The experience leaves both possibilities open.

If you want to respond to the consultation, the story can help you prepare. It can locate the relevant documents, explain the questions the council is asking, and help turn your concerns into a draft you can edit. It might point out that one of your assumptions is not supported by the plan. But it doesn’t decide your position or manufacture your agreement.

You can then ask to hear when the plan changes in a way that affects the question you were following. For several weeks, you hear nothing. The monitoring continues,

but no material change has met the agreed threshold. When a revised proposal does appear, you get an explanation of the difference, with a route back to the wider story. When the consultation ends, the offer to submit a response closes too.

In this example, acting and leaving space belong together. The system takes on some of the work of finding, comparing and following. The reader retains the work of deciding what matters, what they believe and whether they want to participate. The story can bring something new to your attention without asking you to do anything about it.

The same principle could apply to a consumer investigation that helps someone examine their bill, or a transport story that compares the implications of a timetable change. These are relatively modest uses. They are also meaningful. Someone might leave with a better question, a more accurate understanding, or a practical problem resolved.

Not every story needs to turn into a service. There should still be reporting we read because it is beautiful, disturbing, funny or difficult to place. A story about a life very different from our own does not need to help us do something. It can simply enlarge our sense of the world. If we replace ‘attention’ with ‘task completion’ as the universal measure of journalism, we will have found another way to make its purpose too small.

In the journalism of tomorrow, there should also be room for encounters we cannot fully explain in advance. A finite edition can include a story chosen by an editor because it deserves to be seen. A story can invite you to explore without deciding where you should end up. A reader can linger on an image, follow a small detail, or stop halfway through. A thought might surface later, while they are doing something else. Not every moment needs to become a signal used to make the next experience more efficient.

What this asks of newsrooms

All this requires room for hesitation on the other side of the story as well.

To make this possible, newsrooms would need tools that help reporters start with a question before producing an answer. Tools that surface conflicting evidence, identify a missing perspective or let someone try an idea without immediately turning it into an article. Sometimes the most useful intervention might be to leave a fragment alone while the reporter works out why it matters.

For a newsroom, this would involve learning to describe the usefulness and limits of its reporting more precisely. What has this story actually established? Which people or decisions might it affect? What can we responsibly help someone understand, compare or do? What would change our assessment? Those questions are editorial work. They cannot simply be passed to a model after publication.

Journalists would need to become more comfortable working with designers and developers on how their reporting could behave more like services. How do you test whether an adapted explanation still preserves the most important facts? How do you recognise when a helpful suggestion becomes a nudge towards a particular outcome? At what point does this turn journalism into activism?

The responsibility continues after publication. A calculation can become outdated. A correction may affect an answer someone has already received. A newsroom offering these capabilities needs someone who can maintain them and a way to withdraw them when their basis no longer holds. By 2030, this could start with a few stories that offer a useful service, and grow only if the benefit is worth the work of maintaining them.

A new sense of value

I do not have a definitive economic model for this more breathable form of journalism. A calmer interface does not make reporting cheap. An acting story can create new costs as well as new value. And the people who would benefit most from help navigating public decisions may have the least money to spend on it.

Still, there is a different value exchange worth exploring. What might people value enough to support, beyond access to more content?

- **Being done.** A product that respects their time enough to end. An unread pile can bring fatigue and guilt, however good the journalism. A useful ending might make a subscription feel like something that helps rather than another obligation.
- **Being able to act.** Stories that help them do something that matters in their lives, with a newsroom accountable for the reporting and the tools it offers.
- **Being surprised.** A trusted source that regularly makes their world larger, not smaller.
- **Being able to trust the silence.** Knowing a story is still being followed without being disturbed until something that matters has changed.

Start with a ma brief

We do not need to wait for an agent that understands everything about us to explore this future. A simple experiment over a couple of months, built around one continuing issue, is enough to get started: a school merger, an energy transition, a hospital reorganisation, a housing plan. Build one story that can help people understand the issue and, if they choose, do something about it.

But before anyone writes a line of code, before the feature list and before the prompt engineering, write a Ma brief. One page. Four questions.

- Where does it end? When does a person know they are done, and how does the product tell them?

- Where does it pause? Which moves should wait, ask twice, or suggest sleeping on it?
- What does it open? What will it offer that nobody asked for, chosen by whom, and labelled how?
- When does it fall silent? What is the threshold below which it says nothing, and how can a person check that the story is still being followed?

Keep the first version small. Offer a few supported ways to explore or apply it. Make it easy to see how your question connects to the wider issue. Test it with a small group that includes people who rarely use the newsroom's products. Compare their experience with the article on its own and the newsroom’s existing alerts.

Did the people testing it understand the issue more accurately? Did they encounter a relevant question they had not considered? Could they take a useful step on their own terms? Could they finish when they had what they needed, without feeling pushed to keep interacting or take a suggested action?

Record the failures and maintenance costs as carefully as the successful uses, then share the findings publicly so other newsrooms can learn what helped, what failed and what it took to maintain.

At the end I would ask everyone who has tried this Ma-experiment a simple question: after spending time with this journalism, do you have more room to think and act for yourself?



CHAPTER III

Rebuilding Newsrooms

07 Gard Steiro
08 Martin Schori
09 Olle Zachrison
10 Sofie Hvitved
11 Sara Inkeri Vardar



07

The noble art of flogging dying horses

Gard Steiro

Editor-in-Chief, VG and SVP News Media, Schibsted

How do you lead a news organisation through a transformation when nobody knows what the future holds? In this essay, Gard Steiro (Editor-in-Chief, VG and SVP News Media, Schibsted) draws on sinking ships, dying horses and his own experience of leading one of the Nordics' largest newsrooms to argue for a culture of constructive paranoia. He offers ten concrete moves that media leaders can make today to prepare their organisations for an uncertain, AI-mediated future.

Let’s start with a grim truth. Many venerable news organizations will, in all likelihood, die an undignified death. The second piece of bad news is that many talented journalists risk ending up in a seemingly endless unemployment line. In a new, AI-mediated reality, there may simply no longer be a need for them.

The hope lies in this: the future of journalism is not predetermined. Independent news media can still play a crucial role in the future ecosystem of information and public debate. But it will require major change. In this essay, I propose what some of those changes could be.

Before we get there, I want to introduce you to mister Wallace Hartley, a man I increasingly find myself identifying with.

Hartley, born in Lancashire in the late 19th century, was the son of a Sunday-school teacher who introduced his local congregation to the hymn Nearer, My God, to Thee. And that is precisely what young Wallace eventually did: he moved closer to his Maker at a speed of 21 knots.

On 10 April 1912, the violinist Hartley signed on to what was then the largest and most advanced means of transport in the world — more astonishing in its day than self-driving cars and high-tech drones are today. Three days later, Hartley stood on deck with the rest of his orchestra. They were wet, cold and frightened, but apparently not ready to put down their instruments. They remained at their posts, entertaining their audience into the night.

Then RMS Titanic sank.

More often than I care to admit, I identify with bandleader Wallace Hartley. Not musically, obviously. But we share a certain fate. A few years ago, not a living soul believed that VG, the successful Norwegian news organization I lead, could sink. Now I find myself standing in the cold, waiting for the ship to list. We have hit an iceberg. Possibly several at once. They have torn open the hull and exposed the fact that our construction was not quite as watertight as we liked to believe. Water is pouring in. We may sink.

But we keep playing! Just as before. And me? I wave the conductor’s baton and urge the newsroom to keep time: “Nearer, my God, to Thee, nearer to Thee!”

Most people in media understand the metaphor. The Titanic represents the proud news organizations of the Western world, some with roots stretching back to the age of the Morse telegraph. These institutions are shaped by strong cultures that have served them well, but those cultures have also cemented working methods and created a mentality that can, at times, border on arrogance.

The iceberg in my analogy can have many names: OpenAI, Meta, Perplexity or Claude. Agentic consumption, liquid content, creators. Take your pick. Every new

competitor, every shift in behaviour, has the potential to tear traditional media business models to pieces. We can already see the tips of the icebergs above the waterline, but none of us knows how far they extend beneath the surface. The only thing we can be reasonably certain of is that they are much larger than what we can see with the naked eye.

Our great challenge is this: a ship the size of the Titanic cannot change course in a second. Unlike smaller vessels, we are not built to navigate quickly and efficiently through treacherous waters. In conditions like these, it is easier to be a start-up defined by curiosity, technological understanding and agility than an establishment with the self-image of the Sun King.

Some will find these reflections provocative, perhaps even frightening, or to be thoughts an editor would be wiser to keep as an inner monologue. I strongly disagree. The most dangerous thing we can do as media leaders is to downplay the challenges facing journalism. That is why I use metaphors like this regularly with the organization I lead. I am open about the fact that despite all our progress, despite reaching half the Norwegian population every day, despite 80 per cent of our users coming directly to us, we could still be dramatically diminished within a few years.

In my view, it is impossible to create a genuine understanding of the need for change in a legacy organization without being brutally honest about the weaknesses in today's journalism and the products we offer our audiences. There is no shortage of people claiming to possess the objective truth about the future of independent news media. I, for one, have limited faith in self-appointed oracles and travelling futurists. I trust the people who are genuinely uncertain — and who dare to admit it. I think uncertainty builds trust, at least in organizations populated by thinking people and critical journalists.

If I had to describe the culture I want to create at VG — the state of mind that I believe gives us the best possible chance of facing the future — I would call it constructive paranoia. We should always be alert, never take our users' trust for granted, and always look for ways to improve — incrementally and radically — but never become paralysed by uncertainty.

No-regret moves to keep the horse trotting

In Norway, we have an old expression about carrying oats to a dead horse. The English equivalent would be flogging a dead horse. It describes projects, initiatives or arguments that are doomed to fail. Faced with an agentic internet, chatbots capable of producing generic news faster than any newsroom, and platforms devouring our business models at roughly the speed of a hummingbird's heartbeat, it is not entirely unreasonable to compare traditional media to horses receiving palliative care.

I take great pride in trying to keep the horse alive. I believe there is a cure. I believe we can gallop again. But I am fairly certain we will not do so in the same form as

before. The problem is that we do not really know what will be expected of news media in the future. We do not know who our competitors will be, or how the information ecosystem will work.

We are in a period of transformation. It may last several years, perhaps less. Either way, we should make every move we can today to prepare ourselves for a very different tomorrow. In a landscape where we have no map and barely recognize the terrain, the smartest approach is to look for no-regret moves. Fortunately, there are quite a few. Below, I list ten suggestions.

1. Understand what users expect

Yes, it sounds obvious, but the truth is that too many newsrooms still fail to navigate according to signals from their audiences. If we are to succeed in the transition to an AI-driven media world, the first commandment is to use our resources as effectively as possible. That requires understanding which problems we are supposed to solve for an audience whose expectations of journalism are constantly changing.

We need quantitative data, but we also need to meet our users face to face, on equal terms, if we want to understand their needs and expectations.

The days of distribution monopolies are over. Soon, so too will the era when audiences came to us simply to receive information they had neither asked for nor particularly needed. We are entering a period in which users will have different reasons for seeking out news media. We are moving from attention to intention. That is a profound change. If we do not understand our users' intentions and needs better, we will fall behind. And if our strategy is simply to continue making the products we think are brilliant and hope that audiences will eventually understand why, we will fail.

2. Structured journalism

The article is not sick. It is stone dead. I do not mean that the article as a format will disappear, but that there is little relevance in journalism today still being planned, produced and structured around a format created for cellulose-based publications that enjoyed their golden age while Mark Zuckerberg was still in diapers.

If we are going to version our content and serve a diverse audience with news in the form each user actually needs, we have to change the way we structure our proprietary data. It must be readable by machines as well as humans. Journalists need to focus on the story, not the article. We have to prepare for a world in which much of the value lies in atoms of journalism — quotes, images, audio, verified facts and data.

I do not know whether, in the future, machines rather than humans will be our most frequent visitors. But I am absolutely convinced that, like our new competitors, we need to offer some form of liquid content. Not exactly as they do, though. We must do

it in a way that preserves our editorial, ethical and professional standards.

3. Culture eats everything

The media world is full of executives who go on study trips or attend conferences where ideas and insights are endlessly recycled, then return home clutching a new tool they promptly order the entire staff to use. My view — hardly a controversial one — is that media companies focus far too much on technology and far too little on creating a shared understanding of the challenges and opportunities ahead.

A new tool will not save the media. Nor will a clever new CMS. What matters is our ability to act quickly. What matters is a culture of innovation. That kind of culture is extremely difficult to create if there is no trace of it to begin with, but it is more important than ever to give people and teams with the ability and willingness to rethink journalism the freedom to experiment.

Highly hierarchical media organizations with a taste for top-down decision-making will struggle to keep up and to create a genuine learning culture. One of the most important things the media industry can do now is invest in education. We need to educate our staff in AI. Because this shift is not primarily about teaching people to use new tools. It is about giving them the knowledge they need to apply new technology intelligently within their own domains. That is how we create real value from AI. It does not help if the people at the top understand what is hitting us if the crew is given neither the time nor the space to prepare.

4. Document every workflow

What actually happens inside a newsroom? How is journalism produced? How is it quality-assured? What is our secret sauce? What distinguishes the content produced by your newsroom, your department or your beat from everyone else’s? In traditional media, surprisingly many of these questions remain unanswered. The knowledge exists, but it is not sufficiently documented.

Without that insight, it is difficult to realize the potential of AI, and it is almost impossible to reprioritize resources. If you do not understand where your resources are being spent, you risk becoming stuck in old methods simply because that is how things have always been done. Insight is a prerequisite for change.

At Schibsted, we have now done the heavy and rather boring work of mapping all our editorial workflows. And for the first time, we have data on every newsroom process across every Schibsted brand. We have assessed time spent, tool support, AI potential and urgency.

Why? Because we now have a much better understanding of where we can free up capacity. And we know where to focus our AI efforts in order to move resources towards unique journalism.

5. Automate as much as you can

There are many important and productive discussions about how journalists should use AI. Different media organizations have different identities and will make different choices. That is entirely natural. Some may want to market themselves as created by humans. Others will choose to use AI much more actively. There is nothing wrong with pursuing different strategies. It comes down to understanding your audience.

But I believe every newsroom will benefit from automating wherever it can. AI gives us the opportunity to improve and streamline almost every part of the journalistic process. We can find more information, process it faster, improve our analysis of complex datasets, version content more effectively and write faster.

Rather than moving cautiously, I think media organizations should push towards the boundary. Automate until it hurts a little. Because only by doing that will we understand the possibilities in front of us. And, more importantly, only then will we be able to free up the most valuable resource in any newsroom: people.

6. Double down on what AI cannot replace

I have said this before, and I am happy to say it again. Adopting AI is not the difficult part. Anyone with a little technical understanding can now build new editorial tools. The much harder task is identifying the journalism we can actually make a living from after the paradigm shift.

We need to understand what each individual media organization does that genuinely matters to its audience, and define the kinds of content that machines cannot replace. The answer will vary from one organization to another. But in my view, it will involve many of the classical virtues of journalism: building source networks, finding information that someone is trying to keep hidden, analysing complex issues, and being out there in the real world alongside the people we serve.

Seen from that perspective, AI could turn out to be extremely positive for news media. It gives us the opportunity to push more of our effort — and more of our journalism — higher up the value chain. That sounds wonderful. In reality, it is difficult. Because this is not only about time. It is also about competence.

Anyone who has spent time in a newsroom knows how much effort and talent it takes to build a good source network. Reporting from a front line is not for everyone. Many of us simply do not have the qualities required to communicate with real credibility and authenticity from difficult situations. One plausible scenario is that there will be fewer journalists in the future. But the very best — those with unique skills and qualities — may have a brighter future than ever.

7. Trust, trust, trust

What will be the most important KPI in the future? Some will say scale and reach. Others will point to the direct relationship with users. I believe trust will become even more important than it is today.

If we lose the trust dimension, convenience wins. And if convenience is all that matters, people might as well ask a chatbot instead of seeking out an independent news organization. But trust is not simply about whether people believe the individual facts we publish. Trust has many drivers. Relevance, for example. The ability — and willingness — to correct mistakes. Adherence to recognized and publicly understood standards of journalistic ethics.

In my view, it is absolutely essential that media organizations become more transparent about our editorial decisions, our ethical trade-offs and our dilemmas. We take it for granted that audiences understand the difference between journalism and all the other information they encounter. We should not.

There is an opportunity here. I believe greater transparency builds trust. And I am convinced that we can build our ethical standards and internal quality-control processes directly into the systems we automate. That could become a genuine differentiator in the new media market.

8. Launch speedboats

When everything changes quickly, there is a danger of becoming paralysed and falling behind. Many of us will face the same Gordian knot. We have loyal users who are very happy with the products we already offer. They do not want radical change. They may be sceptical of new storytelling techniques and artificial intelligence. At the same time, there is a potentially huge new audience expecting an entirely different user experience. They will not use us unless we change radically.

Satisfying both groups at the same time is difficult, perhaps impossible. The danger of changing established products too quickly is that we lose the very people who are paying for our transformation.

The answer is not to sit still. Our existing products have to evolve. But we cannot move too quickly. One answer is to create sandboxes in which we can think from scratch.

At VG, we have chosen another maritime metaphor. We call them speedboats. These are real products in the market, but designed for the next generation of news consumers. There, we can think differently, take bold decisions and experiment without risking the trust of — or simply irritating — the users of our main products.

There is no guarantee that these speedboats will develop into products that grow and survive. That is not necessarily the point. The learning is the point. What we learn from these experiments can gradually be brought back into the mothership once we

have evidence that the changes actually make things better.

9. Don’t get lost

When everything changes quickly and visibility is poor, you have to hold on tightly to the compass. The opportunities for media are almost endless. We can create new products and take journalism in directions we could barely have imagined a few years ago. In that kind of environment, it is easy to drift off course, forget the mission and throw the craft overboard.

Many people misunderstand the core of journalism. They think it is about writing, formats and politics. That is not even half the story. Journalism is about giving citizens the information they need to make informed choices. That is a prerequisite for a functioning democracy.

But journalism is about something else as well. As we navigate the rough seas ahead, we must not forget the people we are here for, the people who depend on us. People and vulnerable groups subjected to abuse or neglect. People carrying stories that need to be told. People who depend on us to be heard, seen and taken seriously.

That part of journalism is unlikely to be replaced by technology. It remains our responsibility. Regardless. And always.

The dream scenario

So what is my dream scenario? I am anything but certain about what the future will bring, but I do have one hope: that the core mission of journalism survives. My dream is that some, if not all, news organizations emerge stronger from this transformation, that we manage to preserve our role as providers of verified, fact-checked information, and that audiences continue to see the value in someone travelling out into the world, observing it and documenting it on their behalf.

The interfaces will be different. Much of what we spend time doing today will be automated. And I am fairly certain that agents will visit us at least as often as humans do.

If we look beyond the immediate challenges,, there is reason for hope. The next technological shift could leave people better informed than they are today. They could gain access to the information they need, in the form they are best able to consume it. There is something profoundly democratising about AIe — if we use it properly.

We in the media cannot complain that technological change makes it harder to maintain the business models that once funded traditional journalism. Our job is to provide a valuable service to people, and help the craft we value so highly remains relevant in the future.

I do not know exactly how we will do that, but I do know one thing: when the hull is leaking, the stupidest thing you can do is sit still. That is how you go down with the ship. We have to do something, and right now, almost anything is better than doing nothing.

As Wallace Hartley might have put it: Keep on playing. But find a new tune.



08

Look beyond the AI tools

Martin Schori
Co-founder, Hint Media

What if the biggest obstacle to journalism's future is not technology, but the way we organise newsrooms today? In this essay, Martin Schori (Co-founder, Hint Media) imagines a newsroom in 2032 where AI handles routine information production while journalists focus on judgement, distinctive voices and relationships with their communities. He argues that real transformation requires rethinking editorial roles, recruitment and incentives, rather than simply introducing new tools into old workflows.

I am sitting in a café with a young reporter who is frustrated.

Normally, I'd say she is a future star: an original thinker, strong voice, a clear sense of what journalism should do, and plenty of management ambition. She is, in other words, precisely the kind of person media organisations always claim they need. Now, but also, increasingly, in the future.

But right now, she isn't the star. On the contrary, she can't even get a job.

Why? Because she does not fit the template. Every time she meets with a hiring editor, she pitches ideas about community relationships, innovative formats, and audience value. And every time, the editor is impressed, knows that he or she should hire her, but has no place for her in their organization.

And also because the editor across the table—whom has many times been me—is simply trying to solve an immediate problem. The editor doesn't want to discuss long-term visions; he or she needs someone to cover the 6:00 AM breaking shift. Preferably yesterday.

The schedule had holes, and the holes always won.

Meanwhile, industry conferences and editorial retreats endlessly echo the same gospel: in the age of artificial intelligence, newsrooms must double down on what makes us uniquely human. We must create deep relationships, build communities, and attract the people who have a distinct voice. Like the reporter at the café. Still, we don't.

This paradox has now reached a tipping point. The technology question in journalism is in many parts of the world more or less settled. Within a year or two, virtually every newsroom will have access to the same baseline AI capabilities. The gap between fast adopters and slow adopters in mature markets is rapidly closing.

When underlying technical capabilities become a commodity, "AI-driven product innovation" ceases to be a competitive moat. We are not there yet, and many news organizations – and markets – have a long way to go before they have adopted AI tools and, even further away, organized themselves around an AI mediated world. But that time will come. And as AI absorbs the mechanics of standard production, speed, and aggregation, the so-called "soft" qualities of journalism—editorial judgement, personal voice, and relational trust—stop being nice to haves. They become the core craft itself.

The next generation of news products will come from newsrooms that have radically reconstructed their organization, culture, and incentives around the things AI cannot replicate and which actually adds value in a completely new world and logic.

To understand what this looks like in practice, we have to look ahead. Let me take you to the newsroom I think we should be building, around six years from now!

The radically reconstructed newsroom

By 2032, the media landscape has drawn a sharp line in the sand, and most newsrooms have finally abandoned the dangerous middle ground.

On one side lies the automated information layer: fast, hyper-personalized, algorithmically synthesized "ambient news" that delivers commodity facts cleanly and instantly to users' devices (whatever they are in 2032). Publishers that manually tried to compete on speed and volume within this layer were, unsurprisingly, not successful.

On the other side are the human-centered editorial media.

Step inside one of these newsrooms on a Tuesday morning in 2032, and you will not find a traditional morning meeting organized around filling different sections on a homepage. You will probably not find separate, isolated "desks" for politics, local news, or crime either.

The newsroom could be organized entirely around audience need clusters: Help Me Understand, Give Me Perspective, and Connect Me to My Community.

Automated background systems monitor public records, transcribe council meetings, flag data anomalies, and maintain routine tracking feeds. The journalists do not touch this layer unless human verification is needed, or ethical questions are to be answered.

A new type of editor

Meet Maya, a Lead Curator & Perspective Editor (the same reporter from the café six years earlier, now leading a team).It is a silly title, but so was the behaviour in the newsrooms of 2026, where we asked ourselves “what should we publish today?” or “what’s the next article from the foreign news department?”.

2032 Maya does none of that. She reviews overnight signals synthesized by background intelligence tools and asks something like:

"What’s noise that we can ignore so our audience doesn't have to be bothered by it?"

"Where is human presence, moral clarity, or live investigation irreplaceable today?"

Her morning is not spent writing three rapid updates to keep an impression counter ticking. She spends two hours on the ground conducting an in-person investigation, followed by co-hosting a closed-loop audio forum with community members affected by a local policy shift. Later, she publishes a deeply voice-driven piece of analysis, evaluating why it matters and who should be held accountable.

The closed forums, the meetups, the Q&As where readers can talk back and journalists actually answer, that is the product, not just an engagement tactic bolted onto the bottom of an article.

Enabling new user behaviours

Take Leila, 54, a preschool teacher in a Stockholm suburb where the local school is about to be closed. In the ambient layer, she can get the decision, the budget figures and a neat summary in a matter of seconds. What she can't get there is someone who sat in the room, asked the politician the follow-up question she would have asked, and then, two weeks later, comes back to the forum and says: here is what has happened since we last talked. That is what she happily pays for in an editorial media product: the feeling that someone is paying attention on her behalf.

Maya is a well known figure in Leila’s community and people trust her. Her primary metrics are something we could call “Trust Density” and “Relational Dept”. Measures of subscriber retention, active user participation, and perceived value.Her 2032 newsroom operates on a simple premise (that probably is too obvious to be stated): AI handles the information supply; humans hold the relational trust that people pay for.

What it will take to get there

To make the 2032 newsroom possible, we have to stop treating AI as an efficiency tool designed to do old work slightly faster.

The promise of AI was always that it would hand time back to journalists. Yet, in most newsrooms today, those freed-up hours are quietly refilled with more output: more rewritten briefs, more social posts, more noise. We have automated the wrong half of the job.

If we want to build news products that endure, three foundational shifts must happen within editorial culture and skill-sets:

1. Stop measuring old values

If an organization’s incentives, KPIs, and editorial praise still flow toward publishing speed and raw click volume, nothing will change. Journalists will simply use AI to produce the same disposable content at double the velocity. Newsrooms must measure what actually creates long-term value: direct audience contribution, retention-per-interaction, and perceived trust.

2. Shift from Gathering Information to Editorial Triage

In an era of synthetic information overload, the primary value of a journalist is no longer gathering raw data; it is knowing what to ignore. Organisations like Fojo—and the broader journalism training ecosystem—were built on upskilling reporters for shifting realities. Today, that upskilling should mean teaching ruthless editorial triage: the capacity to slice through noise, evaluate algorithmic output, and exercise moral and factual judgment.

3. Cultivate voice and relational competence

Generative text is hyper-polished, perfectly balanced, and utterly bland. Human value lies in distinct perspective, lived experience, and intellectual honesty. Furthermore, journalists need skills in facilitation, active listening, and community leadership—turning passive consumers into active participants. As Swedish researcher Emma Frans noted in a different context: AI won't take your job, but someone more fun to be around might. I believe the same goes for products.

Translated to media: AI will not replace journalism, but newsrooms that build unique, authentic human relationships will easily outlast those that don't.

In pursuit of relationship-building journalists and profiles (“the voices” if you would), media leaders need to be mindful of the mismatch between future talents and legacy roles. If you recruit the most original, voice-driven talent in the country and slot them into a templated legacy role built for production output, you will most likely lose them within a year or two—to Substack, to podcast networks, or to platforms where they can actually build something of their own. Every time we pass on or underutilize an unconventional talent because they don't fit an existing shift, we risk funding our future competition.

So what should we do now?

Set KPI:s for trust and relationships. And make the editor-in-chief personally responsible for protecting it the day the 6:00 AM shift needs covering, because that day will come.

Every hour AI hands back to the newsroom should be accounted for. Decide in advance how much of it goes to relationships, presence and investigation, and write it down. If you don't, it will be eaten by more briefs and more social posts, just like today.

Remove one volume metric from the management dashboard and put a relational one in its place. Just one. Then watch what happens to the conversation in the morning meeting.

None of this requires waiting for 2032. It requires media leaders who, the next time someone like the reporter from the café walks in, stops asking whether she can cover

the morning shift, and asks what she would build if we let her.



09

From low adoption to real enablement – a journey

Olle Zachrison

Head of News AI, BBC News and Co-founder, Nordic AI Journalism

What does it take for AI to serve as meaningful, value-creating tooling in a newsroom? In this essay, Olle Zachrison (Head of News AI, BBC News and Co-founder, Nordic AI Journalism) confesses organisational sins committed in the pursuit of scaling AI in public service media, and shares tangible lessons and advice for moving beyond a focus on shiny new things to promoting wide-scale, organisational enablement.

September 2026. It was a Tuesday, just before lunch, when an email landed in my inbox from a colleague with some thoughts about potentially valuable uses of AI in the team's workflows. They asked for help. I asked one of my team members to follow up with a call, both to understand the proposals better and to ask whether they had tried anything with Copilot. 'Copilot, do we have access to that?' they had asked.

At first, I wasn't sure I had heard correctly. Copilot? This is the AI tool with a capital A that was rolled out to all employees more than two years ago. The enterprise solution everyone is expected to learn through a series of training modules. Our natural starting point when superusers demonstrate best practice. That some journalists complain loudly about and others love.

Of course, this isn't about any one colleague. I'm sure they use AI every day in all sorts of ways, even if they don't happen to use that particular tool. But it is an illustrative starting point for reflecting on the state of AI adoption in media organisations. AI is on everyone's lips. But how, when and where AI should be used at scale in editorial work – that is a nut we haven't cracked as an industry.

Like my co-authors, I am convinced that AI technologies – used creatively and responsibly – can reinvent both the audience's news experience and our internal ways of working. But before imagining possible futures, I want to begin with a self-critical reflection on my six years working with AI in journalism, drawing on leadership roles in editorial AI at Swedish Radio and BBC News. As usual, I will exaggerate slightly to emphasise my points. So please bear with me, and I hope you recognise some of these patterns from your own organisations.

Organisational archetypes and leadership blind spots

In all our feverish excitement about the dizzying opportunities of generative AI, I believe we forgot one thing: bringing the vast majority of our colleagues with us. In the past few years, I have come across four distinct types of colleagues when it comes to AI attitudes. A first group consists of those too tied up in daily production to have time to explore the new. A second are grumpily sceptical, as journalists always are. A third category are fierce opponents, although I believe that group is considerably smaller than some claim.

The fourth group is perhaps the most interesting: the pioneers, or 'trailblazers' as they are called in BBC terminology. Truly mind-boggling things were happening among this group almost from day one after the launch of ChatGPT in November 2022. For many, it was simply the next logical step after years of 'hacking' together their own digital solutions. But the best things cobbled up during 2023–24 were rarely brought into the light, at least not beyond their immediate circles. Sometimes this was due to a fear of getting caught using non-sanctioned tools or being accused of 'cheating'. But fundamentally, there was neither a culture nor dedicated time for

sharing those kinds of experiments.

I feel very much complicit in that. During my first few years in genAI, I focused a lot on theoretical possibilities and risks instead of spotting the concrete gems and giving them a bigger stage. Responsibility for practically exploring AI was left to the individual, whether they were at the cutting edge or largely indifferent.

There were attempts. I remember how Swedish Radio reporter Jesper Ericson showed how he had found an unexpected angle for an investigation into how schools in the Swedish county of Kronoberg compared with the national average. A textbook example of how AI can create editorial value. When he presented it at a management conference, people were fascinated. But back then we didn't find an effective way to scale his methods across the wider reporter community.

Before I continue, I want to pause for a second and make one thing clear: in 2026, real progress is being made in broadening the use and integration of AI. Some media organisations are starting to crack it. Maturity is increasing. More people are getting involved. Real value is being delivered. But before we discuss that, let us dwell a little longer on the – and often my own – sins of the past. We will get to the forward-looking part, I promise.

Confessions of professional misconduct

This may sound shocking, even like gross professional misconduct, but in my early days in AI leadership, I very rarely took the time to talk to people further down the organisation about their needs, fears and dilemmas. Things were made worse by my move from an operational management job into a central strategy role, which meant I no longer had daily contact with frontline journalists.

So I rushed around different leadership forums and international conferences, knocked together slides and talked to like-minded people. I don't think I have ever been as inspired and professionally happy as I was during those years. But did I spend enough time trying to understand my own colleagues? Hardly.

So what did we AI optimists miss during the first two or three years after the breakthrough of generative AI? To start, we enthusiastically jumped on the coolest innovation ideas. The shiny new things. 'Now we're going to launch new summary formats' – which the audience hadn't asked for. 'We'll let people ask our archives any question they like' – but doesn't Google do that better? 'Look how easy it is to create illustrations' – oops, that character seems to have six fingers on one hand.

Enthusiasm, improvisation and pioneering spirit rather than strategic focus. Quite understandable. We were sucked into the global super-narrative that the revolution had arrived and the race was on. Boards and management teams shouted: 'Do something!'

The policy trap

Speaking of strategy – or rather the lack of it – many of us in the first phase of AI focused on guidelines. Media organisations spent much of 2023 and 2024 developing policies, all emphasising the importance of keeping a human editor in the loop. New committees and AI councils were established. Some of that principled editorial policy work was absolutely necessary at the time as a cornerstone of responsible AI work. Not least in public service media, where maintaining the audience's trust is everything.

But it took a great deal of time and effort. The internal guidelines in particular started from what we should not do, rather than what value the technologies could add. Nor did they go hand in hand with a clear direction on what was most important to invest in. No wonder journalists were put off when the basic message was 'Be careful!' rather than 'Have a go!'

Guidelines are not strategy. It took me a while to understand that. Nor did it sink in, as I rushed ahead, that strategy also has to be matched with real resources for implementation.

The far-sighted leaders

Some leaders saw the bigger picture earlier than others. When VG's then-Publisher Gard Steiro, at the Nordic AI in Media Summit 2024, was openly paranoid about the industry's limited appetite for change, many attendees were rightly shaken. After all, he was the one who was simultaneously initiating some of the most radical projects.

Swedish Radio (SR) Director-General Cilla Benkö was also early in forcing discussions about strategy and adoption into the executive meetings. 'How do we involve all our newsrooms across Sweden?' she asked with a piercing look. She and Programme Director Sofia Wadensjö Karén asked me to put together a first, simple slide with a forward-looking message. The recipe for success was its simplicity: just three points:

SR will use AI to…

…improve and develop our journalism

…create better user experiences

…improve internal productivity

… throwing spaghetti on the wall

We had a direction.

That didn't stop me and my team from continuing to rely heavily on gut instinct and focusing on flashy things we could boast about at media conferences and on LinkedIn. 'We've started working with few-shot prompting instead of model training';

‘We’ve built a knowledge graph for our metadata extraction’; ‘We’re going to help train a Swedish language model – even one for Sámi’.

Of course, that’s not the whole story. Projects were in fact launched that could create significant value over time, such as automatic transcription of podcasts and an interactive news search. But somehow, we never quite got the execution right. Instead of rallying behind the strategic direction, we kept tinkering around the edges.

It was easier that way. Involving hundreds of journalists takes a lot of time and effort. You have to move at lightning speed in the AI era, we had been told. Blue-sky thinking was more fun than grey reality. Moreover, our colleagues in Product struggled to squeeze all our ingenious proposals into their packed roadmaps. Once again, we were haunted by the curse of the central innovation team: not that much of the spaghetti actually sticks to the wall.

The news factory kept the lights on

Meanwhile, our fantastic journalists kept toiling away in the news factory. And it was hardly business as usual. A dramatic world of wars and political convulsions created an unprecedented 24/7 news flow. An insatiable hunger for new formats across ever more platforms. Tighter finances meant chairs disappearing from newsrooms. And there was still a long way to go on the winding road from analogue to digital media organisations.

‘What, you haven’t started prompting and personalising your GPTs?’; ‘Surely you have taken part in Copilot Chat Prompt Workshops: Go Deeper?’ No wonder colleagues became even more tired when some AI maniac from headquarters appeared on screen with their insinuating questions.

Then there was the issue of middle managers. To put it bluntly: they were forgotten even more than colleagues on the newsroom floor. The first years’ attempts to increase AI engagement effectively bypassed the middle ranks, going straight from the executive team to individual employees

Learning to master new tools, navigating responsible AI principles, keeping up with external trends – all of that was largely left to each individual. At most, the middle manager might be expected to nag people into attending a few mandatory courses. But there was never any expectation for them to develop a plan for integrating AI into their team’s core processes.

Signs of bottom-up change

It is difficult to pinpoint when greater self-awareness began to set in. It was probably a gradual process, like a stalactite forming drop by drop. I had an early moment of illumination in spring 2024, when we asked Swedish Radio employees what they would use a perfect AI assistant for in their everyday work. The question was tucked away in a free-text field at the end of our introductory course on editorial AI.

We had hoped for a few suggestions. Instead, 1,224 (!) ideas came in. Many were considerably smarter than anything we could ever have come up with in our central AI team. The suggestions were clustered using AI. There were loads of practical ideas for journalistic work – research support, editorial feedback and audio processing.

But one other area was by far the biggest: administration. Colleagues at every level were stuck in the routine handling of information: reporting, email management and scheduling. It reinforced the impression that all employees desperately wanted to free up time. Had we focused too much on the editorial side?

The shared ideation process led to several tangible projects, such as a bot for HR questions. But we were still stuck in a mindset of running centralised projects. I had yet to realise that real empowerment might lie in radically decentralising the development of AI solutions.

Adoption challenges on a global scale

Perhaps it wasn't until I joined a truly global media organisation in spring 2025 that I began to see the formidable challenges of scaling AI adoption. If it had been difficult to reach 26 local channels at Swedish Radio with information and activities, I had seen nothing yet. BBC News, with 4,200 journalists around the world, operates across so many time zones, formats, rhythms and production tools that engaging everyone is a Herculean undertaking.

At the same time, I was deeply impressed by the ambition of the BBC projects, often carried out with our R&D department. An in-house model for reversioning local democracy articles. A daily podcast covering 20 football clubs, where AI generates a script based on articles and a synthetic voice presents the content in the right local accent. A bespoke translation tool aimed at 43 language services around the world.

But switching on a new feature for thousands of busy journalists does not mean everyone starts using it. Launch and actual use are two completely different things. Even if we organise mandatory training and lunch-and-learn sessions, it doesn't mean everyone takes part – at least not actively. To make a crude but affectionate generalisation: journalists are news avoiders when it comes to internal information.

The BBC had also, on several occasions, published critical research into how the tech giants' AI assistants distort our journalism. It was a very valuable contribution to the debate about the lack of accuracy and attribution in the AI-mediated ecosystem. But if AI models from California are seen one day as striking at the heart of our mission, it is quite natural that reporters won't rush to embrace Copilot, ChatGPT or Gemini the next.

Investing in newsroom change

My bosses Nathalie Malinarich and Anjali Kapoor began to see this: AI use was not increasing at the pace we wanted. At least not according to the way we were measuring it. We needed a proper strategy, linked to an action plan and greater capacity to implement it. Our growing project portfolio already spanned everything from video reversioning, style guides and newsletter support to AI voices, but now a new strand was added to the News AI team's remit: adoption and scale.

In a short space of time, my hard-working colleagues set up an ambitious superuser programme. When I heard the participants' comments in a workshop in London more pennies started dropping, one by one: 'We don't know what we're allowed to do'; 'AI is almost never discussed in the morning meeting'; 'We want more practical examples of how AI is being used by different teams across the BBC'; 'Copilot is getting better, but it's hard to use for journalism.'

At the same time, I was overwhelmed by the engagement and precise ideas of the 40 participants. Once the magnifying glass was pointed their way, good examples started springing up like mushrooms. A colleague who had prompted their way to a subtitle-reformatting tool with an intuitive interface. A journalist who had scraped and transcribed hundreds of hours of YouTube videos to uncover stories and find people directly affected. An early version of an agentic chain that laid the foundations for an investigation with major impact.

Scale requires sustainable infrastructure

But, symptomatically, many of the most ambitious colleagues had built their tools using home-made solutions. Creative and impressive, but makeshift and not sustainable in the long term. The combined burden of barriers to accessing technology, legal issues, information security and governance made the process cumbersome, particularly because the innovators often had to try to navigate the internal bureaucracy on their own.

In a newsroom, having clear and easily navigable editorial systems is crucial. When I started publishing online in the early 2000s, you almost had to be an engineer to get an article onto the website. Things slowly improved as UX and data flows became more refined. Attempts were made to merge publishing systems as we moved towards a digital-first mindset.

But in the first genAI phase, we found ourselves back in a world of proliferating tools, with expensive licences for new third-party products. Without the appropriate experience, these immensely powerful general-purpose assistants are difficult to use for specific and repetitive tasks in a newsroom. Here I have been impressed by organisations such as Austria's ORF, Denmark's Ekstra Bladet and Finland's Yle, which understood early on that their own AI toolkits were central to broad adoption. A bespoke interface with the company's own branding gave journalists the seal of

approval they needed to feel safe.

Perhaps this is my most important lesson for the next phase: we need to bring AI to journalists, rather than the other way around. Much like news audiences, I believe journalists are generally more interested in the user experience than in the sophistication of the underlying AI models. The technology needs to adapt to their needs and workflows, not force them to adapt to the technology.

The hard work starts now

So how do we equip our organisations for the wave of disruption that has only just begun? First, I hope my retrospective has provided some clues to the challenges we need to address as we move up the AI maturity ladder. Unfortunately, there is no magic wand. Swedish Radio's Lead for Editorial AI, Danina Mahmutovic, has spoken about three elements that must go hand in hand: easy-to-use interfaces, hands-on coaching and continuous skills development.

But the day-to-day slog has to be paired with a visionary discussion about what journalism could become. That open conversation needs to take place – as it does in this anthology – across industry and disciplinary boundaries, and also within our own organisations.

For my own part, I believe that AI's greatest opportunity in journalism is hyper-personalisation, both for news consumers and journalists. In terms of audience value, we have a real chance to offer content and services that are genuinely relevant to individual users. We can finally move away from mass media, which publishes the same content for everyone, towards a highly tailored offering based on interests, intentions, behavioural patterns and demographics. Is this the moment when public service media's eternal dream of true universality can come true?

If we succeed, we can put real human needs at the centre in an entirely new way, with content coming to people rather than consumers having to buy physical products, turn knobs on appliances and navigate complicated apps. The media experiences of the future will evolve seamlessly with people throughout their lives. Access to vital, trustworthy information could increase dramatically.

It may sound utopian, but something similar could also emerge in editorial workflows. AI support could adapt organically to the individual journalist – based on their role, task, abilities and needs. It could become an ever-present personal coach, rather than yet another inflexible standardised system. AI could become a shortcut to the smooth digital workflows we have dreamed about for years but failed to achieve with traditional production tools. A bit like how some developing countries leapfrogged straight to mobile phones without taking the detour via expensive fixed-line telephony.

Glimpse of our possible future

Earlier this year, I caught a glimpse of the hyper-specialised editorial workflow when my colleague Annika Grosser showed me a prototype of a multi-agent system she had built that scraped, summarised and translated information from social media in a dictatorship. This provided material for a potential investigation. The workflow was carried out in a well-designed, bespoke tool – in practice, a vibe-coded CMS for that specific story.

I was dizzy with the possibilities. Imagine every journalist being able to build and use their own tools at every step of the journalistic process: from idea and research to production, publication and traffic optimisation. Instead of all digital development having to pass through procurement processes and battles over priorities, the person who actually has the idea and the right tech can start creating the solution themselves. The code monopoly would finally be broken.

But the flip side of the coin is darker. An AI leader at one of the world's largest news organisations recently told me: 'We're now rolling out Claude to everyone and it's incredibly exciting. But how we're going to make sure everyone follows our journalistic guidelines when they build their personal agents, we haven't got a clue.'

And it's not all sunshine when it comes to the audience's news consumption either. We compete, coexist and collaborate with the most powerful technology companies the world has ever seen. Is there any chance at all that people will be reached by the quality journalism we seek to provide? Will anyone recognise that a particular piece of content comes from us? Will hyper-personalisation on steroids create even more social fragmentation?

It is important to ask ourselves what it is we are trying to protect. No, it is not particular article, video or audio formats that are at the heart of journalism, as my fellow contributors Nikita Roy and Shuwei Fang often eloquently point out. Nor is it the media organisations themselves; more are likely to disappear in the coming years.

But good journalism is characterised by certain values: careful, methodical work with facts, a critical mindset, engaging storytelling, openness to different perspectives, an ambition to include underserved groups, and accountability when we get things wrong.

Journalism cannot be done entirely from behind a computer screen. Out there in the physical world are the talkative sources, eyewitnesses to dramatic events, and those affected – all absolutely essential to making stories come alive and hit you in the gut. If all we do is report ephemeral information, we're toast. As warnings grow about deskilling – the gradual loss of professional skills – I'm less worried about us forgetting how to write than losing our ability to interact with people out in the world.

What needs to change to safeguard the core values

So what do we need to change organisationally to step things up? First, AI must become an integral part of every transformation process, rather than being treated as something separate and alien. That need becomes even clearer as we now introduce concepts such as ‘agentic AI’ without being able to explain what it means for the organisation, or even exactly what the term itself means.

It might be a good idea to start fewer office conversations with the diffuse and loaded term ‘artificial intelligence’ and instead…

…talk about enablement – giving people the chance to engage with new technologies – rather than about adoption, which casts people in a more passive role.

…focus more on local usefulness than on centrally developed, ready-made solutions.

…let go of our fixation on formats and artefacts and talk more about values and mission.

…make the strategic direction clearer rather than placing responsibility on the individual. Here, centralisation is actually needed.

…share more inspiring examples from peers who have created clever solutions.

…connect journalists with specialist expertise elsewhere in the organisation.

It’s the people, stupid!

Speaking of people. For 25 years, I have worked in different media organisations – print, radio, television and online media. At every point, some kind of crisis has loomed over us. First it was falling print revenues and listening numbers, then the threat from social media platforms, and now the rise of AI. Every organisation I have worked in has felt small and weak compared with the giants ‘over there’. All have gone through painful reorganisations.

But amid the turmoil are the people, our colleagues. If you are an inspired and hard-working bunch who want to do something together, you can move mountains. And if you don’t quite make it all the way, at least you’ll have tried, and hopefully had fun along the way.

So don’t wait. Connect with each other and get cracking. And please, feel free to learn from my professional sins whilst I get busy making new mistakes we can talk about next time.



10

Futures friction: a new approach to how media organisations remember, learn and reimagine

Sofie Hvitved

Director of Media Futures & Innovation, Copenhagen Institute for Futures Studies

How might media organisations better build on collective learnings in pursuit of innovation? In this essay, Sofie Hvitved (Director of Media Futures & Innovation, Copenhagen Institute for Futures Studies) explores how second brains can make organisational knowledge accessible across teams, and how futures friction can help teams question and develop their ideas. She explores how AI-supported organisational memory and deliberate challenges to established assumptions can help newsrooms learn from past experiments and imagine futures beyond today's strategic constraints.

Stop for a moment. Step away from your quarterly goals. And ask yourself: Will the formats, products or services that I am investing my time and resources in still be relevant in a few years?

The assumptions behind our choices often seem so obvious that we rarely question whether they still hold. If media organisations only make room for ideas that fit today's strategy, they risk rejecting the very ideas that could help them remain relevant as the surrounding system changes. How can we make room for that kind of reimagining if the way we develop, assess and share ideas remains shaped by the very assumptions we need to question?

In my work, bringing futures-driven innovation closer to media organisations comes down to these two connected questions: how do we make better use of our collective organisational memory, and how do we keep questioning what we know as the logics around us change?

Working with media organisations, it's becoming clear how difficult it is to connect innovation across teams and build on what others have learned. Ideas struggle to survive daily operations or find support beyond the team that developed them. Those that challenge the current business are often measured by metrics that overlook their potential future value. So, ideas get stuck or are squeezed into a different logic than originally intended, while others may never learn what the team tried to achieve or which assumptions shaped its decisions. An unsuccessful experiment can still produce valuable learning if we preserve the context needed to understand what happened. A good idea might not fit the existing project but be gold in another context.

Could connected second brains make learning more accessible across teams, and might futures friction help question our assumptions in light of possible futures?

The weight of history in news organisations

News organisations carry a particular legacy that needs to be considered in how they preserve and share organisational knowledge. Protecting sources, keeping a scoop from competitors and guarding rights are important parts of the work and culture. Those instincts also shape how knowledge moves within the organisation, including outside the newsroom. Competition between desks can make it harder to share what they have learned, and the reasoning behind a project may stay with the people who made it.

At the same time, there is a kind of exhaustion from the ongoing waves of change in media, from print to web, mobile and platforms. Any practice that asks media organisations to share and question what they know must be built with this history in mind: within the boundaries that source protection and rights require, and with as

little extra effort as possible.

Where do you start when people are exhausted by change and used to protecting knowledge? By making it easier to learn from what others have tried and why. Technology could bring that learning into existing workflows, making the organisation's collective knowledge a starting point for futures-driven innovation.

Second brains: shared organisational knowledge layers

One concept I'm interested in exploring in this relation is the idea of a second brain: a way of organising knowledge outside our heads so we can return to it, connect it and use it in new situations. The term was popularised by Tiago Forte as an approach to personal knowledge management, but the ambition to extend our memory through technology has an even longer history. In his 1945 essay 'As We May Think', Vannevar Bush imagined the memex, a desk-like device for storing books, records and communications. People would connect these materials through trails of association that they could revisit and share with others. AI could make it easier to build, maintain and draw on such connections, opening new possibilities for how organisations use their accumulated knowledge.

AI could make this knowledge easier to retrieve and connect across an organisation, helping people examine the evidence and reasoning behind earlier decisions.

It all starts with collecting contextual data: the purpose of an experiment, the audience insights it drew on, what was tried, why a decision was made and where people disagreed. A report saying "user demand was low" tells us little about why but knowing the context can give another team something to work with when that context is accessible and reusable.

Second brains operate at several connected levels: individual second brains that hold employees' own knowledge, team second brains that capture shared experiments and learning, and an organisational second brain that preserves strategic reasoning across projects and helps connect innovation to the organisation's purpose and values.

AI researcher Andrej Karpathy describes a related approach in his LLM Wiki: a collection of linked Markdown files that an AI model writes, organises and updates from source material. People select the sources and guide the work, while the original material remains separate from the AI-generated knowledge.

Making context machine-readable allows people and AI agents to work with it across projects. Sources, dates and the conditions behind a conclusion need to travel with it, as does the distinction between evidence and interpretation. Otherwise, a confident interpretation could become accepted as evidence simply through repeated retrieval. Formats and protocols are emerging to support different parts of this.

This way, context becomes infrastructure for how we remember and learn in organisations.

Futures friction: using futures to inform the present

Futures friction brings futures thinking and foresight into the innovation process. It introduces deliberate friction to test ideas against possible futures, considering changes that might happen in both the short and long term.

A convincing AI-generated answer can make an idea feel settled before we have explored other possibilities. This is the risk of premature convergence. Adding intentional futures friction helps us challenge the assumptions and biases behind that sense of certainty, making room to stop, think and iterate our ideas.

Futures friction builds on a futurist's way of thinking. Thinking in futures means recognising that the conditions we work within could develop in several plausible directions and exploring what those differences would mean for the decisions we make today. It requires us to look beyond current short-term strategies and pay attention to signals that could challenge what we take for granted. Some changes may be incremental; others could fundamentally alter how people access information or whom they trust, which are existential questions for media to pay attention to. At the same time, many needs and behaviours may endure.

The aim is to become better prepared for these different possibilities. What would need to remain true for an idea to make sense? Under what conditions might an overlooked idea become relevant?

This also means staying willing to change our minds as we learn. Short cycles of testing and reflection help us explore these questions in ongoing innovation work. Each experiment should add to the context others can draw on.

Shared context and futures friction need to work together. A second brain helps turn shared context into infrastructure for remembering and learning. Futures friction brings questioning into that infrastructure, helping organisations explore possible futures and recognise what they may need to reimagine.

A reimagined media organisation in 2030

What might this look like in practice? Imagine a media organisation in 2030, where teams draw on connected second brains and use futures friction to question their ideas.

A journalist returns from a walk with a colleague. They have been discussing an idea for a new local information service that could help residents understand how political decisions connect over time and what they mean for their lives. It would bring together the organisation's reporting, public records and knowledge of the community. Developing it would require editorial, product and technology teams to

work together.

The colleagues have chosen to record their conversation and process it on a local device, where it stays private. Afterwards, an AI agent suggests a short outline for the journalist's individual second brain, capturing the idea, their reasoning and the questions they have yet to resolve.

She reviews the outline and quickly adjusts which parts to share with her colleagues in the team. Her private notes and the original recording stay in her individual second brain. The approved outline becomes part of a team second brain, where colleagues can develop the idea together. She controls what moves out of her personal space; contributing an idea does not give the organisation access to everything behind it.

The team likes the idea and draws on the organisational second brain. Within minutes, an agent finds an experiment from four years earlier. Its summary says demand was low, but the original evaluation reveals something more specific: people found the proposed service useful but disliked having to install yet another proprietary app.

This is the value of shared context. The team can see what was tried, why it was judged unsuccessful and whether that conclusion is relevant to the new idea. The organisational second brain also preserves the strategic reasoning behind the earlier project, including its purpose of helping residents understand local decisions.

Now futures friction enters the process. Would the idea make sense if residents increasingly expected information to reach them through their own AI assistants? An agent helps the team explore that possibility alongside an alternative: what if people resisted AI mediation and wanted direct contact with local journalists? Observations that residents have agreed to share give the team reasons to take both possibilities seriously.

The team decides to stress-test several ways of accessing the service. Its purpose remains the same, but the assumptions about how to fulfil it remain open to reimagining. The decision, its reasoning and the unanswered questions enter the team second brain. Learning that could help other projects is also shared through the organisational second brain.

Residents can follow how the service is being developed through the reasoning and evidence the team has chosen to share. They can question its choices and contribute their own experience. Their feedback adds to the shared context behind the project.

A few weeks later, a competitor launches a similar service. The team uses futures friction to revisit its assumptions: if this kind of information becomes widely available, what would make its own service useful? It adjusts the experiment and records what changed and why in the team second brain. That learning also reaches the organisational second brain, where other teams can draw on it.

In this scenario, AI helps collect, connect and retrieve knowledge, leaving more room for thinking, meeting colleagues and spending time with residents. Leadership has protected that time and recognised contributions to collective learning. The same efficiency could have gone to cuts; here, the organisation has chosen to invest it in relationships, on-the-ground journalism and futures-driven innovation.

Connected second brains help the organisation remember and learn. Futures friction keeps that learning open to challenge as conditions change. Together, they give teams a stronger basis for deciding which ideas to fund, test and reimagine.

Many conditions would need to hold for this future to take shape, and it is unlikely to unfold exactly as imagined here. The value of scenarios is to help us explore what could change and what we might need to prepare for, including the choices about ownership, time and resources that make this way of working possible.

Where to start?

Back to 2026. I met a software developer who described working with his second brain as adding ingredients to a huge pot, stirring them around and then delivering. His accumulated knowledge provides the context for each new task, changing how he works even as he is still trying to understand the implications. Media organisations could learn from that approach by working more systematically with context.

Start with one current innovation project. Bring together relevant context: earlier experiments, audience insights, signals of change and the reasoning behind the proposal. Make it machine-readable and preserve its sources. Agree who controls the knowledge, what can be shared and reused, which AI systems can access it and how you retain control if you change providers. Source protection, privacy, copyright and cybersecurity need to shape these choices from the start. What stays local, and what is shared through cloud services? What knowledge can be exported to other systems to avoid knowledge lock-in? Who owns what knowledge? Begin in a controlled environment with limited access and carefully selected material, then assess what can be connected more widely.

Bring futures friction into the work from the beginning. What must remain true for the idea you’re working on to make sense? Test it against contrasting plausible futures, including one where the underlying media logics change. Choose a small experiment and agree what you need to learn, so today’s performance metrics do not dismiss an idea before its future value can be explored. Record what changes in your thinking and give someone responsibility for revisiting the assumptions.

This takes practice, from organising shared context to questioning what AI retrieves and thinking in futures. Leadership needs to protect time for it, recognise contributions across teams and allow learning to change priorities. A second brain will be of little use if people are rewarded for holding knowledge back or the organisation refuses to reconsider its direction.

This also connects to a call for existential innovation: rethinking how media organisations sustain themselves and what they exist to do. For me, shared context and futures friction could help create the conditions for that work, making it easier to learn across teams, question inherited assumptions and give unfamiliar ideas room to develop in a time where existential questions need to be asked.

So, before returning to your quarterly goals, consider which ideas never make it onto the agenda because they challenge the current logic. Give one of them time, resources and the benefit of what others have already learned. Then use futures friction to explore where it could become relevant.

A futures-driven media industry begins with decisions like these, including the decision to take an idea seriously before it fits today's strategic environment.

ABOUT THIS ESSAY

This article was developed with support from my own second brain, using it to question assumptions and explore the reasoning. The intention, arguments and editorial responsibility remain mine.

We are developing a free, downloadable Futures Friction Assistant as part of AI MediaLab Denmark to support futures-driven innovation in media, a collaboration between the H3 AI Lab at the Copenhagen Institute for Futures Studies and Center for AI DMJX, supported by the Google News Initiative.

The Futures Friction Assistant will become freely available worldwide in December 2026 for anyone to use with their own second brain.

Read more about the project here: https://www.cifs.dk/cases-stories/ai-medialab-denmark.



11

Rewiring organisations starts in our heads

Sara Inkeri Vardar

Innovation Manager, JP/Politikens Media Group and Co-lead, Nordic AI Journalism

What if reinventing journalism starts not with new technologies or organisational structures, but with changing how we think? In this essay, Sara Inkeri Vardar (Innovation Manager, JP/Politikens Media Group and Co-lead, Nordic AI Journalism) reflects on the psychological and cultural barriers that prevent media organisations from challenging their own assumptions. Drawing on personal experiences and research into innovation, she argues that curiosity, psychological safety and the courage to ask seemingly stupid questions may be among the most important capabilities for the newsrooms of tomorrow.

Imagine a newsroom in 2036. What comes to mind?

Wait, no, don't imagine it. Because what you imagine will probably look a lot like the newsroom today, just with some future sprinkled on top.

We are conditioned to think in certain ways, so when we are asked to reimagine something, the result may look new and shiny on the outside while the core remains unchanged. When I started writing this essay, I had a nagging feeling that my vision of the future was not particularly innovative, either. Probably because my idea of a media company was shaped by the ones I already knew. In my writing process, I ended up questioning pretty much everything. What is journalism, anyway? How do we make money ? Who am I ? You see, it got quite deep and existential.

At some point, I had to stop and actually get writing. One draft turned into another. Suddenly, I had five, and I had still not figured it out. With all the scenarios and structures I tried to come up with, I was left feeling... meh.

Then one night, it hit me. The whole problem was that I kept trying to come up with A Solution. But maybe what journalism needs the most right now isn't another brilliant idea for the future, but the curiosity to question the assumptions that shape the ideas we come up with in the first place.

Are we solving the right problems?

In a 2025 essay for Prospect, Rasmus Kleis Nielsen made an observation that has stayed with me: The crisis facing journalism is not necessarily experienced as a crisis by the public. While news organisations worry about declining revenues and changing consumption habits, people are getting on with their lives, finding information and answers elsewhere. They do not necessarily feel that anything is missing.

When young people do not read our articles, we start discussing whether we should make them shorter, turn them into videos or build a new feed. We are already looking for new ways to deliver our traditional journalism without necessarily understanding what people are looking for in the first place. What if the problem isn't the format? What if there are information needs that we are not meeting, barriers we haven't noticed or questions we haven't thought to ask?

We might be so caught up in finding solutions to our problems that we have forgotten to investigate theirs.

Of course, journalism has a democratic responsibility that goes beyond giving people what they ask for. We should investigate what matters, even when the public doesn't know that something needs investigating. But that responsibility should not make us any less curious about the people we are trying to serve. I wonder how many ideas for

new journalistic products begin with a genuine understanding of peoples' needs, and how many begin with an assumption about what will make them consume more of what we already produce.

The value of a stupid question

Let me give you an example from my own organisation. A while back, I suggested using AI to let readers adjust the complexity level of a news article. Journalism can be full of complicated language, institutional references and terminology that assumes a certain amount of prior knowledge. If you do not have that knowledge, reading an article can feel like joining a conversation where everyone else already knows what happened before you arrived. So why not let readers of our journalism adjust the language to something they actually understand?

I raised the idea with journalists and immediately faced objections: What happens to nuance? What if AI changes the meaning? Who takes editorial responsibility for the result? All perfectly legitimate questions. But I will admit that I left those conversations feeling a little stupid, almost as if I had suggested something that went against the very principles of journalism. Which was rather frustrating, when I was only trying to make it more accessible.

Accuracy, nuance and editorial responsibility are not details to ignore in the name of innovation. But making journalism more accessible to people who currently struggle to understand it seems like an important journalistic ambition, too. So why were those ambitions treated as if they were in opposition?

Journalists can become so concerned with protecting editorial integrity that they dismiss technological possibilities too quickly. Developers can become so focused on what technology makes possible that they underestimate the journalistic consequences. Both perspectives are important, but neither should be allowed to end the conversation. The risk is that somewhere in the middle of all these perfectly reasonable professional concerns, and in our hesitation to be thought of as asking stupid questions, we forget the user.

The burden of innovation

Thinking differently is exhausting. When I asked a journalist, who also happens to be a friend, to reflect on these questions with me, he replied, "No, please, Sara. Let's not do that. It gives me a headache."

It is so much easier to do what we have always done. We know how it works, what is expected of us and, more or less, what the outcome will be. Questioning the assumptions behind our work means giving up that certainty. Suddenly, even the things we thought we understood become open questions. And who has the energy - or time - ¨ for that when there are deadlines to meet, stories to publish and a hundred other things demanding our attention?

This is what I think we underestimate when we ask people to innovate. We are asking them to do something that is not only professionally challenging, but also mentally draining. To question what they are good at, to risk getting things wrong and to explore problems without knowing whether they will ever find an answer. And to add to the burden, we often expect them to do all of that on top of their everyday work.

In my master thesis research on innovation in established Nordic media organisations, I found that the problem often wasn't a lack of ideas or ambition. It was the difficulty of creating the time, resources and authority to act on them. When new initiatives compete with everyday operations, everyday operations tend to win. And, why wouldn't they? Existing products have customers, deadlines and measurable results. New ideas come with assumptions, not KPIs. We are asking people to create something unpredictable while judging them by how predictably they deliver.

If we are serious about innovation, we have to make room for the effort it takes. We talk about reimagining journalism at conferences and workshops, but then return to our daily deadlines and the products we already know. If we want to do more than improve what already exists, we need to protect time and resources to question the assumptions we build our work around, and experiment with ideas whose outcomes we cannot yet predict. Even if they never lead to a new product, they might change how we think about journalism itself.

Soft skills, really?

I'm contributing to this anthology alongside people whose work I have followed and admired for years. When I was invited, I was excited. But as I started writing, I also became increasingly nervous. Am I innovative enough? Is any of this interesting? And am I really going to write about all the soft stuff while everyone else is imagining extraordinary technological futures?

At one point, I became so uncertain about my ideas that I didn't even want to show anyone what I had written. Which is slightly ironic, considering that I am writing an essay about the importance of asking questions, challenging assumptions, and being willing to get things wrong. Apparently, it is easier to advocate for those things than to practise them yourself.

As you can see, I did end up sharing my thoughts. And I think a big part of the reason was that I felt safe enough to do so, even when I wasn't entirely confident in what I had to say. But it made me think about all the ideas that never get shared. The questions people do not ask because they are afraid of sounding stupid, or because everyone else in the room seems to know better.

My mentor, Lena K. Samuelsson, reminded me that creating that sense of safety is one of the most important responsibilities of a leader. Because if people do not feel comfortable asking questions, challenging assumptions or admitting that they do not

know something, how are we supposed to innovate?

And yet, we often dismiss these things as soft skills, as if they were somehow less important than technological expertise or strategic decisions. But the ability to ask a question without knowing the answer, disagree with someone more experienced or share an idea before it is fully formed can determine which possibilities an organisation ever gets to explore.

Perhaps that is what makes it so difficult. It requires employees to risk being wrong in front of people whose opinions matter to them, and leaders to accept that they do not always know best. But ultimately, it is not just about being willing to listen. It is also about being willing to give up power. Who gets to decide what is worth pursuing, who controls the resources, and who is willing to give up some of that control?

There is a paradox here. The more pressure our industry faces, the more we need to innovate. Yet that same pressure can make us even more protective of our resources, our roles and the influence we already have.

Maybe the things we tend to call soft are actually some of the hardest, and the most important, things to get right as we move towards an AI-mediated future.

The newsroom of 2036 (or at least my attempt)

So, what would actually look like in practice?

I imagine a media organisation where teams of journalists, developers, designers and people with backgrounds we do not traditionally associate with journalism have the time, resources and authority to investigate problems together. Where they are not handed a product brief with a predefined solution to execute, but are encouraged to explore questions that the organisation has not yet figured out how to answer. And where the people we are trying to reach are involved in that exploration, rather than invited in at the end to test whatever we've already decided to build.

My biggest hope is that we will become much better at opening our doors to people who do not look, think or work like the people already inside our newsrooms. Not simply by hiring more diverse talent into the same old structures, but by creating different pathways into journalism and allowing people with different skills and experiences to contribute to its development. If we want people to question the status quo, it seems rather counterproductive to keep recruiting people who have been trained to see the world in the same way.

That also means creating organisations where challenging established ways of working is welcomed rather than treated as an inconvenience. Where people do not have to be brave to ask questions or rely on a mentor to vouch for them when they challenge the status quo. And where leaders actively create the psychological safety that makes those conversations possible, while also giving people the authority to act on their ideas. Because inviting people to think differently means little if we are not

willing to share the power to change things.

Most importantly, I would like to see organisations where we are honest about what makes all of this difficult. About our fear of being wrong, our tendency to protect what we have, and the power we might have to give up.

Rewiring ourselves

After five drafts, several existential crises and a fair amount of self-doubt, I am still not able to tell you exactly what the newsroom of 2036 should look like. But what I do know, is this:

In the midst of extraordinary AI agents running future newsrooms, radically reinvented organisational charts or spectacular new business models, here will still be us, our beautiful yet tricky brains, and the phycological work we need to do to allow our businesses to change.

Read, reflect, act.

Play an active role in shaping what comes next.

FOJ.NORDICAIJOURNALISM.COM

Nordic AI Journalism

Fojo. Medieinstitutet