



November 2025

Agentic AI in Journalism: The New News Agents

Case Studies, Challenges, and Applications

WHITE PAPER



Mo Shehu, PhD

 columncontent.com

Executive summary

Journalism is being shaped by three converging forces: the collapse of traditional business models, the fragmentation of audience behaviour, and the rise of artificial intelligence (AI) systems that can now read, analyse, and act across entire workflows. Over the past decade, newsrooms have faced several crises including shrinking revenue, bankruptcies, layoffs, reduced staff capacity, lowered public trust, and platforms that increasingly sit between publishers and audiences. At the same time, AI moved from simple automation to fully agentic systems capable of gathering information, monitoring signals, drafting structured outputs, and supporting tasks that once consumed hours of human time.

This paper argues that agentic AI represents one of the most significant shifts in newsroom infrastructure since the arrival of digital publishing. It can help restore reporting depth, accelerate research, scale investigative work, expand multilingual coverage, support new business, and make small teams more competitive. But it also introduces new governance challenges—accuracy, bias, privacy, explainability, reputational risk, and trust—that demand deliberate integration rather than passive adoption.

Agentic AI will reshape how journalism is produced, distributed, and consumed. The question for news executives, media strategists, and policymakers is not whether the industry should use these systems, but how to deploy them in ways that strengthen public-interest reporting and protect editorial values and revenues. This report outlines the opportunities, risks, and practical steps required to build a human-AI partnership that can sustain journalism through the next decade.




1. The state of journalism

If you've worked in or around news over the past few years, you already know the ground has been shifting under the industry. It's not one thing, either, but everything at once: money drying up, audiences scattering, platforms rewriting the rules, and technology changing faster than any newsroom can keep up with.

Let's start with the business side, because that's where most of the pressure has built up. Advertising used to bankroll everything. But when Amazon, Alphabet, and Meta now control 55% of the entire global ad market, the scraps left for newsrooms aren't enough to keep most operations stable. Even subscription growth—which spiked during the early pandemic—has cooled off, though a few still believe there's further growth room. Private equity promised to save journalism in the US, but ended up largely hollowing out the industry and creating local news deserts. The result is that a few large publishers survive, while hundreds of small and mid-sized outlets cut teams, shrink coverage, or shut down entirely.

The safety picture hasn't improved either—only worsened. The Committee to Protect Journalists has reported record-high imprisonments and killings in multiple consecutive years. By the end of 2022, at least 363 journalists had been put behind bars, a 20% increase from the year before. And that's before we get to digital harassment: deepfakes, coordinated pile-ons, and legal threats designed to drain a newsroom's time and cash. You don't need to be in a conflict zone to feel the pressure anymore.

Then there's the growing shift in how people increasingly consume news. Younger audiences now treat social media feeds, influencers, podcasts, and explainer videos as their primary entry point, not news homepages. Reuters data shows that weekly use of generative AI for news has doubled—from 3% to 6%—driven mainly by 18–24s who use AI to simplify or summarise complex stories.



Younger readers haven't forsworn the news; they just expect it to come to them through the people, platforms, and AI systems they already use.

This has rewritten the distribution game. Homepages matter less, SEO doesn't guarantee visibility, and audiences increasingly consume news through AI summaries that sit between them and the original source.

Technology inside newsrooms has changed just as much. Generative AI is already part of the news workflow, transcribing long interviews, cleaning up notes, exploring headline options, or pulling together background context in seconds.

Tools like ChatGPT, Perplexity, and newsroom-specific automation platforms are now part of daily work. But full AI-written articles remain the line most editors won't cross, mainly because trust is brittle, and a few high-profile scandals have shown what happens when things go wrong. One mistake can undermine a brand.

So the industry sits in a strange moment: stretched thin and under-resourced, yet expected to move faster and cover more ground than ever. And while the tools have advanced, the operating model hasn't caught up. That tension explains why AI, and specifically agentic AI, has entered the conversation.



2. Why journalism must take AI seriously



Newsrooms now operate with tighter budgets, fewer specialists, and rising expectations from audiences who want context, speed, and accuracy at the same time. This strain shows up in long research cycles, slow verification processes, labour-heavy transcription, growing data volumes, laborious multimedia editing, and constant pressure to publish across multiple platforms—all while reducing churn and raising revenue. These problems are already shaping daily editorial decisions, and it's where AI enters the picture.

As UK-based media strategist [Dietmar Schantin](#) notes, AI's impact operates in layers. At the basic level it removes manual strain—transcription, edits, SEO cleanup. At a deeper level it elevates journalism by improving research, mining documents, and personalising content for different formats and audiences. And at the most transformative layer, it begins to shape entirely new ways people interact with news, shifting from one-way publishing to something closer to an ongoing conversation.

But AI isn't one thing. Different branches of AI affect newsroom work in different ways depending on the model, function, process, or taxonomy, and each one is already showing up in real workflows.

- **Generative models create text**, images, audio, and video. They draft explainers, summarise long reports, rewrite witness statements, produce early graphic mockups, and clean or enhance field recordings.
- **Predictive and decision models** classify, rank, and forecast. They detect topics in community chatter, flag issues gaining traction, identify gaps in coverage, and help editors prioritise story leads.
- **Autonomous systems** run multi-step workflows. They scan large batches of court filings, extract key names and events, match them against earlier cases, and produce structured briefs without manual stitching.
- **Perception models** interpret raw inputs. They transcribe interviews, translate foreign-language clips, detect objects or scenes in citizen footage, and turn scanned PDFs into searchable text.
- **Natural language processing systems** understand, analyse, and structure language. They extract entities from documents, detect sentiment in public comments, segment long transcripts, classify themes in reader messages, and break down large text sets for quicker reporting.
- **Retrieval and grounding systems** search, cluster, and link information. They surface archived context during breaking news, check claims against previous coverage, cluster leaked documents, and anchor new reporting in verified facts.

Recent surveys show how widespread this shift already is. A 2023 [study by JournalismAI](#) of more than 105 newsrooms found that over 75% now use AI tools in at least one part of news gathering, production, or distribution. Another industry survey from WAN-IFRA reported that about half of newsrooms were using generative AI tools, though only around 20% had clear guidelines in place.

A [2025 Trint study](#) added another layer: 42.3% of journalists were using generative AI tools at work without formal approval from their organisations. And on the audience side, Reuters Institute data from 2024 found that only 36% of respondents felt comfortable with news created by humans with AI support, showing the gap between internal adoption and external trust.

But despite growing experimentation, not all journalists are onboard with AI. In Cision's [2025 State of the Media Report](#), 30% of journalists cited the emergence of AI as one of the biggest challenges the industry has faced in the last year—up from 26% in the [2024 report](#). New York-based media consultant, journalist and founder of The Media Copilot, [Pete Pachal](#), validates this unease, arguing that the short-term impact of AI is painful for reporters. He notes that its writing abilities may feel like an existential threat, forcing journalists to reassess what human judgment uniquely adds.

Other experts take a different view. UK-based journalist, editor, and co-founder of Flare Data [Daniel Flatt](#) argues that the greater danger is not experimentation but stagnation. In his view, AI is already replacing parts of the workflow long assumed to be “human-only,” and the real opportunity lies in redirecting resources toward the parts of journalism that remain uniquely human, like discovering news worth reporting.

For journalism, the question isn't whether AI should be used. The work already depends on it. The real question is how to use these systems in ways that expand capacity without weakening standards—especially with agentic AI systems in the mix.



3. What are agentic AI systems?



If you've been following the AI conversation over the past year, you've probably heard the term **agentic AI** thrown around. It sounds technical, but the idea is simple: instead of waiting for a prompt, an agent can take the next step on its own. It can observe, decide, act, and adjust without someone nudging it every five seconds.

Think of it as the difference between a calculator and a colleague. A calculator only answers the question you type in. A colleague might notice a pattern, flag an issue, suggest a better approach, or go fetch the missing data. Agentic AI aims for the latter.

There are a few flavours of agentic AI to understand, not least because much of the literature uses 'autonomous' and 'agentic' interchangeably. In a 2025 essay, Feng, McDonald, and Zhang define five levels of autonomy for AI agents, from operators, collaborators, and consultants to approvers and observers. Some agents can plan, take actions, and refine those actions without human supervision. Others are semi-autonomous, where you set boundaries or goals and they handle the steps in between. The agents we see today usually sit somewhere in the middle: helpful enough to run tasks, but still need a human to keep them aimed in the right direction.

Most modern agents are powered by large language models, but they increasingly combine other abilities. A single agent might read documents, watch a video clip, query a database, and run a workflow tool, all in one chain. We call those multimodal agents. And then there are orchestrators—agents built to coordinate other agents, like an editor coordinating a newsroom.

Under the hood, four pieces make this work: perception (taking in information), decision-making (choosing what to do), memory (remembering context), and execution (actually doing the task). When these pieces are stitched together, the system stops being a passive tool and starts behaving more like a junior assistant with initiative.

Adoption is already spreading. Research groups tracking AI usage show that companies across tech, finance, and media are experimenting with agentic workflows—everything from automated research assistants to agents that summarise meetings or analyse documents. The tools are young, but the direction is clear: more autonomy, more chaining of tasks, and greater ability to act without micromanagement.

Of course, this shift carries risk. Agents still hallucinate, especially when they're allowed to operate without checks. They pick up biases from their training data. Some systems become opaque once you start stacking steps together. And an agent that can act independently also has more room to make the wrong call quickly.

4. How agentic AI can address journalism's challenges

Agentic AI can be the connective tissue inside modern newsrooms, linking reporting, production, distribution, and audience work in ways that reduce pressure on shrinking teams. Instead of treating research, drafting, adaptation, verification, safety monitoring, and audience analysis as isolated tasks, an agent might move through them in sequence.

A system might monitor late-night filings from a city council, flag an unusual budget line, draft a short memo for the morning meeting, generate a newsletter-ready digest, and prepare a mobile version for the social team. That's five steps in a row. Another might watch breaking news channels, pull fresh quotes from live captions, assemble a quick TikTok script built from the reporter's notes and voice, and check incoming footage for obvious signs of manipulation.

Flatt, the editor and data consultant, notes that some publishers he works with are already automating the acquisition and processing of publicly available information end-to-end, using AI systems to sort, extract, and prioritise material before a human editor ever sees it. These experiments show how agentic workflows are beginning to move beyond isolated tasks toward full-chain automation.

Journalism and agentic AI: Use cases

The above flows don't replace editorial judgment, but they open up time for it by removing the repetitive steps that usually slow teams down. This impact becomes obvious when we zoom out and look at newsrooms holistically.

In **reporting and editing**, an agent might sift through thousands of pages of environmental impact documents overnight and leave a reporter with a list of ten pages worth a closer look. It might watch dozens of local WhatsApp groups and surface early chatter about a policy change before it breaks. It could assemble the first-pass framing for a templated story—an earnings update, a weather alert, or a local council vote—so reporters spend their energy on the nuance rather than the structure.

Operationally, agents might help route assignments, track shifting deadlines, or organise contributions from a distributed team of freelancers. A multilingual newsroom might rely on agents to produce fast translation drafts for editors who then refine idioms and tone.

Safety and compliance work also stands to gain. An agent might track the geolocation of a reporter covering a rally, monitor social chatter for threats, or run analysis on submitted images to detect likely manipulation before publication. It might also flag sensitive phrases or legal risks in a draft, giving editors a clearer starting point.

Audience teams could use agents to identify regional variations in interest, generate story versions tailored to different reader groups, schedule posts when engagement tends to spike, or summarise comment-section themes that should inform follow-ups.

There's a **commercial** angle, too. A system could manage the constant flow of inbound emails from agencies and local businesses, sorting genuine opportunities from generic outreach. It might scan available inventory across the site, newsletters, podcasts, or video channels and propose the best placements based on audience fit and upcoming publishing cycles.

Once a deal is agreed, an agent could draft insertion orders, prepare ready-to-send invoices, track campaign dates, and route any questions to the right person on the sales team. None of this replaces relationship-driven work, but it reduces the administrative drag that usually slows the path from inquiry to income.

Schantin, the media strategist, argues that as automation takes over technical steps, the core human parts of journalism become even more valuable. Interviewing, synthesising complexity, shaping narratives, and producing genuinely original reporting sit at the centre of what AI cannot replicate. In his view, the craft of journalism is being distilled to its most intellectual work precisely because agents handle the repetitive load around it.

Use case	What AI agents might do	Example of newsroom workflows
Editorial augmentation	• Surface early signals from large document sets • Build first-pass drafts or explainers • Assemble structured outlines for recurring formats	A city-hall agent scans filings overnight and leaves a reporter a shortlist of anomalies with suggested angles.
Multimedia and platform adaptation	• Produce platform-specific versions of a story • Generate short video recaps from transcripts • Suggest metadata or headline variations	A breaking-news agent turns a reporter's field notes into a mobile summary, a TikTok script, and a newsletter paragraph before dawn.

Use case	What AI agents might do	Example of newsroom workflows
Operational streamlining	• Route assignments and track workflow changes • Coordinate freelancer inputs • Provide instant transcription, translation, and summaries	An editor assigns an agent to manage drafts from three freelancers, ensuring versions are tracked, merged, and summarised for review.
Safety and compliance	• Track reporter locations in risky assignments • Run media authenticity checks • Flag legal, ethical, or privacy-sensitive elements	While a reporter covers a protest, an agent monitors geolocation, scans footage for manipulation, and flags phrases needing legal review.
Audience targeting and feedback loops	• Produce regional or interest-based story variants • Optimise publishing schedules • Summarise audience reactions and questions	A newsletter agent generates Lagos-specific, Nairobi-specific, and Johannesburg-specific versions based on local interest patterns and feedback.
Commercial operations	• Sort and triage inbound ad inquiries • Match campaigns to available inventory • Draft insertion orders and invoices • Track campaign flow • Generate performance reports	An advertising agent screens agency emails, proposes placements based on audience fit, prepares invoices, and routes complex deals to sales.

5. Implementation considerations: Workflows, talent, tools, and standards

Implementing agentic AI is not a plug-and-play decision. It reshapes how work flows through a newsroom, how teams collaborate, and how editorial responsibility is defined. Instead of asking which tool to buy, the smarter question is how the newsroom itself must adapt. Four areas matter most: **workflow design**, **organisational change**, **infrastructure**, and **ethics**.

Workflow design and integration

The first task is mapping out every workflow that an agent might touch. Most newsrooms have undocumented processes held in people's heads, which is exactly where automation fails. You can't transform what you don't fully understand. Teams need to map who does what, when they do it, what they rely on, and what decisions hinge on human judgment. Only then can you decide where an agent fits.

Clear accountability lines follow from that map. Reporters, editors, fact-checkers, and audience leads need to know which steps remain human-only and where an agent is allowed to act. Notification paths matter too. If an agent flags a correction, surfaces a legal concern, or detects a safety risk, the newsroom must already know who receives that alert.

Prompting becomes part of editorial craft. Not in the gimmicky sense of "prompt engineering," but in the practical sense of giving an agent enough structure, constraints, and context to stay on track. A politics desk might maintain prompt templates specifying source hierarchies and phrasing requirements.



A climate desk might define rules for numerical ranges, units, or scientific terminology. This is editorial quality control expressed through instructions.

Feedback loops and explainability keep these systems reliable. Reporters must be able to mark an output as biased, shallow, or misaligned, and the agent should adjust. Teams also need visibility into how an agent made a decision: what it retrieved, what it ignored, and how it ranked its options. Without this, trust erodes and oversight collapses.

Branislava Lovre, a journalist, AI ethics consultant, and co-founder of Almedia and Almpactful, stresses that the biggest risk isn't the technology itself, but uneven access to support. Some colleagues will quietly fall behind simply because they didn't receive enough training, encouragement, or time. In her view, good AI integration starts with clear guidance, steady support, and a culture where everyone feels included as workflows evolve.



“AI becomes most impactful in journalism when shaped by human judgment and ethical intent. A human-centric approach ensures that emerging technologies strengthen truth, accuracy, and public trust.”

— Prof. Devadas Rajaram, Alliance University

Organisational change and talent

Agentic AI shifts job balance rather than replacing jobs. Reporters may spend more time interviewing, analysing, and shaping narratives, and less time on repetitive tasks like transcription or platform adaptation. Editors shift from managing drafts to orchestrating workflows and refining agent constraints.

New roles appear. Some newsrooms will need AI-literate specialists who maintain agents, monitor logs, manage guardrails, and troubleshoot behaviour. Others will appoint AI editors responsible for quality, escalation, and risk review. These roles ensure that automation strengthens rather than weakens standards.

As Schantin puts it, newsrooms can't outsource this shift to engineers. Journalists need to lead the transformation by defining what AI should do, what it must not do, and where editorial judgment remains non-negotiable. The challenge is less technical than intellectual: aligning the system with newsroom values, language standards, and expectations of quality.

Tools, data, and infrastructure

The choice between building and buying determines the newsroom's level of control. Open-source agents offer transparency, tunability, and on-premise deployment, but they require deep technical skill and steady maintenance. Proprietary systems offer convenience and safety features but introduce risks around lock-in, pricing changes, and limited customisation.

Data is the backbone of every AI agent. Newsrooms need clean archives, structured metadata, reliable analytics, and clear data boundaries. Confidential documents, unpublished drafts, and sensitive reporting materials must never flow into external training pipelines. Even small misconfigurations can lead to privacy breaches or loss of source trust.

Governance tools are essential: logging, monitoring, audit trails, and permission systems that track which agent accessed what, when, and why. Transparency mechanisms should make it clear to any editor how a specific output was generated, what sources influenced it, and which parts were human-authored.



“With support from newsroom leadership, AI becomes much easier to adopt. When people have a safe space for training, opportunities to experiment, make mistakes, and improve through open conversations, everyone gets a fair chance to explore and make informed decisions together.”

— Branislava Lovre, Co-founder at Almedia and Alimpactful

Ethics, safety, and standards

Every newsroom must decide what ethical boundaries apply to agentic systems. Bias mitigation is a central concern. Without careful constraints, agents may reproduce harmful stereotypes or skew story selection based on flawed patterns in training data. Safety considerations include ensuring that agents do not surface sensitive information, misuse private data, or push unverified claims.

As Professor Devadas Rajaram of Alliance University notes, AI becomes most impactful in journalism when shaped by human judgment and ethical intent. A human-centric approach ensures that emerging technologies strengthen truth, accuracy, and public trust.

Privacy protections need to be strict. Anything involving reader identities, whistleblower communications, embargoed material, or source documents requires locked-down handling. These systems must not become accidental gateways for exposure.

Finally, transparency—both internally and externally—shapes trust. Teams should know how AI fits into their workflow, and audiences should know when AI meaningfully contributed to a story. Hidden automation creates confusion and suspicion, while open practice builds confidence.

In sum, these decisions shape whether agentic AI becomes a stable part of the newsroom or an unpredictable add-on. Getting the workflows, roles, tools, and safeguards right at the start gives teams the structure they need to use these systems with confidence.



6. When things go wrong



Agentic AI systems move quickly, act independently, and operate across workflows that were previously slow, linear, and human-gated. That speed is useful, but it also means errors can spread before anyone notices. And because newsrooms already rely on automation and AI in distribution, analytics, advertising, and content workflows, we've seen enough real-world cases to know what happens when these systems slip out of alignment.

Accuracy is the easiest place to start. CNET's [experiment](#) with AI-generated financial explainers showed how easily a system can produce text that looks polished but carries errors. Between 2022 and 2023, the outlet published dozens of finance articles built with an internal AI tool.

When outside reporters examined them, they found basic arithmetic errors, misleading examples, and uncited phrasing familiar from other publications. None of it was malicious, but it was the kind of subtle wrongness an automated system can introduce when it doesn't understand the stakes. When the errors surfaced, the newsroom had to issue multiple corrections and defend its editorial processes. The public frustration wasn't just about the mistakes, but about discovering that AI had been used at all. Correction rarely repairs the full damage to trust.

Bias follows a different path but creates similar pressure. In the run-up to the 2020 U.S. election, Facebook's ranking system elevated hyper-partisan political content at a scale the company later acknowledged. These posts spread not because people searched for them, but because the algorithm rewarded material that triggered strong emotional reactions. That skewed the public conversation long before journalists could contextualise or counter it. For newsrooms considering agentic AI that recommends angles, surfaces topics, or highlights trending narratives, a system optimised for engagement can distort editorial judgment unless its incentives are shaped carefully.

Data privacy risks emerge when automated systems gain access to information beyond what teams expect. The Meta Pixel investigation by The Markup revealed that tracking code on hospital websites was sending sensitive health information like appointment details, medication names, and search terms back to Meta. No one visiting those pages knew their data was being captured. Although this wasn't a newsroom case, the mechanism is nearly identical to what can happen when AI tools are wired into content-management systems, archives, analytics dashboards, or document libraries. A misconfigured agent could pull unpublished drafts, internal notes, or confidential source material into a process where it doesn't belong. In a field where source protection is non-negotiable, that kind of leak is unacceptable.

Reputational risk often arrives through the distribution layer. In 2017, YouTube faced global backlash when advertisers discovered their campaigns being placed next to extremist content. The placement decisions were made by automated systems running at massive scale, but the fallout landed squarely on YouTube and the brands involved. Newsrooms face a parallel risk when agents generate social content, select thumbnails, or advertising placements. One mismatched placement can create the impression that a brand endorses a context it never intended to touch. Audiences won't blame the algorithm; they'll blame the newsroom.

Opacity itself can undermine trust, even when nothing goes catastrophically wrong. Agentic systems make decisions in long chains—retrieval, ranking, drafting, refinement, and optimisation. If an editor can't trace how a piece of information entered a story or how an agent reached a conclusion, they can't vouch for the work built on top of it.

Trust is the through-line across all these failures. Research from the Reuters Institute on public attitudes toward the use of AI in journalism in 2024 found that audiences are generally comfortable with AI performing backstage tasks like transcriptions, translations, and summaries, but less comfortable when AI is generating new content.

When CNET's AI-generated stories came to light, some of the strongest public reactions weren't about the numerical errors, but the lack of disclosure.

Together, these examples show that agentic AI doesn't introduce new categories of risk so much as it amplifies existing ones and accelerates their consequences. The lesson is not to avoid the technology—journalists already depend on automated systems in dozens of ways. It's that governance must be built in at the start.



“The biggest risk is to stagnate or do nothing. AI threatens to take out many journalistic processes that have long been considered the preserve of the human.”

— Daniel Flatt, Co-founder of Flare Data

Transparency, clear boundaries, and strong human oversight are the difference between a tool that strengthens reporting and one that compromises it. Newsrooms that prepare for these risks early will be able to use the technology confidently. Those that don't will face the same cycle we've already seen across tech: after-the-fact disclosures, public backlash, damaged trust, lost revenue, and systems rushed back behind guardrails.



“Don't let engineers lead the AI transformation—journalists must. The challenge isn't technological; it's intellectual and linguistic. Success depends on journalists' ability to articulate precisely what AI should do and why.”

— Dietmar Schantin, Co-Founder at AI-Collective UK

7. Conclusion

Journalism has always adapted to whatever tools shaped the public conversation. The printing press, radio, television, the web, and mobile each changed how information moved and who controlled it. Agentic AI is another shift of that scale. It adds a kind of capacity that most newsrooms have not had in decades: the ability to process large volumes of information quickly, to monitor emerging signals around the clock, and to support the parts of reporting that usually get squeezed when resources are tight—all autonomously.

This doesn't replace the work that matters most. AI doesn't interview sources, weigh competing claims, understand a community's history, or carry the responsibility of publishing under a trusted name. What it does is create space for those human parts of the craft by taking on the heavy, repetitive, time-sensitive work that slows teams down. Used well, agentic AI can help newsrooms regain the depth, range, and accuracy that shrinking budgets have eroded.

Pachal, the media consultant and journalist, asserts that AI may even strengthen journalism in the long run. As more people get information through AI systems rather than news sites, he argues that incentives shift toward deeper, more substantive, and more distinctive reporting—the kind that builds engaged, loyal audiences even if overall traffic declines.

But it only works if the integration is intentional. Newsrooms need clear boundaries, transparent policies, and a culture that treats AI as a support system rather than a shortcut. They also need the courage to explain these choices to audiences openly. Trust doesn't vanish when AI enters the workflow; it vanishes when audiences feel kept in the dark.

The bigger point is that journalism is already competing with systems that move faster, personalise better, and operate at global scale. If newsrooms avoid agentic AI entirely, they risk falling further behind those forces.



You can't ignore AI. You can choose not to use it, but it's still a large and growing factor in the information ecosystem. The best way to deal with it is to understand it, and the only way to understand it is to use it.

— Pete Pachal, Founder of The Media Copilot

Agentic AI won't replace journalism, but it can help sustain it. And at a moment when the old models are cracking, agentic AI may be the difference between a profession that survives the next decade and one that fades as the news agents of tomorrow take shape.

Meet the experts



Branislava Lovre is a journalist and AI ethics consultant who specialises in the intersection of media, technology, and responsible AI adoption. She co-founded Almpactful and Almedia, where she guides organisations on ethical AI integration in newsrooms.

Connect: [in linkedin.com/in/branislava-lovre](https://www.linkedin.com/in/branislava-lovre)



Daniel Flatt is a business and finance journalist and editor, and the co-founder of Flare Data. He works on data mapping, AI governance, and due-diligence workflows, with a focus on how AI and datasets reshape financial reporting.

Connect: [in linkedin.com/in/flatt](https://www.linkedin.com/in/flatt)



Devadas Rajaram is a digital journalism educator with over 15 years of experience in mobile journalism, social media storytelling, and generative AI. He is an Assistant Professor at Alliance University and a long-time trainer in emerging media practices.

Connect: [in linkedin.com/in/devadasrajaram](https://www.linkedin.com/in/devadasrajaram)



Dietmar Schantin is a media strategist and advisor with more than two decades of experience helping major news organisations modernise their editorial, product, and commercial operations. He co-founded AI-Collective UK, supporting publishers as they build AI-ready newsrooms.

Connect: [in linkedin.com/in/schantin](https://www.linkedin.com/in/schantin)



Pete Pachal is a media consultant and journalist, and the founder of The Media Copilot, a publication focused on how AI is transforming media and content creation. A former editor and strategist, he advises on AI adoption and digital media strategy across the industry.

Connect: [in linkedin.com/in/peterpachal](https://www.linkedin.com/in/peterpachal)

About the author



Mo Shehu, PhD

CEO @ Column

Mo Shehu is a writer, researcher, and communications leader focused on the future of media, technology, and communications. He leads Column, a research and media agency that supports B2B brands and public organizations with clear, data-driven storytelling. His background includes work in AI policy, digital culture, analytics, and editorial strategy across Africa, the UK, and the US. He lives in the UK.

About Column

Column is a research and media agency that helps organisations communicate complex ideas with clarity. The team produces reports, briefs, and editorial content that turn data into insight. Column supports clients with content, research, and narrative development.

Learn more at:



columncontent.com