THE LONDON SCHOOL OF ECONOMICS AND POLITICAL SCIENCE

AI and the newsroom next door

Experiments and best practices for small publishers

2024 Innovation Challenge

Google News Initiative JournalismAi POLIS Journalism at LSE

About this report

This report gives you 35 great examples of how smaller newsrooms have used AI technologies to improve their journalism, strengthen their organisations and to expand their public mission. The JournalismAI Innovation Challenge, supported by the Google News Initiative programme builds on eight years of work by the LSE JournalismAI project. We have built training and education programmes, a network of thousands of practitioners around the world, practical courses to help newsrooms to understand and implement AI, and programmes that support experimental collaborations. We work across the globe and, apart from English, in Spanish, Portuguese, Arabic and French.

The Innovation Challenge was the next step forward as it allowed us to directly support smaller news organisations to put new AI tools or systems into action. Those are of immediate and longer term value to those newsrooms. But as important, they have given us an invaluable account of how others can follow in their footsteps. We offer some strategic advice in this report and our other publications about how to approach and adopt AI. However, in the end each news organisation is different. How you use AI will depend upon your particular needs, standards, business model and mission.

This report is not an instruction manual, but it should serve to inspire you with a wide range of ways to use AI. The journalism and AI journey is still in its initial phase. However, it is already becoming a routine part of the journalism process. There is no time like the present to get on board with these technologies and to explore what AI can do for you. Not just to help make your current work more efficient and effective, but to open up new possibilities for editorial creativity and impact.

I want to congratulate all the participants in their hard work, commitment and innovative flair. AI is an impressive technology but putting it into practice in the real world can be problematic. I also want to thank my JournalismAI team who supported the news organisations and the mentors and experts who were part of the process. There will be more Innovation Challenges and the wider JournalismAI project will continue. Please check out all our programmes, sign up for our newsletter and get involved.

Professor Charlie Beckett,
Director of Polis, LSE, leader of LSE's JournalismAI Project

Contents

About the JournalismAI Innovation Challenge

Previous research[1] has shown that many, mainly larger and well resourced newsrooms, have benefited from and seen the transformative power of AI, while smaller news organisations often face unique challenges in adopting this technology. So, it is important to lay the foundation in enabling smaller news organisations to experiment with AI technologies. The JournalismAI Innovation Challenge, supported by the Google News Initiative, a nine-month long programme, is aimed at enabling small and medium-sized publishers to experiment, implement and share best practices of artificial intelligence (AI) technologies. It is designed by the JournalismAI team at the London School of Economics and Political Sciences (LSE) and supported by the Google News Initiative (GNI).

For this inaugural Innovation Challenge, 35 news organisations were selected to experiment with AI solutions across three thematic areas:

- Helping fight misinformation and disinformation
- Experimenting with new formats to engage new and existing audiences
- Identifying ways to grow subscriptions and support diverse forms of revenue

Ten ideas received USD $250,000, while 25 ideas received USD $50,000.

To support these organisations in achieving their goals, as outlined in their applications, we approached the first phase with sessions focused on organising and working in interdisciplinary teams. This was followed by sessions on using a product lens to develop a prototype, while another session focused on how to work with budgets. All these sessions were led by experts in the field, who are also part of the JournalismAI community.

The second phase of the programme was the discovery and experimentation phase. The timeline for this phase was split over nine months. Grantees followed a time-boxed "sprint" format where they would build for four weeks, and report milestones in the fifth week to the JournalismAI team. They had access to coaches who guided them and shared expertise on how to further develop their ideas, while at the same time leveraging each other's strengths and support as a cohort.

In this report we provide a complete record, and share the outcomes of these experiments and best practices about using AI technologies for small and medium-sized news organisations.

---

1 Generating Change. 2023. Beckett, C.H., and Yaseen, M.
New Powers, New Responsibilities, 2019. Beckett, C.H.

The cohort

For this inaugural Innovation Challenge, we selected 35 grantees from 22 countries globally: Argentina, Brazil, Colombia, France, Germany, Greece, India, Israel, Jordan, Mongolia, Nigeria, Paraguay, Philippines, Serbia, South Africa, Spain, Switzerland, Tunisia, Turkey, Ukraine, United Kingdom and United States of America.

The grantees are listed below, according to their thematic areas:

Experimenting with new formats across mobile, video and audio

For the new formats thematic area, 12 organisations were selected, with four (4) organisations receiving the $250,000 and eight (8) receiving the $50,000.

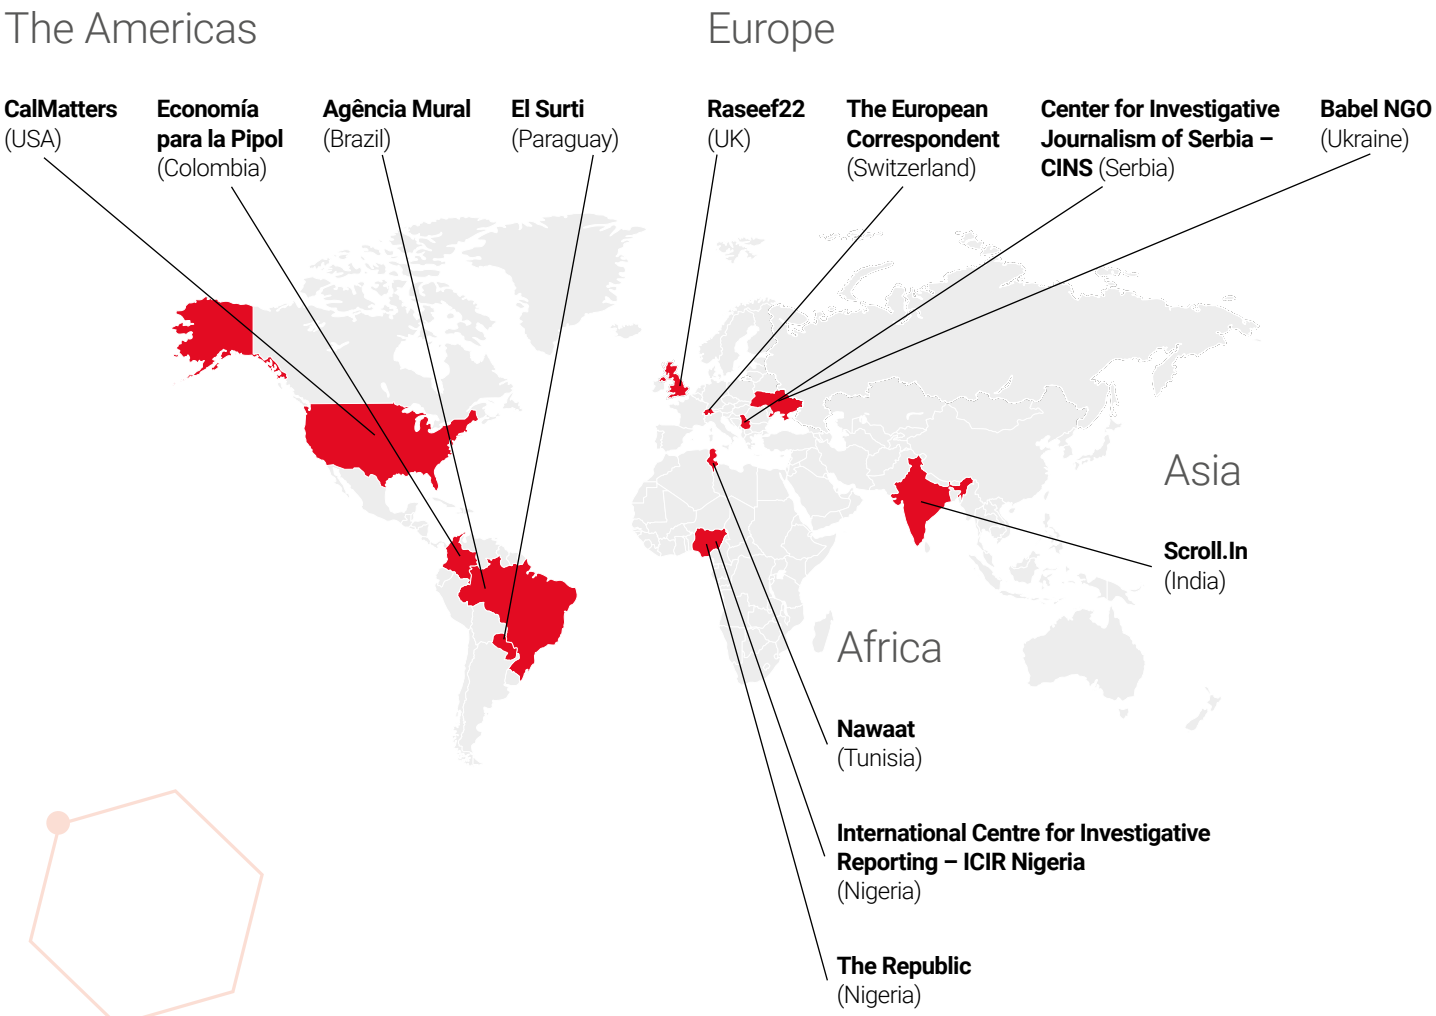

Identifying ways to grow subscriptions and support diverse forms of revenue

We received the least number of applications in this thematic area, and therefore the grant allocation was lower than the other two thematic areas: nine (9) organisations were selected, with two (2) organisations receiving the $250,000 and seven (7) receiving the $50,000.

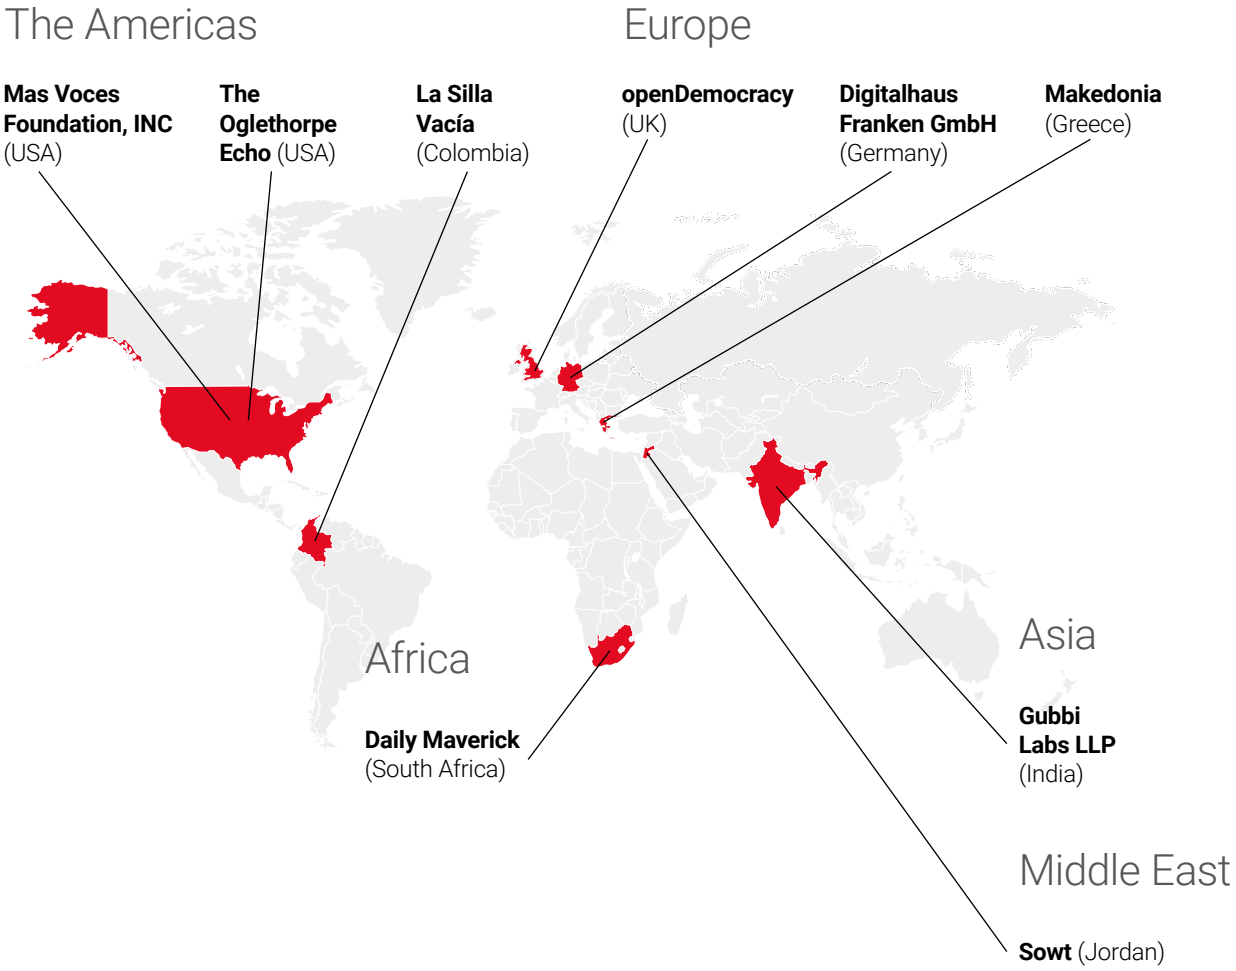

To fight misinformation and disinformation

The fighting misinformation and disinformation thematic area received the most grantees: 14 organisations were selected, with four (4) organisations receiving the $250,000 and 10 receiving the $50,000.

Overview of Chapters

Although the grantees were selected through the three thematic areas, soon after the programme started, we found other threads that linked these organisations and we will present their case studies according to these nine alternative thematic groupings that may be more relevant to a newsroom setting. These alternative thematic groups form the nine (9) chapters of case studies which focus on:

- **AI-powered chatbots**: the biggest chapter of case studies ranging from chatbots for fact-checking to chatbots for public engagement and knowledge sharing. Some of these chatbots are more complex, with added layers of multilingualism, and others linked to third party platforms.
- **Disinformation tracking on social media**: These projects are using AI to monitor, analyse, and archive dis/mis information data on social media.
- **Audience engagement and retention**: These projects are using AI to personalise landing pages, building audience engagement and retention.
- **Content repurposing**: These projects are leveraging AI to repurpose articles into various formats which includes translation and voice recognition, as well as interactive content.
- **Content matching**: The two projects leverage the use of AI for content matching.
- **Credibility and misinformation detection**: Projects in this section are using AI to enhance journalists credibility; providing editing assistance.
- **Newsroom workflow and efficiency**: These projects are using AI to enhance journalistic workflows.
- **Real-time data analysis and community alerts**: These projects use AI to process and analyse complex datasets to provide real-time insights.
- **Government transparency and accountability**: These projects use AI to enhance public access to government information, track policy decisions, and ensure accountability.

In the final chapter of this report, we share some reflections and a summary of strategic learnings we found from these successful newsroom experiments, that will help you create a strategy.

1. AI-Powered Chatbots

Over the past year various publishers have been experimenting with AI-powered chatbots with the goal of increasing reader engagement, and also leveraging archives that have laid dormant for many years. This thinking was not any different from the Innovation Challenge grantees whose projects focused on building AI-powered chatbots. These ranged from chatbots for fact-checking and those used for audience engagement and knowledge sharing. However, these chatbots are more complex than they look on the surface. Some were multilingual chatbots, while others were built for WhatsApp or the respective publications' websites. What are some of the lessons from these organisations?

For the organisation's whose chatbots focused on fact-checking, the thread that connects them all is that AI is an amplifier and not a replacement for human intelligence. In their use cases, AI's role assisted human fact-checkers and enhanced their efficiency, rather than replacing them.

These organisations also have an unwavering commitment to accuracy and trust, as they prioritise rigorous standards for accuracy and public trust in their outputs. They built their bots to specifically pull information from their own database of thoroughly "verified and debunked" articles, and not from general sources, to ensure trustworthiness. This goes hand-in-hand with the shared understanding of the responsible deployment of AI in the newsroom and transparency. That AI is a tool and not a panacea and its use requires transparency, clear principles and realistic expectations. Another unique consideration for success when building AI-powered chatbots, hinges on tailoring AI solutions to specific regional, linguistic and cultural contexts.

A point that has been emphasised in our previous research, are the benefits of collaborating with external partners, whether academic institutions, tech companies or community networks[2].

2 Generating Change, 2023, Beckett, C.H., and Yaseen, M.

All the grantees shared common approaches to their technology stacks, primarily focusing on AI solutions that are customised for local contexts and that can be integrated with widely used platforms like WhatsApp and X (formerly Twitter). Most organisations leveraged generative AI to power their fact-checking efforts. AI models connected via Retrieval Augmented Generation (RAG) to process queries and retrieve information were used by others. To ensure accuracy and trustworthiness, all systems relied on their own verified content or carefully curated external sources, rather than general web searches. Beyond this, publishers also built custom-trained AI models for local languages to address the challenges of integrating them into systems designed for Western languages.

Some of the key lessons learned by these organisations is making sure that platforms/ publishers meet audiences where they are and ensuring accessibility as key to impact. Fact-checking organisations, they need to focus on solving real problems their organisations face rather than simply chasing AI trends. Not every problem can be solved with AI, so it is crucial for organisations to understand what they can and cannot achieve with it.

Over the next couple of pages you will review publisher case studies that focused on AI-powered chatbots from different perspectives, regions and focuses.

1.1 Tama Media: Bringing AI-powered fact-checking to Africa's languages

Tama Media, a Pan-African media outlet that employs African journalists to tell local stories, has developed an innovative AI-powered chatbot designed to combat misinformation while breaking down language barriers. The organisation's approach demonstrates how artificial intelligence can be adapted to serve local contexts, particularly in regions where traditional fact-checking methods struggle to reach vulnerable populations.

The problem: Reaching beyond digital divides

Africa faces a unique challenge in the fight against misinformation. As Tama Media's project lead Moïse Mounkoro explains: "In Africa, there are people who do not or cannot read, and yet, they still use social media platforms where they easily can share fake news without being aware of it". Due to this, traditional fact-checking approaches – articles, videos, and podcasts – aren't effectively reaching these communities.

Tama Media recognised three critical challenges: making fact-checking easier and more engaging; ensuring accessibility for people who primarily communicate in local African languages rather than French or English; and reaching illiterate social media users.

"The idea is to see how AI can contribute to fighting misinformation – making fact-checking easy, fun, and accessible."

Building the solution: AI meets local expertise

Tama Media's response was to build an AI-powered chatbot that users can interact with through voice or text. The chatbot allows users to verify information by simply asking questions like, "I have seen this on social media, is it true?" The system responds with verification and provides multiple sources to support its answers. Crucially, the tool integrates with popular platforms like WhatsApp and Facebook, meeting users where they already are.

What sets this project apart is its hybrid approach. The team built their solution using both ChatGPT and a custom-trained AI model for local African languages, starting with Bambara – a language spoken in Mali and, in related forms, across West Africa, particularly in Côte d'Ivoire, Burkina Faso and Republic of Guinea.

The system maintains quality through careful source curation. "We selected media news outlets that are known and that we are sure don't spread fake news," explains Mounkoro. "We have selected outlets including BBC Africa, RFi, France 24, and local African news outlets."

When users aren't satisfied with responses, human oversight kicks in. "When the user responds negatively, the system sends our newsroom a message that there has been a question and the person is not satisfied. When we answer the question, the user gets an alert that the answer is ready."

Building local capacity

The project's development involved two distinct teams: journalists responsible for data collection and user responses, and a technical team based in Abidjan. The technical team includes specialists in AI, design, and integration – a deliberately local approach that initially faced scepticism.

"I talked to several experts in Paris about this tool and they said that it's too complicated and unfeasible. But then we got in touch with a technical team in Côte d'Ivoire, they acknowledged its complexity but they rose to the challenge."

This local focus proved crucial for creating a tool designed by Africans for African users, despite the technical challenges involved.

Overcoming technical hurdles

The team encountered several significant challenges. Initially, the system took five to 10 minutes to process queries – impractical for social media users expecting instant responses. Through persistent optimisation, they reduced response times to 20-30 seconds.

Accuracy presented another hurdle. The team needed to ensure 80-90 per cent reliability in fact-checking results while working with limited resources as a small, local team.

Perhaps most significantly, integrating African languages into AI systems designed primarily for Western languages required innovative solutions. The team collected extensive data in Bambara – podcasts, dictionaries, and other materials – to train their custom model.

Their solution involves real-time translation: when users submit queries in local languages, the system translates them to English or French for processing, then translates responses back to the original language for delivery.

The opportunities: Partnerships and sustainability

The team plans to offer their AI model to other African newsrooms, enabling similar fact-checking capabilities across the continent. They're also exploring sustainable revenue models through advertising that could support ongoing fact-checker compensation.

They also envision their project as a foundation for broader collaboration. "We hope to expand future partnerships to include non-media organisations in fields affected by misinformation, such as the World Health Organization."

Lessons for newsrooms

Tama Media's experience offers valuable insights for similar initiatives worldwide.

- **Build on existing strengths**: Launch AI projects by starting with your team's existing expertise in areas like fact-checking. This provides a strong, practical foundation for success.
- **Set realistic AI expectations**: Remember that AI is a tool, not a "magic" solution. It requires study and patience. It's crucial to understand its limitations and what problems it can and cannot solve for your organisation.
- **Integrate human and machine efficiency**: Effectively fighting misinformation in diverse languages requires combining AI's efficiency with human expertise and local knowledge. This creates a scalable model that respects cultural and linguistic nuances.
- **Prioritise symbiotic collaboration**: The future of fact-checking is about creating a symbiotic relationship where AI amplifies human capabilities. Don't aim to replace human experts, but rather to use technology to enhance their local knowledge and better serve the community.

Visit our website for the full case study and to learn more about this project.

1.2 **Verify: Using AI to combat misinformation in Syria**

When the Assad regime collapsed in December 2024, Syria's information landscape transformed overnight. Suddenly, regions previously inaccessible to independent media opened up, but with this freedom came an unprecedented surge in misinformation. For Verify, a Syrian fact-checking platform, the challenge was clear: how to handle the flood of verification requests while expanding their coverage across newly accessible territories.

Their answer? An AI-powered WhatsApp bot that could instantly access their database of verified information, allowing their human fact-checkers to focus on what they do best – investigating and debunking false claims.

The problem: Information chaos in a transitional society

"Everyone is trying to use misinformation as a weapon to affect civil society during this transition," explains Ahmad Primo, founder and director of Verify. The platform, which had operated primarily in northern Syria before December 2024, suddenly faced requests from Damascus, southern regions, and areas previously under regime control.

The scale was overwhelming. Verification requests poured in through encrypted messaging platforms – spaces where misinformation spreads unchecked and fact-checkers have no visibility. "The majority of people use encrypted chat groups to share information. We cannot access these spaces to check whether they're sharing verified information or fake news," Primo notes.

This challenge reflects a broader global issue, but in Syria's context, it carries particular urgency. False information about everything from vaccines to political developments can directly impact vulnerable populations including refugees and women. The team realised that traditional fact-checking methods were no longer sufficient.

Building the solution: From database to bot

Verify's approach centred on making their existing database of fact-checked information instantly accessible through AI. "The whole idea is to create a bot that gets information from our website specifically," Primo explains. "When we write articles, we ensure every single word is verified and debunked. The bot summarises this trusted information rather than pulling from Wikipedia or other sources."

The technical implementation uses RAG, connecting their Arabic-language database to an AI model that can respond to queries on WhatsApp – the platform most Syrians rely on for news and information.

Creating this system required bridging the gap between journalism and technology. Verify's internal team consisted of fact-checkers already familiar with AI tools in their daily work. However, they needed external expertise for the technical implementation.

Their first attempt with a technology company failed. "We don't have technical experts in our team, and they didn't understand the fact-checking environment," recalls Rami Magharbeh, the project lead. "We're speaking different languages."

The breakthrough came when they partnered with a company that had previously developed Verify's website. This shared history provided crucial context. "They have the minimum understanding of fact-checking in conflict zones, which involves multiple layers of complexity," Magharbeh explains.

Beyond the core team, Verify expanded their network of "citizen fact-checkers" – volunteers including media students, activists, and journalists now stationed across Syria. These individuals receive training in fact-checking tools and methodologies, becoming guardians of truth in their communities.

Practical solutions for complex challenges

The project faced several significant hurdles, each requiring creative solutions:

- **Language complexity**: With a database primarily in Arabic but users potentially querying in English, Verify developed a workaround. They prepared a "whitelist" of trusted English-language sources like Reuters and The Guardian. When the bot cannot find answers in their Arabic database, it can provide basic information from these pre-approved outlets. "We asked our developer to prepare this whitelist so the bot can help users get answers about general information from trusted sources," Primo explains.
- **Building user trust**: To make the bot feel more human, Verify insisted on natural conversation flows with clear escalation paths. "We challenged ourselves and the technology company to let the bot react as a human, not just as technology," Magharbeh explains. When the bot reaches its limits, it seamlessly transfers users to human fact-checkers, ensuring no query goes unanswered while filtering the volume of direct requests to the team.

The opportunities: Scaling truth in the age of AI

Despite challenges, the team sees transformative potential in their AI implementation. "AI will help us summarise requests and make the process easier for both us and our audience," Primo explains. The system also generates new fact-checking assignments when it encounters questions without answers in the database.

Looking ahead, Verify plans to expand beyond WhatsApp to other platforms like X, creating an ecosystem where verified information is always accessible. They envision a future where AI tools are essential for combating AI-generated misinformation.

"When we use traditional methods to fact-check against AI-generated misinformation, we're already too late," Primo argues. "We need AI to fight disinformation and propaganda campaigns effectively."

Lessons for newsrooms

Verify's experience offers valuable insights for other fact-checking organisations in conflict or transitional contexts:

- **Local context matters**: Generic AI solutions often fail in specific regional contexts. Success requires technology partners who understand both the language and the complexity of working in conflict zones.
- **Human-AI collaboration is key**: Rather than replacing fact-checkers, AI should amplify their capabilities, handling routine queries whilst humans focus on complex investigations.
- **Community engagement strengthens impac**t: Training citizen fact-checkers creates a distributed network of truth guardians, essential in societies recovering from authoritarian control.

As Syria navigates its transition, Verify's AI bot represents more than a technological upgrade – it's an investment in the country's information integrity. In a landscape where false narratives can derail democratic progress, the ability to deliver trusted information instantly may prove as vital as any political reform.

"We believe everyone in our society can affect their community by defending the truth," Primo concludes. "With AI, we're giving them the tools to do exactly that."

Visit our website for the full case study and to learn more about this project.

1.3 El Surti: Building understandable common knowledge

El Surti, a Paraguayan media collective known for its visual and community-driven journalism, has spent over a decade creating content that reflects the realities of underrepresented populations, content that is particularly relevant to speakers of Guarani and its hybrid with Spanish, Jopará. These are primarily oral languages that remain largely invisible to mainstream digital tools and AI systems, due to a lack of training data and linguistic representation in Large Language Models (LLMs).

To tackle this challenge, El Surti launched AI KUAA, an initiative whose name comes from the Guarani word kuaa, meaning "knowledge." The project aims to bridge the gap between AI language models and non-centralised, oral languages and improve the technological representation and usability of Guarani through a three-pronged approach:

- Improving Guarani representation in Mozilla's Common Voice database
- Developing a chatbot capable of understanding Guarani audio inputs
- Designing a toolkit for other media outlets to build similar tools and reach underserved audiences.

The problem: A digital divide for oral languages

The inspiration for AI KUAA stemmed from El Surti's previous project, "EVA," a chatbot about women in prison for micro-trafficking in Paraguay. During the development of EVA, the team discovered a crucial limitation: AI tools and transcription services struggled to recognise and process Guaraní, particularly the hybrid language known as Jopará, which blends Guaraní and Spanish.

As El Surti's Director, Alejandro Valdez Sanabria, explains, "AI KUAA seeks to solve a gap between users of non-centralised oral languages and their recognition in the large languages used by artificial intelligence and the most popular chatbots." This difficulty is exacerbated in digital spaces, as there isn't enough documented data to train AI models in these languages. The team realised that to truly serve their community, they needed to make AI understandable and usable in their native tongue.

Building the solution: Roadmap to prototyping

The project's roadmap is a collaborative and iterative process, balancing technological development with community engagement. The goal is not just to build a chatbot but to create a sustainable and evolving solution that empowers the Guaraní-speaking community.

The team

El Surti assembled a diverse team to tackle the project, to "balance both the technical side and the community side." Sebastián Auyanet is the project manager who also oversees the business model and community-oriented aspects of the project. For editorial and community, they hired a community coordinator who is Guaraní-speaking, Leila Bareiro, while Valdez kept the editorial and general leadership.

The technical team is composed of Sebastián Hacher and Axel Marazzi, conversational and UX designers with prior experience on the EVA project, along with a developer.

Tools and process

The team used a mix of established and open-source tools.

A key part of the project involves training AI to understand Guaraní using Mozilla's Common Voice dataset. This platform allows volunteers to donate their voices, creating a repository of spoken language data. The team organised "mingas" (community gatherings) where people could record their voices to train the dataset, to gather several hours of Guaraní voice data, and increase the dataset's validation percentage.

The conversational flows and chatbot logic are being built using platforms like Voiceflow and Botmaker. These tools are robust enough to handle millions of conversations, providing a strong foundation for the project. They work through an API El Surti built, connecting open-source tools like Transformers and PyTorch to process audio and transcribe it.

A key challenge has been adapting these tools for an oral language. The team is working on a voice-first approach, aiming to create a chatbot that can recognise and transcribe voice messages in Guaraní, a feature that has become more readily available in recent months.

Challenges faced

- **Technical limitations**: Training a model to handle Guaraní and Jopará required starting from scarce datasets. Models had to process hybrid sentences with two languages in the same audio. "The difficulty is to detect, translate separately, and make sense of it all in real time without delay," explains Hacher.
- **Community engagement**: Mozilla's Common Voice platform requires the ability to read Guarani and Spanish to participate, which excluded many fluent Guaraní speakers who were unable to read the language. The team explored new APIs to enable voice-only contributions through WhatsApp.
- **Shifting technologies**: "Dashboards change every week," says Auyanet. Constant updates in voice recognition tools forced the team to repeatedly adjust their workflows.

The opportunities: AI KUAA beyond language recognition

Strengthening the presence of Guaraní in the digital sphere reinforces cultural identity and offers new ways to connect with audiences, especially in low-connectivity areas where WhatsApp is a lifeline. The bot could help deliver hyperlocal, personalised information and to enable collaborative storytelling.

By documenting their methodology and sharing their code, El Surti hopes to support other organisations facing similar challenges. With half of the world's languages underrepresented in AI, the model could have far-reaching applications.

Lessons for newsrooms

- **Community is the foundation**: The project's success is deeply tied to its community-driven approach. By engaging the Guaraní-speaking population in "mingas" to build the dataset, El Surti is not just creating a technical solution but also fostering a sense of belonging and representation. This collaborative model, as Hacher notes, is the "design layer between technology and the community."
- **Iterate and adapt**: The AI landscape is constantly changing, with new tools and models emerging every week. El Surti's team has embraced an iterative approach, constantly testing and fine-tuning their solutions. This flexibility is crucial for developing a product in a rapidly evolving technological environment. The first version of their chatbot, while still in Spanish, was used to sign up subscribers, providing valuable insights into user behavior and community needs.
- **Prioritise narrative over automation**: The team's prior experience with the EVA project taught them the importance of human-centered storytelling. They use AI as a tool for interaction and understanding user intent, but the core narrative remains crafted by human journalists. As Marazzi and Hacher discussed, the challenge is to use these new generative tools to enhance, not replace, human creativity and journalistic integrity.

Visit our website for the full case study and to learn more about this project.

1.4 Economía para la Pipol: how to build an AI chatbot that democratises economics

Economía para la Pipol, a Colombian digital media outlet, partnered with Datasketch to create an AI-powered chatbot that transforms how citizens access and understand economic information. Rather than relying on jargon-heavy traditional channels, the chatbot delivers insights in the accessible language that has become Economía para la Pipol's signature style.

The problem: When economics meets the streets

The outlet emerged in 2021 during social unrest triggered by a controversial tax reform. Many Colombians feared its impact on their finances but struggled to grasp the technical details. "We wanted to bridge the gap between how the government and traditional media explain economics and what people understand from it," says Camila González Olarte, CEO and Co-founder.

Initially, their content on social media answered basic questions and proved highly successful, but these mostly reached urban, internet-savvy audiences. To connect with wider groups, the team launched a website and eventually developed the chatbot.

Building the solution: From social media to structured intelligence

The chatbot was designed around the questions people actually ask. The team mapped recurring queries from social channels, validated them through audience research, and created a database of more than 900 answers. This was enriched with official requests to Colombian government entities, ensuring the responses reflected accurate and current information.

Tools and technical implementation

The chatbot's architecture combines natural language processing with direct access to Colombia's official datasets, particularly from DANE (the National Administrative Department of Statistics) and the Central Bank. This allows it to provide real-time indicators rather than static explanations. Responses are grounded in Colombian terminology and complemented with visuals, graphs, and links to related videos.

The team developed two distinct AI models to power the chatbot's response capabilities:

- **RAG Model (Retrieval-Augmented Generation)**: This model finds and communicates the most relevant text information based on their curated question database. The RAG approach enhances response quality by retrieving specific knowledge from their extensive collection of pre-written answers before generating the final response, ensuring accuracy and relevance.
- **Datasketch's Proprietary Natural Language to SQL Model**: The partnership implements a specialised model that converts natural language questions into SQL queries. This capability allows the chatbot to directly query public databases in real-time, providing up-to-date economic data from official sources.

Building the right team

Around 12 professionals from Economía para la Pipol and Datasketch collaborated on the project, bringing together journalists, data scientists, developers, and designers. This mix of editorial and technical expertise ensured the tool was accurate, user-friendly, and trustworthy.

Navigating complex challenges

The project encountered several significant obstacles that tested both technical capabilities and strategic thinking. One of the most persistent challenges has been accessing official information and connecting new government databases to the chatbot. In Colombia, information access laws require formal legal procedures for many data requests, creating delays and complications. The team must navigate government communication channels and freedom of information processes to obtain the real-time data that makes the chatbot valuable to users.

Another key challenge has been understanding user language patterns. Despite Economía para la Pipol's experience with question-based content, the team discovered that people formulate queries differently in a chatbot environment compared to social media searches. This required continuous refinement of the database and question interpretation algorithms, face-to-face testing and online sessions with audience members to gain insights into how users naturally ask economic questions. These findings led to significant adjustments in their approach.

Finally, the chatbot has to position itself in a market already dominated by ChatGPT, Gemini, and other general-purpose AI tools. Testing, however, revealed its unique value proposition: users appreciated the specialised economic focus and trusted Economía para la Pipol's journalistic approach. They particularly valued the integration of visual content, including graphs and videos from the outlet's social media, which created a multimedia learning experience not available in generic AI tools.

The opportunities: Future opportunities and strategic vision

The chatbot project has revealed several strategic opportunities that extend beyond its immediate functionality.

- **Sustainable revenue models**: With media grants becoming increasingly scarce, the chatbot represents a potential path to financial sustainability. The team is exploring partnerships with organisations that need economic information translated for broader audiences, as well as a WhatsApp version.
- **Expanding the content ecosystem**: A single piece of research can be transformed across multiple platforms and formats, maximising the value of journalistic work by creating content that serves audiences across different preferences and consumption patterns.
- **Deeper audience understanding**: The chatbot development process has provided unprecedented insights into Economía para la Pipol's community. Beyond social media metrics, the team now understands why people seek economic information and what specific concerns drive their questions.

Lessons for newsrooms

- **Technology as a journalistic multiplier**: The chatbot project demonstrates how AI can amplify existing editorial strengths rather than replace journalistic expertise.
- **Partnership success requires shared values**: The collaboration between Economía para la Pipol and DataSketch succeeded because both organisations shared fundamental beliefs about information democracy and accessibility. Technical capabilities alone were insufficient; the partnership needed alignment on mission and values.
- **User research reveals hidden assumptions**: The difference between how people search for information on social media versus how they ask questions in a chatbot requires significant adjustments to content strategy and technical implementation.

Visit our website for the full case study and to learn more about this project.

1.5 **Raseef22: Serving readers by using AI to answer SRHR questions**

The team at Raseef22, a Lebanese independent media organisation, had two goals in mind when they decided to build "Ask Aunty" – an AI-based chatbot. One, was to build a tool where its readers would feel comfortable asking sexual and reproductive health rights (SRHR) questions.

Trusted sources of information in Arabic about SRHR are not easy to access in the Middle-East, according to Rokaya Kamel, Product Lead at Raseef22. Information is fragmented and the ecosystem is broken when it comes to accessing accurate and judgement-free SRHR information.

The second goal was to serve all their readers, but particularly, young women and LGBTQ+ readers in having a safe space to ask sensitive questions, added Kamel.

The problem: Safe access to SRHR information

They wanted what they were building to create a stigma-free, anonymous space that would reduce misinformation around SRHR information generally. They also wanted to focus information-sharing on taboo or legally sensitive topics.

The problem that they were looking to solve was informed by data. Raseef22 found that a third of organic searches from their readers were for articles related to SRHR on their website. They took it as an educated signal that young people were searching for information that they couldn't reach.

"We thought about ourselves and how would we solve a similar problem like that before the internet? We would ask our aunty. In the Arab world and the Middle East, it's very normal for you to ask your aunt before you ask your mother, because she's more culturally grounded," explained Kamel.

The thought led them to creating a chatbot for readers of Raseef22 using only their own and trusted partners' articles and archives. In keeping with the tradition of asking an aunty, they aptly also named the app after it: Ask Aunty.

"It would be easier for people to ask the question and not search for individual articles. We decided to use AI because it would be much easier for people to just ask this question and not wait 24 hours for a response," she added.

Building the solution: A friendly AI chatbot

Interestingly, the team working on the Ask Aunty chatbot comprised four professionals from Raseef22's marketing team: Line Itani, Diana Khanafer, Michel Abi Rashed. It's the third AI project that the team is working on. They also hired an AI consultant, Fawzi Ammache, to support them in their journey.

"We have a little bit of background experience with AI. We were not very intimidated by it but also we were very proud because we know that thousands of people will actually benefit from this project because of how much of a taboo this project handles," said Kamel.

They plan to internally test the chatbot for two months with hopes to release it to subscribers following that, and then to the public.

In terms of technology, the team initially wanted to use Delphi AI or Zapier. However, in the end they decided to build and use a custom GPT which Kamel said proved to be the "perfect" tool for their use case.

Choosing the right fit was not without its challenges, said Kamel. One of them was making the AI systems work in Arabic generally, and more specifically, reconciling the chatbots Egyptian Arabic with a knowledge base that consisted of Modern Standard Arabic content.

"It was very difficult and surprising how little the AI tool is adapted to Arabic. We had this problem that the language was completely butchered and in Ask Aunty. Arabic is in theory supported in these tools, but in practice the tools weren't properly trained for it, so we faced many language-related problems. Thankfully we don't have this problem right now with custom GPT so it's way better dealing with the Arabic language," explained Kamel.

They also faced issues early in the building phase when the data added to the custom GPT was ignored by the tool and it instead scoured the internet for answers.

"The second challenge was that we needed to ensure that chatbot answers were very accurate and very non-judgmental. They had to only use answers from Raseef22's articles and from trusted partners. We had this problem and now it's fixed. We have a closed dataset," she said.

Focusing on editorial controls

Since all questions asked to the chatbot are anonymous, the organisation is also putting into place safeguards to make sure the system is not abused. While there is no such concern from a security standpoint, Kamel said, it can be from an editorial point of view.

"Since it's an SRHR chatbot, people cannot ask a question and turn this chatbot into something that it's not, for example – this is an outlet for me to ask sexual questions, so let's abuse that. We're trying to put editorial guidelines into the chatbot to be able to answer any kind of foul play," she added.

Impact on the organisation

The Ask Aunty chatbot, apart from being a service to readers, is also positively impacting Raseef22's operations. For example, it's helping them identify what's missing from their articles.

"Right now, we're expecting users to ask some questions that will open our eyes to some aspects in SRHR articles that we don't have. Now we can talk to the editors and tell them, 'hey people are asking these questions. Let's make sure to add this to our article so we can answer them accordingly. Let's make our data bigger, let's enrich it so this will help us create more articles and better information,'" explained Kamel.

Future goals for the Ask Aunty chatbot include voice implementation, WhatsApp integration, and potential partnerships with SRHR organisations to use the API. They also plan to expand the database to include culture and history.

Lessons for newsrooms

- **Values over technology**: One of the biggest lessons for Raseef22 that can apply to other newsrooms while building with AI is to always center organisational and editorial values in the technology. Kamel said that their newsroom prioritised these values over just implementing a technology solution. This made a difference to their considerations in building the chatbot whether it was maintaining privacy or inclusivity.
- **Specify editorial guidelines for every project**: The team created editorial guidelines specific to the Ask Aunty project that were different from the guidelines they used for other AI projects. This was to ensure accuracy and keeping in mind the sensitivity of the subject, Kamel said.

Visit our website for the full case study and to learn more about this project.

1.6 Vera Files: Using AI to create SEEK, a GenAI fact-checking search engine

VERA Files, a Philippine media non-profit organisation, identified several bottlenecks during their fact-checking process. One, was in monitoring misinformation. Then came creating a taxonomy of fact-checks since it was a digital-only organisation. Last, was the issue of incentivising their users to engage with their database of over 5,400 fact checks and counting.

From the many use cases, Celine Samson, Head of Online Verification at VERA Files, and her team decided to use AI to improve the number of people engaging with their fact-check database.

"We decided to do this because we already have the dataset available to train the AI on. Also, at the moment we're really on a mission to increase audience engagement with VERA Files," said Samson.

The Problem: Engaging users to explore their database

Samson who is an alumna of the 2022 JournalismAI Academy said that one of the reasons why they decided to participate in the Innovation Challenge was because their team was now primed for it. The team had just welcomed two developers into their organisation apart from the existing reporters and editors.

"I also learned from the Academy that apart from just having journalists, you actually do need people who know how to carry out the tech part of things," she added.

Apart from this, AI was the next option after they tried many qualitative methods to achieve their distribution and engagement goals.

"We've already tried so many other approaches to try to engage with our audience. We've had town halls and focus groups with readers. We've tried to optimise how we engage with people on social media, for example. But we also wanted to try and see what AI could do to help us with achieving these goals," said Samson.

The combination of factors led them to create SEEK, a chatbot search assistant that draws on VERA Files' archive to answer their fact-checking queries or those on misinformation trends. It offers users a "Quick Answer" and "Think Deeper" setting to allow them to customise the level of detail they'd like in responses.

Building the solution: Aiding users in quicker fact-checks, enabling discovery

SEEK was inspired by other newsrooms' and fact-checking organisations' experiments in this space, said Samson.

"One of our main inspirations was Fátima of Aos Fatos where we actually talked to their director for innovation, Bruno Favero, in preparing for SEEK. We asked questions like how did they develop Fátima? What were their considerations? How do they make sure that the answers Fátima comes up with are accurate?" Apart from this, they also used Encyclopedia Britannica's chatbot and the San Francisco Chronicles' Kamala Harris news assistant as an inspiration to create SEEK.

For user research, Samson shared that their research involved speaking with teachers, students, and disinformation researchers, all of whom exhibited varying degrees of trust in AI systems and had experience using them. They found that most users employed AI to launch research or to check grammar.

"The overwhelming feeling when we do our fact-checking training is sometimes people tend to get fed up especially when they have to keep fact-checking their parents, for example, who keep on falling for scams. That also became a target market I would say, the 'fed up family fact-checkers', how do we make their lives easier?" she explained.

Key priorities identified by users for SEEK's development included speed, clarity, transparency, accuracy, and localised language. Users also indicated they would not use the tool if the interface was too complex, if it was paid, if data breaches occurred, or if it failed to answer their questions.

In further research, they also found that most Filipinos were already using the AI chatbots, and concluded that the learning curve to use their product would not be too steep.

Technology stack

The technology underlying SEEK involved the use of LangChain for the chatbot and then custom building the other features for the AI tool, for example, using the RAG system. This came with its own challenges.

"None of us has had experience building an AI tool from scratch. Our developers had some experience working with LangChain, but not in the capacity of creating a search assistant or a chatbot. It was something they had to learn from the bottom up. Also, we used the RAG. So that was something that the tech team had to learn and then they cascaded it to the rest of the team so that we were all on the same page.

Samson also shared that working towards creating SEEK alongside the full-time duties of their job proved to be challenging as well. In the end, she had to prioritise and re-prioritise tasks in order to reach their goals, both for the search assistant, as well as for their non-grant related tasks.

Impact on the organisation

The impact on VERA Files as an organisation has been positive, shared Samson, with the entire team "proud and excited" to test the tool. It also led them to work substantially on their AI policy.

"Some members were surprised and didn't know you could use AI that way for journalism and then to see something that was created by VERA Files itself not by a partner, that really impressed them. Another impact was we already had a skeleton of an AI policy but creating SEEK also actually led us to review the AI policy to make it fit the direction we want to go in the future," explained Samson.

SEEK was beta tested in September 2025, involving 77 people and receiving 39 feedback evaluations. It scored high in question comprehension, answer quality, natural language performance, and information satisfaction. Users particularly appreciated the design, citation of sources, and the "quick answers" versus "think deeper" settings, shared Samson.

Future iterations of SEEK would likely involve using a larger dataset. They would also explore partnerships with mainstream news organisations.

Lessons for the newsroom

- **Prioritise the problem, not the tool**: "Don't put the cart before the horse," says Samson. Identify specific journalistic problems first, and then determine if AI is the most appropriate solution, rather than simply looking for a way to use AI.
- **Resist AI pressure and understand limitations**: Don't feel forced to adopt AI just because it's a popular topic. It's crucial to understand what AI cannot do and ensure you use it appropriately to achieve a concrete goal.
- **Conduct deep user research**: Cast a wide net in defining and detailing your target groups to thoroughly validate your product idea. For example, VERA Files included teachers, researchers, and students to ensure their work met a broad set of community needs.

Visit our website for the full case study and to learn more about this project.

1.7 Verified News Network: Building trust in Indigenous communities with AI

In an era of widespread misinformation, Verified News Network (VNN) embarked on an ambitious project: to combat the spread of false information within Indigenous communities in northeast Oklahoma, USA.

The problem: A flood of misinformation

For seven years, VNN has served Indigenous communities, witnessing firsthand the damaging impact of misinformation and disinformation. Brittany Harlow Tidwell, VNN Co-founder, says, "We found that one of the biggest issues and challenges facing our community members is misinformation and disinformation, particularly relating to their communities."

This problem is exacerbated by historical exploitation, where information, resources, and even land have been extracted from Indigenous communities. "This is a centuries-long issue in Indigenous communities," explains Harlow Tidwell, adding an additional layer of complexity to building trust and introducing new technologies.

Furthermore, misinformation often serves to benefit external interests, with Harlow Tidwell noting, "Even really on the state level there is a lot of misinformation being weaponised to benefit the state and harm our tribes here in Oklahoma."

The rise of AI presented both a threat and an opportunity: to create a tool that could counter this tide of false narratives.

Building the solution: The Indian Country Chatbot

VNN conceived the "Indian Country Chatbot," an AI tool designed to provide accurate information sourced exclusively from Indigenous entities. These sources included tribal governments, native-owned companies, and native-led non-profit organisations.

The team prioritised cultural respect and data sovereignty throughout the development process. As Co-founder, Kelly Tidwell emphasises, "Everything we learned is closed source. Additionally, we got explicit permission from the different people that we got the data from."

This commitment to ethical data handling, while making the process more challenging, was paramount. "It made our job a lot harder," Harlow Tidwell admits, "but it was very important to the entire process of being culturally respectful."

Instead of building a chatbot platform from scratch, the team opted for chatbot.com due to its speed, affordability, and closed library functionality. However, the development was not without its hurdles. A significant challenge arose from the limited number of contributors willing to share data for AI training, as many Indigenous organisations lacked AI policies or were outright resistant to the technology. This necessitated a shift towards original content creation, with the team developing "fact sheets" to train the chatbot.

The team consisted of the two co-founders, alongside two reporters who assisted with research and writing, and a part-time web developer. Project management was facilitated through Coda, an online operations platform.

Tidwell describes their process: "Whenever an unknown chat came in, we would load it into Coda to find the answers. Then we'd go to our reporters, who would point out any misspellings or areas needing cleanup. Once proofed, we'd plug it back into the chatbot for training."

This iterative approach ensured accuracy and quality.

The opportunities: Fostering dialogue and understanding

The Indian Country Chatbot, despite its initial challenges, has opened up new avenues for engagement and learning. The team sees opportunities to further develop the chatbot, focusing on community feedback sessions to understand local needs and concerns regarding AI. Harlow Tidwell states, "The next step is really to have some community feedback sessions to learn more about our communities."

Another critical opportunity lies in collaborating with academic institutions. While some initial engagement occurred, further dialogue and the development of robust data-sharing policies are necessary to secure broader acceptance and contributions from the academic community. Ultimately, the chatbot can serve as a catalyst for awareness and education about AI within tribal communities.

As Tidwell notes, "working on tribal communities raises awareness on the nature of AI and chatbots, which is critical because awareness is the big driver."

Lessons for newsrooms

VNN's journey offers valuable lessons for newsrooms looking to implement AI solutions, especially when collaborating with marginalised communities.

- **Integrate community throughout development**: Don't just inform the community. Actively prioritise and integrate community perspectives across the entire AI solution building process. Involvement helps ensure the tool meets their needs.
- **Value and correct assumptions**: Understand that even with prior experience, your assumptions about a community's acceptance of new technology may be incorrect. Be open to having initial assumptions proven wrong and value that learning experience.
- **Address historical context and concerns**: When collaborating with marginalised communities, you must cultivate a deeper understanding of their historical concerns. Be prepared to address resistance related to issues like the environmental impact of AI or data privacy, which can become significant sticking points.

Visit our website for the full case study and to learn more about this project.

> "...we got explicit permission from the different people that we got the data from."

2. AI for Disinformation Tracking on Social Media

Where disinformation used to take flight from certain offline spaces such as from politicians interviews, it has quickly migrated to and evolved on social media. A key problem for organisations in the fight against disinformation is the virality that social media affords it. In other words, disinformation quickly goes viral before it can be debunked. Manual efforts to quell this quickly spreading fire, can often be insufficient and resource intensive. This is where organisations decided to turn to using AI.

The organisations tackling this problem during the JournalismAI Innovation Challenge were Género e Número (Brazil), Science Feedback in the UK, CCIJ (USA) and Newtral (Spain).

One of the key lessons from these organisations lay in the power of interdisciplinary collaborations to build AI tools. It helped them build tools that staff would use in-practice and intuitively. For example, editorial staff and fact-checkers were key to providing the right input, the ethical lens, and for some, even annotating for them. Meanwhile technical teams provided the technical guardrails for the tools and technical know-how within this context.

All the organisations believe that key fuel to fight disinformation lies in sharing their tools and lessons with their peers and networks. The organisations plan on making their AI tools available for those outside their organisations to use, like for example, other International Fact-Checking Network (IFCN) members, journalists, researchers, and policymakers.

These publishers also used iterative approaches to building their products, where they refined their paths, testing, technologies, and discarding anything that didn't work quickly.

2.1 Gênero e Número: Understanding anti-gender discourse on YouTube: a feminist approach to AI-powered monitoring

Gênero e Número, a Brazilian media outlet dedicated to covering gender, race, and inequality, joined forces with Novelo Data to tackle one of today's most pervasive threats: anti-gender disinformation. Their project seeks to map and understand how narratives targeting women and LGBTQ+ people spread on YouTube, especially during politically sensitive moments such as elections.

For the first time, Género e Número is using AI to expand its capacity to monitor and analyse online content. The partnership combines the newsroom's editorial expertise with Novelo's technical know-how, creating a platform designed not just for internal use but also to support researchers, journalists, and activists with accessible, data-driven insights.

The problem: Mapping the opaque disinformation flows

The idea grew out of Gênero e Número's 2019 investigation *O Reino Sagrado da Desinformação*, which exposed how political, religious, and media actors coordinated disinformation campaigns on X. That investigation relied on painstaking manual analysis. With YouTube now central to Brazil's digital debate, and with elections on the horizon, the team knew manual methods would no longer suffice.

"We realised we needed to better understand how these narratives circulate, who is targeted, what formats generate the most engagement, and how they reinforce each other," said Vitória Régia da Silva, Executive Director of Gênero e Número.

Unlike explicit hate speech, anti-gender discourse often hides behind humour, cultural references, or visual cues. Detecting and categorising these subtleties at scale was the challenge that led the team to adopt AI.

Building the solution: From concept to prototype

The collaboration followed an iterative path. Novelo Data led the technical side, testing and discarding early prototypes before arriving at a stable system. Each step helped refine the project's scope.

The workflow now runs through a modular pipeline hosted on Amazon Web Services: extracting metadata, transcribing audio, analysing transcripts, and classifying results. An open-source Python library developed by Novelo simplifies access to YouTube's API. On the frontend, a JavaScript interface built from Figma designs makes the platform usable for non-technical collaborators.

"We designed an integrated pipeline that transforms videos into audio, transcribes them, and then classifies their content using AI models," explained Guilherme Felitti, Founder of Novelo Data. "It's modular and scalable, which means we can adapt quickly."

Collaboration across disciplines

The project brought together two distinct but complementary worlds. From Gênero e Número, da Silva coordinated the editorial vision, joined by researchers and a designer shaping usability. From Novelo, Felitti oversaw data extraction and architecture, supported by specialists in classification, backend, and frontend development.

Bridging the technical – editorial divide was crucial. Felitti, who has a background in journalism, emphasised the need for open communication: "Technology can be intimidating. We made it clear from the start: ask us anything. We'll explain it as many times as needed."

For da Silva, this sense of alignment went beyond workflow: "We've worked with technical partners before, and it's often a struggle. But with Novelo, we felt aligned — not just on workflow, but on values."

Challenges along the way

- **Narrative complexity**: Anti-gender discourse is rarely explicit. The team had to build a framework to detect indirect strategies, identifying three key dimensions — themes, strategies, and subjects — to structure the analysis.
- **Project scope**: Initial ambitions proved too broad. "We don't do this because it's easy — we do it because we thought it would be easy," Felitti joked. The team decided to prioritise an MVP that "does one thing very well."
- **Workflow coordination**: With editorial and technical contributors working in parallel, coordination became as challenging as the technology itself.

The opportunities: What comes next for the team

For Gênero e Número, the project marks a new chapter: its first experiment with AI applied to feminist and anti-racist journalism. “This opens the door for us to integrate AI into journalism in Latin America,” said da Silva.

They also launched two products: the Gênero e Número AI Policy, and “Anti-Gender Discourses on Social Media: A Practical Guide to Recognize, Investigate, and Contextualize Anti-Gender Narratives” — an unprecedented publication that brings together concepts, references, and strategies for those seeking to understand how anti-gender narratives operate and how to investigate them in digital environments.

Looking forward, the team envisions a self-sustaining platform with features like keyword search, richer data visualisations, and granular filters. With elections approaching, the potential impact is clear: equipping journalists and civil society with better tools to track harmful narratives.

Lessons for newsrooms

- **AI needs human framing**: Algorithms alone do not explain disinformation. Editorial framing remains essential to give meaning to the data.
- **Disinformation is strategic**: The project reinforces that anti-gender narratives are not accidental. Patterns show coordination, timing, and intent.
- **Platforms fall short**: By surfacing harmful content, the project also underscores the responsibility of platforms to act.

Visit our website for the full case study and to learn more about this project.

> “Technology can be intimidating. We made it clear from the start: ask us anything. We’ll explain it as many times as needed.”

2.2 Newtral: Using AI to tackle disinformation on social media

Newtral is a journalism, fact-checking, and data verification organisation based in Spain. They have experimented with AI for many years, specifically to detect disinformation in political discourse. However, they soon noticed that the channels used for political disinformation were evolving. They were no longer coming from politicians' interviews and other political spaces alone, but had migrated to social media. They found that a lot of disinformation went viral on social media before it was debunked causing a chain reaction of a misinformed public.

The problem: Debunking disinformation on social media

One of the key issues that they faced while verifying and fact-checking social media data was the amount of disinformation that was prevalent on semi-private platforms like Telegram. The magnitude of disinformation was so high that manual-fact checking would have proved to be an endless task for their organisation.

"For example, here in Spain we had a natural catastrophe in Valencia in October 2024. There was a lot of disinformation around this specific event especially on Telegram. Our fact checkers were able to tackle some of the disinformation by following some channels. However, the [disinformation] wave was so big that many people started believing some information that was not real. So we wanted to help avoid such situations before it grew even greater," said Sara Estevez, NLP and Prompt Engineer at Newtral.

To help their team with such situations, Newtral engineers decided to tap into using AI technologies and create a tool they've named FactFlow AI.

Building the solution: How FactFlow AI works

FactFlow AI is designed for use by fact-checking organisations, including Newtral and others, to accelerate disinformation detection on Telegram. The FactFlow AI dashboard allows organisations to access and monitor potential disinformation, verify content, and also view what's already been verified, by showing relevant narratives. It also allows fact-checkers to customise and monitor different channels. Estevez explains that the idea in being able to customise the tool is to make it "modular" and give fact-checkers "an environment they can actually use on a daily basis."

FactFlow AI is currently in its internal testing phase at Newtral and will be ready for other fact-checkers in the next few months. In its internal testing and user feedback phase, the team received several requests for several features.

"This was a feedback given by the team because we initially were just shown the messages from Telegram which were potential disinformation. However, while it was very useful to them, they also wanted to know if some of these messages are already fact checked. The fact-checkers also annotated data for the AI models and so feedback from fact-checkers was always a part of the process. The fine tuning we did on the different models was actually really scoped under fact-checkers' criteria," added Estevez.

The opportunities: Iterating on the technology

FactFlow applies AI for three main purposes: 1) automated identification of suspicious channels; 2) detection of potential disinformation messages; 3) grouping related content

"The reason is that we wanted our models to have reasoning capabilities and we wanted them to be generative, and training that type of model from the beginning was not within our scope. We wanted to just fine-tune an already existing one in order to make it specific to our task," said Estevez.

They started with using a Microsoft model which was Phi 3.5 mini instruct. "It was working fine but in the last months a new generation of open-source Qwen models appeared and they appear to be the best on performance over the different dashboards," explained Estevez.

This led them to use the Qwen model since it provided them better results. "We are not using any closed source model like OpenAI here," she added. The tech stack included using multiple structural databases.

"We have the Telegram data entered on one big database then we are cleaning that data and putting it in another one [database]. We are then generating embeddings of that data so that we can do the claim-matching. We have different databases with the different data that we need," she explained.

The team and the challenges they faced

The journey was not without its challenges, including ensuring that the AI model performed "sufficiently well" without bias, performing as much annotation as possible for the model to perform correctly, and defining which exact Telegram channels to acquire data from.

"Selecting these channels at the beginning was a bit difficult and the hard work was done by fact checkers in selecting which channels have potential disinformation, that helped us," she added.

The team working on FactFlow AI comprised three parts. The first part was an AI team which developed the AI models, fine-tuned and also selected data for annotation. The second, consisted of fact-checkers for data annotation, testing, and feedback.

"They were the ones actually confirming that the process followed was done correctly and on fact-checking terms," explained Estevez.

Apart from this, they also received support from the software development team and a project manager, Diana Cid, who managed the entire FactFlow AI project. One of the ways they ensured all different departments coordinated and collaborated were by establishing strict deadlines and regular meetings to track progress.

“We had meetings in order to make sure that the timelines were followed as it was crucial for us. So I think that is where we put most of our efforts in trying to have meetings like every week or two weeks and putting big efforts to try to fit into the rest of our daily schedule,” added Estevez.

A challenge they also had to tackle was multilingual support. FactFlow detects potential disinformation based on the presence of “disinformation patterns”, in other words, linguistic cues commonly used when spreading disinformation. As the original LLM was trained in Spanish its performance will be better for similar languages like Italian or Portuguese than in, for example, Russian or Armenian. They look forward to collaboration with other fact-checkers from other countries to improve FactFlow capabilities. They also want to include a debunking capability to FactFlow AI functionality in the near future.

Lessons for newsrooms

Newtral’s successful implementation from data collection to disinformation detection offers several lessons for others looking to explore this space.

- **Iterate and pivot quickly**: Be prepared to explore different AI models and technology stacks and don’t hesitate to pivot when you find a solution that yields better results. Fast iteration is key to successful AI implementation.
- **Capitalise on interdisciplinary collaboration**: Foster strong collaboration within your organisation. This collaboration helped Newtral in two ways: they created tools that their users (fact-checkers) would actually engage with and make use of. It also enabled them to create more ethical AI tools considering how input and support from fact-checkers themselves centered how the tool would function.

Visit our website for the full case study and to learn more about this project.

2.3 Science Feedback: Building an AI system to combat climate misinformation

Science Feedback, a fact-checking organisation founded by climate scientists, has spent over a decade combating scientific misinformation across digital platforms. This French-based organisation, known for its rigorous evidence-based approach and network of scientific experts, has become a leading voice in verifying climate claims on social media and online news outlets.

Yet despite their extensive experience, Science Feedback recognised a critical blind spot in the fact-checking community's coverage: traditional broadcast media. This realisation led to the development of Climate Safeguards, an ambitious AI-powered project in partnership with Data For Good and QuotaClimat to monitor and analyse climate misinformation across French TV and radio news programmes.

The problem: Unmonitored misinformation at scale

While social media misinformation has received significant attention, traditional broadcast media – despite reaching millions daily – remained largely unchecked. Science Feedback identified this critical gap.

"French TV news programmes reach about two million viewers daily – the same influence as a viral social media claim," explains Charles Terroille, the project lead. "Yet we had no idea how much misinformation existed in these trusted channels or what form it took."

The scale proved daunting: 35,000 two-minute segments weekly made manual fact-checking impossible. Meanwhile, partner organisation QuotaClimat observed that regulatory responses to reported misinformation took months, allowing false narratives to spread unchallenged.

Building the solution: AI as a filter, not a replacement

The team's approach was deliberately measured. Rather than replacing human fact-checkers, they designed AI to act as an intelligent filter, identifying segments likely containing misinformation for human review.

"The goal wasn't to automate fact-checking, but to empower it," notes Terroille. "We needed to maintain credibility, and that requires human expertise verifying every detection."

The technical architecture combines a GPT-4o mini fine-tuned model that scores transcripts for potential misinformation, Whisper to enhance transcript quality for flagged segments, Label Studio for human annotation, and Metabase for visualisation and trend analysis.

Crucially, the team leveraged existing expertise. Their data scientist, Charlotte Samson, had previously developed a TikTok misinformation detection tool, while Data For Good provided infrastructure expertise. This in-house capability avoided outsourcing pitfalls, where technical teams often lack journalistic sensibility.

Learning through iteration

The project began with extensive manual analysis. “We fact-checked hundreds of cases manually to understand what we were dealing with,” recalls Terroille. This groundwork revealed that broadcast misinformation differs significantly from social media content.

“A model performing brilliantly on social media simply wouldn’t work for TV and radio,” explains Samson. “The speaking patterns, the transcript quality, the context – everything was different.”

The team discovered that comprehensive coverage mattered more than speed. As scientists, they wanted “the full picture, not just scattered detections” to make credible claims about misinformation trends.

Navigating unexpected challenges

Summer 2025 brought an unanticipated test. Climate debates intensified in France, and detected misinformation quadrupled overnight. “We went from about 100 detections monthly to 400, and it stayed that way throughout the summer,” Terroille recalls. “We never anticipated that workload.”

The team faced an ethical paradox: using energy-intensive AI to combat climate misinformation. Their solution was to filter for climate content first, then apply AI selectively. They’re in the process of developing lighter open-source alternatives, though these currently sacrifice some accuracy.

Expanding impact beyond France

The project’s replicable framework has enabled rapid international expansion – already operational in Brazil, with Poland and Spain launching soon. “Countries with smaller budgets can implement it,” explains Terroille. “Everything’s open-source except data costs and human coordination.”

Early findings reveal surprising geographic variations: Brazil shows significantly less broadcast climate misinformation than France. “This nuance is crucial,” Terroille emphasises. “Misinformation isn’t uniform across contexts.”

Creating new possibilities

Beyond detection, the system could one day enable real-time journalist support during interviews, and cross-country misinformation analysis. “We could provide channels with the top repeated false claims and dedicated fact-checks,” suggests Terroille.

As one of the few large-scale analyses of broadcast media misinformation, the project advances scientific understanding of how misinformation operates across media ecosystems while serving stakeholders – media, regulators, and the public.

Lessons for newsrooms

The Climate Safeguards project shows that responsible AI implementation requires:

- **Human-in-the-loop design**: AI amplifies human expertise rather than replacing it.
- **Ethical resource use**: Targeted application minimising environmental impact.
- **Contextual adaptation**: Solutions must account for medium-specific characteristics.
- **Sustainable architecture**: Replicable, efficient design enables broader adoption.
- **Comprehensive approach**: Understanding phenomena requires systematic, complete analysis.

As misinformation tactics evolve and amplification tools become more accessible, Science Feedback’s approach offers a template for combating false narratives while maintaining fact-checking’s human credibility. Their success scaling across countries proves that responsible, targeted AI can strengthen journalism’s defence against misinformation without compromising ethical principles.

Visit our website for the full case study and to learn more about this project.

> The goal wasn’t to automate fact-checking, but to empower it. We needed to maintain credibility, and that requires human expertise verifying every detection.

2.4 CCIJ: Unmasking disinformation during elections

The proliferation of disinformation poses a significant threat to democratic processes, particularly during elections. Newsrooms, especially in regions like Sub-Saharan Africa, often find themselves overwhelmed by the sheer volume of false information, lacking the real-time tools and resources to effectively combat it. This challenge was the driving force behind the Center for Collaborative Investigative Journalism's (CCIJ) development of ElectionWatch, an innovative AI-powered system designed to empower journalists in their fight against election-related disinformation.

The problem: Navigating a sea of election disinformation

CCIJ's journey to ElectionWatch began with a year-long investigation into Nigeria's 2023 presidential elections. This extensive work exposed a chaotic landscape of data manipulation, including doctored ballot boxes and fraudulent image uploads. What became clear was the immense amount of manual effort required to sift through this "craziness in the data," says Nelly Kalu, CCIJ's Editorial Project and Product Manager. The team recognised that much of this labor could be automated and that their experience offered valuable lessons for other journalists.

CCIJ's ElectionWatch originated from an investigation into Nigeria's 2023 elections, revealing widespread data manipulation. This highlighted the need for automation to manage the sheer amount of anomalous data". The core problem CCIJ aimed to solve was assisting journalists with disinformation investigations during elections. They expanded the definition of disinformation to include "malinformation" and malicious intent. Their goal for ElectionWatch was to strengthen election reporting and protect democracy from information weaponisation by tracking traditional falsehoods combined with malicious intent.

Building the solution: ElectionWatch takes shape

The solution to this pressing problem was an AI-powered system that functions as an AI agent. CCIJ deliberately chose this approach to avoid the risk of hallucination and misinterpretation, ensuring that the tool delivers precise, fact-based analysis. ElectionWatch analyses disinformation, maps out the involved actors and their networks, and creates timelines of events. This empowers journalists to quickly grasp narratives, identify key players, understand the chronology, and ultimately produce impactful reporting that counters disinformation. It shifts the focus from basic fact-checking to a comprehensive analysis of disinformation's real-time effects.

After securing a grant from JournalismAI, CCIJ embarked on the development journey. They began with extensive consultations, discussing the tool's design and data gathering strategies during elections. They met with various organisations, which had amassed a wealth of data from the Nigerian elections. These conversations helped define the scope of data collection, with a particular focus on scraping difficult-to-access platforms like Telegram and TikTok.

CCIJ then hired a UX/UI designer to map out the user journey. Through user research, their Minimum Viable Product (MVP) addressed 11 out of 13 identified pain points and needs of their target users: fact-checkers, journalists in Sub-Saharan Africa and small newsrooms globally. Our vision for ElectionWatch is to be a global tool used not only in the global south but newsrooms globally who may need its service. The emphasis was on creating a simple, step-by-step tool that provides useful analysis. Visual appeal was also crucial; CCIJ wanted ElectionWatch to be a "companion" that engaged journalists, making their work feel less like an additional burden.

Collaborations were vital. Through the extensive JournalismAI community, they connected with experts who guided the brainstorming and contributed to data analysis from TikTok and Telegram. For the DevOps aspect, CCIJ partnered with WebWorks, an African team whose understanding of regional nuances and data complexities, coupled with their willingness to follow CCIJ's lead, made them an ideal choice, says Kalu.

The technical build used Python, MongoDB Atlas, and Google Cloud services for the core stack. Machine learning and AI involved the Google ADK, Open Router, spaCy, scikit-learn, and various transformers to pre-train LLMs for disinformation classification, tokenisation, and entity/relationship recognition. Collaboration tools included GitHub, Slack, WhatsApp, Discord, Notion, and Basecamp. The front-end and visualisation relied on HTML, JavaScript, D3.js, and Plotly. Testing and quality assurance used pytest, Postman, and Cypress.

A significant aspect of the development was the integration of local languages. Recognising the importance of nuance, CCIJ hired journalist-translators for Igbo, Hausa, Yoruba, and Nigerian Pidgin English. This ensured the LLMs were trained to detect these specific linguistic nuances, a challenging but crucial step for the Kenyan WebWorks team. The team also worked with a couple volunteers to crack the TikTok and Telegram challenge. One of the volunteers came from Factiverse.

CCIJ structured their team to work in an agile manner, establishing a four-month build timeline. The CCIJ team focused on product and project management, data analytics, editing, and data development, while the WebWorks team provided technical expertise, including a technical project manager, technical lead, API data engineers, machine learning engineers and interns, backend and frontend developers, and designers. Editorial leads handled narrative framing, translation, and reporting, with a QA lead and community testers ensuring quality. The team's agility and commitment were paramount, especially given the rapid development cycle.

The opportunities: Expanding ElectionWatch's reach and impact

The future holds exciting opportunities for ElectionWatch. CCIJ's immediate goals include enabling real-time scraping and live fact-checking. This means highlighting existing fact-checks around disinformation, rather than generating new content, to avoid the risks of hallucination. They also aim to strengthen the governance and ethical frameworks surrounding the tool and enhance its real-time interaction capabilities with journalists.

A key development is the creation of a robust "data engine" – a cleaner, more deliverable, and referenceable data backbone for newsrooms worldwide. This engine will incorporate metadata, timestamps, and geotags. CCIJ also plans to develop an AI newsroom kit, evolving their methodology into a comprehensive "how-to" guide and offering training for newsrooms on using ElectionWatch.

While mindful of the risks, CCIJ is exploring the potential for generative AI at the very last stage, only for summarising completed analysis, emphasising their commitment to avoiding unrequested or hallucinated content. The ultimate goal is to transform ElectionWatch from a data analyst collector into a full-fledged, automated tool that robustly supports journalists.

Lessons for newsrooms

- **Embrace humility and iteration**: Approach data, disinformation, and machine learning with humility and a willingness to learn. Adopt an iterative approach, unafraid of mistakes, by accepting to "Let it be ugly... Let's see that it's not working" to speed up development.
- **Recognise linguistic nuance "translators"**: Acknowledge the crucial need for specialised experts ("translators") to bridge the gap between technical teams and local context to accurately capture linguistic nuances in research and data.
- **Cultivate active inquisitiveness**: Product managers must be actively inquisitive, striving to understand why things fail even in areas outside their primary expertise. Attention to detail is vital for effective problem-solving.
- **Maximise budget for future funding**: Newsrooms, especially those with limited resources, should maximise the demonstrated potential of their current investment to clearly justify and secure future funding.
- **Foster partner and team belief**: Belief in the product among partners (like WebWorks and Factiverse) is essential for deep investment. Small teams should leverage all available tools, learning resources, and partnerships to compete and create impactful tools, proving that a small team can build a global product.

Visit our website for the full case study and to learn more about this project.

3. AI for Audience Engagement and Retention

The organisations mentioned in this section have used AI in a wide gamut of audience-centred use cases from retention, to acquisition, and most interestingly, to continually engage their communities.

Makedonia, a legacy, local news organisation in Greece, turned to AI-powered personalisation and push notifications to engage its readers in the digital age. Daily Maverick, an independent South African newsroom, built a 360° revenue strategy powered by AI-tools and with a focus on community engagement at its heart. Más Voces, on the other hand, built a tool informed by its community's needs to understand Cuba's informal markets. Sowt in Jordan leveraged AI to boost subscriptions by analysing customer behaviour, personalising marketing and predicting churn. Lastly, in Germany, Digitalhaus Franken, launched an ambitious project to reduce subscriber churn, deepen loyalty, and engagement.

One of the clear shifts that these grantees exemplify is the growing use of AI to support audience engagement initiatives and not just audience development. While audience development refers mostly to user acquisition, engagement refers to building long-lasting, meaningful relationships with existing audiences while having the potential to attract new ones.

Some of the main lessons from these organisations are to reject a one-size-fits-all approach to audiences. Others like Más Voces spoke about the challenges of coordination while working in exile. Daily Maverick and Digitalhaus Franken shared the importance of working across siloes in interdisciplinary teams and the project management required for such scenarios. Many of the organisations also said that it was important to have deep technical expertise while building out complex engagement tools.

3.1 Makedonia: Reviving regional journalism through personalised community engagement

Makedonia, a traditional regional newspaper in Greece with nearly a century of history, faced a crisis familiar to local media outlets worldwide: how to remain relevant in an era where social media delivers instant neighbourhood news and younger generations increasingly avoid traditional journalism. Their innovative response – an AI-powered personalised newsletter system – offers valuable lessons for regional media grappling with similar challenges.

The problem: When local news loses its local advantage

Regional and local media across Europe are grappling with an existential challenge. The challenge is particularly acute for outlets like Makedonia, which historically filled a clear gap by covering local neighbourhood stories that major media ignored. As Nikolaos Panagiotou, the project's lead researcher, explains: "Two decades ago it was very obvious that the local or regional [news organisations] were actually covering a gap regarding what is happening in my neighbourhood. But what about now that everything can be easily reached through social media?"

The challenges are multifaceted. "AI is considered to challenge the viability of media, but most importantly it presents a major challenge in how to integrate AI into current newsroom day-to-day operations," Panagiotou notes. "If this is a major issue for big media, you can imagine that for regional and local media, it is an even bigger problem, because they lack the resources, the knowledge, and the relevant staff."

Andreas Panagopoulos, the technical lead on the project, adds another dimension: "Journalistic staff at Makedonia is quite small, and most of them lack training not only in AI but also in digital tools."

Building the solution: Finding opportunity in a crowded market

The project emerged when Makedonia's management decided to relaunch their website and increase newsletter subscriptions as a pathway to print subscriptions. However, the team quickly realised this required more than traditional digital marketing approaches.

The breakthrough came from recognising that personalisation could address multiple challenges simultaneously. Research revealed a gap in the market: "Our research showed that very few outlets are using this system. Perhaps major international media, but certainly nothing in Greece, and nothing in the broader region – the Balkans, Turkey – for personalised newsletters."

The team identified AI's potential to automate routine tasks while enhancing rather than replacing journalistic work, leveraging AI to deliver "specialised content to targeted audiences".

At the heart of the project is the development of an advanced Personalised Newsletters and Push Notifications (PNN) system. By analysing user behaviour and preferences, the team expects this AI-powered tool to enable Makedonia to deliver content that is tailored to each reader's interests. The goal is simple yet powerful: to offer a more meaningful, personalised news experience that builds loyalty and trust over time.

Building the technical foundation

The tool Makedonia built combines technical innovation with editorial oversight. The infrastructure integrates ChatGPT for content analysis, Mailchimp for distribution, and a custom WordPress-based content management system.

When users interact with newsletter content and indicate preferences, AI provides feedback to ensure the next newsletter is delivered "without human intervention". However, the system includes a crucial editorial safeguard: "We're working on a special feature where every newsletter will include an 'editor's pick' – two or three articles – to avoid the filter bubble effect and give users what journalists feel they need to know, not just what they want to know," Panagopoulos explains.

The technical implementation required integrating university AI experts with web developers to create optimal user experiences for both readers and journalists. The team planned to test multiple LLMs, starting with ChatGPT, but expanding to include Gemini and possibly others.

Bridging the newsroom divide

The project's greatest challenge wasn't technical – it was cultural. "What was more challenging was integrating the AI experts with the journalists – getting AI experts to understand journalists' needs and, on the other side, getting journalists to understand what the AI experts were doing," Panagopoulos notes.

The team organised four comprehensive seminars to address this gap. "Journalists have fears: 'We're going to lose our jobs and AI is going to replace us.' So we had to start from the beginning, explaining the goal of this specific project and how AI can be their assistant, not something that will take their jobs."

Success required creating common vocabulary through "a small glossary so everyone could understand what we mean when we talk about LLMs, what an LLM is, how LLMs work", and using appropriate benchmarks. The team learned to avoid inappropriate comparisons: "If you present journalists at Makedonia with the New York Times example, they'll say 'this is the New York Times' and immediately reject it."

The opportunities: Scaling beyond Thessaloniki

The Makedonia project represents more than a single outlet's digital transformation – it offers a potential template for regional media across Southern and Eastern Europe. "We think that this tool will be an opportunity for expanding to other media outlets, probably to have a startup on that," Panagopoulos explains.

The team sees expansion possibilities once "the algorithm is trained from everyday use". Future enhancements include testing headlines to optimise newsletter open rates and extending personalisation to push notifications. The broader implications extend beyond technology: "Newsletter subscribers and users are the more loyal audience for every website all over the world."

Lessons for newsrooms

Makedonia's approach suggests three critical insights for regional media:

- **Prioritise easy adoption**: Develop solutions that are designed to be "easily introduced and adopted" by newsroom staff, minimising disruption and training overhead.
- **Offer implementation support**: Provide comprehensive support for implementation and change management, not just the technical tools themselves, to ensure successful integration.
- **Reject "one size fits all" models**: Avoid generic solutions and instead reject **approaches that ignore vital local contexts and unique regional needs.**
- **Balance AI with editorial control**: Create a sustainable model by leveraging AI for efficiency while firmly maintaining human editorial control. This allows outlets to expand reach without compromising their local mission.
- **Use AI for personalised local value**: View AI-powered personalisation as an opportunity to create appealing content that helps preserve the local focus, particularly the distinctive value of regional journalism, as traditional subscription models face limitations.

Visit our website for the full case study and to learn more about this project.

3.2 Daily Maverick: Using AI to boost its membership strategy

Choosing a voluntary membership model in reader revenue comes with a plethora of problems. Daily Maverick, the South African newsroom, has faced many of them head-on. The three key components of a membership funnel include acquisition, engagement, and retention. The Daily Maverick team found that while they had strategies in place for all three components of the funnel, with a growing audience, these strategies required consistent iterations.

The problem: Moving from reader interest to action

There was a significant drop-off from people clicking to convert and those who actually completed their membership signup. While the global industry standard for paywalls is 11 per cent, Daily Maverick's voluntary membership was sitting at 3.5 per cent, according to Fran Beighton, Head of Growth at Daily Maverick.

"When we dug a little deeper, we saw that often our messaging was so emotional by the time they [readers] got to our standard landing page, something didn't resonate with them and they failed to complete (the signup)," explained Beighton.

Equally important to them was the focus on community to anchor their membership, said Beighton. "Community can not be a one-way street; Daily Maverick has benefited from our members so much but they need to also feel how the community can change their lives and businesses," added Beighton.

Building the solution: A 360° view of membership

Keeping this in mind, the team came up with an idea for a holistic membership strategy, aided by AI, to target the three components of a membership strategy, and with a renewed focus on community and engagement.

It led to optimising the landing page for acquisition, maximising engagement by creating the Daily Maverick Connect platform, and an Impact Tracker tool to demonstrate the impact of their journalism. The team behind the Revenue 360 project included the reader revenue team, audience development team, head of operations, and the editor-in-chief, said Beighton.

First, the landing page

A crucial element of the entire puzzle was in optimising their landing page for conversions. "Even with a voluntary membership, we had an added issue. There's no impetus to absolutely 100 per cent signup. It's not like it's an e-commerce store where you want to buy a pair of trainers and you find other offers or testimonials. That element of urgency for the potential member isn't there. So, we have to be really specific and careful about what is on the landing page and what's just a waste of time," said Beighton.

They added heatmaps to the page to find out and discovered that 75 per cent of the people didn't go below the fold. "What's really interesting there, is that it showed us where we need to concentrate the information, where we need to add behavioral nudges, to explain more, while still keeping it clean and not overwhelming," added Beighton.

They changed the landing page from a static page that encompassed all ideas and started delivering tailored messages to specific audience targets. Initial testing showed a 55 per cent increase in conversion. It also led to ongoing iterations and improvements.

The team built a custom A/B testing module specifically for this, which intercepts traffic to ab.dailymaverick.co.za and then redirects to set URLs based on the parameters set, with a tracking component. For the landing pages themselves, they integrated the pages into their Wordpress/Woocommerce setup to enable integration into their Maverick Insider system, and included Hotjar and Heap tracking.

Next, fostering community connections

The team also recognised the importance of deepening meaningful engagement with their community to eventually help in more acquisition and retention, said Beighton. This is where the Daily Maverick Connect social networking platform comes in.

"We want our members to be able to network with each other, whether that's finding new clients or new jobs or a new book club. That could exist by itself, but the purpose behind it is to engage with each other and with Daily Maverick," explained Beighton. The tool, which was in beta-testing with staff at the time of this interview, will be first opened to an initial closed group of members before wider expansion.

The biggest challenge they faced while creating this platform was in continually engaging users without it becoming addictive, intrusive, or overwhelming. Beighton confirmed that the team has many safeguards in place to avoid toxicity, specifically by training an AI bot to prevent misuse and misbehaviour. The tech stack for the Connect platform includes using Discourse and Docker.

"We are using ModerationAPI's AI model which we tuned based on its response to test inputs in order to encourage the behaviour on the site that we want to promote," said the team. "The AI Moderation tool we have implemented, will allow us to moderate comments at scale, with minimal human intervention, by passing them through custom prompts and training the system on what we regard as acceptable and non-acceptable comments. This will allow DMC to perform optimally, and conversations to happen on the platform without a slow and cumbersome moderation regime," they added.

Last, impact for retention

The team also created an impact tracker tool, ImpactEngine, that uses an LLM trained on past stories to scan parliamentary notes (Hansard) and identify how Daily Maverick's journalism effects change in parliamentary meetings. This tool serves as a research asset for journalists, allowing them to track the progression of stories, find experts, and design their reporting with intentional impact. It also generates reports on the effect of specific teams' work which is then shared with readers to demonstrate the organisational impact.

"We didn't anticipate how useful it would be for the very construction of journalism," said Beighton. The tool will also be used to cover local committee meetings, which Beighton says is crucial for a "resource-constrained" industry without enough people to cover towns and municipalities with "boots-on-the-ground".

The technology behind the ImpactEngine include SvelteKit, Ollama (Mxbai-embed-large) OpenAI (GPT 4.1), MongoDB and Qdrant. Most of the "AI" logic is built around vectorisation, so LLMs are not particularly important for this project, shared the team. They selected GPT 4.1 as it has a large context window and is good at creating the types of reports that ImpactEngine outputs. The biggest challenge they faced was in terms of time taken to ingest data cleanly into the system, well-chunked, and with sufficient metadata to be useful.

Lessons for newsrooms

- **AI empowers smaller newsrooms**: Recognise the significant possibilities AI offers to smaller news organisations, specifically its power to fast-track project development and overall technology adoption.
- **Value interdisciplinary teams**: Consider it a "huge mistake" to approach AI projects without an interdisciplinary team. Since the first users of internal-facing tools are journalists and staff, breaking down silos and collaborating across the board is essential for success and user-friendly design.

Visit our website for the full case study and to learn more about this project.

3.3 Más Voces: Building on market needs' shoulders

Imagine an ecosystem in which one of the main leads that individuals need to make decisions is missing. Or, well, maybe not missing but scattered and distributed in so many pieces that putting it together takes an immense amount of time and effort, eroding the possibilities to actually act in favour of your interests.

Más Voces developed elTOQUE Markets as an AI-powered solution to help stakeholders from within and outside Cuba understand the island's complex informal economy, particularly currency rates – a key challenge for readers of their main project, elTOQUE. This subscription-based data platform provides structured, real-time information on exchange rates, remittance costs, and food prices sourced from fragmented and unregulated channels, transforming what was initially identified as a content gap into a new potential revenue stream.

Operating in one of the most opaque information environments in the region, the project leverages artificial intelligence to address a fundamental lack of economic transparency while also testing a path to financial sustainability for a newsroom in exile.

The problem: Why elTOQUE Markets was born

In Cuba, the informal economy isn't a peripheral phenomenon: it's a foundational system that structures the interactions at every level. Due to the government's refusal to publish transparent economic data and the near-total absence of a functioning formal exchange market, economic actors depend on informal systems to survive and understand which is the most efficient current exchange rate.

From currency exchange and remittances to food distribution, unregulated systems underpin daily life. But reliable, real-time information about these markets is scattered, unstructured, and hard to access, especially for entrepreneurs, diaspora investors, or anyone trying to make informed decisions.

Building the solution: From tracker to product

The roots of elTOQUE Markets lie in the outlet's earlier project: a daily informal exchange rate tracker. Building on that success, the team designed a more comprehensive platform aimed at a wider range of economic indicators, with the potential to be monetised.

The project began with audience and market research to validate user needs, followed by the design of the platform's information architecture. With this foundation, the team developed a MVP in the form of a paid newsletter, while also building the CRM and integrating payment systems to support its operations. These steps laid the groundwork for scaling toward the launch of the full web platform.

Tools and technologies

To collect and process the unstructured data behind Cuba's informal economy, the team combined:

- Data capture bots deployed in WhatsApp, Telegram, and Messenger groups
- A NLTK-based model to extract and clean raw textual data on currency exchange offers
- LLMs like ChatGPT to interpret slang, typos, and evolving informal terms
- A CRM system to manage subscriptions and users
- MongoDB with separate collections to organise exchange data, product listings, and remittance info
- And a global vector database on top of MongoDB that is used for semantic searching and unify structured and unstructured data via enhanced microservice APIs.

The team behind the platform

The team is spread across Cuba, Spain, Chile, Ecuador, and the U.S., working under conditions shaped by exile and funding instability. Key roles include a product manager, a lead developer, AI engineers and data specialists, and CRM and frontend/backend developers.

Challenges faced

- **Structuring the unstructured**: Cuba's informal markets exist primarily in private or semi-private digital spaces, like Telegram and Facebook groups. The language used in these spaces is informal, inconsistent, and constantly evolving – requiring ongoing fine-tuning of AI models.

  It also involves a great deal of normalisation efforts to uniform concepts: "There's no universal way people spell the word 'Zelle' in these groups. It might be 'cel', 'celle', 'celly', or something else entirely," explains Nieves.

  The team also had to adapt to frequent changes in where and how economic information appears, as groups open, close, or shift platforms.
- **Payment restrictions due to the embargo**: USA sanctions block financial platforms like Stripe and PayPal from processing Cuba-related transactions. To bypass this, the team relied on QvaPay, a peer-to-peer cryptocurrency marketplace, which allows users in Cuba to purchase prepaid cards using local methods – a workaround that remains fragile.
- **Remittance data still requires human input**: Because many remittance providers operate informally – often just via a phone number or chat group – their data is inconsistent. The team still relies on some manual inputs for these sources, later structured by AI.

- **Operating from exile**: Developing this tool while in exile brings challenges around coordination, safety, and impact. Despite limited resources and the loss of half the team due to funding freezes, the project continued. Nieves assures: "This experiment has kept us going. It gave us hope when we had almost nothing left."

The opportunities and learnings

The project offers a replicable model for data in opaque economies, showing how methods used in Cuba, such as collecting unstructured data from messaging apps and social networks and structuring it with AI, can be adapted to other countries with large informal markets, like Venezuela or Argentina, to generate actionable economic information.

At the same time, it represents a new mindset for exiled media. For Más Voces, the initiative is not just a tool but a shift in sustainability strategy: "Most exile media rely on grants. But funding is drying up. If we can make this work, it could be our way to survive."

Lessons for newsrooms

- **AI can unlock public value in opaque environments**: By structuring unregulated data, newsrooms can inform both local and international actors – from businesses to policymakers.
- **Tech solutions must adapt to local languages and practices**: Slang, typos, and informal channels challenge standard AI tools. Continuous training and human oversight are critical.
- **Media innovation must include sustainability**: elTOQUE Markets is as much a product experiment as a business one. It shows how media can build revenue by solving real user needs – even in restrictive contexts.

Visit our website for the full case study and to learn more about this project.

> This experiment has kept us going. It gave us hope when we had almost nothing left.

3.4 Digitalhaus Franken: Using AI to predict subscriber churn

Digitalhaus Franken, a German digital media house, faced a challenge familiar to many publishers navigating the shift from print to digital: while they successfully attracted new digital subscribers, keeping them proved difficult. This led to Project PULSE, an AI-powered initiative to predict and prevent subscriber churn.

The problem: High acquisition, poor retention

As digital subscriptions became increasingly central to their business model, Digitalhaus Franken discovered a troubling pattern. “The beginning of our funnel worked quite well – we had lots of conversions,” explains Lisa Riech, the project lead. “But we struggled to keep subscribers long-term.”

The company recognised a fundamental truth about subscription businesses: retaining existing customers costs significantly less than acquiring new ones. They’d addressed retention before, but never with AI’s predictive capabilities. “This was our first time using prediction models,” Riech notes. “We wanted to identify customers at risk of churning before they even realised they wanted to cancel.”

Building the solution: From consensus to action

Unlike many organisations hesitant about AI adoption, Digitalhaus Franken found immediate buy-in across the company. The proliferation of various generative AI tools had already familiarised staff with AI’s potential.

“Everyone was on board,” recalls Riech. “We all had the main aim of delivering an MVP, but equally important was learning how to manage machine learning projects.”

The team distinguished early between what they called “classic AI” (machine learning and deep learning) and generative AI, helping set appropriate expectations for this prediction-focused project.

The new workflow

Project PULSE short for Predictive User Loyalty System for Engagement, unfolded across several carefully planned milestones. The first phase, lasting until mid-April, focused entirely on data exploration and integration. The team needed to understand what data they had, which systems housed it, and what would actually improve their model’s predictions.

“We learned that throwing all available data into a machine learning model doesn’t necessarily produce better results,” Riech explains. “You need to experiment with different data combinations.”

The team tested three machine learning models: XGBoost, Random Forest, and Support Vector Machine. After extensive experimentation, they selected XGBoost for its superior explainability. “With XGBoost, we can examine predictions for each customer ID (anonymised) and see how changing one value affects their churn risk,” Riech notes.

To make the predictive insights actionable, the team decided to visualise both the input data and the calculated churn probabilities in a Power BI dashboard. This approach allows the entire customer base to be segmented according to their churn risk, ranging from “very low” to “very high.” In addition to current predictions, the dashboard also presents historical data, illustrating how customer engagement – such as pages read or time spent – along with churn probability, has evolved over time. This perspective is crucial for evaluating the effectiveness of marketing initiatives aimed at reducing churn, enabling the team to measure impact and adjust strategies based on real results.

The technical team remained deliberately small: a project lead, data manager, two data engineers, and the head of AI, supplemented by external partners who provided two data scientists and one data engineer. This external expertise proved essential, handling approximately 50-60 per cent of the technical work, particularly the model selection and experimentation phases.

What worked well

Communication, often a stumbling block in technical projects, functioned smoothly. Weekly meetings with external partners created accountability and momentum. “They asked questions every week, and we needed good answers,” Riech recalls. “This pushed us to stay prepared and move quickly.”

The decision to maintain a small, technically focused core team proved crucial. Initially, more people from different departments were included, but the team soon recognised the efficiency gains from keeping the group homogeneous. “When a project is predominantly technical, ensuring clear communication with non-technical colleagues involves considerable effort in translating technical concepts. Working with a fully technical team has allowed for faster execution and more efficient workflows,” Riech observes.

Data quality, an anticipated challenge, proved sufficient despite initial concerns about not having “millions of customer records”.

Building sustainable capabilities

Crucially, the external partners built a foundation the internal team could maintain independently. “They established the model, and afterwards we set up the initial clusters and values. Now we just feed in new data – it’s a ten-minute process,” explains Riech. “We only need external help if we want to fundamentally change the model.”

This sustainability extends to the technical infrastructure. The internal data engineering team handled all data formatting and provision throughout the project, ensuring they could continue operations post-launch.

The opportunities: A roadmap for media organisations

Project PULSE succeeded in creating a functional churn prediction system, but as Riech emphasises: "The project provides a good foundation, but now the real work with customers begins."

The marketing team must now leverage these predictions to craft targeted retention campaigns.

For media organisations considering similar projects, Digitalhaus Franken's experience offers an encouraging message: with clear objectives, the right expertise, and strong project management, AI-powered subscriber retention is achievable – even without millions in customer data or massive internal AI teams.

Lessons for newsrooms

The project revealed three key insights for Digitalhaus Franken and other media organisations considering similar initiatives:

- **Need dedicated data science capabilities**: "We have plenty of people with data engineering and management backgrounds, but we're actually missing the data science component," Riech acknowledges. "If you want to pursue prediction AI projects on a large scale, you need this expertise in-house."
- **Tap "classic AI" for prediction**: Machine learning offers untapped potential beyond generative AI's current dominance. "There's significant value in 'classic AI' for prediction models that we're not fully utilising yet," notes Riech.
- **Prioritise exceptional project management**: Clear roles, defined timelines, and focused teams matter more than having cutting-edge technology.

Visit our website for the full case study and to learn more about this project.

3.5 Sowt: Transforming data into actionable insights

In today's fast-paced digital landscape, publishers face a significant challenge: transforming vast amounts of raw data into actionable insights. Many newsrooms collect extensive data on audience engagement, content performance, and operational metrics, yet often lack the necessary tools and processes to effectively analyse this information. This deficiency can lead to missed opportunities, inefficient resource allocation, and an inability to make data-driven strategic decisions.

The core problem lies not in the scarcity of data, but in the struggle to move beyond mere collection and actively leverage it to inform and enhance every aspect of their operations.

The problem: Bridging the data insight gap

Newsrooms today often grapple with an overload of data, struggling to convert it into meaningful insights. They frequently possess extensive audience engagement metrics, content performance data, and operational figures but lack the necessary tools and processes for effective analysis.

Sowt Media, a podcast company in Jordan, faced this challenge directly. As Head of Marketing, Diala Otaibi highlighted a key difficulty: "Coming up with decisions specifically for marketing campaigns to be able to measure our efficiency in terms of getting subscribers."

Traditional campaign evaluation involved a "two-day period where we would need to wait to measure the effectiveness of the campaign," says Otaibi. Resulting in "a lot of wasted time on campaign understanding and resources as well," Otaibi added. This manual, time-intensive approach underscored the critical need for a more efficient and data-driven solution.

Building the solution: Unlocking key benefits and results

Sowt's strategy began with robust data modeling, creating a unified and comprehensive view from previously siloed datasets. This involved identifying critical data sources, defining relationships, and ensuring data quality and consistency.

Otaibi explained: "We collected all the data and we integrated them through BigQuery into Looker."

This integrated data included "Google ads, podcast analytics, subscription data, audience demographics," all consolidated into "one main dashboard and one main tool," she said.

The team then focused on transforming this modeled data into actionable insights via interactive dashboards. This collaborative process involved engaging directly with Sowt stakeholders to gather business requirements and pinpoint critical metrics and Key Performance Indicators (KPIs) for visualisation.

Prototypes were developed for review and feedback, allowing for an iterative process that ensured user-friendly dashboards aligned with stakeholder workflows. The final dashboards, meticulously built in Looker Studio, featured tailored visualisations for specific business needs. These interactive reports empower users to explore trends, drill down into data points, and gain deeper insights. Rigorous testing for data accuracy and full functionality preceded delivery, guaranteeing reliability.

At the core of the solution was an "AI model that provides accurate performance predictions within a 24-48 hour campaign launch," according to Otaibi. This "prediction engine" allows Sowt to "forecast how my marketing campaign is going to go in the next week, or two weeks ahead."

Otaibi says the tool also analyses "the video title, if this title is doing well or not doing well, based on previous historical data" and "categorises the audience into segments or into countries and therefore it gives you the subscription impact as well." Crucially, it provides a "recommendation," advising whether to "keep the campaign" based on specific objectives.

The opportunities

The implementation of these interactive dashboards opened up a wealth of opportunities for Sowt. They now possess a clear, real-time understanding of their audience's behaviour and content performance. This translates into several key advantages:

Their improved content strategy, driven by data on audience resonance, will lead to increased engagement, and a more loyal audience. This deep understanding of demographics, preferences, and consumption patterns will enable personalised content delivery and targeted marketing.

Optimised operational efficiency will result from data on production costs and content performance, allowing Sowt to identify inefficiencies and save costs. This faster optimisation facilitates immediate decisions without manual hassle, leading to a significant increase in marketing ROI.

Sowt Media is also looking ahead, with plans to "continue growing." Future aspirations include "creative content analysis," and exploring "multi-channel attributions and more competitive intelligence features."

Lessons for newsrooms

Sowt's journey offers valuable lessons for newsrooms looking to leverage data effectively.

- **Prioritise cross-organisational collaboration**: Involve stakeholders from across the organisation from the outset, understanding their specific needs and pain points to build solutions that deliver real value.
- **Define business-first requirements**: Start by clearly defining critical metrics and KPIs that matter most to your business objectives, ensuring dashboards are relevant and actionable.
- **Embrace iteration and continuous feedback**: Develop prototypes and seek continuous input to refine solutions and meet evolving needs.
- **Invest in the right tools**: Select powerful, flexible, and user-friendly data visualisation tools like Looker Studio, aligning them with your technical capabilities and business requirements.
- **Focus on actionable insights and accuracy**: Present data in a way that enables users to draw meaningful conclusions and take specific actions. Ensure data accuracy and reliability through rigorous testing, as user trust is essential for informed decisions.
- **Foster a data-driven culture**: Actively work to value, understand, and use data to inform daily operations and strategic planning across the entire newsroom.

By following these principles, newsrooms can emulate Sowt's success in transforming data into a powerful engine for growth, efficiency, and deeper audience engagement.

Visit our website for the full case study and to learn more about this project.

> At the core of the solution was an AI model that provides accurate performance predictions within a 24-48 hour campaign launch.

4. AI for Content Repurposing

These organisations have demonstrated that news organisations are globally evolving into technological innovators by leveraging AI to repurpose content into various formats. The Republic, and the International Centre for Investigative Reporting (ICIR) built tools that integrated indigenous languages often overlooked by AI platforms in language translation and voice recognition in Africa. Scroll, and Nawaat built tools for interactive content repurposing, while Babel focused on AI avatars.

In building these tools, the organisations had cross-cutting challenges. These ranged from concerns about the costs of AI tools and the long-term viability of their projects. Some of the organisations are considering selling their products to other newsrooms. Coupled with this is the concern of infrastructure limitations. For example, The Republic experienced a basic infrastructure gap such as custom engineering for payment processing in Nigeria, and poor internet connectivity in some regions, which caused additional barriers to building the tool.

Some of the challenges they experienced were within their own organisations. Newsrooms encountered internal resistance from staff who were concerned about job displacements, and a “conservative culture” regarding new technologies. Another common thread across many of the grantees was in seamless communication and coordination between editorial teams and technical staff. Some of these organisations needed to collaborate with external technical staff in order to build their tools, many had also built custom technical stacks or adapted existing tools towards their respective unique challenges.

4.1 Babel: Sustaining media operations during crisis

For small newsrooms operating in crisis zones, maintaining consistent content production becomes a matter of survival. The Ukrainian digital media outlet Babel discovered this reality firsthand when their hypothetical emergency scenario became an urgent operational challenge, leading them to pioneer AI-assisted content creation as a solution for staff shortages.

The problem: The challenge of sustaining journalism during national emergency

In early 2024, Ukrainian media outlet Babel's hypothetical grant scenario became reality: what happens when journalists are mobilised for military service?

"We outlined a hypothetical situation in our submission and then it happened", the project lead Svitlana Moskalenko recalls. Historian Serhiy Pyvovarov – whose expertise is central to Babel's YouTube series That's the History – was called up for military duty, leaving the newsroom scrambling to maintain their production schedule with reduced capacity.

So how did Babel leverage AI voice technology to sustain their content production during a national emergency? Here are valuable lessons for small media organisations operating under extraordinary circumstances.

Building the solution: From crisis to innovation

The newsroom's initial vision was ambitious: create AI-powered video content featuring avatars to maintain consistent programming even with reduced staff. "We knew the product we wanted to see. It was in the video format", explains Moskalenko. Their approach was comprehensive, assembling a toolkit that included:

- Synthesia for avatar creation
- Replit for coding and automation
- Canva for photo and video editing
- Paddle for captions
- Eleven Labs for voice generation.

However, the team quickly discovered that not all AI tools are created equal for journalistic purposes.

The Synthesia roadblock: When AI meets editorial reality

The project hit its first obstacle when Synthesia, their chosen avatar platform, refused to process their content due to restrictions on historical and political themes.

This forced a crucial pivot from video avatars to AI-generated voice narration – a decision that proved both challenging and enlightening.

"We need to keep up with production. If we do not publish, we lose our audience," the team realised. The pressure to maintain audience engagement drove rapid adaptation.

Building the team: Skills for an AI-powered newsroom

One surprising discovery was that implementing AI didn't require hiring new technical specialists. "The good thing about AI tools is that you don't need extra skills. You just need to be a little sharp to make it work," notes Moskalenko. The team structure remained largely editorial:

- **Editor-in-chief**: Provided backup support and helped with text editing
- **Designers**: Adapted their visual skills to work with AI tools
- **Two editors**: One for voice, one for video
- **Programmer**: Handled coding and automation needs
- **Remote contributor:** The mobilised staff member who could dedicate limited hours.

"At its core, it's still editorial work. We didn't add anything beyond what's already part of our daily newsroom agenda," Moskalenko emphasised, highlighting that AI served to enhance, not replace, traditional journalism practices.

Time, tone, and technology: Unexpected challenges

While the technology worked, implementation brought unforeseen complications:

- **Time investment**: "We thought AI tools would be easier but they're really time consuming," the team discovered. A single piece required over a full day of editing: "Because of all the names and dates, it has to be generated repeatedly until all the stresses are correct."
- **Tonal mismatches**: The AI's limitations in matching tone to content created jarring contrasts. "We have feedback that the content is so sad and the voice is so cheerful. So, do something with your host. He's too cheerful for sad things," audience members complained. This challenge was later addressed when ElevenLabs released an update that smoothed out tonal mismatches.
- **Financial sustainability**: "Most of these programmes aren't cheap," acknowledges the project lead. As the grant ends, the newsroom faces hard questions about viability: "We're thinking about changing the project architecture to make it cheaper and possibly continue without donor support."

The media literacy crisis: An unexpected discovery

The most surprising finding was that audiences couldn't distinguish AI-generated voices from real ones, even with clear labeling. "People don't realise the voice is artificial. We include a notice about AI use, but no one reads it," the team noted.

This highlights broader awareness challenges. As Moskalenko notes: "Media literacy needs to evolve. We assumed everyone could tell the difference [between AI-generated and real content]."

The opportunities: Promise amid technology-literacy concerns

Despite challenges, the newsroom sees potential in their pioneering work. Positive audience feedback validated their experimental approach – "We've received good feedback, and honestly, we expected worse."

However, concerns remain about the gap between advancing technology and public understanding. "So many tools are evolving quickly, but only a small percentage actually use them," they observed.

Lessons for newsrooms

- **Technology evolves faster than implementation**: "The pace of technological development is unbelievable. It changes from month-to-month," the team observed. Tools that were ineffective at the start of the project had become viable options just a few months later.
- **AI as agent, not specialist**: "AI applications, they're more like agents as opposed to things that require heavy technical knowledge", making them accessible to traditional newsrooms without extensive retraining.
- **Community knowledge gaps**: The lack of established knowledge-sharing practices within the journalism community proved challenging: "Most newsrooms tend to work in a fairly isolated way, and there's no real system for sharing experience within the community. "We didn't really have anyone in the country we could ask about their experience." This isolation made planning and troubleshooting more difficult.
- **Conservative culture as a barrier**: "Journalists tend to be conservative, and it's hard to get them to work with new technology," Moskalenko noted, reflecting the broader industry resistance to change.

Visit our website for the full case study and to learn more about this project.

4.2 The Republic: Building AI for African Voices

When Nigerian digital media platform The Republic struggled to find text-to-speech tools that could pronounce African names correctly or capture the cadence of local languages, they decided to build their own solution. The result is Minim, an AI-powered platform that promises to revolutionise how African stories are told – and heard.

“If you’re telling a Nigerian story, you don’t want to be listening to it in an American accent,” explains Wale Lawal, Founder of The Republic. “You want to listen to it in a Nigerian accent.”

This simple observation sparked a journey that would see a media company transform into a tech innovator, tackling one of the most pressing challenges in African digital media: the underrepresentation of indigenous languages and accents in AI tools.

From student journal to tech pioneer

Founded as a student publication in 2018, The Republic gained prominence during Nigeria’s 2020 #EndSARS protests, earning coverage from Al Jazeera. But as a subscription-based platform, they discovered that basic infrastructure didn’t exist in Nigeria. “We couldn’t accept subscriptions in Naira because nothing existed for that,” Lawal recalls. “It could take us up to 60 days to find and license images, which slowed our ability to tell the compelling stories needed to drive subscriptions in the first place.”

The Republic turned obstacles into opportunities, winning a 2022 Google News Initiative award to build Atlas, a media licensing platform. Now, they’re applying the same innovative approach to audio with Minim.

The problem: A billion voices unheard

The numbers tell a stark story. While millions of Africans speak at least one indigenous language, these languages remain dramatically underrepresented in digital media. Existing text-to-speech tools like ElevenLabs offer high-quality audio but lack African accents, creating a disconnect between content and audience.

“Media is something that people get very attached to,” Lawal notes. “Can we get a regional accent? Can we get a regional female voice telling this story? Can we tell it in English? Can we tell it in Swahili?”

The challenge extends beyond accents. Reader engagement across digital news platforms is declining, with newsrooms struggling to create immersive content. For long-form journalism – The Republic’s specialty – readers increasingly want audio options for consuming content during commutes or other activities.

Building the solution: Four languages, multiple voices

Minim's MVP aims to support four languages: English in Nigerian accent, Pidgin English, Hausa, and Swahili. The platform offers two pre-loaded voice profiles (male and female) but also allows users to train the model on their own voices.

The technical stack reflects the complexity of the challenge:

- Python for AI model refinement
- Hugging Face for hosting
- The Urroman toolkit for text processing
- Google Cloud for computing power
- Node.js and Next.js for backend and frontend development.

The team structure reveals how media companies must evolve to embrace AI. Beyond traditional editorial roles, The Republic assembled software engineers, an external AI specialist, UI/UX designers, and a project manager to coordinate between technical and editorial teams.

Navigating the human-AI bridge

Perhaps the most significant challenge wasn't technical but human. "When we started focusing on AI, we got resistance within the team," Lawal admits. "People asked, 'Is this going to replace us? Are there going to be cuts?'"

The Republic addressed these concerns by positioning AI as a productivity tool rather than a replacement for creativity. However, coordinating between journalists and AI engineers proved complex. The team appointed two internal staff members – a project manager and a full-stack engineer – to bridge this gap.

"The AI engineer is very concerned with just building," Lawal explains. "But that's not where the project ends. We need to get feedback from audiences, organise focus group sessions to sense-check Minim's features."

The scarcity of AI talent in Nigeria added another layer of complexity. With limited local expertise available, sourcing and verifying local AI consultants cost The Republic up to three months of "building time", and, once resolved, created dependencies that highlighted the urgent need to develop in-house AI capabilities.

The opportunities: Ethical AI leadership from the Global South

Minim is more than a technical solution – it's ethical AI development in practice. By integrating with Atlas, The Republic's licensing system, it creates a marketplace where voice creators monetise their work while newsrooms access authentic voices.

"This will affect any Africa-focused newsroom that wants to create more enriching media," Lawal says.

The opportunities: new revenue streams, enhanced reader engagement, and linguistic diversity preservation. Minim shows how Global South media can lead ethical AI development. This Nigerian newsroom proves that technology can amplify rather than replace authentic voices.

Lessons for newsrooms

The Republic's journey offers crucial insights for media organisations venturing into AI development:

- **Bridge the communication gap early**: Successful AI integration requires dedicated liaisons between technical and editorial teams. The Republic learned this lesson after struggling with coordination: having translators who understand both worlds is essential from day one.
- **Address AI anxiety transparently**: When staff worried about job losses, The Republic positioned AI as a productivity tool, not a creativity replacement. Clear communication about AI's role helps maintain team morale.
- **Plan for in-house expertise**: Relying on external AI consultants proved challenging. "If we're going to scale this platform, we're going to need internal AI skills," Lawal reflects.
- **Lead with trust**: Newsrooms have an advantage: journalistic ethics. "We don't think 'move fast and break things.' The kinds of questions we have around ethics are not questions a tech company would ask," Lawal observes. "We're thinking about sensitivity checks, focus groups, and ethical considerations because it's embedded in journalistic practice.

Visit our website for the full case study and to learn more about this project.

4.3 **Nawaat: AI platform aids journalists and readers overcome historical amnesia**

Nawaat, an independent Tunisian newsroom, has been in existence for over two decades. During this time, it has witnessed many historic political events and movements, including a dictatorship and the 2011 Tunisian revolution. The team responded to the call of bringing a long-held dream to life – to make the organisation's vast historical archive across the English, French, and Arabic languages accessible to all, with the support of AI.

The problem: Providing utility to its archival data

While the idea to implement this project with AI is relatively more recent, Nawaat is no stranger to working at the cutting edge of technology. In fact the idea for Nawaat AI dates back to 2011, when they had decided to create a tool they called the "Time Machine" to understand the links between different entities – people, persons, events etc. At the time, they used MySQL and Wordpress to create taxonomies since the overall AI technology landscape wasn't mature during the period. They still use wordpress as the basic framework for nawaat.org and taxonomies to structure content. Things changed post 2020 with the rise of AI and they revisited the idea with renewed enthusiasm.

"We started even before getting the grant to build certain capabilities on AWS and testing some LLMs. So we've seen that the idea that was discussed and dreamt of 13 years ago is now really taking shape and facilitated with the use of AI," said Sami Ben Gharbia, Co-founder of Nawaat.

While some implementation was taking place, this grant helped them break resource constraints like being able to hire the right people and developing the core AI infrastructure, said Houssem Hajlaoui, DevOps Engineer at Nawaat.

> Nawaat is literally the only place where you basically have the history of the country.

Building the solution: Enabling discovery and record-keeping

Nawaat AI taps into the newsroom's archives and works twofold – both internally and also externally for its readers. It helps the journalists in the team to easily discover and access articles related to what they're researching or writing even if they have only just joined the organisation. It also helps its readers engage with historical coverage of an issue.

The overall tool allows readers to navigate a certain topic chronologically through history, or through the evolution of the theme itself. For example, this could be about how women's rights evolved in Tunisia in the last 20 years and important milestones.

"It helps in content discovery and it resolves the issue of staff needing to know what Nawaat covered during the last 20 plus years. The other part is for our users, our readers, to understand the evolution of our coverage around a certain topic or certain person or certain region," said Ben Gharbia.

It includes an advanced search feature with a possibility to ask follow up questions, a Time Machine feature, and a timeline. In the chat feature, users can ask a question and receive a well-structured answer with references to Nawaat's previous coverage. The tool also allows them to convert this into a timeline for quick and easy comprehension. The third, Time Machine feature, summarises a particular period that the user selects to reflect the crucial news moments during that period.

A greater and more crucial need for Nawaat AI's Time Machine feature lies in the context of press freedom and freedom of speech in Tunisia, which is leaning towards a dictatorship.

"Nawaat AI is literally your primary source and this is not only a credit to Nawaat and the work that's been done for 20 plus years, but this is also because it's in a dictatorship. If I am to research, for instance, women's rights or feminism in the United States, I have 10 billion sources, because to the extent it's a democracy. But in the context of a country that was a dictatorship all the way to 2011 and lived through a democratic transition for 10 years, it's not that obvious. Now it's unfortunately relapsing back. Nawaat is literally the only place where you basically have the history of the country," explains Hajlaoui.

To the team, it's not just archival content they need to organise, but serves a higher mission of recording the history and struggles that their country has been undergoing over time. It is a way to overcome historical amnesia, and educate those "who didn't live through it all", said Hajlaoui.

The opportunities: Breaking siloes by working in interdisciplinary teams

The overall strength of staff at Nawaat stands at about 16. However, the team building out the tool itself was small and interdisciplinary covering development (infrastructure, design, frontend), user engagement, and editorial input. The technical team comprised external professionals including for frontend design, backend assistance, and AI/infrastructure development. Meanwhile, internal newsroom staff provided essential feedback and testing.

One challenge that the team kept in mind while developing the tool was the accuracy and relevance of articles that the system fetches. They also have to ensure that all content produced is in line with Nawaat's editorial charter and principles.

Extensive research and optimisation, including providing the entire Nawaat archive as input and the editorial charter as guidelines, led to a very high rate of accuracy, said Hajlaoui. Nawaat AI aims for 95-98 per cent accuracy. Still, human validation is required for content published on the main website, added Hajlaoui.

Other challenges the team faced included deciding which AI model to choose, what system to use while retrieving content, and, to avoid being too ambitious with their goals.

Lessons for newsrooms

- **Plan for long-term financial sustainability**: Develop a strategic plan to optimise and cover significant infrastructure costs (servers, storage, maintenance, AI processing) that will arise after the initial grant period to ensure the project's long-term survival.
- **Explore product monetisation**: Consider selling the product or service to other smaller, independent newsrooms in the region as a potential strategy to generate revenue and support ongoing operations.
- **Focus on innovation over competition**: Prioritise innovation as a survival strategy to navigate difficult operational periods or industry transitions, rather than innovating solely to gain a competitive edge in the market.

Visit our website for the full case study and to learn more about this project.

4.4 ICIR: Building Nigeria's culturally aware transcription tool

The International Centre for Investigative Reporting (ICIR) is an independent, non-profit newsroom in Nigeria committed to strengthening accountability through investigative journalism.

The problem they wanted to tackle was a persistent challenge in Nigerian journalism: the lack of AI tools that understand and serve local languages.

Their solution, NativeAI for Newsrooms, a culturally aware AI-powered tool that transcribes audio-visual content into text and translates it into Nigeria's three major local languages: Hausa, Igbo, and Yoruba.

The problem: When foreign AI tools don't speak your language

For Nigerian journalists, transcription has long been a tedious burden. At ICIR, reporters spent hours manually transcribing interviews, press briefings, and reports – often at the expense of investigative work. Existing tools offered little relief: Whisper, an open-source model developed by OpenAI, required technical fine-tuning before use, while Google Cloud Speech-to-Text was both complex and costly.

Translation posed its own challenges. Many existing models frequently fell short in accuracy, particularly when it came to capturing the nuances of the Nigerian English accent and inflections, and often produced inaccurate outputs in local languages.

"Accent differences are a major challenge," explained Chukwudi Iwuoha, ICIR's Senior Programmes Officer. "That was why we needed a tool that could adapt to Nigerian languages and accents."

But the problem extended beyond journalists. Nigeria's deaf community, supported by growing efforts to promote inclusion and expand access to education, continues to face barriers in accessing transcriptions of audio-visual contents. At the same time, many grassroots media outlets publishing in local languages lack reliable tools to produce accurate translations.

These frustrations sparked a conversation within ICIR about building a solution that was local, simpler to use, and free.

Building the solution: An African-born AI tool

ICIR set out to build something different – an African-led, culturally aware AI transcription and translation tool. They needed to train their AI model to understand and adapt to the unique characteristics of the Nigerian and African accent. With support from the JournalismAI Innovation Challenge, they developed NativeAI for Newsrooms, designed specifically to:

- Transcribe audio-visual content automatically
- Adapt to Nigerian accents and speech patterns
- Accurately translate English transcriptions into Hausa, Igbo, and Yoruba
- Support inclusivity by making content accessible to the deaf community
- Promote the use of local languages in communication, research and teaching.

"Our solution serves newsrooms, organisations, and individuals seeking efficient transcription services," said Iwuoha. "It also supports the deaf community, who may require translation of audio-visuals into readable texts that help bridge gaps to access and bring more people into national conversations."

From proposal to a working prototype

The journey from idea to prototype was collaborative all the way from the start. The project kicked off with the entire ICIR team aligned on a shared vision. "It was all hands on deck from the very beginning," recalled Iwuoha. "Everyone contributed – from the proposal stage to the final build. We deliberated extensively to ensure the project matched exactly what we needed."

The team was cross-functional, bringing together machine learning and data engineers and a full-stack developer, alongside journalists, editors, and fact-checkers who ensured the outputs met editorial standards. Partnerships were essential, as noted by Iwuoha:

"We wanted this to be African-built, for Africa. So we engaged not just our newsroom but external engineers and partners who could provide the expertise and datasets needed." The ICIR team sketched out a roadmap.

A key milestone was model selection. After testing several options, and with input from external engineers, the team chose Whisper ASR, an open-source speech recognition model, because of its flexibility, speaker differentiation and potential to be fine-tuned for Nigerian English accents and reduce background noise. For translations, they adopted M2M-100, Meta's multilingual model, which could render accurate outputs in Hausa, Igbo, and Yoruba while handling diacritics and code-switching to preserve meaning.

To ensure reliability, the models were deployed on Google Cloud with HTTPS protocols for secure data handling. Journalists played a central role, uploading interviews then providing feedback that guided improvements in accent recognition and translation quality. The NativeAI model currently achieves approximately 90 per cent accuracy in transcription.

The roadblocks that tested the ICIR team

Like any innovation project, NativeAI faced its share of hurdles – the biggest one being transcription of large files. Initially, the tool struggled to transcribe files longer than 20 minutes, sometimes taking over an hour to process, which defeated the purpose of efficiency. This was resolved after data engineers fixed a backend–frontend communication error. Now, shorter files (5-15 minutes) transcribe in just 1-2 minutes, while hour-long recordings take 5-7 minutes. Translation is even faster, averaging about 3 minutes regardless of file length.

Other key challenges included:

- **Data scarcity**: Reliable, structured datasets in Nigerian languages were limited, requiring the team to spend considerable time cleaning and preparing data.
- **Noisy environments**: Street interviews and crowded press events continue to affect transcription accuracy.
- **Accent diversity**: While the model performs well with Nigerian English, it still struggles with uncommon words and regional name variations.
- **Connectivity**: Poor internet in some regions disrupts performance, reflecting infrastructure challenges beyond the tool itself.

The opportunity: Expanding NativeAI across Africa

NativeAI may have started as a prototype, but its story doesn't end there. ICIR plans to refine its accuracy, expand its language coverage, and explore partnerships with other newsrooms and institutions. "This model will not stop in Nigeria. It was made in Africa, by Africa, and for Africa. Expanding to include other widely spoken languages, like isiZulu [a language spoken in South Africa], would make it truly continental and allow other countries to benefit as well" said Iwuoha.

Lessons for newsrooms

From ICIR's experience, three key takeaways stand out:

- **Localisation is everything**: AI tools become truly valuable when adapted to local realities – from languages and accents to the specific needs of journalists in their context.
- **Collaboration is key**: NativeAI came to life because ICIR combined the skills of journalists, data engineers, and external partners. The blend of newsroom insight and technical expertise was essential to solving the problem.
- **AI is an aid, not a replacement**: Even the best models need human oversight. Journalists remain indispensable in verifying accuracy, providing context, and upholding ethical standards.

Visit our website for the full case study and to learn more about this project.

4.5 Scroll: Factivo AI, a force-multiplier for multi-format journalism

Scroll, an independent digital newsroom in India, saw early on that a single article format could not serve every audience need. The team had been experimenting with off-the-shelf tools that turned text stories into video summaries since 2016 but had failed to find a tool that understood journalistic values. This gap led them to build Factivo 1.0, a multilingual article-to-MP4 tool designed specifically for newsrooms to distribute short videos with high-fidelity-to-source on social platforms.

The tool proved its value during India's 2024 general elections, when Factivo-generated videos were piloted on Instagram. The newsroom saw a sharp spike in both followers and engagement, with audiences responding positively to the videos, which were explicitly labelled "*Made with AI*" and "*Verified by Scroll's Editors*."

The breakthrough: Providing information to users at the most opportune time

But Factivo 1.0 also revealed a bigger opportunity. If a text article could be turned into a video, why not also into interfaces like audio streams, calculators, decision trees, or context sliders? For Scroll, this opened the possibility of creating richer, more interactive experiences for users within their own website rather than relying on third-party social platforms.

"A text article may not always be the best way to deliver news. Some stories work better in other forms. Our challenge was: how do we give users the information they most need, in the form they most want, at the time they need it?" said Sannuta Raghu, Head of AI Lab at Scroll.

That insight became the foundation for Factivo 2.0, developed over the last nine months as part of the JournalismAI Innovation Challenge. This multi-modal versioning tool takes a single piece of journalism and transforms it into multiple formats tailored to a user's needs.

The problem: One size does not fit all

A traditional news report on India's Union Budget or on gold prices answers the essential questions – what happened, who said what, and when. While this context is crucial, readers often also want to know something more immediate: *what does this mean for me?*

For example, alongside a budget story, some readers may want only the Finance Minister's soundbites. Others might prefer a calculator that shows their exact tax liability. With a gold price update, users may want to instantly see how much they could buy or sell gold for that day. Scroll wanted to not only explain events in context, but also give users tools to directly understand their personal impact. Repurposing every story into such tailored formats manually was impossible for a small newsroom. AI offered a way to bridge this gap at scale.

Building the solution: One story, many hats

Currently live as a public sandbox, a user is able to access a news report as summary sliders, timelines, mind maps, decision trees, calculators and expanders, at the click of a button. Each form-to-form versioning of Scroll's articles with Factivo 2.0 is executed with large language models (LLMs). It is dynamic, based on verified news reporting and is designed to have high fidelity to source.

In addition to these formats and their contextual variations, Scroll also developed a lo-fi news stream player and workflow. The idea was to provide an audio stream of news that users could listen to while performing routine chores – like folding clothes or washing dishes, if they wanted to. The audio is neutral-toned, non-alarmist and has an optional lo-fi background track.

The lo-fi news stream is borrowed from Scroll's news feed, The Latest. Every article published on The Latest, is summarised, voiced and sent to the playlist. The workflow is designed such that every story is linked to the previous one in the stream – making it seem like a long radio programme. "This format is still a work in progress but the proof of concept has given us confidence."

Under the hood: How vectorisation helped

One of the breakthroughs of Factivo 2.0 was vectorising Scroll's news archive. In simple terms, vectorising means turning every article – and even every sentence – into a numerical "fingerprint" that captures its meaning. These fingerprints live in a database, where similar ideas, events, or entities can be instantly found and connected. Unlike keyword search, which only looks for exact words, vectorisation understands the context and relationships between pieces of information.

Vectorisation helped Scroll unlock their archive and extract facts from temporally and semantically similar events. "If we have produced 20 reports on a particular flood in 2025, the timeline or expander or mindmap format is able to accurately extract facts from these 20 reports, and not from reports about other floods, or a previous flood in the same geography. Our users had been asking for deeper engagement with our archives, and this is a great way to meaningfully deliver that," Raghu said.

Team structure and challenges faced

The team that is building Factivo is organised in three layers: Raghu and a machine learning engineer are building, experimenting, testing, debugging and iterating everyday. This forms the first layer. Scroll's chief product officer, chief architect and editor-in-chief come in to provide feedback and guardrails. They are constantly updated on Slack everyday. This forms the second layer. The third layer is the designers and editorial testers, who come in as and when required. Scroll's team is completely remote and works across two timezones.

Though working async has not yet proved to be a challenge, the team has faced several challenges while building Factivo.

Adapting the formats to different Indian languages continues to be a problem. Indian languages are codified as accurately as English yet. "How do you deliver the calculator dynamically with Hindi, for example? This is what we've been testing and the results are usually not precise. So we're not able to automate it," she added.

Another significant challenge was ensuring temporal precision in archival retrieval to ensure context remains relevant to the specific story – a problem they worked on for three months.

Versioning of quotes, especially complex or nested quotes was another challenge while working with formats that needed rewriting (like the complexity slider or the lo-fi news stream). By default, many quotes used to get summarised as statements and that needed specific prompting to be corrected.

Lessons for newsrooms

- **AI as a force multiplier**: Recognise that AI can be a force multiplier, unlocking potential and capabilities that would be very difficult for a small, resource-strapped team to achieve otherwise.
- **Ambitious products with tiny teams**: Believe that ambitious, high-quality products can be built by a tiny, committed, and skilled team. "I am not an engineer, I am a journalist – I have learnt everything on the job. From basic coding to using the same jargon as our engineer. And the same goes for him: He has understood how to teach journalism to models, " Raghu said.
- **Deeply integrate editorial input**: Avoid building in silos. Work deeply with editors and journalists because the real challenge in using AI is not the technical coding, but the deliberate design required to teach a model how to preserve the original editorial intent of journalistic work.

Visit our website for the full case study and to learn more about this project.

5. AI for Content-Matching

Two publishers worked on tools that leveraged the use of AI for content-matching. Maltida.es built an AI assistant for content clustering, debunking and expert recommendation, while openDemocracy built a platform for content-tagging in order for news organisations across the world to find partners with whom to share their newsletter real estate with.

What is evident in both organisations is that they leveraged AI to solve precise, labour-intensive problems within their newsroom operations, aiming to enhance efficiency and reach. They emphasised starting with identified user needs and real problems, rather than simply applying technology as an end in itself. A lesson for other publishers is to start with a problem, not the technology, and to question whether AI is the best solution for the problem.

Small publishers should also leverage their ability to be agile and adaptable when it comes to development. They must prioritise continuous testing, learning, and significant adjustments based on real-world feedback and emerging challenges.

Significant hurdles both organisations encountered included substantial technical and practical requirements, particularly concerning data quality, model performance and integration. They also emphasised the need and importance for interdisciplinary teams with a blend of journalistic, technical and data science skills for successful AI project implementation.

5.1 Maldita.es: Building intelligence into fact-checking workflows

Spanish fact-checking organisation, Fundación Maldita.es, processes hundreds of fact-checking requests daily through its community-driven chatbot system. As the organisation has grown since 2018, so has the complexity of managing vast amounts of data – from articles and expert databases to multimedia submissions requiring verification. The grant allowed them to tackle these scaling challenges by introducing AI into specific workflow points, focusing on modular integration that supports journalists rather than replacing them.

The problem: Understanding the user's needs

The inspiration for Maldita's AI assistant emerged from recognising a fundamental gap between the organisation's technological capabilities and those of disinformation spreaders. "We realised that our languages and our codes were a bit behind in comparison to those who spread disinformation," explains Pablo Pérez Benavente, the project's lead engineer.

The team discovered multimodal models capable of processing text, images, and voice messages – exactly the type of content flooding Maldita's fact-checking pipelines. But the primary motivation remained centred on supporting their journalists.

As knowledge became fragmented across the organisation, when writing a new article, fact-checkers often had to search for the previously published articles of Maldita in Google, using a not very efficient in house search engine or rely on colleagues who wrote the articles, which is challenging, especially to newcomers. At the same time, Maldita's methodology requires that every article is supported by experts as sources. Maldita has a database that has been improved over the years but it was hard for fact-checkers to look for the exact experts they need when they are investigating a content. The AI assistant aims to streamline this retrieval process both for the articles and the experts, and eventually serve the broader fact-checking community through Maldita's technology stack.

Building the solution: From user research to prototype

The development followed a systematic approach across multiple work packages, each addressing different aspects of the fact-checking workflow to identify specific pain points where AI could provide meaningful assistance. The methodology began with extensive user research, including one-on-one interviews and focus groups.

The team then collaborated with Maldita's design team to visualise feasible solutions, creating both minimum viable scenarios and ideal development targets.

The roadmap involved iterative testing with end users throughout the development process. “We would do some testing usually with the end users, and then once the testing is done, either iteration listing out next steps or ideally deployment,” Pérez Benavente explains. Insights from impact coordinators and social media specialists, not just fact-checkers, underscored the value of broad research.

Technical architecture and tools

The team used Coder for collaborative development, allowing engineers to work independently whilst sharing APIs. Hugging Face provided access to open-source models, whilst OpenAI’s API offered additional capabilities.

The first work package developed a system for matching images, videos, texts, and voice messages. This “simple” task proved complex, with interconnected modules affecting each other’s performance. The second work package built an AI-powered search engine for Maldita’s article database and experts database, a cornerstone for other Maldita’s tools like grounded chatbots and narrative analysis. The second work package included a new “Analyze with AI” function that, when implemented, will suggest to the fact-checkers both published articles of Maldita and possible experts to be contacted as sources.

Assembling the right expertise

The project required skills that differed significantly from Maldita’s previous AI implementations, as this project demanded more quantifiable, metrics-driven approaches requiring “core data science type of process” skills. The team needed expertise in planning different tests, defining test conditions, and identifying data variability combinations to ensure comprehensive evaluation.

Key roles included:

- An archivist with deep knowledge of Maldita’s tagging schemes, labels, and data registration processes
- An AI engineer focused on model development and integration
- A backend engineer handling software development for API implementation.

Having a global perspective on technical development proved vital to reduce misunderstandings across roles.

Navigating implementation challenges

Open-source model integration brought obstacles. Although Spanish is widely represented, frontier models often exist only in English, with performance in Spanish “clearly subpar,” notes Pérez Benavente.

Beyond language issues, the clean, structured data used for model training differed dramatically from "the collection of highly varied, messy and unbalanced datasets", like the screenshots and extremely long forwarded texts that Maldita processes daily. This forced "a modular system using different models for different kinds of data, affecting both usage and efficiency."

Performance evaluation was another hurdle. Full dataset testing proved unfeasible, and inconsistent documentation of test conditions made it hard to distinguish dataset issues from model shortcomings.

The opportunities and future directions

The AI-powered search engine serves as a foundation for future developments, including chatbots grounded in Maldita's article database and systems that combine articles with narrative analysis to identify debunked information.

The modular matching system also offers potential for expansion, now that its strengths and weaknesses are better understood.

The project also revealed internal coordination opportunities: multiple teams had unknowingly worked on related database improvements, showing the value of cross-team communication. Time constraints imposed by the grant also proved beneficial, pushing the team to prioritise effectively.

Lessons for newsrooms

- **Start with systematic user research across departments**: Maldita's approach of conducting extensive interviews and focus groups with various newsroom roles proved essential for identifying real needs rather than assumed problems.
- **Plan for the complexity of modular AI systems**: What appears simple often involves intricate, interconnected components where failure in one module affects the entire system. Careful conceptualisation at the beginning, identifying potential weak points, and understanding module dependencies can prevent significant delays and performance issues.
- **Document everything, especially testing conditions**: Without proper documentation, it becomes impossible to distinguish between dataset variations and actual model performance.

Visit our website for the full case study and to learn more about this project.

5.2 openDemocracy: Building bridges between newsrooms through an AI-powered newsletter exchange

When social media traffic becomes unpredictable and search algorithms keep changing, email newsletters have emerged as a vital lifeline for news organisations. But manually coordinating content swaps between publications is time-consuming work that many understaffed newsrooms simply can't manage. openDemocracy's solution? Build an AI-powered platform that automates the entire process.

The problem: Manual content sharing takes too much time

openDemocracy, an independent global media platform covering democracy and human rights, had been running a successful but labour-intensive programme: manually arranging content swaps with other organisations to cross-promote their newsletters.

"We would reach out to people and there'd be a back and forth. It was very laborious," explains the project lead, Matthew Linares, who serves as both journalist and coder at openDemocracy. "In the context of a challenging distribution environment where social media and search traffic is precarious, newsletters have become one of the key elements for journalistic organisations to maintain that link with readers."

The manual process worked – organisations would share each other's content with sign-up links, helping everyone grow their subscriber base and increase donations. But it simply didn't scale.

Building the solution: CopySwap's AI-powered matching system

CopySwap was designed to automate this "virtuous circle" of audience sharing. The original vision was straightforward: organisations would upload content they're happy to share, and AI would suggest the best matches for each publication's newsletter.

"The idea was that AI would suggest appropriate content to match newsletters from the pool of submitted content," Linares explains. "Users would submit their content and receive AI-enhanced suggestions of what would suit their newsletter."

Building the platform: Tools and team

The development team kept things lean and focused, building on CiviCRM as the foundation – an open-source CRM that was already managing their donations and email campaigns – and adding a JavaScript app for additional functionality. Their technical stack initially included Quadrant as a vector database for content matching and OpenRouter for accessing various LLMs, with a deliberate focus on using "relatively lightweight LLMs" to minimise environmental impact. The team composition was similarly streamlined, consisting of a project lead who was a journalist/editor with coding skills, an audience manager, a multimedia producer, and a freelance developer.

"I'm perhaps unusual in that I directly code and understand the technical architectures whilst being a journalist and editor," notes Linares. "I bridge the two worlds, which simplifies things."

The pivot: When AI became the content creator

Seven months into development with only five active users, the team discovered organisations weren't creating enough content for the AI system to function effectively. "Users were struggling to create multiple messages on top of their other work," Linares admits.

So they flipped their approach. Instead of matching existing content, AI would generate the messages. "We're deploying a function that takes their social media feed and uses an LLM to create suitable messages," explains Linares. "It's like having a new marketing team member that transforms existing output for different platforms."

Ethical considerations: Keeping AI sustainable

Unlike many AI projects that prioritise functionality over sustainability, openDemocracy made environmental concerns central to their development process.

"We're very aware that AI has a large environmental footprint," Linares emphasises. "We think it's critical to minimise LLM usage, pick the right models, use the smallest footprint models, and use them only where needed."

This meant researching carbon footprints and ensuring human oversight remained paramount. "If something doesn't need to be AI, then don't introduce it," he advises.

Human-centred design: AI as assistant, not replacement

CopySwap treats AI as a helpful assistant rather than an autonomous system. When generating content suggestions, the platform presents both raw social media posts and LLM-generated alternatives, clearly labelled.

"The human gets the final sign-off. It's really just a helper," Linares explains. "We make it easy for users to accept it, take inspiration from it, or reject it entirely."

Future developments might include allowing users to add prompts or keywords, putting them "even more in the loop" to guide the AI toward their specific needs.

The opportunities: Scaling beyond journalism

Despite challenges, the team sees significant potential. The platform could expand beyond journalism to include campaigning and research organisations, potentially supporting fee-based models for sustainability.

The international dimension also offers promise. The team has begun attracting users from different regions including South Africa, and their global focus aligns naturally with their multilingual publishing approach and international content strategy.

Lessons for newsrooms

- **Question the need for AI**: Resist the "hype cycle" and clearly ask if AI is truly needed for your problem. Be sure there aren't cheaper or better-tested alternatives; sometimes, "boring technology" is the most desirable solution.
- **Prioritise the problem, not the tech**: CopySwap identified a real need – helping small news organisations grow audiences and revenue through newsletter partnerships – then thoughtfully applied AI where it genuinely helped, pivoting when user behaviour revealed a better application.
- **Be ready to pivot**: Successful implementation requires understanding user behavior and being willing to pivot when the original vision or application of the technology doesn't match reality.
- **Focus on user needs and ethics**: Success isn't about using the most advanced technology; it's about deeply understanding user needs, maintaining ethical standards, and being flexible enough to adjust your approach.

Visit our website for the full case study and to learn more about this project.

6. AI for Credibility and Misinformation Detection

In an era where information travels faster than verification, the ability to detect misinformation and establish credibility has become fundamental to preserving democratic discourse and public trust. The challenge of distinguishing reliable information from manipulation requires innovative approaches that combine technological sophistication with deep journalistic expertise, particularly as traditional gatekeeping mechanisms struggle to keep pace with the volume and velocity of digital content.

The organisations in this section demonstrate diverse yet complementary approaches to this challenge. Shomrim's investigative newsroom has developed an AI-powered fact-checking tool that identifies approximately 30 different types of content flaws, from omitted factual details to unsubstantiated claims, by deconstructing journalism's fundamental elements to reveal often invisible manipulation. The Nest Center for Journalism Innovation and Development in Mongolia has created the country's first AI-powered fact-checking and sentiment analysis system, integrated into their Pluma.media platform. Their tool flags content based on probability thresholds of misinformation while also analysing sentiment across news sources, enabling readers to contextualise whether they're encountering outlier perspectives or mainstream consensus. The European Correspondent tackles the complexity of verifying and contextualising information across linguistic and cultural boundaries through their AI companion editor Vika, ensuring stories remain accurate and relevant for audiences spanning from Estonia to Portugal.

These organisations share critical implementation challenges: calibrating AI autonomy to prevent hallucinations while maintaining effectiveness, overcoming cultural resistance and trust deficits that limit data access and collaboration, and designing human-centred systems that augment rather than replace journalistic judgement. Their experiences reveal that successful AI integration requires not just technical sophistication but careful attention to psychological barriers, ethical considerations, and the fundamental principle of keeping humans at the centre of editorial decision-making.

6.1 **Shomrim: Teaching AI to spot hidden bias**

When journalists at Shomrim, an Israeli investigative newsroom, began developing their AI-powered fact-checking tool, they faced an unexpected philosophical challenge: how do you teach a machine to recognise something that even experienced journalists often miss?

After hours of debate about what constitutes a fact and deconstructing journalism's five Ws, the team created a sophisticated AI system that identifies approximately 30 different types of flaws across four categories: omitted factual details, statements without context, linguistic issues, and unsubstantiated claims. As Michael Levi, Shomrim's Head of Data Journalism, explains: "When you deconstruct these elements, you begin to see the invisible manipulation that happens, often unintentionally, in news articles."

The problem: When emotions replace facts

The tool emerged from observing troubling patterns in Israeli journalism that proved global. "While the Israeli media rhetoric, like the people themselves, tends to be informal and warm," Levi notes. "We discovered this isn't just an Israeli problem – it's simply more nuanced in other markets."

The core issue? News outlets, driven by engagement metrics, use emotional content to compensate for missing facts. "Every news outlet is a business," Levi explains. "To make money, they need engagement. To get engagement, they evoke emotions – usually negative ones."

This creates "cognitive gap-filling" – readers unconsciously supplementing missing information. "Readers aren't even aware they're filling in missing information," says Doron Sela, Shomrim's Chief Operating Officer. The team drew inspiration from Aristotle's concept of enthymeme – omitting information to make audiences believe something was stated without the speaker taking responsibility. In modern journalism, this ancient manipulation technique has found new life.

Building the solution: Strategic sacrifices and unexpected lessons

Shomrim assembled a compact team: a product manager, project supervisor, and three data scientists. Budget constraints meant strategic sacrifices. "We decided UX/UI design wasn't essential initially," Levi recalls. "Our target audience is professional journalists – they don't need fancy interfaces."

This proved problematic. "We learned it's not a 'nice to have' but necessary," Sela admits. "Even professional users need intuitive interfaces."

The system uses OpenAI's GPT-4.0 Mini as its primary language model, though the team tested Claude and Gemini extensively. The infrastructure relies on standard tools: AWS for storage and JSON for parsing articles. Importantly, they used standard AI models without custom databases or RAG, allowing the LLM to analyse articles purely based on its training.

The challenge of AI freedom

The project's biggest technical challenge involved finding the right balance of AI autonomy. "Initially, we gave the LLM very strict instructions," Levi explains. "Then we realised it could parse articles into factual units on its own, very effectively."

However, giving the AI too much freedom led to what Levi diplomatically calls "less accurate output" – essentially, hallucinations. The team found themselves in a philosophical dilemma: "Sometimes you wonder, maybe the LLM is right and I'm wrong. There's this paradoxical feeling when working with AI."

The solution required months of back-and-forth calibration. "We're now at a balanced point where we have specific instructions that still give room for the LLM to work its magic," Levi says.

Cultural resistance and the Zuckerberg example

The most significant challenge was internal resistance. When an article about Zuckerberg and Trump omitted Zuckerberg's CEO role, journalists dismissed it as nitpicking.

"Everyone said, 'everybody knows who Mark Zuckerberg is,'" Levi recalls. "But omitting his professional context creates an ominous atmosphere. It becomes personal – just 'Zuckerberg,' not 'Facebook CEO."

This epitomises the "invisible manipulation" the tool exposes. Though initially sceptical, consistent flagging awakened what Levi calls "dormant critical muscles." Even explaining the system's logic began changing how journalists read articles.

The opportunities: Building outward from Shomrim's newsroom

The team envisions a tiered approach to impact, beginning with enhancing their own newsroom's journalism through better sources, better articles, and increased efficiency. From there, they plan to expand to other newsrooms and journalism schools. As Levi notes, "It's about training future journalists to do better." The ultimate vision is a browser extension for general readers, though this presents the challenge of bridging the gap between public expectations for simple verdicts and the tool's sophisticated, nuanced analysis. "The general public wants a simple good/bad verdict, but that's not what our tool provides," Sela cautions, highlighting the difficulty of democratising technology designed for professional critical analysis.

The Future of “cyborg journalism”

The project represents what Levi calls “cyborg journalism” – humans and machines continuously train each other. The team is implementing feedback mechanisms for journalists to dispute the AI’s findings, creating a dialogue that improves both.

“This is a mental cyborg, a professional augmentation that makes you aware of things and trains you to see them yourself,” Levi explains.

This positions AI not as journalism’s replacement but its enhancement. As Sela emphasises: “We’re showing AI can improve journalism without replacing it – not in editing or writing, but in sharpening our critical faculties.”

By teaching machines to see what humans miss, this newsroom is pioneering a new approach to journalistic integrity in the age of AI.

Lessons for newsrooms

- **Expect cultural resistance**: The primary resistance to AI won’t be from the technology itself, but from newsroom culture. When AI reveals uncomfortable truths about existing journalistic practices, initial dismissal should be anticipated.
- **Prioritise user experience (UX/UI)**: Don’t underestimate the importance of design, even for professional users. Treating an intuitive interface as optional for internal tools is problematic; even experts require intuitive interfaces to meaningfully engage with AI feedback.
- **Strive for balance in AI freedom**: Successful implementation requires careful balance. Too much freedom for the AI can lead to hallucinations, while too little stifles new insights. The right balance allows the AI to identify human-missed patterns while maintaining accuracy.
- **Prepare for professional discomfort**: Understand that the most valuable AI applications will challenge professional blind spots. This inherent discomfort is precisely where the tool’s true value lies.

Visit our website for the full case study and to learn more about this project.

6.2 **Nest Center: Building Mongolia's First AI-Powered Fact-Checking System**

The Nest Center for Journalism Innovation and Development, a pioneering NGO in Mongolia, confronted a critical challenge in combating disinformation while empowering independent journalism. This organisation, which has operated for five years building innovative media platforms, has emerged as a crucial player in Mongolia's media landscape through its commitment to serving underrepresented communities, its development of freelance journalism infrastructure, and its evidence-based approach to fighting false information in a media ecosystem dominated by political and business interests.

The problem: A media ecosystem under strain

Mongolia's media landscape presents a unique paradox. With 460 registered media organisations serving just 3.5 million people, the country has one of the world's highest per-capita media densities. Yet this apparent abundance masks a troubling reality: 75 per cent of major outlets are owned by politicians or businesses, seemingly functioning more as propaganda tools than independent news sources.

"When you look at the market size and ownership structure, it becomes clear that most media organisations aren't run as businesses but as influence tools," explains Dulamkhorloo Baatar, the project lead. With 80 per cent of news coverage focused on parliament and 60 per cent of the population concentrated in the capital Ulaanbaatar, vast swathes of Mongolian society – particularly rural communities and marginalised groups – remain voiceless in mainstream media.

This concentration of coverage creates a perfect storm for disinformation. Without adequate editorial oversight for independent journalists or proper fact-checking resources for smaller outlets, false information spreads unchecked through Mongolia's media ecosystem.

Building the solution: AI-Powered tools for truth

The Nest Center's response was ambitious: develop Mongolia's first AI-powered fact-checking and sentiment analysis system, integrated into their existing Pluma.media platform – a space designed for freelance journalists to reach underserved audiences.

The project introduced two key innovations. First, an AI-assisted fact-checking tool that scans incoming content for potential inaccuracies, drawing on a database of 4,400 fact-checks conducted over five years. "We identified five key features where compromised information typically appears," Baatar notes. "If content shows a 20 per cent probability of being false, it's flagged for human review. At 60 per cent, readers see an immediate warning that fact-checkers are investigating, which is then updated with the verdict of the fact-check, clearly labelling compromised information as either false, missing context or misleading."

The second tool analyses sentiment across news sources, showing how different outlets cover the same story. “Readers can see whether they’re reading an outlier perspective or mainstream consensus,” Baatar explains. Early users report this increases engagement, encouraging deeper reading.

Building the solution with limited resources

Creating AI tools for Mongolian language presented unique challenges. As a “low-resource language” with limited digital presence, Mongolian lacks the vast datasets that power AI development in English or Chinese. “We have 100 years of newspaper history that isn’t digitised,” Baatar notes, “and no government plans for a national language model.”

The team built most components from scratch, choosing a one-shot model for sentiment analysis while developing a custom fact-checking assistance system. “We couldn’t use readily available tools because everything needed localisation. We literally compiled lists of words that commonly appear in problematic content,” explains Byambajargal Ayushjav, a member of the technical team.

The platform’s complexity demanded a full-stack approach with careful integration across separate systems for journalists, readers, fact-checkers, and sources. Finding qualified engineers proved challenging, particularly those fluent enough in English to work with AI documentation while understanding Mongolian language nuances.

Trust: The unexpected barrier

Perhaps the biggest challenge wasn’t technical but cultural. “We planned to onboard 50 news organisations but ended up with seven,” said Baatar. “Even though we’ve operated for five years, asking for access to archives required executive-level decisions. The level of trust among Mongolian media organisations simply wasn’t there.”

This forced an ethical approach to development. Rather than scraping data without permission, the team pursued collaborative agreements. “We promised partner newsrooms they could use our trained model for their own journalism tools. We wanted this to be the baby of Mongolia’s entire media ecosystem, not just ours.”

Psychological barriers also emerged. “Journalists still tell us they’re afraid of AI, worried it will replace them,” Baatar reveals. This fear, combined with language barriers and limited AI literacy, created resistance even among potential beneficiaries.

The opportunities: AI literacy and niche audiences

Despite challenges, the project is reshaping Mongolia's media landscape. The sentiment analysis tool unexpectedly increased reader engagement, with time spent on articles tripling when users could explore different perspectives. A collaboration with another local organisation added an AI chatbot, the Voyager AI widget, that suggests related stories, further deepening engagement.

More significantly, the project catalysed Mongolia's first community of AI-literate journalists. "Nine months ago, these conversations weren't comfortable. Now we have a group that can discuss AI, ask questions, and solve problems together," Baatar notes.

The platform also enables journalists to serve niche audiences economically unviable for traditional newsrooms. "A newsroom might not survive serving 10,000 people with disabilities, but for an individual journalist, that's a substantial audience," Baatar explains.

Lessons for newsrooms

- **Set ethical standards early**: In emerging AI markets, recognise that you are "breaking the ice," and it is your responsibility to proactively set clear ethical standards for AI collaboration in the country.
- **Prioritise ethics over commercial speed**: Emphasise that ethical considerations must precede commercial ones. "Going slower with proper collaboration is better than setting exploitative precedents," Baatar emphasises. This patient approach, prioritising human rights over profits, offers a model for other developing nations navigating AI adoption.
- **Local understanding trumps resources**: The Nest Center's journey demonstrates that meaningful AI innovation doesn't require Silicon Valley resources – just clear purpose, ethical commitment, and deep understanding of local needs.

Visit our website for the full case study and to learn more about this project.

> We couldn't use readily available tools because everything needed localisation.

6.3 The European Correspondent: An AI companion editor for its journalists

The European Correspondent is a pan-European, independent newsroom that covers news stories from all over Europe. Covering the continent comes with its challenges for their newsroom, namely in trying to ensure that its journalism is useful and relevant for all of its readers.

“Why should someone from Estonia need to know what is happening in Portugal? How do you tell a story from Norway that is important for someone from Italy? We don’t have the resources in terms of editors to manage that many stories, so a lot of times, you have to do the work on your own,” said Philippe Kramer, Co-founder at The European Correspondent.

The problem: Ambitious, but resource-constrained

At the same time, they realized that a lot of mistakes that journalists make in their stories are repeated by their peers in their newsroom. While the journalists may often realise that they’ve made these mistakes, they may need a “subtle ping” to help them notice that they’ve repeated the error, he added.

“Every so often we have a text. In the beginning of the text, we make a promise that we will explain xyz and then we will actually not explain maybe one aspect of xyz that we promised our readers we would at the start of the story. That can happen because the day is hectic or resources are scarce,” said Kramer.

While many newsrooms currently use AI to write texts, Kramer said their AI works more like an editor, leaving comments in a document and through that “increasing the number of conscious decisions they make”. This led to creating Vika, an AI based companion editor that suggests comments and edits to reporters’ drafts, similar to what editors would do.

Building the solution: AI to assist and contextualise editing

The team building Vika consisted of Kramer, the editor-in-chief, and a developer. The remaining newsroom took part occasionally in brainstorming sessions.

Kramer and his team were convinced that the only way the tool would work for their organisation would be if it mimicked the existing writing environment. The entire editor works alongside the organisation’s CMS. Vika’s primary capability lies in its “analyse” function which provides editorial suggestions and commentary throughout the draft of the article. Kramer demoed the tool to us in one article with multiple suggestions from Vika.

Some of the use cases are quite simple. “In a sentence where there is a verb that’s quite boring, for example, it will ask if you want to turn it into something else. In another paragraph which you may have maintained as a placeholder, but forgot to remove, it would suggest that you add text or remove it,” he said.

It works through a system of what they call “mini AI agents”, where different agents conduct different functions that can then be performed across the draft document. For example, one mini AI agent checks for style and tone, another one that checks if the article meets The European Correspondent’s audience needs model, and yet another that simply checks the title of articles. Journalists in the newsroom can also custom create their own mini agents using Vika to detect different things in their texts according to their requirements.

“At the core of everything that we do is helping the journalists write the story and the editors who write it, make a greater number of conscious decisions of what they want to look at. Vika directs them to what they should potentially look at and then they can use their true expertise and let it shine,” explained Kramer.

Despite the use of the term “AI agents”, Vika does not actually use agents at all. In fact, Kramer said it was a “very long, very structured” prompt that led to its creation. They mainly used Gemini for building Vika, however Kramer noted that the structure could easily be adopted to other AI models.

“When it came to AI we realised we didn’t need something very fancy or autonomous. The true innovation for us came in making it really convenient for this one use case that we have,” he added.

The opportunities: Building a seamless system for their users

There was a great amount of testing that went into creating and structuring Vika, according to Kramer. They took inspiration from other text editors that already exist including for structure and logic. They didn’t have to reinvent the wheel and just had to learn from it, he added.

The feedback for Vika has been positive, he explained. They first launched Vika without the AI and just as a writing environment so that the journalists could learn to use it intuitively. This also helped them tune out any bugs that were associated with it. Later, the AI elements were added.

“What has really been interesting is how the newsroom started building a passion for it. They started coming up with their own suggestions on what they would like to see in it, about what should be added. They appreciated how they didn’t have to perform crazy document management; that everything is built into the system,” said Kramer.

Creating Vika also had its share of challenges. One of them was in learning to prioritise the most important features. “Our project was really ambitious. Knowing where to stop and what’s good enough so that it works – was challenging,” said Kramer. They were also worried that, despite the positive feedback, their team of journalists would not use it and would circumvent the system being built, he said.

One of the bigger challenges they face is in ensuring the relevance of Vika’s suggestions. “If you have ten suggestions, and eight of them are not relevant, then you start getting annoyed with the AI and will stop using the tool pretty quickly. So the AI needs to know when to stop suggesting and when it’s not relevant anymore,” he added.

To solve this, they built a two-layered prioritisation system within Vika. This system ensures that smaller editing issues like punctuation errors are not suggested right at the beginning. Instead, things like structural changes to text are suggested first. To help with accuracy, they also have an internal feedback system tied to the suggestions. Moreover, since the AI doesn’t rewrite text and only highlights potential issues for journalists to rewrite at their discretion, most bias problems that come with an AI model can be avoided, said Kramer.

“It’s not just human-in-the-loop, but the human is the person doing the task. That solves the problem of hallucinations which may not be captured in time,” he added.

Lessons for newsrooms

To Kramer and his team, design and limitations are some of the key factors to consider while using AI.

- **Design for precision through limitation**: Avoid focusing on the “magic element” of AI (like simply generating text). Instead, limiting how you apply the AI makes it more precise and therefore truly useful for newsroom tasks.
- **Focus on stronger decision-making**: “If you want to use AI in your newsrooms, it helps in being able to plan and demonstrate that using AI will enable your journalists to make stronger and more impactful decisions. This is a powerful lesson for us.”

Visit our website for the full case study and to learn more about this project.

7. AI for Newsrooms Workflow and Efficiency

In newsrooms where resources are increasingly constrained and audience expectations continue to rise, the integration of AI into daily editorial workflows has shifted from experiment to necessity. The challenge isn't just about using new AI technology, but about integrating tools that maintain journalistic standards while actually making newsrooms more efficient.

The organisations in this section demonstrate pragmatic approaches to workflow transformation through carefully designed AI integration. Chequeado in Argentina developed Asistente IA, a lightweight Chrome extension that packages a growing prompt library directly into journalists' existing Google Docs workflow, prioritising accessibility and regional relevance for fact-checkers across Spanish- and Portuguese-speaking Latin America. La Silla Vacía in Colombia created a centralised hub of modular AI assistants addressing specific editorial needs – from a style-editing tool aligned with their manual to automated newsletter builders – each designed to save time while maintaining editorial consistency. Finally, Gubbi Labs in India created an AI-based news workflow improvement tool which transforms research articles into engaging stories for science communication.

These organisations encountered similar implementation challenges: the tension between development timelines and newsroom urgency, the complexity of designing intuitive interfaces without dedicated UX expertise, and the delicate process of determining which journalistic tasks could be automated without compromising quality. Their experiences reveal that successful workflow integration requires not just technical solutions but organisational patience – accepting that capabilities evolve rapidly, that regional collaboration demands flexibility, and that real efficiency improvements come from iterative refinement based on actual newsroom usage rather than theoretical possibilities.

7.1 Chequeado: Building AI accessibility for fact-checkers

Chequeado's Asistente IA is a lightweight assistant designed to make AI practical for fact checkers and journalists in Argentina. It works as a Chrome extension and is easily integrated within the Google Docs workflow reporters already use. The tool packages a growing prompt library into an interface that is easy to adopt, adaptable to each newsroom, and oriented to collaboration. As part of the JournalismAI Innovation Challenge, the team prioritised accessibility, editorial usefulness, and regional relevance.

The problem: Bringing AI to where the user needs it

The project responded to two recurring pressures. First, the gap between the promise of AI and what newsrooms could adopt in practice. Second, the need to automate certain fact-checking tasks without sacrificing rigour. "The main focus was how to make AI accessible to journalists, especially fact checkers," said Joaquín Saralegui, engineer and product manager.

The team also aimed to support peers across Spanish- and Portuguese-speaking Latin America, where resources are limited but needs are shared.

Building the solution: Roadmap to prototyping

The team combined user research with technical iteration. Reporters were interviewed and regular check-ins identified concrete tasks worth solving. "We discovered that journalists really needed to chat with the solution," noted data journalist Ignacio Ferreiro, shaping the interface.

They prototyped a plug-in that launched prompt-driven tasks directly in Google Docs, then iterated with journalists to refine output quality and editorial fit. A demo showed how users could generate a social thread from an article, draft a newsroom-ready summary, and scan a speech for checkable claims, all within a single session. Tasks can be tweaked, reordered, or shared so improvements propagate to colleagues.

Tools and technologies

The stack combined familiar web technologies with newsroom-ready integration:

- Backend and core services in Python and Django
- Google Chrome extension interface to bring functionality into existing workflows
- A configurable prompt engine, with context and per task knowledge bases
- LLM access via API keys controlled by each organisation
- An open repository with more than 15 prompts in three languages, covering fact checking, social publishing, writing aids, and productivity tasks

The local open version works out of the box with an API key. Organisational features that require shared backend services are available in Chequeado's managed deployment.

The team and the skills involved

Asistente IA was built by a cross functional group: engineers and developers, a product manager and designer, and Chequeado's newsroom. The wider organisation supported logistics and infrastructure. Crucially, developers and journalists worked in tight loops to judge model outputs by editorial standards. As Saralegui put it, without that "synergy between the two it's really difficult to evaluate how good you're doing," which risks shipping a tool misaligned with newsroom quality.

Challenges they encountered

- **Aligning tempos**: The fundamental difference in working pace between development teams and journalists created initial friction. "The development team takes like six months to develop a product and the journalists are working by the minute," Saralegui observes.
- **Regional collaboration**: Whilst essential for impact, regional coordination proved to be consistently challenging. Coordinating across multiple organisations, from multiple countries in similar situations of limited resources requires extraordinary patience and flexibility.
- **Finding the automation frontier**: Perhaps most significantly, determining which journalistic tasks could genuinely be automated required extensive experimentation. This uncertainty led to unexpected discoveries about AI's capabilities in newsroom contexts.
- **Moving targets in model capability**: The rapidly evolving nature of AI technology added another layer of complexity and the need of constant reevaluation. "The thing you tried on month one and didn't work might do on month eight," Saralegui explains.

The opportunities built from a collaborative approach

The project has revealed significant potential for reducing repetitive, non-creative tasks in newsrooms whilst preserving the essential human elements of journalism. “We discover a great potential for this tool and the kind of approach to reduce these repetitive non-creative tasks that happen in the newsroom,” Saralegui confirms.

Regional collaboration has strengthened participating organisations’ AI capabilities while creating knowledge-sharing networks.

Regional collaboration raised AI capacities across participating organisations and fostered knowledge-sharing networks. The tool also proved relevant beyond fact-checking, since many solutions apply to broader journalistic contexts.

Lessons for newsrooms

- **Integration trumps innovation**: Rather than creating a standalone AI tool, Chequeado’s Google Chrome extension succeeds, because it meets journalists where they already work in existing workflows, removing adoption barriers that plague many AI solutions.
- **Collaboration between technical and editorial teams is essential**: Without this synergy, AI tools risk producing technically impressive but journalistically irrelevant results.
- **Embrace the experimental nature of AI development**: The rapid evolution of AI capabilities means that failed experiments may become viable solutions months later. Maintaining a backlog of unsuccessful attempts allows teams to revisit solutions as technology advances.

Visit our website for the full case study and to learn more about this project.

> “The thing you tried on month one and didn’t work might do on month eight.”

7.2 La Silla Vacía: Building a custom AI tools hub

La Silla Vacía, a Colombian political outlet known for editorial independence and investigative reporting, developed a centralised hub of AI-powered tools to improve newsroom efficiency and audience engagement. The hub hosts modular assistants tailored to specific editorial needs, such as:

- A style-editing tool ("FranBot"), aligned with La Silla's manual but adaptable to other style guides
- An assistant to build X threads
- A tool for drafting the daily news round-up ("Duerma Informado")
- A newsletter builder for exclusive "Superamigos" members.

All aim to save time, ensure consistency, and strengthen audience connection.

The problem: Behind the idea

The idea for a centralised hub emerged during the team's AI Lab discovery sessions, after realising that individual tools delivered in isolation often failed to integrate smoothly into editorial workflows. Multiple platforms, logins, and configurations created friction, and journalists were less likely to adopt tools that felt fragmented or required significant effort to learn.

"We saw that even when a tool worked well, it risked being forgotten or ignored if it wasn't easily accessible," said Karen De la Hoz, Product Manager at La Silla Vacía and project manager of this initiative.

Each of the tools of the Hub are purposefully selected, and addresses a concrete problem identified through user research and internal feedback. As De la Hoz explains, "Creating X threads was a pain point. One editor said our assistant does 85 per cent of the work now."

Building the solution: On the way to prototyping

The foundations for the hub were laid two years earlier, when La Silla successfully migrated its CMS with support from the same product and tech team. That project helped establish internal trust and effective communication channels with the newsroom, enabling smoother collaboration on AI integration.

Initial problem identification was participatory. Editors were invited to submit ideas, and the tech team also approached specific desks with targeted offers to collaborate. Each concept was validated with senior editors to ensure sufficient newsroom alignment and long-term utility.

The team worked in weekly sprints, with cross-functional collaboration at the centre. They iterated quickly through prototypes and conducted daily stand-ups to identify blockers and adjust priorities. Feedback loops were built into every stage.

Tools used

The project combines several LLMs and development platforms:

- **LLMs tested**: ChatGPT (including GPT-4, GPT-5 and Opus), Anthropic Claude (Sonnet and Opus), Google Gemini, and DeepSeek.
- **Back-end and infrastructure**: MongoDB, Vercel, Google Cloud Platform, SupaBase, WordPress API.
- **Prototyping**: V0 from Vercel was particularly helpful early on for rapid front-end deployment, although its quality declined after a later update.

The team conducted multiple prompt engineering cycles and attempted fine-tuning four times for one of the tools, though with limited success. They prioritised flexibility in tool integration, switching models or APIs depending on performance and cost constraints.

Team and skills

The project was led by a hybrid team with backgrounds in journalism, technology, and product development and included a product manager and editorial bridge, with a background in journalism, an external collaborator with experience in chatbots and civic tech, full-stack and a backend developer. The team also included a feedback coordinator and user liaison during the project implementation, after discovering how relevant it was.

The challenges in the way

- **Reassessing and restructuring the hub**: Halfway through the development process, the team decided to rebuild the hub's architecture to better support external users. While this change added technical complexity and delayed timelines, it was necessary to future-proof the product.
- **UX design limitations**: No one on the team was a UX specialist. Eventually, they organised a "Design Critic" session, studied its documentation, and revised their approach in a space facilitated by their consultant. This helped improve usability, but the process highlighted the need for better design resources in future iterations.
- **Feedback volume and complexity**: The team underestimated feedback volume and learned that collecting input without acting on it would frustrate users. "We decided we could only ask for feedback if we had the capacity to act on it," said De la Hoz. They created a dedicated feedback owner role and established a lightweight system using WhatsApp voice notes, which are transcribed into actionable documents, enabling quick adaptation while reducing user friction.

The opportunities: What's ahead

The team sees strong potential to scale the hub to other newsrooms in Latin America and beyond. "We want to figure out how to make this technically and economically viable for others," De la Hoz said.

Internally, the tools are already being used regularly. FranBot, the daily round-up tool, and the X thread assistant have all been integrated into newsroom routines. Their continued use by a sceptical and demanding editorial team serves as a strong indicator of impact.

Lessons for newsrooms

- **Feedback systems are as important as the tools**: Without mechanisms to gather and act on user input, even the best tools will fail.
- **Adoption requires trust and experience**: Past wins like the CMS migration helped build internal credibility. Journalists were more willing to collaborate on AI tools once they saw tangible benefits and trusted the team's intentions.
- **AI projects need cross-disciplinary integration**: Success depends on merging technical development with product thinking, newsroom collaboration, and user-centred design. "A tech team alone cannot build the right tools. You need newsroom buy-in from day one."

Visit our website for the full case study and to learn more about this project.

> We saw that even when a tool worked well, it risked being forgotten or ignored if it wasn't easily accessible… 'A tech team alone cannot build the right tools. You need newsroom buy-in from day one.'

7.3 The Oglethorpe Echo: Empowering local news through AI efficiency

Small, local newsrooms and academic journalism partnerships in the United States are grappling with significant challenges, from financial strain and staff shortages to the rise of "news deserts." These issues severely limit their ability to provide essential local coverage and uphold democratic functions. The team at The Oglethorpe Echo partnered with YESEO, an AI-powered tool to address these critical inefficiencies.

The problem: Doing more with less in "news deserts"

Financial pressures, amongst other challenges, have led to widespread closures, creating "news deserts" with no local journalism. As Amanda Bright, instructor and assistant editor of The Echo noted, "Most local newspapers like the small ones in the US have one or two reporters tops and they're covering these huge areas and they're just not able to serve people the way they need to be in a functioning democracy." This severely limits their ability to serve communities effectively and uphold a functioning democracy. The core issue boils down to doing more with less, as the appetite for news remains high while resources dwindle. These newsrooms struggle to generate sufficient revenue and audience growth, partly because they spend an inordinate amount of time on repetitive tasks.

Another critical inefficiency, particularly within news-academic partnerships, is the high turnover rate. Students and new employees frequently need extensive onboarding to understand their beat, community, and key local figures. This process is time-consuming, often delaying actual journalistic work and relying heavily on editors or ad-hoc internet searches. Without a systematic way to quickly get new staff up to speed, newsrooms lose valuable time and resources, hindering their ability to produce timely and in-depth reporting.

Building the solution: Creating efficiencies in news production

Recognising these challenges, the team behind YESEO aimed to leverage AI to create newsroom efficiencies. The initial concept, spearheaded by YESEO founder, Ryan Restivo, built upon a pre-existing Slack-based AI tool that integrated with Google Trends. This tool helped with tasks like headline writing and social media messaging. However, the grant funding from JournalismAI allowed the team to deeply invest, grow, and specifically target the app's capabilities based on direct feedback from local newsrooms. The goal was to understand what truly helped them increase revenue, subscriptions, and efficiency.

The project focused on The Oglethorpe Echo, a news-academic partnership in Georgia, as its beta test site. Early surveys and qualitative conversations with other newsrooms and academic partnerships informed the development. The solution centered on addressing two primary inefficiencies. First, YESEO streamlined content production for various platforms. It assisted in generating headlines, social media posts, and newsletter content, preventing reporters from starting from a blank page for every new output. This significantly saved time, especially on days when creativity was low or a fresh approach was needed for basic production tasks post-reporting.

Second, and unexpectedly, the project developed a crucial onboarding and backgrounding solution. Recognising the high churn in academic newsrooms, YESEO was designed to help new students and reporters quickly learn about people, places, trends, and topics within the community. Instead of solely relying on editors or random online searches, the Slack-based app allowed users to input names or organisations. The AI then performed an LLM analysis, providing information like contact details, past quotes, an analysis of perspectives on key issues, and a broader understanding of individuals' viewpoints. For example, it could quickly summarise what a county commissioner has said about solar farms, including quotes and historical context. This systematic approach drastically reduced the time needed for new staff to get up to speed, freeing them to focus on journalism sooner.

Restivo emphasised that the data used for these profiles and analyses within YESEO comes exclusively from The Echo's own published articles. This ensures that the information is authoritative and relevant to the local context, particularly in rural communities where public figures may not have extensive online profiles elsewhere. The tool also allows newsroom staff to add notes and update information, creating a continually growing and accessible knowledge base that can be shared across different classes and reporting periods.

The opportunities

The impact of YESEO extended beyond mere efficiency. The project demonstrated tangible positive outcomes in audience growth and revenue generation. The increased efficiencies in digital product creation and the enhanced context in reporting allowed The Echo to move faster. This freedom enabled them to add a new "in-depth storytelling" beat, supporting two new reporters in producing long-form, deeper dives. They also significantly ramped up video production, successfully monetising it through digital ads.

Metrics revealed an increase in digital ad sales, subscriptions, and overall analytics. While multiple factors contribute to such metrics, these sharp increases were unprecedented for The Echo over the past four years, indicating a correlation with the implementation of YESEO. The project underscored that improved internal efficiencies directly translate into more resources for audience engagement and revenue-generating activities.

The demand for practical, ethical AI tools is evident. As Restivo and Bright presented YESEO at various conferences and press associations, the consistent feedback was, "We want that," says Bright. This highlights a widespread interest in AI that moves beyond philosophical discussions to deliver tangible, useful solutions for newsrooms. The tool's integration into Slack, a platform already used by many newsrooms, further enhances its appeal by providing a multi-dimensional solution within a familiar environment, minimising the learning curve and promoting collaboration.

Lessons for newsrooms

- **Prioritise transparency and ethics**: Ensure proactive communication about AI usage and emphasise human oversight to build crucial trust within the community.
- **Address specific newsroom needs**: AI implementation must focus on specific, real-world pain points within newsroom workflows, rather than creating tools for technology's sake. Deep engagement with staff is essential to find the most impactful support areas.
- **Foster flexibility and experimentation**: Cultivate a culture of adaptability and be willing to change course as new needs emerge. Encourage staff participation in AI product development.
- **Leverage existing local knowledge**: Utilise organisational knowledge, such as newsroom archives, to ensure AI outputs are accurate, relevant, and tailored to the local context, which is key for credibility.

Visit our website for the full case study and to learn more about this project.

> "Most local newspapers like the small ones in the US have one or two reporters tops and they're covering these huge areas and they're just not able to serve people the way they need to be in a functioning democracy."

7.4 Gubbi Labs: Using AI to improve story production workflows

Gubbi Labs is an Indian research-based news organisation that has been communicating science since 2014. Their typical day-to-day workflow for writing stories includes perusing research papers from all domains of science, engineering, and humanities, understanding them, and writing news stories based on leads that they generate. The entire routine, including writing a draft, used to take the small team of six anywhere between a week and two weeks to complete.

The problem: Time-consuming manual workflows

The organisation's prior manual workflow involved significant efforts to skim hundreds of sources to identify newsworthy pieces and a multi-step editing and publishing process.

With the arrival of GPT 2 and later, ChatGPT into the AI scene, they realised that they could potentially tap into AI to help them with initial drafts and improve turnaround times. "Clearly, that is where I guess using LLMs became sort of the go-to solution for us. In the sense, to really have greater and quicker turnaround times that will help also in getting more throughput out," said HS Sudhira, Director at Gubbi Labs.

This led to creating the AI-based news workflow improvement tool, Babbler.

Building the solution: Roadmap to prototyping

After being selected as a grantee for the Innovation Challenge, the team zeroed in on certain aspects of the workflow to automate using AI. One of the first, was to identify and create a "newsworthiness index".

"We have come up with our own matrix to kind of see through them and assign an index to say what could be newsworthy. So now instead of skimming through thousands of papers every month, which from India alone is anywhere around two-and-a-half to three-and-a-half thousand papers, we now know we just have to pick up the top 50 from the newsworthiness index across different topics," said Sudhira.

Using AI tools in this instance significantly reduced the time and effort involved in picking research papers to read. Where the team used to spend hours, they now spend about an hour overall in a week.

The opportunities: An answer that has many applications

Using LLMs to solve the newsworthiness issue also lent itself to other applications within their workflow. For example, they built a pipeline using engineered prompts to generate story drafts and social media collateral.

“Something that we have been very cognisant of is that we have made that at different touchpoints we are bringing in a sort of human-in-the-loop. So our editors will still play a very critical role in basically making sure the outputs are good, reasonable or any changes to be quickly edited,” added Sudhira.

The user interface for the tool is designed to be intuitive especially for editors.

The dashboard allows editors to view and sort papers by newsworthiness index, institution, or journal, then select a paper to generate a story and social media collateral with pre-derived prompts that can be modified. To tailor the system to the type of editor using it, it has two modes: the what-you-see-is-what-you-get aspect for immediate revisions, and advanced settings for experienced users to tweak models.

Technology stack

The newsworthiness index was developed by analysing Google News for relevance including freshness, authoritativeness, relevance, and context. They identified eight broad topics of social relevance (health, ecology, environment, climate change, technology) to categorise research articles. The index also incorporates a ranking system based on the institution and impact factor of the journals where papers are published.

“So higher impact factor journals would have a greater weightage for instance or an institution with a higher ranking would have a greater weightage. We arrived at sort of a combination of these weightages and topics and then linked it with Google News’ API to see which are trending. Based on that we derived the newsworthiness index,” said Sudhira.

The backend uses Python as the core language, with a SQL database for journal papers and Python for API requests, newsworthiness calculations, and story generation. The frontend is built using React and Material UI.

Team and challenges faced

The team assembled to work on the project included consulting researchers to test the newsworthiness index, and a UI/UX consultant for design feedback. Apart from this, their internal team of editors also took part in reviewing and evaluating the generated content.

During the initial phase, the team faced several challenges.

"Some of the things that we had realised was that when you parse a PDF (to LLMs) it would kind of simply take all the text in and not just the main text that is required because at times in a research paper you have a lot of other information like journal name, page numbers, keywords, institutional affiliation, author names and all of it," explained Sudhira.

They also faced initial resistance from some members of the newsroom to apply AI keeping in mind the quality and issues of bias. However, they addressed this by keeping human-in-the-loop systems as a part of the tool.

Lessons for newsrooms

- **Benefits of time-boxed "sprints"**: Embrace the use of short product development cycles (like four-week sprints). This approach helps in clearly conceiving the end goal and achieving it through smaller, focused milestones.
- **Sticking to focus and goals**: Using sprints helps the team maintain focus on the immediate task while ensuring that each milestone fits the larger goal of arriving at the final tool.

The team aims to build their revenue and subscription strategy by making Babbler, available commercially for academic institutions and other research institutions.

Visit our website for the full case study and to learn more about this project.

8. AI for Real-Time Data Analysis and Community Alerts

Only two publishers worked on this thematic area of using AI to effectively deploy real-time data analysis and community alerts, particularly for underrepresented or vulnerable communities. Both organisations and initiatives are based in Brazil.

Agência Mural's tool called the Local Climate Alert focused on providing real-time environmental risk alerts, such as for air quality and floods, to underrepresented communities in metropolitan areas. Fogo Cruzado on the other hand built VANIA, a conversational chatbot designed to make Latin America's largest gun violence database accessible to Brazilian communities, beyond just numbers.

Their work highlights the potential of AI to transform complex data into actionable, accessible information that is delivered directly to those who need it most. Both these initiatives underscore common themes and critical insights into the work they are doing and experimentation.

Both initiatives prioritise the delivery of information through channels and platforms that their target audiences already use and are familiar with. They meet their communities in the spaces they congregate and converse in, such as on WhatsApp and other conversational platforms. Their work is community-centric. They address information gaps by leveraging AI to bridge gaps in understanding. They also provide accessible and verified information that can empower communities to make informed decisions and hold authorities accountable.

Their case studies highlight the practical challenges of AI implementation in a resource-constrained environment.

8.1 Agência Mural: Using AI to deliver service journalism where people are

Agência Mural is a nonprofit newsroom in Brazil with its reporting focused on underrepresented communities in the metropolitan areas of São Paulo. With Local Climate Alert, the team is applying AI to service journalism, sending concise, location-aware alerts about climate risks directly to people's phones and inboxes. As Brazil prepares to host a major climate summit, the project positions local audiences to make safer, faster decisions.

The problem: Being usefully fast

The impetus is rooted in lived experience. "All of us are from the peripheries," says Vagner de Alencar, director of journalism at Agência Mural. "We know what is happening in these places." Mural's network of more than 60 local correspondents consistently reports that climate impacts are felt first and hardest in the margins. The team wanted to move beyond one-off reports to provide actionable, real-time information in the channels people actually use. "Our audience wants information directly through WhatsApp," notes de Alencar. "They want to know about air quality or floods right now."

The project also builds on prior WhatsApp experiments, including a pandemic-era micro-podcast sent as short audio messages and Papo Reto no Zap, a community initiative across five neighbourhoods. Those experiences validated WhatsApp as a primary interface for news and services.

Building the solution: Roadmap to prototyping

The team followed a pragmatic path from concept to pilot:

- **User research and scoping**: They ran surveys and interviews to learn what people wanted, how often, and in which formats. The priority was clear, short alerts about immediate risks, delivered where people already converse.
- **Data sourcing and cleaning**: Rather than rely on third-party APIs, the team identified open public databases relevant to weather, floods and air quality. Early tests with generative text "had too much hallucination," so they designed a pipeline focused on extraction, cleaning and concise templated outputs.
- **Technical partnership**: Recognising a gap in-house, Mural partnered with Agência Tatu, a newsroom in Brazil's northeast with data and automation experience. "We put together our strengths. Without them it would be almost impossible," says de Alencar.

- **Minimum viable alerting**: They launched a WhatsApp Community and an email list for five diverse communities across the capital region. People register once and receive periodic alerts, plus links to a project web page that hosts related service content.
- **Live community launch**: True to Mural's ethos, the pilot launched with an in-person event featuring exhibitions, conversations with residents and participation from local correspondents. "AI is online, but our essence is being connected in person," de Alencar reflects. "We exist for these communities. AI is a means to serve them better."

The tools and technologies involved in Local Climate Alert

The team intentionally deprioritised free-form LLM writing in favour of structured templates to avoid factual drift. "We now have a process that extracts and cleans the data," says de Alencar. To build and make Local Climate Alert work, they relied on different tools:

- Open public datasets for air quality, rainfall and flood risk
- A data pipeline to extract, normalise and validate fields before publishing
- Lightweight templating to generate clear, non-hallucinatory alert copy
- WhatsApp Communities and email for distribution
- A web hub to host evergreen service information and updates
- Visual and UX support to clearly label what is AI-assisted versus human-authored.

The team was assembled to place community needs at the centre of the development process, with technical expertise provided primarily by Agência Tatu, Mural's partner in the initiative. Within Mural, an editorial lead ensures the service is framed in ways that remain relevant to local audiences, while a tech and business lead coordinates the build and partner integration. Audience engagement is strengthened by a community coordinator who works with more than 60 correspondents to root the alerts in local realities. Research is carried out by a dedicated role focused on surveys and interviews with residents and specialists. A visual and web designer contributes to clarity, accessibility, and proper labelling across channels. Additionally, data and automation support is provided via Agência Tatu, bringing Python and interface development expertise to the project.

Challenges they encountered

- **Limits of WhatsApp Communities**: Communities are discoverable but less direct than one-to-one messaging. “Our dream is to send as a pop-up directly to people’s numbers,” says de Alencar, but business messaging at scale costs money and is priced in dollars. That makes full automation difficult for a nonprofit.
- **Avoiding hallucinations**: Early experiments with LLM-written alerts produced inaccuracies. The team switched to a constrained, template-driven approach backed by rigorous data cleaning and editorial checks.
- **Data fragmentation and scarcity**: Useful public data exists but is scattered and uneven. Normalising formats, ensuring freshness and matching signals to specific neighbourhoods require ongoing maintenance.
- **Resourcing specialised roles**: It took time to hire for community management and audience development, roles that proved essential to adoption and trust.

The opportunities: To multiply the impact and drive revenue

Replicability is a central insight. “We created something completely replicable,” says de Alencar. The alerting “robot structure” can be pointed to other useful datasets, such as job opportunities or public services. The model also opens pathways to sustainability by deepening local reach and potentially attracting place-based advertising on Mural’s site. Above all, the project has unlocked organisational capacity to produce service journalism at scale, not only as special projects.

Lessons for newsrooms

- **Meet people where they are**: Design for the channels your audience already uses. For Mural’s readers, WhatsApp is the front door to timely, trusted service information.
- **Constrain AI to protect accuracy**: Templated outputs on top of clean, validated data reduce errors and build trust.
- **Pair tech with presence**: Digital alerts travel fast, but legitimacy is earned face-to-face. In-person launches and ongoing community contact anchor the service in real needs and local knowledge.

Visit our website for the full case study and to learn more about this project.

8.2 Fogo Cruzado: Giving a voice to gun violence numbers

Fogo Cruzado, a Brazilian organisation that monitors armed violence, has built one of the most extensive datasets on shootings in Latin America. Since 2016, it has recorded more than 63,000 incidents across Rio de Janeiro, Recife, Salvador and Belém. Yet data alone is not enough to create impact. To address the challenge of accessibility, the organisation developed VANIA, a chatbot powered by artificial intelligence that allows citizens to ask simple questions and get real-time information on gun violence in their communities.

The problem: Shedding light on the violence

Gun violence is a daily reality in Brazil's major cities, particularly in poor and marginalised neighbourhoods. In 2024 alone, 2,922 people were killed by police forces in the metropolitan regions covered by Fogo Cruzado. But the lack of accessible data weakens public debate and prevents communities from demanding accountability.

"We have the biggest database on armed violence in Latin America, but despite its public importance, we still face the major challenge of making it accessible," explains Diogo Santos, communication and innovation manager at Fogo Cruzado.

Building the solution: Roadmap to prototyping

The initial goal was simple: allow citizens to query Fogo Cruzado's database through a conversational interface. Early prototypes were tested in Portuguese and English, offering suggestions of possible questions to guide users. The tool enables follow-up queries on the same topic, simulating a natural conversation.

What seemed straightforward quickly revealed technical complexities. The team had to adapt the LLM to Fogo Cruzado's unique taxonomy, including terms like "police massacre" (three or more civilians killed in an operation) or "drive-by attack". Training the system to understand this specialised vocabulary was a critical step.

The team, technologies and the skills behind the project

VANIA combines a LLM with Fogo Cruzado's API, which feeds real-time information from the database. The chatbot checks each query against the database and validates the result before responding. The interface was designed for simplicity, showing users what kind of questions can be asked and how to refine them. A bilingual version is currently being tested, with Portuguese as the primary language.

The project was led by Fogo Cruzado's communication and innovation team, working closely with developers and AI specialists. Their expertise in database management and violence monitoring was crucial to adapting the model to their dataset. The team also relied on community knowledge to refine terminology and ensure that the language used by the chatbot matched the reality of affected populations.

Developing VANIA highlighted several challenges:

- **Specialised vocabulary**: Training the LLM to recognise and explain local terms and categories required significant effort.
- **Cost barriers**: As Santos notes, "How do we democratise the use of AI and apply it to real problems when cost is a big barrier?" Tools priced in US dollars remain inaccessible to many Latin American organisations.
- **Naming and cultural fit**: The project was originally called BulletData, but the team quickly realised it would not resonate locally. They opted for VANIA, short for Violência Armada em Números + IA (Arms violence in numbers + AI), a name that is both meaningful and familiar to Brazilians.

The opportunities: Transparent and evidence-based conversations on violence

For Fogo Cruzado, VANIA is more than a chatbot: it is an entry point to a future ecosystem of tools for transparency and advocacy. By making gun violence data accessible, it empowers citizens to stay informed, organise collectively, and pressure authorities to act. The tool also has potential for journalists and policymakers, who can use it to identify trends and design evidence-based responses. Beyond Brazil, the model could inspire other contexts where violence is poorly documented and misinformation thrives.

Lessons for newsrooms

- **Accessibility matters as much as data**: Collecting information is not enough. To empower communities, data must be translated into usable, everyday formats.
- **Local language is key**: Training AI systems on context-specific terminology ensures accuracy and cultural relevance.
- **AI can support civic mobilisation**: Tools like VANIA show how AI can be harnessed not only for efficiency, but to strengthen democracy and social accountability.

Visit our website for the full case study and to learn more about this project.

9. AI for Government Transparency and Accountability

This ninth and final section explores how three distinct publishers leveraged AI to address critical challenges in fact-checking, investigative journalism and holding governments accountable. The Centre for Investigative Journalism of Serbia (CINS), Full Fact in the United Kingdom (UK) and CalMatters in the USA, each embarked on custom AI initiatives that reveal valuable insights into the practicalities, benefits and hurdles of integrating advanced technology for government transparency and accountability.

These three case studies collectively highlight several critical insights. Even for small newsrooms: AI should be leveraged as a tool for efficiency and augmentation, and not a replacement, especially when journalists need to process and analyse vast amounts of datasets quickly that would be manually impossible. Given current limitations of 100 per cent reliability, human oversight and editorial judgment remain indispensable to ensure accuracy. Often overlooked publishers also need to bridge communication gaps between technical and editorial teams.

The three organisations largely used common tech stacks such as LLMs, specifically ChatGPT, Google's Gemini, and OpenAI tools, which reveal a reliance on such advanced models for tasks like summarisation, question generation and text processing. Both CINS and CalMatters invested in developing their own tailored AI tools or models to meet their specific needs, especially when off-the-shelf solutions proved insufficient.

Despite facing various challenges, the initiatives showed significant potential and need for custom AI solutions to scale and foster collaboration. For example, CINS sees universal application for its tool across the Balkans, due to linguistic similarities, while CalMatters plans to expand its project to other states in the US. Full Fact's work aims to benefit smaller fact-checking organisations globally and advance research in the field.

In a nutshell, these next three case studies collectively highlight a vision where AI is a powerful ally in enhancing democratic accountability and public trust through deeper, faster, and more comprehensive journalism.

9.1 CINS: Powering data-driven investigations with AI

The Centre for Investigative Journalism of Serbia (CINS) faced a challenge many investigative outlets encounter: how to verify government claims backed by vast amounts of data when traditional manual processing simply isn't feasible. Their solution – a custom AI-powered tool – represents their first major step into artificial intelligence.

The problem: Debunking claims with mountains of data

CINS observed a surge of pro-government news coverage claiming dramatic reductions in healthcare waiting lists for major operations. The 11-person team, with seven working in the newsroom, suspected these claims were misleading.

"We thought that this cannot be true because it's seemingly impossible for them to get thousands of people off the lists in such a short period of time," explains CINS Director Milica Šarić. "So we started wondering, how can we debunk this?"

Their traditional workflow – requesting thousands of documents from Serbian institutions, scanning them, performing Optical Character Recognition (OCR), and manually processing everything – was becoming unsustainable. "We have done everything manually. We have a long history of collecting huge amounts of data," Šarić notes. The JournalismAI Innovation Challenge, offered a timely opportunity to pilot a different approach.

Building the solution: When off-the-shelf AI isn't enough

Initially, CINS attempted to use ChatGPT with scraped data, but quickly discovered its limitations. "It couldn't answer all our needs. It couldn't collect and analyse all the data that we needed," the team found.

This led them to develop a custom solution with an external AI specialist who understood media requirements. "We decided to make something on our own – our AI specialist building our own tool that actually can bring in many Excel files, process them through SQL database imports and queries, and then provide responses like GPT when you put in a prompt."

The resulting tool functions like a specialised ChatGPT that can process multiple Excel files and respond to data queries.

The communication barrier

One of the most significant obstacles to building the tool wasn't technical but human. CINS discovered that the gap between journalists and developers poses a fundamental challenge for media organisations, particularly in the Balkans.

"That's a huge problem... finding appropriate specialists and tech people who will be able to understand our language, understand what we need," Šarić explains. This affects the entire region: "Other media organisations that I know in Serbia encounter the same issue."

Their solution was pragmatic: working with their AI specialist to develop a shared communication approach that would make him more attuned to journalistic sensibility. "We worked with our AI specialist on the way that he needs to talk to us. We now have a new language between us."

Obstruction by design

According to CINS, Serbian authorities deliberately obstruct data access. "Sometimes institutions print everything and then scan it back in a way that obscures the text. They're trying in every possible way to make it hard for us to use digital tools." The team also notes that the Freedom of Information Act has become increasingly unreliable.

These institutional barriers add another layer of complexity to AI implementation, as even the most sophisticated tools cannot process deliberately corrupted or withheld data. The team's custom solution can only be effective when institutions actually provide usable information or when such data is available online.

What they've learnt

The project revealed that AI integration requires organisational transformation. "We don't have a skilled enough organisation. So at some point we should bring on board more technical people, more digitally skilled individuals. But we also need to support and upskill our current staff as part of this journey,," Šarić explains.

The editor-in-chief is now pursuing a master's degree in digital innovation management in the UK, planning to drive the transition to a modern digitally-oriented newsroom.

The opportunities: Enhanced efficiency and digital transformation

CINS views AI as a tool for efficiency rather than replacement. The solution processes data faster than manual methods and promises cost reductions by eliminating the need for additional staff for routine processing.

"This is an opportunity for us to produce more in a faster and more digital way instead of doing everything manually but significantly slower," Šarić notes. However, they maintain editorial oversight and recognise that successful implementation requires sustained technical support.

Regional potential

The tool has universal applications for data journalism. "This tool can serve anybody who works with large sets of data. Anyone engaged with data journalism or investigative journalism can benefit from it."

The team is enthusiastic about knowledge sharing through masterclasses and conferences, viewing their experience as a foundation for broader regional AI adoption in journalism.

Lessons for newsrooms

- **Bridge the communication gap**: Recognise that the biggest hurdle is often the communication gap between journalists and technical specialists. Also, aim to bridge the gap between journalists and the AI solutions themselves, reducing the reliance on expensive technical intermediaries.
- **Build custom tools for specific needs**: Consider building custom tools instead of simply adopting existing solutions. This approach helps ensure the technology precisely addresses journalism's unique requirements.

For newsrooms considering similar projects, three factors proved critical: finding technical specialists who understand media, preparing for organisational change beyond tool adoption, and treating AI as a way to enhance rather than replace human judgment.

Visit our website for the full case study and to learn more about this project.

9.2 CalMatters: Using custom-AI to enhance its political accountability tool, Digital Democracy

The Digital Democracy project had its origins a little over 10 years ago in the state of California, USA. A senior leader from San Luis Obispo, who had completed his service in the State Senate, was disillusioned by how much impact money had in state government. To solve this, he partnered with the California Polytechnic University to build Digital Democracy, a tool that would give Californians access to information about their State legislature. CalMatters, a nonprofit newsroom that was launched in 2015 later adopted this framework and built upon previous efforts.

“We took their tool which made all this information available to the public and started building the capability to turn that information into journalism to generate story ideas,” said Neil Chase, CEO of CalMatters.

The ultimate goal was to make the information more accessible to journalists and citizens by identifying speakers, entities, and financial flows.

The problem: Expanding accountability capabilities

Once the team saw the tremendous interest from other States for a similar tool, they decided to expand the number of regions that had access to it. The JournalismAI Innovation Challenge grant was one opportunity to help them do so.

“It was the boost we needed. We knew we wanted to do the multi-state expansion. We didn’t have a lot of support for it. It has helped us, I think, to see this as a specifically defined project that we can make major progress on this year and do something we think is really good that we hope is going to set an example for other folks. We’re getting interest in this from other levels of jurisdictions and even other countries now,” said Chase.

During the course of JournalismAI Innovation Challenge, the team worked on expanding the tool to Hawaii. They partnered with another nonprofit newsroom, the Honolulu Civil Beat to do so.

“Hawaii has different types of data that are more readily available. They’re kind of comparable to California in the sense that there’s a lot of stuff that’s publicly available on their websites or databases that we could acquire via the help of our Civil Beat partners,” said Ramsey Isler, Director of Special Projects at CalMatters.

Building the solution: AI to decode public meetings

The tool works both externally for citizens and internally for journalists. Once a citizen navigates to the Digital Democracy platform, they will be able to search the database for things like every word spoken in public hearings, money donated to politicians, bills introduced and even research legislators and their activities. It works differently for CalMatters' journalists who are using the system.

"The part that we're not exposing to the public is what I equate to the kind of thing that's in a reporter's notebook. They are notes. They are not facts. They're not things we can publish. And so we draw a line between data and making that as accessible as we can to people and the conclusions or the ideas that the AI is identifying in the tip sheets where it says this is an anomaly," said Chase.

Eventually, after discussion with the editors, reporters use the tips as potential leads for new stories.

The entire process involves heavy transcription work to log words from the meetings, and this is where AI comes in. The team is using LLMs from OpenAI to run through the data and text of bills and transcripts for the tip sheet, said Isler.

"There's a lot of not just raw data, raw structured and unstructured data, but also just plain old text that we need to go through, and the LLMs are really good at that," he added.

They are also working with chat systems that use custom-built AI models. As they move on to more robust transcription systems, they also plan to partner with a tech startup called Cast Insights which focuses on video transcription using several models including Gemini, DeepSeek and OpenAI tools.

The team

The team is led by David Lesher, a longtime California political journalist who co-founded CalMatters and served as its CEO and editor in the early years. He long admired the work done by the team that built the original Digital Democracy, led by CalPoly professor Foaad Khosmood, and worked with that team to design this new approach for version 2.0. In addition to Chase and Isler, the team includes engineering staff on board from CalPoly along with graduate or undergraduate student employees who have the skill sets required for the project.

"We've got a technical team of engineers who just focus on the tool set that allows us to not only get the videos and transcribe them, but also we have an internal tool for managing the workflows and managing who's working on what and assigning tasks to our transcribers," said Isler.

Meanwhile, the students are mainly transcribers who also verify AI transcription and assign statements to actual speakers since that was a component that AI was not able to do when they started the project. Identification of speakers is crucial so that they may be matched with the organisations they represent, allowing users to track the influence of donations from those organisations.

The opportunities: Boosting reporters' nose for news

The impact of the Digital Democracy tool on their newsroom has been significant, according to Chase, with the tool enabling multiple reporters to use its data to write stories about political influence and funding.

"We've built this in a way that is accomplishing the goal of actually identifying not just who the speaker is, but the entities and the money flow and being able to say confidently that this action that was taken by the State was funded by money from these interest groups," said Chase.

"I think we knew it would work, but seeing it actually show up in numerous stories lately has been tremendous. We've now trained 60 plus reporters outside of our staff around California to use this and they're starting to do a few stories," he added.

Lessons for newsrooms

- **Value external expert networks**: One of the big takeaways from being a part of the Innovation Challenge and developing the product further was the ability to tap into the network that JournalismAI offers and explore news ideas through it, said Isler.

  "It is really useful to have another kind of outside force that is also technical and more expert in some of these topics that we can kind of just ask random questions and get some feedback," he said.
- **Prioritise public mission and access**: Maintain a core focus on the primary mission (eg, explaining the state to its citizens) and prioritise public access and informing people over achieving commercial success alone.

Looking towards the future, Chase said CalMatters' primary mission remains to explain California to Californians and use information to improve the state, prioritising public access and informing people over commercial success alone. They aim to expand to multiple states where feasible, but not all 50, and also plan to adapt the approach to the city and county levels to provide local government information.

Editorially, their goals include enhancing the tool for California, serving local journalism, and commercially, generating revenue to support the nonprofit.

Visit our website for the full case study and to learn more about this project.

9.3 Full Fact: Using GenAI to track government promises

Full Fact, the UK's leading independent fact-checking charity, faced a daunting challenge after the 2024 general election: how to sustainably track approximately 300 pledges made by the Labour government throughout an entire parliamentary term.

This fact-checking organisation, known for its rigorous approach to verifying claims in public debate, recognised that traditional manual tracking methods would prove unsustainable over the five-year parliamentary cycle.

The problem: Sustaining momentum beyond the election cycle

When Full Fact began planning their Government Tracker, they spoke to other organisations who had attempted similar projects in the UK. The feedback was consistent: tracking government pledges proved extremely resource-intensive.

"Speaking to other organisations who'd tried this before, they told us it simply took up too much time and resources to keep going," explains Thom Callan-Riley, Full Fact's Delivery Manager. "You'd see these snippets of tracking that would start strong but then drop off. Yet the most useful time for this information is right before the next election."

The scale of the challenge was immense. Full Fact's editorial team identified around 300 trackable pledges in the Labour manifesto. Writing up what each pledge means, providing context, and continuously monitoring progress would require enormous human resources.

Louisa Wania, Full Fact's Fundraising Manager, frames the broader mission: "Our work aims to increase public trust in institutions and politics. We want to provide a trustworthy source people can turn to, allowing them to see in real time how promises turn into action – or don't."

> We could have built a chatbot quickly if we weren't fussed about accuracy. But we can't put anything on our website unless we're 100 per cent sure users are accessing verifiable information.

Building the solution: The genesis

Full Fact wasn't starting from scratch. The organisation had been developing AI tools for fact-checking for nearly a decade, primarily to help fact-checkers monitor large volumes of media content and identify claims worth checking.

"We already use AI to ingest hundreds of thousands of sentences daily across newspapers, TV, radio, podcasts, and social media," says David Corney, Senior Data Scientist on Full Fact's AI team. "But we don't use AI to do the actual fact-checking – it's not very good at that. Instead, it helps our experts work more efficiently."

Combined with their existing AI capabilities and partnerships – particularly with a research team at Cambridge University, a world-renowned expert in AI for fact-checking – they saw potential for a new approach.

The new workflow

Full Fact and Cambridge University developed a two-pronged AI solution. First, they use Google's Gemini to generate comprehensive questions that help fact-checkers explain each pledge in context. "The AI is quite good at capturing how most people would interpret certain terms and phrases," Corney explains. "It reminds our experts that while they understand the details, most people might interpret terms like 'inflation' differently."

Second, they created an automated monitoring system using ChatGPT to track pledge updates, searching trusted sources and creating timelines of officials' statements about specific pledges. The search process is iterative: the initial results are used to generate a series of more specific queries to find more sources that can provide more details that might otherwise be missed. This can discover, for example, a ministerial announcement that has an impact on the delivery of a pledge but doesn't directly mention the pledge in question.

"The AI analyses articles and determines which are most relevant to each pledge," Corney notes. "It then summarises findings as bullet points that our fact-checkers can quickly review."

Crucially, human oversight remains central. Fact-checkers review all AI outputs before publication, maintaining Full Fact's reputation for accuracy.

Progress and challenges

By July 2024, Full Fact and Cambridge University had improved their model's accuracy from 50 per cent to 80 per cent, however, that 20 per cent gap remains significant.

"80 per cent is incredible, but it's not 100 per cent," Wania emphasises. "Unlike other organisations that might accept some level of AI hallucination, that's simply not a risk we're willing to take."

The team also discovered that updating existing pledge pages proved easier than creating new ones – a finding that shifted their development priorities. They adapted by working on multiple project phases simultaneously rather than sequentially.

The opportunities: Creating a model for trustworthy AI in journalism

Full Fact's approach exemplifies responsible AI deployment in journalism. They have clear principles: AI assists but never replaces human judgment, all outputs include source links for verification, and transparency about AI use is paramount.

"We want to teach healthy scepticism about public information," Callan-Riley explains. "Users shouldn't just trust our output – they should click through to see the actual sources and evidence."

This commitment extends to advancing AI literacy more broadly. Full Fact sees their tools as opportunities to demonstrate AI's genuine capabilities while dispelling hype about what it cannot do.

Lessons for newsrooms

Full Fact's experience offers valuable lessons for newsrooms considering AI adoption:

- **Start with existing workflows**: Rather than creating new processes, Full Fact enhanced their established fact-checking methods with AI support.
- **Maintain rigorous standards**: "We could have built a chatbot quickly if we weren't fussed about accuracy," Wania admits. "But we can't put anything on our website unless we're 100 per cent sure users are accessing verifiable information."
- **Iterate based on feedback**: The team discovered their AI initially captured irrelevant local news stories. "The AI would find a county council hiring 17 teachers and think it was relevant to national education pledges," Corney recalls. "Our fact-checkers immediately flagged this wasn't what we needed."
- **Bridge technical and editorial teams effectively**: Callan-Riley's role as project manager proved crucial. "I ask clarifying questions that help open up understanding between technical and non-technical teams," he says. "It's about creating space where everyone can contribute meaningfully."
- **Partner strategically**: Working with Cambridge University provided additional computational resources and academic expertise.

The Full Fact team emphasises focusing on real problems rather than chasing AI trends. "Assess the challenges your organisation already faces," Wania advises. "AI is great, but it's not magic. It requires study, dedication, and patience."

Visit our website for the full case study and to learn more about this project.

AI implementation in small and medium-sized newsrooms: Principles, practices and potential

You've read 35 case studies of small and medium sized publishers from around the world, who, with a small budget were able to produce exceptional MVP's of their ideas. In this section we summarise the learnings and core principles for effective building and implementation strategies, sustainability considerations, challenges and opportunities for integrating AI into newsroom operations.

The emphasis from these organisations, many who have very small teams, is a strategy that leads with a human-centric approach. This approach is focused on AI that augments journalistic capabilities to enhance efficiency and support the goal of public value, while grappling with the daily complexities of ever-changing technologies and ethical responsibilities.

We first begin with the basics of AI implementation in newsrooms. It is crucial to begin by identifying existing problems rather than simply adopting technology for its own sake. The focus should be on human-AI collaboration, where AI augments human expertise for more routine tasks, allowing journalists to concentrate on complex investigations, editorial framing, and ethical oversight.

Publishers need to prioritise their organisational and editorial values, along with ethical considerations such as privacy and inclusivity. Furthermore, local context is vital, as generic AI solutions often fail to account for specific regional, linguistic, and cultural nuances. Coupled with this, publishers need to conduct extensive user research, including interviews and focus groups across various newsroom roles in order to uncover real needs and validate product ideas. Below is a summary of the lessons learned from these 35 grantees for other small and medium-sized publishers around the world.

On building and implementing AI tools

Plan carefully from the start. Even simple-seeming projects can have complex, interconnected parts. To prevent delays, you must carefully conceptualise the project, identify potential weak points, and understand how the different components depend on one another.

Establish clear editorial guidelines. Every AI project needs to fit within the organisation's editorial guidelines. This is especially important for sensitive or critical topics, as it helps ensure accuracy and sensitivity. Without clear rules, an AI tool could produce inaccurate or biased results or in situations where the tool is used in relation to the public such as personalisation.

Work in cross-disciplinary teams. Breaking down silos is a must. Bring together teams that have different skills, like technical development, product design, and newsroom experience. These cross-disciplinary teams are key to creating tools that are both technically sound and genuinely useful for journalists. If you do not have all the required skill-sets within your organisation, find partners to work and collaborate with.

Continuously test and fine-tune solutions. The AI field is always changing, which means you can't just build a tool and forget about it. You need to be ready to adapt based on how journalists or your audiences actually use them. This process is essential for making sure the tool stays useful and accurate over time.

Integrate tools seamlessly into workflows. As obvious as this sounds, for journalists to actually use a new AI tool, it has to fit naturally into their day-to-day work. Tools are more likely to be adopted when they seamlessly integrate into existing workflows, meeting journalists where they already work.

Document everything. To understand a tool's performance, you need to document everything, especially the conditions under which you test it. This helps you differentiate between variations in your data and the actual performance of the model. Comprehensive documentation is critical for effective troubleshooting and improvement.

Challenges and considerations

Many of the publishers on the Innovation Challenge confronted substantial obstacles spanning technological precision, data management, localisation, and organisational adoption. These challenges were frequently rooted in the mismatch between standard AI capabilities and the complex nature of real world journalistic inputs and workflows. The common key challenges experienced by these publishers included:

AI accuracy and handling complex inputs. Publishers struggled to prevent AI from producing hallucinations and inaccuracies, leading one organisation to abandon LLM written alerts in favour of a constrained, template driven approach backed by editorial checks. Some models had difficulty parsing unstructured journalistic content, such as PDFs, where they would take in all text, including extraneous information like journal names and page numbers, and not just the main content.

Data scarcity and system performance. Finding reliable, structured data was a persistent issue, particularly for projects involving Nigerian languages, requiring significant time for data cleaning and preparation. In other instance, teams working with under-represented languages faced similar challenges of limited datasets, leading them to re-strategise by engaging local communities to collaboratively collect and enrich data to improve their system performance. System performance was also a hurdle; one system initially took five to 10 minutes to process queries, which was deemed impractical for users.

Localisation, cultural fit, and cost barriers. The high cost of specialised AI tools priced in US dollars remains a barrier, making them inaccessible to many smaller organisations. Beyond cost, ensuring cultural resonance was necessary; one project changed its name to VANIA (Violence Armada em Números + IA) because the original name would not resonate locally in Brazil. In terms of function, dynamic features did not work precisely when adapting formats to different Indian languages.

Design, ethical, and financial hurdles: Teams often discovered they lacked internal UX expertise halfway through development, necessitating external consultations and significant restructuring of the product architecture. Ethical considerations were central for some, who made efforts to minimise LLM usage due to the large environmental footprint of AI, choosing the smallest-footprint models. Financial sustainability was a critical concern, with teams thinking about changing project architecture to make it cheaper to continue without donor support.

Opportunities

A core thread across the publishers is their commitment to ethical AI deployment, ensuring that technology assists human judgement rather than replacing it. Their tools are driving significant efficiencies in data processing, reducing costs, and automating repetitive tasks, which allows reporters to focus on in-depth journalism and increase content production. Furthermore, many of these tools demonstrate AI's potential to generate new revenue streams, foster deeper audience engagement through personalised content and chatbots, and build AI literacy among both journalists and the wider community.

Many AI applications are becoming more accessible "agents" that don't require heavy technical knowledge, making them viable for any newsroom. Individual journalists can experiment and adopt tools that suit them, though they should tell the wider team what they are doing. Thoughtful use of AI can empower journalists to make stronger and more impactful decisions. AI-powered prediction models can be used to craft targeted subscriber retention campaigns, even for newsrooms without massive internal AI teams.

Enhanced efficiency and digital transformation. Many organisations view AI primarily as a tool for increased efficiency, allowing newsrooms to produce more content faster and reduce costs. The implementation of YESEO at The Oglethorpe Echo led to increased efficiencies in digital product creation and reporting, which allowed them to move faster. This improvement in internal efficiencies translated directly into more resources for audience engagement. The tool developed by CINS and partners revealed significant potential for reducing repetitive, non-creative tasks in newsrooms. In another newsroom, Sowt anticipates optimised operational efficiency resulting from data analysis on production costs and content performance, helping them to identify inefficiencies and save costs.

Improving journalistic quality and context. AI is widely leveraged to improve the depth, accuracy, and trustworthiness of reporting. Full Fact exemplifies the creation of a model for trustworthy AI in journalism by ensuring AI only assists human judgment, outputs include source links, and transparency is important. Nawaat AI achieved a very high rate of accuracy (95-98 per cent) by using their entire archive and editorial charter as guidelines during optimisation.

Financial sustainability and revenue generation. Projects frequently open up new pathways to sustainability, revenue, and economic viability. The Oglethorpe Echo demonstrated tangible positive outcomes in audience growth and revenue generation via the YESEO project, specifically noting an increase in digital ad sales and subscriptions. The Republic's project, Minim, is more than a technical solution; it's ethical AI development that creates a marketplace where voice creators monetise their work, showcasing new revenue streams for the newsroom. Economía para la Pipol views its chatbot as a potential path to financial sustainability, exploring partnerships that require economic information translated for broader audiences.

Scaling, replicability, and regional adoption. A number of the organisations see their projects as templates for broader adoption, particularly within regions facing similar challenges. Agencia Mural created an alerting "robot structure" that is described as "completely replicable," meaning it can be pointed to other useful datasets, such as job opportunities or public services. The Makedonia project offers a potential template for regional media across Southern and Eastern Europe. Verify plans to expand its fact-checking capabilities by offering its AI model to other newsrooms in the Middle East. El Surti aims to support other organisations facing similar challenges by documenting their methodology and sharing their code for the AI KUAA project.

Audience engagement and AI literacy. AI tools are being used to deepen audience understanding, increase engagement, and foster greater AI literacy. In Mongolia, a sentiment analysis tool unexpectedly increased reader engagement, with time spent on articles tripling. This project also catalysed Mongolia's first community of AI-literate journalists. The chatbot development process provided Economía para la Pipol with unprecedented insights into their community, moving beyond social media metrics to understand why people seek economic information.

Sustainability and growth

One of the pre-requisites for participating on this Innovation Challenge was for grantees to have thought through their project's long-term survival, especially after this initial grant funding ends. Publishers needed to plan for financial sustainability. What could this look like? Publishers should find ways to lower operational costs, like optimising their infrastructure, and develop solid revenue models. A number of the grantees are considering selling their products to other newsrooms. Here are other strategies to consider.

Innovate with a purpose. Innovation should not just be about competing in the market. Instead, it should be driven by the practical need to improve your operations in challenging environments. By focusing on what you truly need, you can create more valuable and resilient solutions.

Maximise your archives. Newsrooms have a powerful resource they can leverage: their archives. By using AI, you can unlock the hidden potential in historical content, drawing out contextual meaning and capitalising on this valuable source of information.

Talk to your audience. Ask your audience for feedback on your work, especially on any AI tools that are public facing, but also ask them what other new services they would value.

Conclusion

The JournalismAI Innovation Challenge led to tangible solutions for the grantees with most of them sharing that they have achieved significant results or breakthroughs with their projects. A majority of grantees (96 per cent) **achieved a working prototype or tool** within the stipulated 9-month period, thanks to their participation in the programme. Many (70 per cent) also **improved internal workflows** and improved or **built AI literacy** within their teams. To a relatively smaller extent, 39 per cent **reached new audiences**.

Some of the outcomes as a consequence of these achievements were:

- It helped them serve their communities and users better
- It helped them in launching new products
- It helped them improve newsroom-wide AI literacy and understanding its applications.

The underlying theme and message from these grantees from small and medium organisations is that AI, and the funding to experiment with it, helps them overcome capacity constraints and the invisible barriers that it brings with it.

Acknowledgements

The JournalismAI Innovation Challenge, supported by the Google News Initiative was led by Polis Director Charlie Beckett and Project Manager and Team Leader Tshepo Tshabalala. The editorial responsibility for the content of this report lies solely with the JournalismAI team.

Additional JournalismAI team members who worked on this programme are: Senior Programme Manager Lakshmi Sivadas, Programme Officer's Igor Celov and Zena Achieng Onyango, as well as Programme Manager and Engagement Lead Ana Paula Valacco. We gratefully acknowledge the administrative support provided by the London School of Economics and Political Science Enterprise Ltd. for this programme.

Thank you to all the selection committee members as well as Ezra Eeman, who was the advisor for this programme. Thank you to the instructors who led sessions during the kick-off week.

This programme would not have been possible without the support of the Google News Initiative. Special thanks to GNI's Matt Cooke and Etan Horowitz for their vital work to make this happen.

Last but not least, we want to thank all the publishers who took the time to submit their applications to this programme as well as the inaugural 35 grantees featured in this report. Your input and participation was invaluable to putting this programme and report together.

Contributors

Tshepo Tshabalala
Project Manager and Team Leader, JournalismAI

Lakshmi Sivadas
Senior Programme Manager, JournalismAI

Igor Celov
Programme Officer, JournalismAI

Ana Paula Valacco
Programme Manager and Engagement Lead, JournalismAI

Zena Achieng Onyango
Programme Officer, JournalismAI

THE LONDON SCHOOL OF ECONOMICS AND POLITICAL SCIENCE

Get involved

Sign up for our newsletter if you want to receive updates about our activities and to meet the community.

If you have any questions about the project, or if you want to be involved in future JournalismAI initiatives, do not hesitate to get in touch with Tshepo Tshabalala at **t.h.tshabalala@lse.ac.uk**

bsky.app/profile/polislse.bsky.social

PolisLSE

JournalismAI, Polis
Department of Media and Communications
The London School of Economics and Political Science
Houghton Street
London WC2A 2AE

Google News Initiative JournalismAi POLIS Journalism at LSE