FOJO CASEBOOK · 11/2026



CASEBOOK ON AI FOR MEDIA

From Experiment to Capability

Armenia, Belarus in exile, Moldova and Ukraine: what newsrooms built with AI, and the choices ahead

Linnæus University | Fojo.

Abstract

This casebook examines how newsrooms in Armenia, Moldova, Ukraine and Belarusian exile are moving from individual experiments with artificial intelligence toward systems embedded in everyday editorial work. It connects the global AI landscape with the practical conditions facing smaller media organizations: limited funding and technical capacity, uneven support for local languages, changing routes to audiences, and the demands of journalism under war, repression and displacement. Its central question is what turns an accessible AI tool into a useful, accountable newsroom capability.

The book combines industry research, regional surveys, published implementation accounts and first-hand newsroom questionnaires. Four detailed cases follow different stages of the production process. Civic Radar AI at Moldova's BAStv monitors official sources and prioritizes potential local stories. SmartFlow at Armenia's Ampop Media collects and clusters coverage before preparing summaries for editorial review. Ukraine's 18000 platform adapts published reporting for social channels within its content management system. BGmedia uses a digital presenter to produce video while reducing conventional production requirements and limiting journalists' on-camera exposure. International comparisons extend these cases into archive search, verification, audience engagement, internal operations and reader revenue.

Across the cases, the strongest pattern is that value comes from redesigning a repeated workflow around a clear editorial need. Collection, data preparation, integration, scheduling and maintenance often matter as much as model choice. Human responsibility remains central: journalists assess relevance, inspect original evidence, correct language and decide what can be published. Because several projects are pilots or early systems, the book distinguishes intended benefits and newsroom-reported experiences from measured outcomes. Regional surveys and the initiative database provide context, rather than a representative census of adoption across all four settings.

The regional evidence also shows why a common technology produces different priorities. Armenian examples emphasize language infrastructure and local context. Moldovan evidence highlights a gap between uptake and governance. Ukrainian newsrooms combine widespread routine use with constraints on resources, skills and infrastructure. Belarusian outlets in exile place particular weight on continuity, security and access to audiences. These differences make local knowledge, editorial rules and trusted source material essential to implementation.

The casebook concludes with choices about open-weight models, shared infrastructure, transparent AI use and publisher-controlled access to verified archives, including through the Model Context Protocol. A phased plan links an initial audit and baseline to a measured pilot, then to decisions about maintenance, expansion or closure. The practical argument is to invest in capabilities that endure beyond any model subscription: reliable workflows, portable archives, language resources, staff skills and accountable editorial judgment.

Under the general editorship by Oleg Khomenok
Responsible editor Fojo:
Veronika Menjoun
Published by Fojo Media Institute, Linnaeus University, Kalmar, Sweden 2026
with financial support from the Swedish Institute

Inside the casebook

A practical guide to AI implementation in Armenia, Belarusian media in exile, Moldova and Ukraine.

Abstract . . . 2

Part 1. Where Newsrooms Stand With AI . . . 5
1. The technology itself: commercial giants, open alternatives, and a Chinese layer . . . 6
2. How newsrooms are actually using it . . . 8
3. What AI is actually changing: business, content, distribution, audience . . . 9
4. Zooming in: Eastern Europe as the sharpest version of all of this . . . 13

Part 2. What Newsrooms Actually Built . . . 20
1. Finding the story before it disappears into the official record . . . 21
2. Turning a flood of coverage into the beginning of journalism . . . 23
3. Distribution is editorial work too . . . 26
4. Making video when conventional production is impossible . . . 29
5. The newsroom behind the published work . . . 31
6. Where the money is . . . 32
7. What changes when a newsroom builds instead of experimenting . . . 33
Sources for the “Elsewhere” examples . . . 34
Postscript. The punk newsroom . . . 35

Part 3. The Choices Ahead . . . 36
1. What's already settled, and what's still up for grabs . . . 37
2. Six forces shaping the next few years . . . 37
3. Open source as the resource-constrained newsroom's way out . . . 38
4. Four scenarios . . . 39
5. The MCP bet . . . 40
6. Country trajectories . . . 41
7. What to do in three phases: 3 months, 6 months, and 1 year . . . 43
Closing . . . 43

Appendix. Sources index . . . 44
1. Landscape & research . . . 45
2. Newsrooms & regional sources . . . 47
3. Implementation cases . . . 50
4. Strategy, distribution & trust . . . 53

How to read this book

Part 1 maps the landscape. Part 2 follows four newsroom systems through the work they do. Part 3 turns the evidence into choices and a phased plan.

THE IDEA

The newsroom proposition and who it is for.

THE PROBLEM

The constraint the tool is designed to address.

HOW IT WORKS

The workflow, the technology and the human checks.

WHAT CHANGED

How testing or implementation reshaped the original plan.

WHAT IT CHANGES

The claimed value, remaining work and limits of the evidence.

ELSEWHERE

Related experiments to adapt to local conditions.



1

Where Newsrooms Stand With AI

The technology, the business, and the regional landscape

Two years ago, "AI in the newsroom" mostly meant one thing: a journalist typing a prompt into ChatGPT. That description no longer holds. By the autumn of 2026, AI sits inside CMS systems, translation pipelines, ad-sales dashboards, comment-moderation queues, and increasingly, the way readers ask a chatbot for the news instead of opening an app or a website at all.

This chapter draws on data checked on 2 November 2026. First, what the technology itself looks like now: commercial giants, open alternatives, and a fast-growing Chinese layer most editors have barely heard of. Then, how newsrooms around the world are actually using it. Then what it's doing to four things editors care about most: the business, the content, the distribution, and the audience's trust. It ends by zooming in on the media in the region this text is about: Armenia, Belarus in exile, Moldova, and Ukraine, where nearly every global pressure described here shows up at once, often in its sharpest form.

1. The technology itself: commercial giants, open alternatives, and a Chinese layer

For most of 2023 and 2024, "which AI should we use" had an easy answer for most newsrooms: whichever commercial chatbot was best that month, usually from OpenAI, Google, or Anthropic. That's no longer the whole picture, and understanding why matters for every decision that follows in this book.

Commercial models, the ones you pay a subscription or per-use fee for, from companies like OpenAI, Anthropic, and Google, remain the default for most newsrooms, especially for complex work: deep research, multi-step reasoning, long-document analysis, and anything where a newsroom wants a company on the other end of the contract to hold accountable if something goes wrong.

Open-weight models are the second tier, and they've changed the economics of the whole field. "Open-weight" means the company that built the model publishes the underlying files, so anyone with the right hardware can download and run it themselves, instead of renting access through an app or API. Three years ago, open models lagged well behind commercial ones in quality. That gap collapsed dramatically in 2025 (for a stretch, an open model briefly matched the best commercial reasoning models on independent benchmarks) before widening again slightly in 2026 as commercial labs pushed further into complex, multi-step reasoning. But on the tasks that make up most of daily journalism, drafting, summarizing, tagging, transcribing, translating, writing small pieces of code for a scraper or a CMS integration, open models are now good enough for the overwhelming majority of newsroom use cases, at a fraction of the cost. Getting equivalent output quality is roughly 50 times cheaper today than it was three years ago, and open-weight releases are the main reason: each time a strong open model launched, it forced commercial prices down behind it.

That progress comes with real concerns attached, and the AI companies building these models don't fully agree on how serious they are. Once a company publishes a model's weights, it can't take them back. There's no way to patch a flaw the way a hosted app gets updated overnight, no way to see who is running a copy or cut off a bad actor, and any safety guardrails built into the model can be stripped out by anyone with the technical skill to do it, permanently. Anthropic's own public position on this, laid out by CEO Dario Amodei, argues that the sharper risk isn't openness itself so much as the chance that an authoritarian government builds a more capable model than anyone else and uses it for military advantage or domestic surveillance, a concern that holds whether or not that model's weights ever get published. Microsoft, in a

letter signed by dozens of AI companies including Google, Meta, and OpenAI, acknowledges the same loss-of-control problem directly, while arguing that restricting open models broadly would backfire, since the people defending computer systems against attacks need access to AI at least as capable as what attackers already have. A researcher at the R Street Institute made a related point with a concrete example from mid-2026: when the US government temporarily cut off worldwide access to one of Anthropic's own frontier models over cybersecurity concerns, independent researchers reproduced comparable results within weeks by combining several older, publicly available open models instead, evidence that restricting one specific model doesn't contain a capability so much as decide who can use it easily and who has to work harder for the same result. Reporting from the BBC on the rapid rise of Chinese open-weight models points to a more immediate question for any newsroom actually choosing one: a model trained under a different country's information and censorship rules can carry those rules' blind spots and omissions into its answers, which is worth weighing alongside price and capability, not after them.

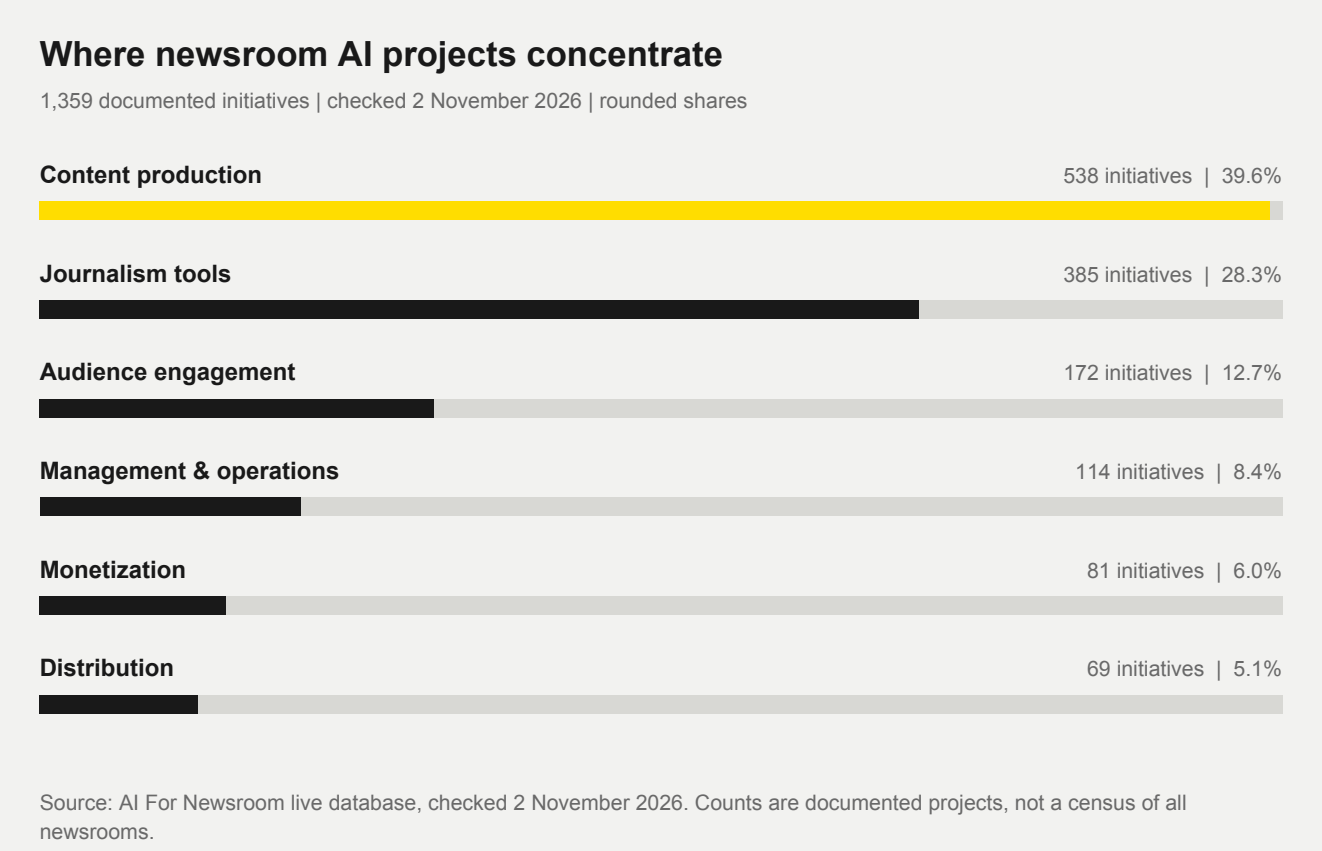
A Chinese layer most editors don't know exists. Here is the part that surprises people outside the industry: measured by how much AI computing actually gets used day to day, Chinese-built open models, from labs like DeepSeek, Alibaba (Qwen), Tencent, and MiniMax, now account for a very large share of global usage, particularly for coding and automated, repetitive tasks. They're cheap, openly downloadable, and improving fast. This isn't purely a Chinese story, though. American companies have followed the same open-weight playbook: OpenAI now publishes its own open-weight models alongside its commercial ones, and so does Google. Open-weight has become a mainstream deployment choice across the whole industry, not a niche alternative.

Even so, Chinese labs building some of the most widely used open models in the world is a different fact from Chinese newsrooms using them in documented, public ways. The collection behind this chapter has a handful of cases from China against well over a hundred from the United States alone, which is likely a visibility and language gap in what gets reported internationally rather than evidence that nothing is happening. Who builds a technology and who's been documented using it are two different questions, and the gap between them tends to be widest exactly where this book is looking, outside the big English- and Western European-language markets.

The last thing worth knowing about the technology layer: open models are easy to try and hard to run well. Across the industry, most newsrooms and companies have experimented with an open model, but only about half of those experiments make it into daily, reliable production use, compared to roughly two-thirds for commercial ones. The gap isn't quality. It's the unglamorous work of keeping a self-run system secure, maintained, and monitored, exactly the kind of technical capacity small and mid-sized newsrooms tend to lack. That gap will matter a great deal later in this book.

2. How newsrooms are actually using it

As of 2 November 2026, the live AI For Newsroom database groups 1,359 initiatives into six broad uses. Content production is the largest category (538 projects, 39.6%), followed by journalism tools such as fact-checking and document analysis (385, 28.3%). Audience engagement accounts for 172 projects (12.7%); management and operations, 114 (8.4%); monetization, 81 (6.0%); and distribution, 69 (5.1%). Production still dominates, but more than one in four documented projects help journalists find, test, or organize evidence.



Sources: AI For Newsroom database

A few concrete examples make this more real than the percentages do.

- **The New York Times** has built its generative-AI use around a small set of standing principles: AI supports the newsroom's mission rather than replacing it, every application gets human review, and its use is disclosed. In practice that means AI helps categorize images, suggest related articles, draft headlines for review, and summarize long documents, not write the actual story. Sources: New York Times AI principles
- **Reuters** has built what it calls an "AI-forward" newsroom around internal assistants, not internal authors: coding tools that help data journalists debug analysis scripts, systems that digest hundred-page documents and surface the key numbers, and publishing-desk checkers that flag likely errors before a story goes out. Reuters pulled back an AI-generated story-summary feature after noticing problems with

attribution and outdated facts, a reminder that newsrooms are still actively correcting course, not just rolling AI out and hoping for the best. Sources: Reuters AI-forward newsroom

- **VG**, Norway's largest news site, built a suite of six AI tools covering transcription, statistics analysis, and finding connections between people and events in its reporting. Transcription, in particular, shows up as the single most common starting point for AI adoption almost everywhere in the world, because it's the task every newsroom already does by hand, all day, at real cost. Sources: VG case
- **Nikkei** and **Süddeutsche Zeitung** have pushed further into audience-facing tools: Nikkei on engagement and personalized delivery, Süddeutsche Zeitung building a chatbot around EU elections and treating AI as a permanent layer in its product, not a one-off experiment. Sources: Nikkei · Süddeutsche Zeitung
- On the money side, **Observador** in Portugal built an AI subscription concierge aimed at deepening reader relationships rather than chasing ad growth; **Ringier** in Switzerland combines reader-revenue and advertising data into a single customer-value model; and in India, **Sakal Media Group** uses AI-driven optical character recognition to make its advertising sales operation more efficient, a much less glamorous use case than a chatbot, and a good illustration of how much of this technology's real value shows up in the boring parts of the business. Sources: Observador

3. What AI is actually changing: business, content, distribution, audience

Business

The clearest business effect so far is on costs, not revenue. Drafting, tagging, transcription, and translation have gotten dramatically cheaper, and that mainly helps smaller and mid-sized newsrooms that could never have afforded a large in-house tech team. There's a hidden cost on the other side of that ledger, though: running any of this reliably and securely, with proper oversight, is itself a real, ongoing expense, and it's the reason so many pilots never turn into permanent tools.

On the revenue side, monetization remains one of AI's least-solved problems in journalism. It accounts for only 6.0% of the live collection, and the most interesting experiments - Observador's subscription concierge, Ringier's blended customer-value model, the Washington Post testing more flexible paths to a subscription, and the New York Times positioning itself as subscription infrastructure other publishers can plug into - all treat AI as a new lever on an old business model rather than a wholly new one.

Content creation

Across virtually every serious newsroom in the collection, the pattern is the same: AI is used to prepare material for a journalist, transcribing an interview, summarizing a long document, suggesting a first-draft headline, converting a written piece into narrated audio or a short video, and a human reviews it before it goes out. That's true at the New York Times, at Reuters, and at VG. The tasks AI has genuinely absorbed are the ones that were mostly mechanical to begin with.

Not every attempt goes well. In Denmark, the outlet Politiken published an AI-generated fact box alongside a story that turned out to be full of errors, a small, specific, entirely avoidable failure that did real damage to reader confidence in that one piece, and a reminder of exactly why human review before publication has become close to a universal rule rather than a suggestion.

Distribution

For decades, distribution for a newsroom meant one thing: get found on Google, get clicked, earn the ad impression or the loyal reader. That model is now under pressure from two directions, one well documented and one still forming, plus a third, more direct approach that almost no small newsroom has caught up to yet.

The scale of the shift. Similarweb's *2026 Generative AI Landscape* report puts real numbers behind what editors have been sensing anecdotally. Generative AI platforms drew an average of 9.5 billion monthly web visits between June 2025 and May 2026, up 70% year over year, alongside 4.4 billion app downloads. Source: Similarweb, 2026 metrics. AI chatbots as a web category grew unique visitors by 57% over the same period, faster than search, social, or e-commerce. And in that same window, one major web category actually shrank: News, down 5%. Every other large category Similarweb tracked grew. News was the exception. Harry Clarkson-Bennett, a partner at DJB Strategies and formerly of The Telegraph, put it plainly in the report: publishers can no longer expect their traditional channels to return the same results, and need to build audience engagement and loyalty by other means.

Zero-click search. Google’s AI Overviews answer more queries on the results page instead of sending users to publishers. Similarweb estimates that the share of US searches showing an AI Overview rose from about 15% in early 2025 to over 40% by May 2026, while use of Google’s conversational AI Mode also grew. For outlets reliant on search traffic and advertising, fewer clicks threaten revenue, especially where subscriptions provide little cushion.

Publishers are trying to improve their chances of being cited through generative engine optimization (GEO) or answer engine optimization (AEO). These strategies build on search visibility: many AI search tools retrieve current web material. But low-quality, mass-produced AI copy can undermine both search rankings and AI visibility. Sources: Similarweb report

ChatGPT answers with live web citations rose from about 1% in June 2025 to 6.8% in May 2026. Citation rates varied by topic: 22.6% for travel, 6.0% for media and entertainment, and 4.8% for education, suggesting that AI systems do not seek sources equally across subjects.

Across all topics, News & Publishers ranked among the two largest source categories cited by AI systems, alongside reviews and user-generated content, ahead of official brand sites, search results, or marketplaces. Journalism is already a major source of evidence for broad AI answers, even when news questions prompt fewer live searches.

Citations do not always bring readers to the cited article. Similarweb reports that 65% of cited URLs were two or three folders deep, while 58.8% of AI referral traffic landed on homepages. URL depth is not a classification of article type. Clear navigation and a useful first-visit experience matter for readers arriving without context.

AI referrals to news grew 115.6% year over year, but more slowly than referrals to fashion (278.2%) and beauty (312.5%). Large marketplaces still attract far more AI-driven traffic than publishers, suggesting that broad platforms have an advantage in this emerging channel.

The report's analysts recommend tracking share of voice, not citation counts alone: an AI mention can build awareness without a click. Similarweb's current report page associates an AI mention with a 2.5-times greater chance of a site visit. That association should not be read as proof that a mention caused the visit.

Large publishers are treating AI discovery as a business channel. Ahrefs identifies Future plc as the most AI-visible publisher, with TechRadar among ChatGPT's most cited domains. In 2026, Future launched Future Optic to sell AI-visibility services, including a Samsung campaign. "Liquid content" restructures articles so AI can quote them accurately while preserving richer experiences for people; The Economist is testing this approach. Google News Initiative programs have also involved publishers such as Kathimerini and Punch.

Other publishers are restricting access: Reuters and Time whitelist selected AI crawlers, favoring controlled, negotiated access over blanket openness. GEO and AEO can improve visibility, but neither guarantees traffic or revenue.

AI platforms are monetizing this attention too. Similarweb's 2026 comparison reports ads in 26% of US desktop ChatGPT chats, up from 14% in May 2026. That measure is ad penetration in a particular channel, rather than the share of every ChatGPT conversation worldwide. Publishers compete for visibility alongside advertisers within AI answers. Sources: Similarweb, 2026 landscape report · Similarweb, metrics and methodology

Getting queried: MCP, still the underrated one. Most newsrooms haven't caught up to this part yet, and it may end up mattering more than GEO and AEO combined.

GEO and AEO are, at their core, about earning a mention inside somebody else's product, trying to influence how a general-purpose AI system talks about an outlet, on that system's terms. A newer approach called MCP (Model Context Protocol) works differently: it's an open technical standard that lets a reader's AI assistant query a newsroom's own verified archive directly, in real time, rather than relying on what the AI happened to learn during training or find in a web search. Think of it less like search engine optimization and more like plugging a newsroom directly into the machinery an AI assistant runs on, becoming a source it can actually call, instead of hoping it remembers you correctly.

AI For Newsroom's reporting on this exact question frames it as a genuinely different kind of pillar for how a newsroom shows up in the AI era: GEO and AEO compete for a mention inside someone else's chat layer; MCP becomes part of how that chat layer answers in the first place. The project ran its own direct test of the idea. After standing up an MCP server, the assistant Claude began appending things like a newsletter subscription prompt to relevant answers, not because it was told to every time, but because the server had simply made that context available for the model to draw on. That's a new kind of distribution: a newsroom's own call-to-action, surfacing inside somebody else's conversation with a reader, without buying an ad placement for it.

Real examples are starting to appear across very different kinds of outlets:

Gazzetta built its own MCP server, reasoning that AI models are quickly becoming the actual interface through which many readers get information at all, and that this matters most precisely where the open web

itself is degraded, blocked, or hard to reach, whether by censorship, shutdowns, or simple neglect. If a newsroom hasn't made itself queryable, an AI assistant asked about that country or topic answers anyway, just from older, compressed, unattributed training data instead of the newsroom's actual, current reporting. Connecting to it from Claude takes three steps (adding a custom connector and pasting a URL), and the same server works from ChatGPT, Claude Code, Cursor, and other tools that support the standard. Sources: Gazzetta MCP · Guardian archive MCP An independent developer built a public MCP server over **The Guardian's** 1.9-million-article archive, letting anyone ask an AI assistant to find the paper's top stories from a specific day, track how it covered a topic over years, analyze one reporter's body of work, or surface related pieces, entirely through conversation, no coding required. It's a vivid demonstration of how much value sits dormant in almost any newsroom's own archive once it's made queryable this way: not a new story to write, just the existing reporting, unlocked. **Reuters** has built its own MCP server letting subscribers pull Reuters reporting they already pay for directly into their own AI-assisted work, a sanctioned alternative to the uncontrolled scraping most newsrooms are trying to fight off, and a step that, as trade coverage has argued, starts to turn a publisher that does this well into something closer to an infrastructure provider, not just a content supplier. Sources: Reuters MCP **Time** has gone further still, converting its entire published archive into a format built specifically for AI agents to read directly, instead of scraping and guessing. **The Economist** is testing a narrower version of the same idea, starting with lower-risk material like its own marketing copy rather than its core paywalled journalism. And this isn't confined to one company's technology: a related, compatible standard called WebMCP, co-developed by Google and Microsoft, is already being tested by at least one other major publisher, aimed at letting a website hand structured information straight to an AI agent instead of being scraped for it. Sources: Economist ChatGPT app For scale: across the wider tech industry, well over 5,000 active MCP servers already existed by the most recent count available, evidence that this is becoming basic infrastructure far beyond journalism, and that news media has so far been a slow follower rather than an early mover.

This matters most for the newsrooms in this book because the two paths cost very different amounts. Getting good at GEO and AEO takes a growth or SEO specialist most small newsrooms simply don't have. Standing up a first MCP server, by contrast, is closer to a solo, weeks-long technical project than a department-level investment. For a newsroom without the budget to compete for AI's attention the expensive way, this is close to the cheapest real seat at the table currently on offer. And one piece of strategic advice is worth carrying forward regardless of which future eventually wins for how publishers get paid for AI use, direct micropayments, blanket licensing, or coalition-style infrastructure deals like the ones Reuters or Germany's dpa wire service are already exploring: every one of those futures starts from the same place, a newsroom's own content, organized and queryable on its own terms. Building that layer is worth doing now, before the shape of the deals is settled.

Audience behavior is shifting in parallel. The Reuters Institute's *Generative AI and News Report 2026*, published on 29 October, surveyed nine countries in August-September. Weekly AI use for news reached 9% in the six countries tracked since 2024, up from 6% in 2025; across all nine countries it was 8%. These are different survey bases, not interchangeable global estimates. They show a growing conversational route to news alongside established channels. Source: Reuters Institute, Generative AI and News Report 2026

Communication with the audience, and trust

The latest Reuters Institute evidence puts numbers on the audience's caution. Across nine countries, 12% were comfortable with news made entirely by AI, rising to 19% for AI-led news with some human oversight. Translation and spelling or grammar assistance were more acceptable. The report, published on 29 October 2026, finds that attitudes toward newsroom AI remain broadly cautious despite increased use of AI tools. These findings concern the surveyed countries; none of this book's four settings was included. Source: Reuters Institute, Generative AI and News Report 2026

None of this means audiences reject AI in journalism outright. What it means is that they are watching closely for the specific failure modes that erode trust fastest: an error that reaches print because no one checked it, as with Politiken's fact box, or a lack of clarity about what was AI-assisted and what wasn't. The newsrooms handling this best, the New York Times's published principles are a clear example, treat disclosure and human review not as a legal safeguard but as the mechanism by which trust survives the transition. The lesson for every newsroom in this book, especially the smaller and more resource-constrained ones: the appearance of cutting corners with AI costs more in reader trust than the time saved is usually worth.

4. Zooming in: Eastern Europe as the sharpest version of all of this

The pressures described above - lower production costs, zero-click distribution, the race for AI visibility, the opportunity in machine-readable archives, audience trust, and uncertain business models - are especially visible in these four media environments. War, exile, limited language support, and fast-moving public policy shape adoption in different ways.

The evidence is uneven across the four settings. Ukraine and Moldova have recent newsroom surveys; Belarusian independent media are best understood through documented cases shaped by exile; Armenia has a growing record of newsroom and governance examples but no representative usage survey. The profiles below keep those evidence limits visible.

Ukraine: adoption under pressure, with uneven capacity

Ukraine shows the most developed and urgent adoption in the region, shaped by war. The Kyiv Independent uses anonymized reader questions to identify gaps in its war coverage; investigative outlets such as Nikcenter use data tools to examine procurement and officials' declarations; and Rayon.in.ua's Raya assistant supports daily newsroom operations. Other documented examples include draft-audit and pitch tools at Euromaidan Press and repeatable multimedia workflows at Gre4ka.info. Suspilne's 2026 strategy adds a sector-wide dimension, linking internal use to public-interest journalism and collaboration with other media. These are leading examples, not evidence that the whole sector has become technically mature. They show how experimentation can become an owned process when teams invest in maintenance, shared learning, and verification. Sources: Suspilne strategy · Nikcenter reporting

The Calliope image-authenticity bot developed by Mykolaiv center for Investigative reporting, lets journalists submit an original image through Telegram and receive a report in Ukrainian and English with

model findings, disagreements, limitations, and a suggested next verification step. Access may require administrator approval. Treat the report as a preliminary assessment, not proof of authenticity or manipulation; it should inform further verification.

A 2025 DW Akademie mapping based on interviews with more than 20 outlets across Ukraine complements the later survey: it documents broad use of commercial, general-purpose tools, but also gaps in regional fit, integration, funding, and privacy awareness. The authors call for practical privacy training, clearer guidance, and locally tailored tools. Because the study focused on early adopters and experimenters, it maps use cases rather than measuring how common they are across the sector. Source: DW Akademie mapping

The Media Development Foundation surveyed 40 local and niche outlets in May-June 2026: 23 small, eight mid-sized, and nine larger newsrooms. The sample leaned toward Kyiv and included no outlets from several western and southern regions or occupied territories, so it is a useful snapshot rather than a representative census. Seventy percent were online-only. Source: Media Development Foundation survey

AI use was concentrated in routine text work: 85% used it to correct, restructure, shorten, or headline copy, while about 10% used it for research or standards checks. Translation (47.5%) and image generation (about half) were also common. Only two of the 40 outlets used AI to verify their own reporting against source material; most AI-assisted fact-checking meant finding a source for a journalist to inspect.

The survey points to a capacity gap beneath adoption. Only 52% had basic digital infrastructure, and reported technical capacity was much stronger among those outlets. More advanced systems - such as source-monitoring agents and reader-facing bots - were limited to two outlets in the sample.

Governance and staffing lagged: 57% had discussed AI strategically, but 27.5% had a written policy, none had a documented AI strategy, and only 22.5% had a designated lead. Money (60%) and skills (55%) were leading barriers. The report names training as the top need (85%), ahead of tool access (60%) and subscription funding (42%).

The report's separate literature review, covering 236 relevant academic sources from 2023 through early 2026, breaks down which barriers to AI adoption come up most often in published research generally, and one entry is worth flagging specifically for this book: language and localization barriers, difficulty with less-common languages and regional accents, and the deeper question of preserving linguistic diversity and authentic local voice in generated content, account for 4.6% of the barriers researchers identified. It's the smallest share among the eight barrier categories, but it's also the one this book's Armenian and Belarusian chapters are built around, a reminder that what shows up as a minor line item in a global literature review can be the central constraint for a specific language community.

The 40-outlet survey is the quantitative backbone, but the same research effort also compiled desk-research case studies of individual Ukrainian newsrooms, and they put faces on the aggregate numbers. Tsukr, an outlet team, uses Claude for internal team management, with an editor quoted saying the tool "knows more about me and my team than I can remember myself." **24 Kanal** trained an AI system to voice its own news bulletins in a synthetic version of its presenters' voices. **Gre4ka.info** built a one-click pipeline that turns written stories into songs and short videos, entertainment framing wrapped around real reporting rather than a news product in its own right. **Texty.org.ua** went the opposite direction and trained AI specifically to flag emotionally manipulative language in coverage of shelling and the war, after finding that a stock ChatGPT setup flagged almost any mention of an attack as manipulation, a useful reminder

that off-the-shelf models can need real retraining before they're safe to use on a country's most sensitive subject. And IT Logs, a technology outlet, uses AI directly in its news-writing pipeline to raise reporter output. None of these are outliers in the sense of being unrepresentative, they're exactly the "two out of forty" advanced adopters the survey's own numbers point to, just now with names and specifics attached.

The financial picture behind Ukraine's governance gap has its own texture worth adding. Only 42% of surveyed outlets budget anything at all for paid AI tools, and even within that group the sums are small: 12.5% of all outlets spend $5 to $50 a month, 15% spend $50 to $200, and a generous budget, meaning anything above that, describes just three of the forty newsrooms surveyed. Governance splits into two different questions with two different answers: 60% of outlets say they've set some kind of ethical rule for AI use, often just an informal internal limit or a decision to disclose AI use to readers, but only 38% have anything like a policy for labeling AI-generated content specifically, and just 27.5% have a written AI-use policy of any kind. The rules that exist tend to be the easiest ones to agree on in a hallway conversation, not the ones that would require sitting down and writing something binding.

What ties the Ukrainian numbers together, in the end, is a training gap the outlets identify themselves: 85% named training as their top need, well ahead of tool access (60%) or subscription funding (42%), and a clear majority say the training they've had so far didn't go deep enough, one outlet asked bluntly not to be shown "what ChatGPT is" again but something built around their actual editorial workflow. A pattern in the data suggests why the gap matters over time rather than just today: outlets that started experimenting with AI a year or two ago now use it far more broadly and systematically than recent adopters, meaning today's divide between early movers and everyone else, the Tsukr and Rayon.in.ua of the sector, is more likely to widen than close without training aimed specifically at the newsrooms still starting out.

Belarusian media in exile: resilience under pressure

Belarus, or more precisely Belarusian independent media in exile, shows what AI adoption looks like when the newsroom's basic survival is the first priority. **Zerkalo**, reporting from outside the country under constant cyberattack, built an AI proofreader directly into its publishing system, a modest, practical tool, but a real one. In this environment, the technical staff and time that might go toward AI experimentation is just as often consumed defending the newsroom's own servers from being taken offline. Gazzetta’s experience points to a broader risk for Belarusian media in exile: when the open web is blocked, throttled, or degraded, AI models may become a key way for readers to access information. A newsroom MCP server could help keep verified reporting accessible through those systems. For an exile newsroom whose main site is already mirrored, VPN-routed, and periodically knocked offline, making its verified reporting directly queryable by an AI assistant isn't a growth tactic. It's closer to a second, harder-to-block front door for a Belarusian reader trying to reach real reporting about their own country.

Moldova: everyday use ahead of governance

The Independent Journalism Center’s 2025 assessment surveyed 101 media institutions in April-May 2025, alongside 16 interviews and three focus groups. Of all institutions, 67% used AI in professional work, 26% did not, and 7% were unsure or did not answer. Among the 68 AI-using institutions, 12% reported a functioning policy and 88% wanted staff training. Separately, 60% of all respondents cited insufficient funds as a barrier to staff training. The findings describe adoption advancing faster than governance and training capacity. Source: IJC needs assessment, 2025, pp. 25-26 and 31-33

National policy and newsroom practice are developing on separate tracks. Moldova’s AI strategy links public policy to EU alignment and funding. At outlet level, #diez discloses AI assistance with translation, transcription, subtitles, and summaries while keeping final review with journalists. NewsMaker discloses AI-generated images and treats AI-assisted verification cautiously. Observatorul de Nord has experimented with virtual presenters and, in 2026, began developing an editorial assistant with Internews and NarativAI. These are governed experiments, not autonomous publication.

The Independent Journalism Center has since acted directly on its own finding. On December 23, 2025, it published a model editorial policy on AI use that Moldovan newsrooms can adapt rather than write from scratch, built around human editorial responsibility, accuracy, impartiality, public interest, transparency, disclosure, and misinformation prevention, treating AI systems explicitly as tools rather than independent creators of journalistic content. It has paired that with hands-on training: at the tenth annual Mass-Media Forum in Chișinău that same month, roughly 230 journalists, editors, and media managers attended, and Deutsche Welle's AI editorial lead ran a dedicated session demonstrating practical AI tools for daily newsroom work, with transparent use and clear internal rules as the explicit message. That push has an obvious target: on the technical-staff side of the same survey, AI use was the single most requested training topic (50%), ahead of even website modernization, and 44% of respondents said the same for creative staff, more than any other subject. The headline finding, wide informal adoption paired with almost no governance, matches the Ukrainian data closely enough to suggest it isn't a coincidence specific to either country.

Moldova: use runs ahead of governance

IJC / Magenta, 2025 | uptake among all 101 institutions; follow-up questions use the bases below.



67% use AI / 26% do not / 7% unsure or no answer (N=101)

What they use it for



Depth and paid tools

Use: 54% to some extent; 37% to a small extent; 9% to a large extent.

Tools: 43% free only; 49% free and paid; 9% paid only. N=68.

Rules lag behind use

12% have a clear, working policy

28% drafting; 44% none, needed;
15% none, not needed; 1% unsure.
88% want training. N=68.

Why non-users hold off

Selected reasons among non-users (N=26):

No perceived need	35%
Do not know how	23%
No opportunity to try	8%
Distrust	8%
Cost	4%

77% would like to adopt AI eventually.

Source: IJC needs assessment (2025), pp. 31-33. Uptake: N=101. Tasks, depth, paid tools, policy and training: N=68 AI users. Barriers: N=26 non-users. Task choices allow multiple answers; percentages are independently rounded.

AI's role in Moldova isn't limited to something newsrooms do to their own workflows, either. During the 2025 election period, Moldovan media operated inside an increasingly sophisticated information environment involving fake accounts, cloned media portals, and synthetic content, and a March 2026 report from the Legal Resources Center of Moldova documented more sophisticated attacks on independent media that year, including AI used to manipulate the image and identity of journalists in fraudulent contexts. That gives AI a dual role in the country at once: production technology inside the newsroom, and a technology used to attack and impersonate the newsroom from outside it, which is exactly why disclosure, authentication, and verification carry more weight in Moldova than they might in a market facing less pressure on the second front.

Armenia: local-language tools and newsroom governance

Armenia's documented adoption is commercial and incremental, with local-language infrastructure and newsroom governance developing alongside individual tools. The country has no comparable national usage count, so the profile below draws on named newsroom cases, research interviews, and public policy developments.

The public record includes a range of newsroom cases. News.am rewrites and adapts articles with AI; Hetq generates illustrations, works through large datasets, and simplifies complex topics; Infocom has used AI to formulate interview questions and generate a conceptual image; and ARKA's in-house "AI editor" role oversees a system that processes and summarizes incoming wire information, reportedly cutting processing time roughly in half. Ampop's SmartFlow collects selected Armenian news feeds, groups related coverage, and prepares preliminary summaries in the outlet's explanatory style. CivilNet's fact-checking team has used AI to research claims and check image and video authenticity. By February 2026, Media.am described transcription, data analysis, headline suggestions, and preliminary scripting as routine practices in some Armenian newsrooms, with a smaller leading group using automated subtitles and synthetic voiceovers.

A 2024 Freedom House policy brief, based on interviews with 15 Armenian media professionals, provides an earlier qualitative baseline. It identifies limited Armenian-language support and uneven journalist skills as adoption barriers, and records uncertainty about when to disclose AI assistance. The brief notes that audiences are more comfortable with behind-the-scenes uses than with realistic AI-generated photos or video, and recommends clear explanations of AI's role, human verification, and newsroom-specific guidelines developed with journalists. Source: Freedom House policy brief

Some of the language infrastructure behind that shift is homegrown. **Virtual Solutions**, operating as VS.AM, built what it describes as the first free public AI platform for the Armenian language, with proprietary speech models trained specifically on Armenian speech patterns and dialects that the company says reach 96 to 98% recognition accuracy, well above the general multilingual models a newsroom would otherwise default to. Its media-facing tools cover automatic subtitling, text-to-speech for reading out bulletins, and voice cloning that lets an outlet keep a familiar presenter's voice on synthetic narration. The platform's most aggressive product, Armenian AI Journalist, goes further than anything documented inside an actual newsroom: it finds a topic, writes a script, and renders a finished short video for TikTok, Reels, or Shorts in fifteen to thirty seconds, with no human in the loop until after the video already exists. It's marketed at businesses and content creators rather than news organizations specifically, and no Armenian newsroom has been documented publishing its output as journalism, but it's a real, working example of autonomous AI content production sitting immediately adjacent to the news ecosystem, which makes the "assistive, not autonomous" line elsewhere in this chapter a description of what newsrooms have chosen to do with the tools available, not a description of what the tools themselves are capable of.

Governance has followed a visible arc rather than appearing all at once, and it's arguably the most mature of the four countries in this book on that front specifically, even without a usage survey to match it against. Armenia's government approved a concept and action plan against disinformation back in December 2023, but it treated AI only in passing, a brief mention of deepfakes and bots among many disinformation vectors, with no real detail on Armenian-language risks or who should build detection tools for whom. The **Media Ethics Observatory**, Armenia's self-regulatory body uniting more than 80 outlets and organizations, filled part of that gap in 2024 with a shared framework of ethical principles for AI use. **CivilNet** has since gone

further than either, publishing its own dedicated AI Ethics Code as a supplement to Armenia's general media ethics code: it requires transparency about where AI was used, labeling of AI-generated content, rigorous human fact-checking before anything runs, protection against discrimination and of sensitive sources, and periodic review of the rules themselves, stating outright that AI-generated material must be verified by a journalist and that the newsroom stays responsible for its accuracy regardless of what produced the first draft. Regulation caught up to that newsroom-level standard in April 2026, when Armenia's National Assembly adopted amendments requiring broadcasters to label AI-generated content on air, turning what had been a voluntary editorial choice at outlets like CivilNet into a compliance obligation for the whole sector.

That regulatory urgency tracks a threat that has already arrived, not one still theoretical. CivilNet's fact-checking team has repeatedly had to defend its own brand from being weaponized: a 2026 investigation traced a fake video, built with a cloned CivilNet logo and a synthetic voice, fabricating a news broadcast about an abortion ban that never happened, while earlier cases involved a fabricated charity-theft claim circulated against the prime minister's wife and a fake story alleging a HIV diagnosis for the prime minister himself, both dressed up with CivilNet's branding to borrow its credibility. Some of the infrastructure built to catch this kind of thing is now purpose-built rather than improvised case by case: **MediaScoper**, an OSINT monitoring console built for fact-checkers and threat analysts, pulls in traditional news sites alongside hundreds of Telegram channels, runs hybrid keyword-and-semantic search against them, and flags abnormal spikes in coordinated narrative activity close to real time, exactly the kind of tooling that CivilNet's manual detective work implies the wider information space needs at scale. As in Moldova, AI in Armenia is doing double duty: a production tool inside the newsroom, and a weapon aimed at the newsroom's own name from outside it. None of the newsroom-facing examples add up to autonomous AI journalism, either. The available evidence points overwhelmingly toward an assistive model, AI drafts and a journalist edits, verifies, and publishes, and the Armenian-language limitation flagged earlier in this chapter cuts especially deep here: with comparatively little Armenian-language text to train on, generated copy can still read stiff, literal, or culturally flat, which is exactly why CivilNet's insistence on a human edit before publication matters more in Armenian than it would in a bigger language.

What the country evidence shows

Taken together, the country evidence points to different gaps. Ukraine’s survey shows broad use concentrated in routine tasks, alongside limited infrastructure, uneven governance, and a strong demand for training. Moldova also reports widespread day-to-day use but few functioning policies, while newsroom and sector initiatives are beginning to address that gap. Belarusian independent outlets in exile prioritize resilience and access under pressure. Armenia has no representative usage estimate, but documented cases show language-specific tools and governance taking shape. Across all four, adoption is being driven by accessible tools; local capacity, language support, verification, and clear editorial responsibility determine whether that use can be sustained.

The next part moves from this national context to the newsroom systems themselves: how they work in daily practice and what other outlets might adapt.



2

What Newsrooms Actually Built

Four core cases, a wider field of examples, and the lessons inside the workflow

The four main cases in this chapter come from Armenia, Moldova, Ukraine, and Belarusian media in exile, and although they all involve AI, they have very little else in common with one another, since each sits at a different point in the life of a story. **Civic Radar AI** at Moldova's BAStv watches official sources and tries to identify the documents that matter locally. **SmartFlow** at Armenia's Ampop Media collects and groups coverage before preparing a preliminary summary for a journalist. The **18000 AI Content Distribution Platform** in Ukraine turns a finished article into versions for Telegram and Facebook and lets a journalist schedule the one they approve. BGmedia, a Belarusian outlet working from exile, uses a **Digital presenter** to turn reporting into video without a studio or an on-camera journalist.

Taken together, they resemble a newsroom production line, and in a small regional newsroom that line is usually run by very few people. This is a large part of the reason AI is spreading through smaller newsrooms even when they have no strategy, no developer, and very little spare money: the work has grown faster than the staff, and the technology can take over some of the mechanical tasks that surround it.

The strongest projects here did not begin with the question, “How can we use AI?” They began with a newsroom complaint: this takes too long; we keep missing this; we cannot afford the usual way of doing it; our readers need this in a form we do not have the staff to make.

The four main cases are pilots or early systems, described mostly from the newsrooms' own accounts. Passages marked **Elsewhere.** describe what other newsrooms have tried on a similar problem, and they are best read as prompts for ideas, since a project that worked in San Francisco or Bern has not necessarily been tested in Chișinău or Yerevan.

1. Finding the story before it disappears into the official record

Local accountability reporting often begins in a place that seems designed to make it difficult. A municipal authority uploads a document late in the day, a regulator adds a PDF to an archive without any notification, and the information is public in the narrowest technical sense while someone still has to know where to look, come back often enough to notice a change, and judge whether it matters. For a small newsroom this monitoring takes real time, and if nobody does it the first public trace of an important decision can go unnoticed.

CASE 01 / MOLDOVA / BASTV / PILOT

Civic Radar AI at BAStv

THE IDEA

BAStv's Moldovan pilot is a web service that monitors official sources, identifies documents of public importance, estimates their relevance to the region, and prepares a short summary and news draft for a journalist to check. It is meant to help before a reporter has a story, by handing the newsroom a prioritized lead together with the original evidence.

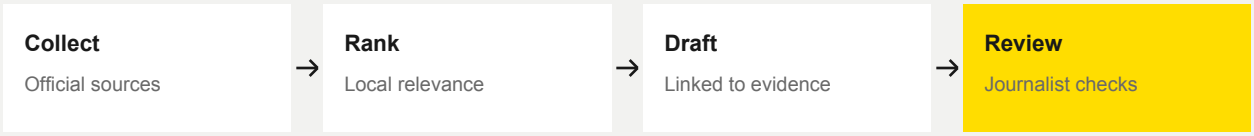
THE PROBLEM

Official information in Moldova is spread across many systems. Some sites offer RSS feeds or sitemaps, others have to be navigated through archive sections or rendered in a browser, and the material may arrive as HTML, a PDF, a DOCX attachment, or an entry in the official gazette, Monitorul Oficial, which needed special handling. Each task is manageable on its own, and the difficulty lies in repeating all of them continuously while telling a routine notice from something that affects the audience.

HOW IT WORKS

The pilot combines a database, an API, and an internal interface. Parsers collect material from official sites, and Google Gemini 2.5 Flash and Pro, accessed through Google Cloud, handle analysis, summaries, and drafts. Cloud credits made those models affordable for the pilot.

CASE 01 / HOW THE WORK MOVES



Workflow summary from the case text. Yellow marks the human editorial decision.

The team's reliability decisions say more about the project than its choice of model. Civic Radar keeps a link to the original source and records the versions of the AI processing so a journalist can see what the system did, and when the software cannot extract a document's text reliably it sends nothing to the model, on the principle that a broken extraction should produce no draft at all rather than a confident draft built on incomplete evidence. The journalist checks the source, decides whether the assessment of importance is right, finds context, calls people, and determines what can be published.

WHAT CHANGED

Testing changed the team's view of what the newsroom needed. A parser that collects everything only replaces one overwhelming list with another, and what BAStv wanted was prioritization: which item matters to the region, which to open first, and what line of reporting might follow.

WHAT IT CHANGES

There is not yet enough before-and-after evidence to claim hours saved or stories gained, so Civic Radar should be described as a pilot. Its value can be measured once it is in regular use, by counting how many suggestions journalists reject, how many become published work, and how often it finds something the newsroom would otherwise have missed. One lesson is already visible, which is that for local media relevance has to be built into the technical work from the start. A system that summarizes a ministry document perfectly but cannot say why one paragraph matters in a particular district has automated the reading and left the journalism undone.

ELSEWHERE / The closest counterpart is the *San Francisco Chronicle*'s public-meeting monitor, in which more than 200 scrapers check government platforms across nine Bay Area counties for new recordings. The audio is transcribed on a self-hosted copy of OpenAI's open-source Whisper model and summarized with a commercial model, reporters set keyword alerts, and the newsroom reports dozens of stories it would otherwise have missed. Its own summary of the division of labor is that the system tells a reporter where to listen and the reporter makes the calls. Two features would carry over to a small newsroom: the most repetitive step runs cheaply on hardware the newsroom controls, and reporters define what counts as relevant from the first day. Chalkbeat, after a pilot of its own, moved to a third-party platform called LocalLens to follow about 80 school districts, treating the summaries as tips and checking every quote before publication.

2. Turning a flood of coverage into the beginning of journalism

Official sources are only one of the streams a newsroom watches. Ten outlets may report the same announcement with slightly different facts. Ampop Media, an Armenian outlet built around explanatory and data journalism, is developing SmartFlow to cut the time spent moving through that stream.

CASE 02 / ARMENIA / AMPOP MEDIA / EARLY SYSTEM

SmartFlow at Ampop Media

THE IDEA

SmartFlow is an internal pipeline and dashboard that collects articles from selected Armenian media feeds, groups related coverage into topic clusters, and produces preliminary summaries in Ampop's neutral style. A journalist reviews the result, checks the underlying material, and decides whether there is an Ampop story to write. The tool is not meant to fill the site with rewritten news, since Ampop's value comes from explanation, and its purpose is to shorten the path to the point where that work can begin.

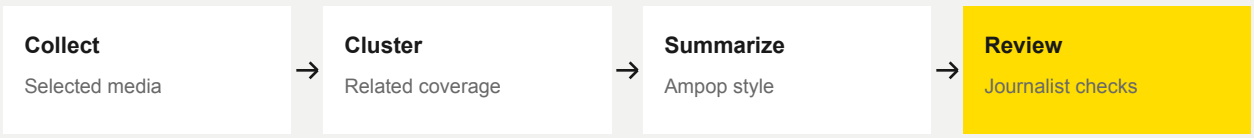
THE PROBLEM

Monitoring several local outlets takes a surprising share of a reporter's day, and most of it is mechanical. SmartFlow automates collection, clustering, and the first draft, and leaves the judgment about meaning with the journalist, because the system may not know that two organizations with similar names are politically different, or that a small change matters because of what happened six months earlier.

HOW IT WORKS

The newsroom's description names Python scrapers for collection, Go services for backend logic and data operations, and the OpenAI API with prompts adapted for Armenian-language processing, tone moderation, and summarization. The stack matters less than the way responsibilities are divided: scrapers retrieve material, backend services store and organize it, the model works on text that has already been selected, and the dashboard gives the journalist a place to inspect the result.

CASE 02 / HOW THE WORK MOVES



Workflow summary from the case text. Yellow marks the human editorial decision.

WHAT CHANGED

Ampop originally considered adding government portals and regulatory databases to the media sites. Early testing showed how costly that would be, because Armenian official sites follow no consistent structure and each one could need its own extraction logic and continuing maintenance. The team narrowed the MVP to digital newsrooms and concentrated on stabilizing the pipeline and making the dashboard clear enough for daily use, which reduced the ambition and improved the chances that the product would actually work. A first useful version usually does one job over a bounded set of material, whereas a broad demonstration of everything tends to end up trusted by nobody.

WHAT IT CHANGES

SmartFlow was scheduled to enter live editorial testing in September 2026. Public material reviewed for this edition does not yet provide results from that phase. The test should answer questions no benchmark can: how often the clusters are correct, how much correction Armenian copy needs, and whether the tool helps Ampop produce better explanatory work or only helps it reach the same stories sooner. The newsroom has already drawn its main lesson: local context decides the outcome. Armenian morphology and the habits of the country's media shape both the prompting and the parsing. A summary can be grammatically plausible while flattening the politics or sounding unlike the outlet, which is why human review in SmartFlow is far more than a ceremonial final click.

ELSEWHERE / VERDAD's two-stage design applies directly to a pipeline like this one: a cheap step can decide which articles belong together, and the expensive model can be reserved for the summary a journalist is going to open. Tamedia in Switzerland shows where a trusted collection layer can lead. Its newsletter tool scrapes municipal websites, tags the material, and pulls in official listings, so that an editor can select and check a community newsletter within minutes, and the company aims to produce over 100 hyperlocal articles a week with municipalities sponsoring the newsletters. Ampop's audience is national, which makes the product itself a poor model, but the underlying idea, that a collection and clustering layer can support more than one editorial product, carries over well.

Small-language infrastructure changes what a newsroom can build

THE IDEA

Global AI services often support major languages better than smaller ones. In Armenia, VS.AM addresses that gap with Armenian speech-to-text, text-to-speech, subtitles, and OCR. The company reports 96-98% speech-recognition accuracy on good audio; this is its own claim, and newsroom use still requires checking transcripts and names. This language infrastructure is distinct from Ampop’s SmartFlow, which monitors Armenian media sources and prepares summaries for editorial review. Sources: VS.AM

VS.AM is an attempt to build that missing layer locally. It offers Armenian speech-to-text, text-to-speech, subtitles, OCR, article-to-audio conversion, voice cloning, and custom voice training, and says its speech recognition reaches 96% to 98% accuracy with good audio. Those figures are the company's own claims, but the range of products shows what becomes possible once a small language has its own technical foundation. A newsroom can transcribe an interview, add audio to an article, search a scanned document,

or train a voice that pronounces local names correctly, and audio in particular reaches visually impaired readers and diaspora audiences.

ELSEWHERE / OCR deserves more attention than it usually gets. The *South China Morning Post* digitized a century of microfilm into structured text for internal search, syndication, and external licensing, and Tempo in Indonesia is bringing together more than 50 years of material as the base for an archive assistant. Both began by getting old material into a form a machine can read, which is the stage an outside company is least likely to build for Armenian or Belarusian.

Ukrainian newsrooms show a related pattern. Texty.org.ua had to train its own system to detect manipulative language in war reporting after a general ChatGPT setup treated almost any account of an attack as manipulation, and Tsukr built more than ten task-specific Claude projects for text development, SEO, social introductions, and donor reporting. Models are becoming increasingly interchangeable, whereas local language, newsroom memory, trusted sources, and a workflow people will actually use cannot be swapped so easily.

3. Distribution is editorial work too

A finished article may need a Telegram post, a Facebook version, an Instagram caption, a newsletter introduction, a video script, and metadata, and each platform rewards a different kind of opening. This work is often described as promotion, as though it happened outside journalism, yet in a small newsroom journalists do it themselves and it influences which facts reach people, since a technically correct post can still fail if it buries the reason a reader should care.

CASE 03 / UKRAINE / 18000 / CMS WORKFLOW

The 18000 AI Content Distribution Platform

THE IDEA

18000, run by the Cherkasy Information Agency 18000 in Ukraine, integrated its platform into the outlet's content management system. It analyzes a published article and proposes several versions of a post for Telegram and Facebook, and a journalist reviews them, edits one if needed, and approves it for immediate or scheduled publication. The platform cannot publish without human confirmation.

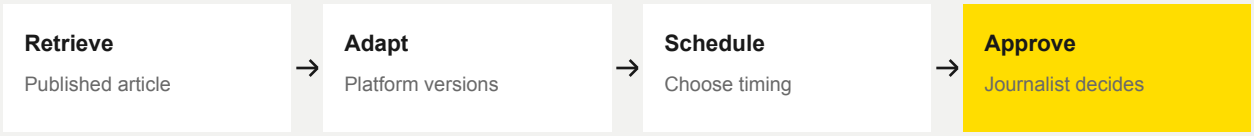
THE PROBLEM

Local outlets have small staffs but face distribution demands similar to those of much larger organizations. Copying the same headline and link to every platform is quick and ineffective, while writing a separate post for each network takes time that could go into reporting. The platform draws on the outlet's editorial guidance and examples, so the suggestions sound like 18000.

HOW IT WORKS

The newsroom built a custom WordPress plugin in PHP, which puts the tool inside the system journalists already use. Its description names Claude Haiku 4.5 and Sonnet 5 for article analysis and post generation, with connections to Telegram and Facebook for scheduling. Model versions will change, but the system sits where the work happens. A separate chatbot would ask the journalist to copy the article, find a prompt, move the result elsewhere, and schedule it, and those small steps together can cancel out the time saved.

CASE 03 / HOW THE WORK MOVES



Workflow summary from the case text. Yellow marks the human editorial decision.

WHAT CHANGED

The first version generated posts. During development the team saw that suggestions alone left much of the task untouched, because journalists also had to decide when each post should appear, and adding scheduling turned a text generator into a basic distribution platform. The model handles the visible AI task, while most of the value comes from what surrounds it: retrieving the right article, applying the right instruction, recording approval, and passing the decision to the next system.

WHAT IT CHANGES

The decision to publish stays with a person, who remains responsible for accuracy, wording, and whether a post goes out at all. A factually accurate statement can still make a weak post, so the system has to be taught the type of story, the outlet's voice, and the difference between Telegram and Facebook through examples and repeated feedback. The immediate gain is time, which adds up over a month when several stories a day go to several platforms, and the larger gain is consistency, since reporting that has already cost money to produce has a better chance of reaching the people it was meant for.

ELSEWHERE / A Slovak media group built a similar tool into its editor. Its post generator offers three suggestions after the journalist chooses the length and the kind of interaction the post should invite, whether discussion, reading, or sharing, which amounts to asking what the post is for before writing it, the same thing the 18000 team found the system needed to learn. A smaller case shows how little is needed to begin: two Northwestern students built a Social Card Generator with Claude for the *Ocala Gazette*, in which the AI designed the layout and writes none of the content. For a newsroom not yet ready to let a model near its words, this marks a useful boundary, because AI can improve the packaging of a story without touching what it says.

One piece of reporting can become several audience products

Rayon.in.ua, a Ukrainian network of nine thematic and 52 district sites, uses AI to turn recorded meetings and original video into text, which gives its sites more original material for less mechanical work. MykVisti has built CMS buttons for SEO descriptions, URL slugs, translations, and headline options, which together remove a layer of repeated formatting from every story. The Philadelphia Inquirer used AI assistance to launch newsletters in four nearby areas it had not covered heavily and planned eight more with two additional staff, which shows software changing the economics of coverage enough to make it plausible without replacing the people responsible for it.

The Kyiv Independent uses AI for audio versions of articles and a chatbot that helps readers buy memberships, which reflects the idea running through this section, that the job continues after the article is written.

ELSEWHERE / Tamedia's newsletters follow the same logic, with an editor checking the material, production taking minutes, and communities too small for a dedicated reporter receiving a reliable local digest, and a district network like Rayon.in.ua could ask which parts of that model would work for it.

4. Making video when conventional production is impossible

The usual argument for synthetic presenters is efficiency, since an avatar needs no studio, camera setup, lighting, or presenter available at the right hour. BGmedia's case includes all of that, but its most serious constraint has little to do with money. Belarusian independent media work in exile, reporting for people inside a country where journalism can bring punishment to relatives who remain there as well as to the reporter, and appearing on camera can carry a risk that no production budget explains.

CASE 04 / BELARUS IN EXILE / BGMEDIA / VIDEO

BG Avatar at BGmedia

THE IDEA

BGmedia turns journalistic texts into short presenter-led videos using a custom digital avatar modeled on editor-in-chief Victor Marchuk, a familiar public figure for the outlet's audience. The newsroom assembles generated presenter segments in ordinary editing software into thematic videos of up to ten minutes. The familiar face may help viewers recognize the format, although the available material does not measure its effect on reach or popularity.

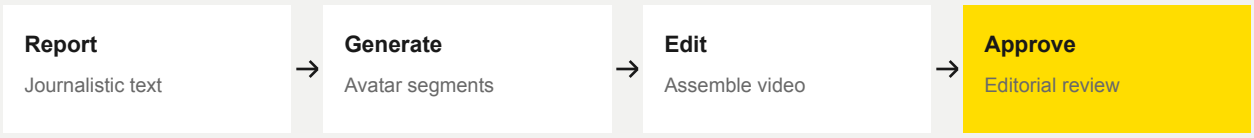
THE PROBLEM

The outlet is a small team without a studio, specialized equipment, or a full-time camera operator, while video platforms remain an important route to Belarusian audiences, and some of its journalists need to stay off camera to protect family in Belarus. A digital presenter lowers the threshold for producing video and separates the public face of a report from the identity of the journalist behind it.

HOW IT WORKS

The newsroom uses HeyGen to generate the avatar footage and video-editing software to assemble finished reports. The stack is deliberately simple, because BGmedia built neither a model nor a platform and put together a repeatable process around a commercial service. Implementing AI does not have to involve custom software, and a newsroom has done it when it has redesigned a real workflow, decided who is responsible, established what gets checked, and keeps using the result.

CASE 04 / HOW THE WORK MOVES



Workflow summary from the case text. Yellow marks the human editorial decision.

AUDIENCE RESPONSE

BGmedia reports generally positive reactions, with only a small number of comments over roughly twenty months pointing out that the presenter was generated. This is useful evidence, though viewers who fail to notice a synthetic presenter have not necessarily accepted one. The public account does not say how the avatar is labeled, whose likeness and voice it uses, or what the audience has been told, and since audiences react differently to news they know is AI-generated, a newsroom that hides the synthetic nature of its presenter is betting that the difference will never be tested.

WHAT IT CHANGES

The newsroom can produce presenter-led video without a studio, make more formats with the same team, and keep a journalist's face out of the finished report when visibility would be dangerous. Production still costs editorial labor, since a journalist has to report and write the text and someone has to check the delivery, edit the segments, and approve the result.

ELSEWHERE / Synthetic presenters have been given very different jobs. Grupo Fórmula in Mexico launched NAT, which began with good-news stories and moved toward entertainment to build a following, and later introduced SOFI for political news. NAT passed 13,000 Instagram followers, and SOFI's videos have reached hundreds of thousands of views and sometimes more than a million. Sharjah Media City used an AI presenter under human supervision as an emergency substitute when an anchor arrived late. Of the three purposes, only personal safety is a reason a newsroom cannot simply put a journalist on camera, and it is also the case that most needs a clear disclosure policy.

The same format can produce a different response

Observatorul de Nord in Soroca, Moldova, had a more mixed experience. Its head of video said viewers could tell the presenter was virtual, yet the tool proved useful during elections, when information arrived late and had to become video quickly. An avatar clearly has no single value for every newsroom, offering security and basic capacity in one place and speed for a narrow class of updates in another, so the real question is which problem it solves well enough to justify losing a human presence. The outlet has since written its own AI guidance.

The opportunity comes with a danger. CivilNet in Armenia has had its branding and the look of a news broadcast copied in fabricated videos, and the Kyiv Independent has investigated deepfakes aimed at

narratives about Ukrainian mobilization, so the same technology that lets an exile newsroom protect a journalist also lets an attacker borrow the identity of a trusted outlet. Disclosure, authenticated source material, and content-provenance standards such as C2PA, adopted by Ukraine's public broadcaster Suspilne, belong to the production process itself.

5. The newsroom behind the published work

Some of the most valuable AI work never reaches the audience, because it happens in meetings, archives, analytics, grant reports, and moderation queues. **Rayon.in.ua** uses Fireflies to record meetings and strategy sessions and produce transcripts and protocols, which for a network spread across several dozen cities works as institutional memory, so that a decision no longer lives only in the notebook of whoever attended the call. **Tsukr** uses Claude for team management, with bots that analyze the performance of social posts and an internal dashboard built through AI-assisted coding, helping a small team see patterns in its own work.

Daily Maverick in South Africa built ImpactEngine, which searches parliamentary records to find where the outlet's reporting may have influenced political discussion. The team found that choosing a language model was the easy part, and that the hard work lay in ingesting the material cleanly, dividing it into sections, and attaching enough metadata for the system to retrieve the right passage. Models look impressive in a demonstration, and it is the quality of the data preparation that decides whether a newsroom can trust the result on an ordinary Tuesday afternoon. Such a system will also find patterns that mean nothing, and a journalist still has to decide what a possible relationship means.

ELSEWHERE / Tempo's archive assistant is the most direct example of newsroom memory as a product. It combines a digital asset management system holding 50 years of photos, articles, video, and audio with a retrieval-based chatbot that answers only from the verified archive, links back to the originals, and is designed to decline when it has no source. Its rule of answering only from what has been verified is the one Civic Radar follows when it refuses to draft from a document it cannot read. An archive organized this way sits one step away from being queryable by a reader's own AI assistant through a standard such as MCP, as the public MCP server built over the Guardian's archive shows, and the groundwork is the same cleaning Daily Maverick described.

Moderation belongs in this group as well. **ModerAItor**, an Armenian project, analyzes context and intent in Armenian, Russian, and Georgian instead of relying on banned-word lists, which matters because irony, quotation, abuse, and reporting about abuse can all contain the same words. The *New York Times* uses Google's Perspective tool to score comments for toxicity so that moderators can prioritize their queue, and says this allowed it to triple the number of articles open for comments. Scoring systems built for English perform unevenly in Armenian, Belarusian, and Romanian, so a newsroom should test one against its own comment history before trusting the scores.

What these internal tools have in common is that they generate almost no content and instead help a newsroom remember, compare, prioritize, and notice.

6. Where the money is

AI has produced clearer cost savings than new revenue. In the database checked on 2 November 2026, monetization accounts for 81 of 1,359 documented initiatives, or 6.0%; only distribution is smaller, with 69. The cases below are therefore closer to careful bets than to proven models. The most convincing of them do not sell AI-generated content and combine the newsroom's data, reporting, or audience knowledge with a service that helps a reader do something.

Saving money is the first business model. All four central cases rest on a cost argument, though a claim that a tool saves time means little unless the newsroom has recorded the old process, the new one, and the work that remains, since a bad draft that takes fifteen minutes to repair may save almost nothing. A saving also raises the question of what the recovered time is used for, whether more original reporting, a newsletter, or a product to sell, or simply a workload that was already too heavy.

Converting more readers. Daily Maverick's Revenue 360 project started from a membership problem: readers showed interest but left before completing the signup, and conversion stood at 3.5% against a paywall benchmark of 11% cited by the team. Heatmaps and audience behavior informed a redesigned landing experience with messages tailored to different groups, and the first test produced a reported 55% increase in conversion. AI was only one part of that result, alongside design, tracking, and testing, which fits a general pattern in which revenue follows from a system that connects audience evidence to better decisions and rarely from adding a chatbot. Tamedia's newsletters add a second logic by seeking income from sponsoring municipalities as well as readers, which suits communities too small for a subscription business.

Selling a useful service. TX Tax, from Hearst's Texas newspapers, lets homeowners find comparable properties in official county data, organize evidence, and generate a personalized summary and talking points for a property-tax protest, and journalists and developers worked with appraisal districts to understand the records. The product costs $69 before discounts, and Hearst described the first Houston version as a major subscription driver before expanding it to Austin, Dallas, and San Antonio. Readers pay for the ability to act, which a model-written article could not give them, and what protects the product is the local records, the newsroom's grasp of the process, and the trust that it exists to help. The **Armenian corruption-data portal** points toward similar services built on declarations, procurement, budgets, property, or court records, some of which should stay free while others could support membership or a paid report. Archives are the other asset newsrooms rarely sell, and the **SCMP** and **Tempo** both name licensing as a goal, though neither has published results.

A subscription conversation for every reader. Observador in Portugal is testing a Subscription Concierge that contacts readers individually by SMS or WhatsApp, drawing on their plan, billing status, and reading habits. It has identified roughly 30,000 active subscribers who might upgrade and a similar number of former subscribers it might win back, more conversations than a human sales team could hold, and it began with a supervised test of 50 to 200 subscribers compared against its telemarketing operation. The conversation is the visible part, and most of the product lies in identity, permission, payment, data protection, measurement, and the ability of a person to step in.

Proving that journalism had an effect. ImpactEngine also serves revenue by showing members what their money supports, for example when an investigation is discussed in parliament or gives another journalist a starting point. Journalism can contribute to a change without causing it alone, so the record

needs checking, but a checked record of impact can strengthen grant reporting and membership communication.

Taken together, these examples suggest that producing more generic content more cheaply is unlikely to give anyone a lasting advantage, and that the clearer path to money runs through trusted local data, a difficult public process, a direct audience relationship, or a body of reporting the newsroom uniquely owns.

7. What changes when a newsroom builds instead of experimenting

In all four central cases the first idea changed along the way. Ampop dropped difficult official sources to stabilize its pipeline, BAStv found that ranking documents by local significance mattered more than collecting them, 18000 added scheduling because generated posts solved only half of the problem, and BGmedia joined a ready-made avatar service to ordinary editing where a custom build would have added nothing. Several lessons are common to all of them.

Start with a repeated bottleneck. Small improvements add up only on tasks that come around often.

Put the tool where the work occurs. A function inside WordPress will be used more consistently than a prompt saved beside six browser tabs.

Treat local context as a system requirement. Armenian morphology, Moldovan regional relevance, the difference between Telegram and Facebook, and the security conditions of exile journalism decide whether a product is useful at all.

Name the human decision. Someone has to open the original document, check the translation, approve the scheduled post, and be able to stop the avatar video, and a reliable workflow gives each of those decisions an owner. Surveys in Ukraine and Moldova suggest how rarely this is written down.

Choose the evidence before declaring success. A discovery tool should count useful leads and missed sources, a distribution tool should track time, reach, and errors, and a subscription product should track conversion, cancellations, and cost per transaction. Without a baseline, enthusiasm ends up as the only metric.

Decide what to build and what to borrow. Chalkbeat piloted its own tool and then adopted a mature one once it knew what it needed, while Ampop and BAStv built their own pipelines because nothing on the market understood Armenian media or Moldovan official sources. Where the hard part is generic, buying makes sense, and where it is local, the newsroom's own knowledge becomes the product.

Match the model to the job. VERDAD's cheap-then-strong design and the *Chronicle*'s self-hosted transcription both reserve the expensive tier for the places where a mistake costs the most, and most newsroom tasks do not need the newest model.

Expect maintenance. Official sites, APIs, and models change and staff leave, so keeping a tool reliable can cost more than building it, and someone should own that job before a pilot becomes infrastructure.

Protect the assets that grow more valuable with use. A model can be replaced, whereas the verified archive, source lists, pronunciation dictionary, local-language examples, corrections, and audience history cannot. Own the memory, not the model.

Ideas worth borrowing

- **Civic Radar (BAStv):** let reporters set their own keyword alerts, as the *Chronicle* does, and run transcription on an open model the newsroom controls.
- **Corruption-data portal (Armenia):** add fuzzy name matching across declarations, following the *Financial Times* pipeline.
- **SmartFlow (Ampop):** cluster with a cheap model, summarize with a stronger one, and use the collection layer as the base for more than one product.
- **18000 platform:** ask the journalist what a post is for before generating it.
- **BG Avatar (BGmedia):** publish a plain statement of when and why a synthetic presenter is used.
- **Archives (all four countries):** start with OCR and clean metadata, then build a verified-answers-only assistant like Tempo's, and consider an MCP server after that.

Sources for the “Elsewhere” examples

These examples come from the AI For Newsroom database, a search sample and not a complete census.

- San Francisco Chronicle, Public Meeting Monitoring System: awards.journalists.org
- Chalkbeat, LocalLens: Nieman Lab
- Financial Times, Story Finding Pipeline: The Audiencers
- Newtral, FactFlow AI: Newtral
- Tamedia, hyperlocal newsletters: Innovate Local
- Grupo Fórmula, NAT and SOFI: Reuters Institute
- Sharjah Media City: City, University of London study
- Ocala Gazette, Social Card Generator: Generative AI in the Newsroom
- Tempo, Tempo Assistant: Tempo.co
- South China Morning Post, Vintage archive: SCMP Inside the Wonton
- The New York Times, Perspective moderation: Google blog
- Guardian MCP Server, VERDAD, Slovak editor tools, and Civio: AI For Newsroom initiatives database

Postscript. The punk newsroom

Punk did not begin because the music industry discovered an underserved market. In the mid-1970s, bands, designers, club promoters, fanzine makers, and audiences assembled a scene from what they could reach. Malcolm McLaren and Vivienne Westwood, working from their shop at 430 King's Road, were part of that unruly network, not its sole authors. The useful lesson is not the mythology around any one figure. It is the method: learn enough to make something, put it in front of people, revise it in public, and build the missing infrastructure while you go.

AI gives newsrooms a similar opening. It does not supply the purpose or the courage, and it does not remove the work of security, maintenance, or editorial judgment. What it changes is the price of a first attempt. A journalist who could never commission a monitoring platform can prototype one. A local newsroom can make a transcription bot that understands its language, connect a public database to a usable search page, or build a small distribution tool around its own publishing routine.

That does not mean every newsroom should recreate every service it buys. Payroll, payment processing, security, and mature collaboration systems are hard for good reasons. But the software-as-a-service era taught organizations to adapt their work to somebody else's product. AI-assisted development makes the opposite question practical again: what is the smallest tool that fits our exact need, our data, our language, and our editorial rules? In many cases the answer will still be "buy." In more cases than before, it can be "build," "adapt," or "share with three neighboring newsrooms."

This is also a route back to technical and data independence. The durable asset is not the model subscription. It is the cleaned archive, the source map, the pronunciation list, the corrections history, the audience relationship, and the knowledge of how the newsroom actually works. Keep those assets portable and local where the risk demands it. Make vendors replaceable. Treat the model as a component, not the institution's memory.

The first punk records were rough, but roughness was not their point. They proved that the gate was not locked. The equivalent newsroom project may begin with the spreadsheet rebuilt every month, the archive nobody can search, the audio nobody has time to transcribe, the membership question nobody can answer quickly, or the public notice everyone discovers too late. Look around the newsroom. Find the drag, the blind spot, or the unrealized idea. Build the smallest useful response. Measure it. Improve it. Then decide what else has become possible.



3

The Choices Ahead

Infrastructure, independence, and a practical route from experiment to capability

Part 1 was a snapshot. This chapter looks one year ahead: what is already settled in AI and journalism, what remains open, and what newsrooms in Armenia, Belarus, Moldova, and Ukraine can do now.

1. What's already settled, and what's still up for grabs

Some things are no longer in question. Drafting, transcription, translation, and headline support - with a person reviewing the output before publication - have moved from novelty to routine practice. Reuters Institute's *Journalism, Media, and Technology Trends and Predictions 2026* treats this as background rather than news, while industry surveys regularly find AI in a clear majority of newsrooms. The AI For Newsroom collection, checked on 2 November 2026, tells a more precise story about documented projects: of 1,359 initiatives, 538 concern content production and 385 concern journalism tools. Revenue remains much less mature. Monetization represents 81 initiatives, or 6.0%, and many of those are licensing agreements or improvements to an existing subscription business rather than a new source of income created by AI alone.

What's still genuinely open is almost everything downstream of that. How much zero-click search actually costs publishers long-term. Whether GEO and AEO become a permanent discipline or a passing scramble. Whether agentic AI, AI systems that plan and execute multi-step tasks rather than just answering a single prompt, turns out to be a real capability shift for reporting or another round of task-focused tools that quietly underdeliver. Reuters Institute's separate piece gathering forecasts from 17 experts across the BBC, the Wall Street Journal, Süddeutsche Zeitung, Semafor, and other newsrooms captures this tension directly: the same report that describes task-focused AI tools as a "strategic dead-end" in their early form also predicts that reasoning-capable agents will start handling real newsgathering, investigation, and fact-checking work within the year. Both things can be true. The technology moved past the disappointing phase; what newsrooms do with the more capable version is still being figured out in real time.

2. Six forces shaping the next few years

Cost keeps falling, and open weights keep driving it. The economics described in Part 1, roughly a 50-times drop in the cost of GPT-4-class output over three years, mostly driven by open-weight releases, isn't finished. Every newsroom in this book, particularly the smaller and donor-funded ones, gets more room to work with each time that floor drops again.

Distribution keeps moving from search results toward AI agents themselves. FT Strategies describes a shift toward "context engines": systems that serve a newsroom's content to its own products while managing outside AI agents that arrive to retrieve it. By September 2026, Reuters had connected its MCP server to CuttingRoom's browser-based video editor, allowing licensed footage to be found, edited, captioned, and reformatted through plain-language requests. The Associated Press and Reuters were also leading work on a shared language for context to move between tools across the production chain. This is a different job from conventional SEO, and almost none of the newsrooms in this book has built the capacity for it yet.

GEO and AEO are becoming a real discipline, unevenly. Part 1 already covered the mechanics. The forward-looking piece of it is that this expertise is concentrated in large, well-resourced publishers first,

exactly the pattern that leaves small-language and donor-funded newsrooms furthest behind unless something changes deliberately.

Agentic AI is arriving for real use, not just for demos. Column's white paper frames 2026 as the year agentic systems move from simple automation into autonomous, multi-step action inside newsrooms, and lists governance, ethics, and public trust as the open questions that come with it, not capability. Reuters Institute's forecasting piece adds a specific prediction worth sitting with: 2026 as the year newsrooms start investing in the infrastructure and training to use these tools strategically, rather than trying them once and shelving them.

Regulation is converging faster than any single newsroom can track, and not only in Europe. Moldova's national AI strategy, tied explicitly to eventual EU AI Act compliance, was covered in Part 1. It isn't an isolated case. In Brazil, the digital rights group Foxglove has pushed regulators to investigate Google's AI systems over the use of news content, and in the United States, the trade group IAB has drafted legislation specifically targeting AI scraping. None of these efforts are coordinated with each other. All of them are responding to the same underlying problem: AI systems using journalism as raw material without a clear, enforceable agreement about compensation or credit.

A trust and verification layer is being built in parallel with all of the above, and it's arriving as policy and labeling rather than as a single silver-bullet technology. A UK think tank has proposed that AI-generated news should carry something like a nutrition label. Süddeutsche Zeitung has published its own internal framework for labeling AI involvement in a story, called *Designing Clarity*. The Local Media Association and Trusting News ran a fresh reader survey in the US and packaged the underlying questions into a reusable "AI Trust Kit" any newsroom can run with its own audience, concluding that trust gets built through disclosure, visible human judgment, and ongoing conversation with readers rather than through a single policy document. The most concrete version of this trend inside the four countries this book covers is Suspilne, Ukraine's public broadcaster, which has adopted C2PA content-authenticity verification technology, the same content-provenance standard being built into cameras and editing software worldwide, specifically so that its own footage can be traced and verified in an information environment already saturated with fabricated war content.

3. Open source as the resource-constrained newsroom's way out

The strongest case for open-weight AI in a resource-constrained newsroom is control. A hosted chatbot sends material to a provider’s systems under that provider’s contract and retention rules. A model that runs on infrastructure the newsroom controls can keep unpublished investigations, source material, and internal archives inside a chosen security boundary. That difference matters in any newsroom and becomes decisive where surveillance, legal pressure, or attacks on infrastructure are routine.

Self-hosting is not free. The model may carry no license fee, but hardware, electricity, updates, access controls, monitoring, and staff time still cost money. Nor is an open model automatically safer: a badly maintained local system can expose sensitive material just as surely as a careless cloud workflow. The practical advantage is choice. A newsroom can use a commercial service for low-risk work, a locally run

model for sensitive or high-volume tasks, and switch components without rebuilding the archive or workflow around them.

This is why chasing the newest frontier model is usually the wrong objective. Transcription, translation, tagging, document classification, and first-pass drafting rarely require the most capable or expensive system available. The harder work is organizational: deciding which material may enter which tool, who checks the output, what happens when the tool fails, and who maintains the workflow six months after the grant or pilot ends.

For smaller outlets, the most credible route may be shared infrastructure rather than solitary self-hosting. One technically capable team can maintain a secure deployment, speech model, archive index, or monitoring service for several newsrooms. Costs and specialist work are pooled; each newsroom keeps authority over its sources and publication decisions. The model layer can change without taking the institutional memory with it.

Small-language markets make the logic especially clear. Armenian speech and text tools exist because local builders treated the language as infrastructure rather than waiting for a global vendor to make it a priority. The same opportunity exists for Belarusian and Romanian-language Moldovan media. The durable asset is not the model itself. It is the newsroom’s verified archive, language resources, corrections, sources, and relationship with readers. Own the memory, not the model. That idea runs through the rest of this chapter, especially the MCP section below.

4. Four scenarios

Infrastructure normalization. AI settles into being routine, unremarkable plumbing: transcription, tagging, drafting-with-review, the way spellcheck or a CMS became invisible after a few years. Roles stabilize around something like ARKA's in-house "AI editor" position, adopted at outlets from Clarín in Argentina to newsrooms far outside this book's four countries. This is the least dramatic scenario and, on the evidence gathered here, currently the most common one.

Platform capture. Zero-click search and slow, uneven progress on GEO, AEO, and licensing leave small-language, donor-funded, or exile outlets effectively invisible to the AI systems increasingly mediating how readers find news. Part 1's finding that News was the one major web category to shrink while every other category grew during the same period AI usage exploded is the leading indicator for this scenario, and a whitepaper from the Center for Journalism and Liberty on privately controlled data marketplaces lays out the deeper structural risk: publishers who don't build their own negotiating position end up inside a marketplace someone else designed, on someone else's terms. This is the scenario most outlets in this book are currently most exposed to.

Sovereign and collective resilience. Regional newsrooms and their funders respond to that exposure by pooling resources: shared small-language infrastructure building on the Podcastle precedent, a shared, professionally-run open-weight deployment, a shared MCP or licensing negotiating body. This is the scenario the recommendations in Part 1, and the open-source argument above, are explicitly trying to make more likely, not a prediction so much as a target.

Weaponized information environment. Disinformation volume keeps outpacing verification capacity even as AI-assisted fact-checking tools improve at the same time, and Ukraine and Belarus in particular are

pushed into a permanent, exhausting verification footing. Suspilne's adoption of C2PA is an early countermeasure already being built inside this exact scenario, not a hypothetical response to it. Worth holding onto here: several serious voices in this space, including the media critic Matt Pearce and the researcher Felix Simon in separate recent essays, have pushed back on the idea that fixing AI's own accuracy and trust problems is journalism's job to solve. Newsrooms can build defenses. They didn't create the flood, and treating verification as now a core part of the job is different from treating it as journalism's debt to repay.

5. The MCP bet

Getting cited inside an AI answer, the GEO and AEO work covered in Part 1, means competing for a mention inside somebody else's product. MCP, the Model Context Protocol, is a different kind of move: it lets a reader's AI assistant query a newsroom's own verified archive directly, on the newsroom's terms, rather than relying on whatever the AI happened to learn in training or find through a general web search.

The FT Strategies piece on AI engines describes where this is heading next: context engines that don't just serve a newsroom's own content to its own products, but manage a growing traffic of outside AI agents arriving to retrieve that content, and increasingly, verify that those agents are paying for what they take rather than simply scraping it. The *Newsroom Robots* conversation frames the strategic stakes in blunter terms: when an AI agent can hand a reader a full briefing without a single visit to a news site, competing with that agent for the click is a losing game. Building for it, so the agent's briefing is drawing on a specific newsroom's verified reporting rather than generic training data, is the more promising position. Whoever does that well, as that conversation puts it, ends up owning the experience, not just the underlying content.

Real, working examples remain thin but are no longer hypothetical. Three realistic paths exist for a donor-funded outlet in this book's four countries. One newsroom, well-positioned technically, could pilot an MCP server alone; the Kyiv Independent, given its existing use of AI tools for audience research, is the strongest candidate. A donor or press-freedom organization could fund one shared MCP service used by several outlets at once, spreading the cost and the security burden the way the open-source section above argues for. Or the region could simply wait and adopt whatever standard the largest global publishers settle on eighteen to twenty-four months from now, arriving late with no negotiating leverage of its own. The Cloudflare-commissioned whitepaper from the Center for Journalism and Liberty makes the case for urgency directly: the shape of these data marketplaces is being decided now, while publishers, especially small ones, still have some say in it. The shared regional path is the one worth actually building toward.

6. Country trajectories

Armenia

Current direction. Armenia's adoption is likely to deepen through Armenian-language infrastructure, but the scale of newsroom use remains unmeasured: there is no representative survey. Documented cases are task-specific. ARKA's in-house "AI editor" processes and summarizes incoming wire material; Podcastle offers Armenian text-to-speech, while VS.AM provides speech recognition, subtitles, and related language tools. These examples show local capacity taking shape, not widespread deployment. Governance is moving ahead of measurement: CivilNet has a newsroom AI code requiring disclosure, human verification, privacy safeguards, and editorial accountability. Amendments adopted in April 2026 require audiovisual broadcasters to label AI-generated synthetic content on air; the requirement does not cover every online article. Sources: Media.am newsroom analysis · Podcastle · VS.AM · CivilNet AI Ethics Code · Armenian government announcement on labeling law

What could shift it. The most useful next investment is shared Armenian-language capacity for text, translation, search, and transcription, extending beyond audio. Common datasets and evaluation checks for dialects, names, and factual accuracy would let small outlets assess tools before they depend on them. Grants should support newsroom pilots, training, and maintenance, with requirements for clear disclosure and safe handling of source material. Without affordable, interoperable options, larger organizations may shape the market while smaller outlets remain dependent on vendors and language systems they cannot adapt or audit.

Belarusian media in exile

Current direction. For independent Belarusian outlets working from exile, continuity and secure access are likely to remain more urgent than broad automation. In 2026, DDoS attacks hit at least six exiled Belarusian news sites and a journalists' association; Zerkalo reported that its technical staff had to divert time from product development to protecting its infrastructure. At the same time, Zerkalo's AI proofreader shows a narrow workflow becoming part of routine production: the tool is integrated into its CMS, corrections are reviewed by people, and iterative changes raised the share of accepted suggestions from about 56% to 80%. The newsroom says it improved weekend coverage and redistributed work, but did not replace proofreaders or remove human responsibility. Sources: Committee to Protect Journalists on DDoS attacks · Zerkalo AI proofreader case study

What could shift it. Funding should pair AI adoption with digital security: threat modeling, resilient hosting and backups, incident response, and tools that can be tested without exposing unpublished or sensitive reporting to unsuitable services. Shared or locally controlled systems could also help preserve access to verified journalism when platforms or vendors change their terms. Those are resilience measures, not a claim that any single tool can bypass state blocking or prevent attacks.

Moldova

Current direction. Moldova has broad informal uptake but little newsroom governance. The 2025 Independent Journalism Center assessment found that 67% of participating media institutions used AI in professional work, while only 12% of those users reported clear, functioning policies; 88% of users wanted staff training, and many outlets lacked the budget for it. The government's 2024 White Book sets a strategic direction for responsible AI and data governance aligned with EU standards, but it is a policy framework rather than a newsroom implementation plan. At outlet level, #diez discloses AI assistance in selected workflows, NewsMaker has published rules for AI-generated images and verification, and Observatorul de Nord is developing an editorial assistant with Internews and NarativAI. Election-related impersonation and synthetic disinformation add pressure for authentication and transparent use. Sources: IJC media needs assessment · Moldova White Book on AI and Data Governance · Observatorul de Nord and NarativAI project · Associated Press on AI-driven election disinformation

What could shift it. The practical step is to turn isolated pilots and the White Book's policy direction into shared newsroom capacity. The IJC's model editorial policy gives outlets a starting point; a targeted grant and training program could help them adapt it, select tools, build staff skills, and fund maintenance. Shared procurement and clear requirements for privacy, disclosure, verification, and human editorial control would reduce duplication while supporting EU-aligned practice. Funding should reach local and small outlets as well as national organizations, since they face the largest gaps in staff time and technical support. Sources: IJC model AI editorial policy

Ukraine

Current direction. Ukraine is moving from widespread experimentation toward more routine use and possible sector-wide coordination, but the evidence also shows a capacity gap. In the Media Development Foundation's 2026 survey of 40 local and niche outlets, 85% used AI for routine text editing, 47.5% for translation, and about half for image generation; only two outlets used it to check their own reporting against source material. Just 27.5% had a written AI policy, while 85% named training as their top need. Suspilne's two-year strategy, approved in August 2026, creates a department for innovation and AI, sets out tools for planning, verification, transcription, archives, and audience products, and calls for editorial, ethical, security, and reputational risk reviews. Its AI Alliance may help connect this institutional work to other Ukrainian media. Sources: Media Development Foundation survey · Suspilne AI strategy

What could shift it. The next step is to make local experiments maintainable and shareable. Regional outlets are already using AI for administrative routines and generated announcements; investigative teams use data tools to examine procurement and officials' declarations. But the 40-outlet survey is a snapshot, not a census, and routine uptake alone does not establish reliable verification or long-term savings. Training, written rules, security review, and stable funding for maintenance should accompany new pilots. Suspilne's alliance offers a route for sharing practices, while the Commission on Journalistic Ethics' guidance keeps responsibility with journalists and editors. Collective investment could help regional newsrooms build common infrastructure and respond to the information-security threats facing Ukrainian media. Sources: Ethics Commission recommendations · regional workflow reporting · Nikcenter AI and data reporting · business and traffic reporting

7. What to do in three phases: 3 months, 6 months, and 1 year

From experiment to capability

3 months

Map use. Name an owner. Set a baseline.

6 months

Pilot one task. Write rules. Verify outputs.

1 year

Measure. Maintain or stop. Share learning.

First three months. Document current AI use, name someone accountable, and set a baseline by checking whether AI systems describe the outlet and its work accurately. Review the newsroom's existing ethics and disclosure rules before choosing a new tool.

By six months. Choose one high-volume task - often transcription or translation - and run a measured pilot. Write a short policy that keeps publication decisions with editors, requires verification of AI-generated facts and citations, and bars unreviewed automatic publication. Disclose material AI-generated text and visuals; avoid AI-generated images to illustrate current news. Check outputs for bias and plagiarism, and keep personal or sensitive information out of prompts. The Ministry's media guidelines also call for lawful use, systematic risk assessment, transparency about AI's purpose and the nature of its output, data protection, human oversight, and clear accountability. The Commission on Journalistic Ethics recommendations and Suspilne's strategy offer Ukrainian starting points. Commission recommendations · Suspilne AI strategy · Ministry media guidelines · Harvard Business Review on AI bias

By one year. Compare the pilot with its baseline, decide whether to maintain or stop it, and share what was learned. If the newsroom has a verified, organized archive and technical support, pilot an MCP server alone or through a shared regional service. Fund maintenance and training, not only the initial build.

Closing

A recent essay in the newsletter *(Re)Structured News* raises a question worth sitting with rather than answering too quickly: as new AI-driven workflows take shape inside newsrooms, will they simply be faster versions of the jobs that already exist, or something genuinely different, built around whatever specific skills turn out to still be worth having? Nobody running a newsroom in Yerevan, Chișinău, Kyiv, or in Belarusian exile knows the answer yet. What connects all four is that the pressure driving AI adoption there was never really about the technology. It was commercial in Armenia, existential in Belarus, tied to political transition in Moldova, and driven by war in Ukraine. Whatever the newsroom of the future turns out to be, in any of these four countries, the parts of the job that stay recognizably human, deciding what matters, showing up under repression or shelling, insisting on a language and an identity no model was built to care about, are likely to be the parts that were never really the technology's decision to make in the first place.



A

Sources index

Research, newsroom accounts and practical references, organized by purpose. Select a source title to open it.

SOURCES / 01

Landscape & research

Core datasets, industry trends and regional evidence.

1.01
AI FOR NEWSROOM
Industry Report
Industry analysis. Its July 2026 editorial snapshot is distinct from the live initiative counts used in this edition.

1.02
AI FOR NEWSROOM
Initiatives Database
Live initiative totals and category shares, checked on 2 November 2026; also the discovery layer for implementation cases.

1.03
SIMILARWEB
2026 Generative AI Landscape
AI-platform growth, zero-click search, citation behavior, referrals, and AI visibility.

1.04
REUTERS INSTITUTE
Journalism, Media, and Technology Trends and Predictions 2026
Industry adoption, strategy, and the transition from experimentation to routine newsroom use.

1.05
REUTERS INSTITUTE
How Will AI Reshape News in 2026? Forecasts from 17 Experts
Forecasts on agents, newsgathering, infrastructure, and strategic investment.

1.06
REUTERS INSTITUTE
Emerging Uses of AI Chatbots for News
Audience use of chatbots for news and the implications for publishers.

1.07 REUTERS INSTITUTE

Digital News Report 2026: Executive Summary

Audience behavior, trust, and changing routes to news.

1.08 MEDIA DEVELOPMENT FOUNDATION

AI in Ukrainian Local Media Research

Survey evidence on adoption, budgets, governance, training needs, and newsroom practice in Ukraine.

1.09 INDEPENDENT JOURNALISM CENTER, MOLDOVA

Media Institutions' Needs Assessment

Survey and interview evidence on AI adoption, policy, skills, budgets, and institutional pressure in Moldova.

1.10 FREEDOM HOUSE

AI Use in Armenian Media: Policy Brief

Regional policy and practice context for AI use in Armenian media.

1.11 MEDIA.AM

Armenian Media in the Age of Artificial Intelligence

Examples of routine and advanced AI use in Armenian newsrooms.

1.12 FREEDOM OF INFORMATION CENTER OF ARMENIA

National Disinformation Concept and Action Plan

Armenia's policy context for disinformation, deepfakes, and institutional response.

1.13 REUTERS INSTITUTE

Generative AI and News Report 2026

Latest audience survey, published 29 October 2026; news use and comfort with newsroom AI, with survey bases stated.

1.14 SIMILARWEB

Is AI Replacing Search? What the 2026 Numbers Actually Reveal

Verified 2026 app-download totals and US desktop ChatGPT ad penetration, with measurement limits.

SOURCES / 02

Newsrooms & regional sources

First-hand case accounts, local tools, editorial policies and regional reporting.

The first four records are unpublished questionnaires supplied for this casebook. All other linked titles open the public source supplied in the draft.

2.01

AMPOP MEDIA

SmartFlow questionnaire response

Unpublished first-hand newsroom questionnaire supplied for this casebook, 2026.

2.02

18000

AI Content Distribution Platform questionnaire response

Unpublished first-hand newsroom questionnaire supplied for this casebook, 2026.

2.03

BASTV

Civic Radar questionnaire response

Unpublished first-hand newsroom questionnaire supplied for this casebook, 2026.

2.04

BGMEDIA

BG Avatar questionnaire response

Unpublished first-hand newsroom questionnaire supplied for this casebook, 2026.

2.05

Armenian Corruption Data Portal

Working public-data product discussed as an example of locally useful newsroom infrastructure.

2.06

VS.AM

Armenian AI Platform

Armenian speech-to-text, text-to-speech, OCR, subtitles, voice tools, and related language infrastructure.

2.07

MediaScoper

Armenian-language and regional media monitoring, semantic search, and narrative detection.

2.08
J-INDEX
Armenia
Country-level journalism and information-environment reference.

2.09
Media Ethics Observatory of Armenia
Self-regulatory framework and shared principles for Armenian media.

2.10
CIVILNET
AI Ethics Code
Newsroom rules on transparency, human responsibility, verification, discrimination, and source protection.

2.11
NATIONAL ASSEMBLY OF ARMENIA
Amendments Requiring Labels for AI-Generated Broadcast Content
Legal requirement used in the Armenia governance section.

2.12
CIVILNET FACT CHECK
Fabricated Abortion-Ban Broadcast Using CivilNet Branding
Example of synthetic media impersonating a trusted outlet.

2.13
CIVILNET FACT CHECK
Fabricated HIV Story
Example of false reporting packaged with newsroom branding.

2.14
CIVILNET FACT CHECK
False Charity-Theft Claim
Example of disinformation borrowing CivilNet's identity.

2.15
KYIV INDEPENDENT
How Deepfakes Manipulate Ukraine's Mobilization Narrative
Evidence for the threat side of synthetic media in Ukraine.

2.16
Rayon.in.ua and Raya
Hyperlocal newsroom workflow, internal memory, and AI implementation.

2.17
DW AKADEMIE
Mapping Ukrainian Media Outlets' AI Use
Cases including Tsukr and other Ukrainian regional publishers.

2.18 NARATIVAI

A Local Moldovan Newsroom Turns Independence into Innovation

Observatorul de Nord's editorial-assistant work with Internews and NarativAI.

2.19 INDEPENDENT JOURNALISM CENTER, MOLDOVA

Model Policy for Ethical and Transparent AI Use

Governance template and newsroom training in Moldova.

2.20 PRESS COUNCIL OF MOLDOVA

Code of Ethics for Journalists

Ethical baseline for editorial responsibility and public trust.

2.21 COMMISSION ON JOURNALISTIC ETHICS OF UKRAINE

Recommendations on AI in journalism

Editorial responsibility and guidance for AI-assisted publications.

2.22 MINISTRY OF DIGITAL TRANSFORMATION OF UKRAINE

Guidelines for AI use in media

Lawful use, transparency, data protection and human oversight.

2.23 MYKOLAIV CENTER FOR INVESTIGATIVE REPORTING

NikCenter

Investigative reporting, public records and regional data work.

2.24 MYKOLAIV CENTER FOR INVESTIGATIVE REPORTING

Calliope image-verification bot

Telegram tool for a preliminary image-authenticity assessment.

2.25 DETECTOR MEDIA

How Ukrainian media integrate AI into their work

Reported examples from local and national newsrooms in Ukraine.

2.26 SUSPILNE UKRAINE

Artificial intelligence strategy

The public broadcaster's institutional approach to AI implementation.

SOURCES / 03

Implementation cases

Original accounts of newsroom projects, traced through AI For Newsroom.

3.01 SAN FRANCISCO CHRONICLE

Public Meeting Monitoring System (Online News Association awards entry)

Original case reference surfaced by AI For Newsroom.

3.02 CHALKBEAT

LocalLens (Nieman Lab)

Original reporting linked by AI For Newsroom.

3.03 FINANCIAL TIMES

AI-Powered Story-Finding Pipeline (The Audiencers)

Original case article linked by AI For Newsroom.

3.04 NEWTRAL

FactFlow AI

Newtral’s own account of its Telegram misinformation tool; linked by AI For Newsroom.

3.05 TAMEDIA

Hyperlocal Newsletters (Innovate Local)

Original project case linked by AI For Newsroom.

3.06 GRUPO FÓRMULA

NAT and SOFI (Reuters Institute)

Original reporting on synthetic presenters linked by AI For Newsroom.

3.07 SHARJAH MEDIA CITY

AI Virtual Presenter (City, University of London study)

Academic source linked by AI For Newsroom.

3.08 OCALA GAZETTE

Social Card Generator

Builder's case account linked by AI For Newsroom.

3.09 TEMPO

Archive-to-Knowledge Assistant (JournalismAI)

Detailed project source linked by AI For Newsroom.

3.10 TEMPO

JournalismAI Innovation Challenge Grant

Tempo's own announcement and project context.

3.11 SOUTH CHINA MORNING POST

Vintage Archive

SCMP's account of archive digitization and reuse; linked by AI For Newsroom.

3.12 THE NEW YORK TIMES

Perspective Comment Moderation (Google)

Original case account linked by AI For Newsroom.

3.13 DAILY MAVERICK

Revenue 360 (JournalismAI)

Original project page linked by AI For Newsroom.

3.14 DAILY MAVERICK

Membership and ImpactEngine (WAN-IFRA)

Follow-up source on membership experimentation and impact measurement.

3.15 OBSERVADOR

Subscription AI Concierge (JournalismAI)

Original project account linked by AI For Newsroom.

3.16 HEARST NEWSPAPERS

TX Tax

Hearst's expansion announcement linked by AI For Newsroom.

3.17 HOUSTON CHRONICLE

TX Tax Tool

Public product referenced by the AI For Newsroom case.

3.18 PHILADELPHIA INQUIRER

AI-Assisted Newsletters

External coverage exposed by the AI For Newsroom case record.

3.19 THE GUARDIAN

MCP Server (Nieman Lab)

Original case article linked by AI For Newsroom.

3.20 PROJECT VERDAD

Political Misinformation in Spanish-Language Radio

Primary article exposed by the AI For Newsroom initiative record.

3.21 PUBLIC DATA WORKS

VERDAD Project Organization

Organization page exposed by the AI For Newsroom initiative record.

3.22 NEWS AND MEDIA HOLDING, SLOVAKIA

Social Media Post Generator

The AI For Newsroom record exposes the publisher’s site but no separate public case article.

3.23 CIVIO

JournalismAI Innovation Grants

Grant-program source exposed by the AI For Newsroom record.

3.24 FUTURE PLC

Future Optic / GEO as a Service (Press Gazette)

External source surfaced through AI For Newsroom for the publisher’s AI-visibility business.

3.25 THE ECONOMIST

ChatGPT Graphs App (Nieman Lab)

External article identified in the AI For Newsroom record.

SOURCES / 04

Strategy, distribution & trust

Infrastructure, audience relationships, governance and the future of news.

4.01
FT STRATEGIES
Liquid Content
How publishers can structure reporting for both human readers and AI retrieval.

4.02
FT STRATEGIES
AI Engines and Context Engines
Distribution through agents, context engines, licensing, and publisher-controlled access.

4.03
REUTERS
Reuters MCP
Primary product source for licensed Reuters content delivered through model-context infrastructure.

4.04
CENTER FOR JOURNALISM AND LIBERTY
Joint White Paper on Independent Journalism
Publisher bargaining power and privately controlled AI data marketplaces.

4.05
CENTER FOR JOURNALISM AND LIBERTY / CLOUDFLARE
White Paper PDF
Full white paper supporting the platform-capture and collective-infrastructure discussion.

4.06
COLUMN
Agentic AI and Journalism Research
Research landing page for the agentic-newsroom white paper cited in Part 3.

4.07
NEWSROOM ROBOTS
Archive
Industry conversations on agents, newsroom products, and distribution beyond the click.

4.08
LOCAL MEDIA ASSOCIATION AND TRUSTING NEWS
Audience Survey on Newsroom AI
Reader attitudes toward disclosure, judgment, and trust.

4.09
TRUSTING NEWS
AI Trust Kit
Reusable audience-research and transparency materials.

4.10
SÜDDEUTSCHE ZEITUNG
Designing Clarity: Labeling AI in News
Practical framework for explaining AI involvement to readers.

4.11
SUSPILNE
C2PA Authenticity Verification Pilot
Primary newsroom source for content-provenance implementation in Ukraine.

4.12
C2PA
Technical Specifications
Authoritative standard for content provenance and authenticity.

4.13
RESTRUCTURED NEWS
It’s the Economics, Stupid
Analysis of AI distribution and the economics of publishing.

4.14
RESTRUCTURED NEWS
We Have Met the Enemy and He Is Us
Industry commentary on strategic choices around AI and journalism.

4.15 MATT PEARCE

It’s Not Journalism’s Job to Make AI Trustworthy

Critical argument about journalism’s responsibility for AI accuracy and trust.

4.16 FELIX SIMON

Six Provocations About AI in News

Critical framing of AI adoption and newsroom strategy.

4.17 FELIX SIMON

Should Journalists Use AI for Writing?

Evidence and argument about audience attitudes and editorial use.

4.18 AXIOS

Anthropic, Open Weights, and National-Security Risk

Source for Dario Amodei’s public position on open-weight risks.

4.19 MICROSOFT

Open-Weight Models and Responsible Access

Industry position on security, innovation, and access to open-weight models.

4.20 R STREET INSTITUTE

Open-Source AI and Cybersecurity

Policy analysis supporting the open-model security discussion.

4.21

Jor-MCP

Example of model-context infrastructure built for journalism.

4.22

Gazzetta

Project homepage for a newsroom-oriented MCP and AI-distribution experiment.