

THE END OF THE HYPE CYCLE: AI AS MORE THAN A NOVELTY

AI Launchpad EMEA 2025

CONTENTS

03 Introduction

05 Introduction to the AI Launchpad Program

08 From Experimentation to Impact

Building the Foundation for AI

Optimising Editorial Workflows

Pushing the Frontier

51 Conclusion

55 Appendix - Additional publisher case studies

INTRODUCTION

As we continue to explore what AI can do for news organisations in the media industry, the focus and topic have profoundly shifted from understanding and discovering the new technology through scattered experimentation to carefully planning, integrating and managing a portfolio of active initiatives. At the same time, models and tools are becoming more refined and easier to use, making approaching and using AI easier than ever before.

In conferences and roundtables on the topic, both sides of the argument around AI have moved to a more moderate, considered position. There is no longer so much a question of “if” AI will become part of the daily operations of news organisations, but publishers are asking “in what ways”, “when”, and “how profoundly.” On the other hand, AI is no longer regarded as a “silver bullet” that can solve every problem. What we are left with is a deep understanding that AI is an important tool to improve and support automating existing processes, and open new possibilities to increase revenues and engagement.

In this new era of AI as more of a business tool than a novelty, the central question becomes how to measure return on investment (ROI). ROI remains challenging, especially for early-stage or enabling use cases, but the effort to do so is indicative of the broader shift above. The costs of buying AI tooling

and building custom features and integrations need to translate into clear value, whether operational, editorial or commercial. Having a clear approach to ROI is an important enabler of the organisational adoption of AI.¹

Discussions around what value AI creates for news organisations fall into two broad categories:

- AI as an efficiency driver: many organisations use AI to reduce time spent on repetitive or administrative tasks (including metadata tagging, content translation or summarisation). These efficiency gains allow teams to reallocate time to core editorial work, helping to maintain quality while working at greater scale or speed.
- AI as an incremental revenue and engagement driver: others are exploring how AI can deepen

audience relationships or support new business models. This includes increasing engagement with existing readers, reaching underserved demographics with more tailored formats, or introducing monetisable products powered by generative AI.

Publishers are operating in a difficult environment characterised by economic and geopolitical uncertainty and pressure on existing business models. Subscription models are maturing and slowing, with a stagnant pool of 18% people paying for news in wealthier countries.² To continue growing their audience, publishers need to adapt to evolving behaviours, such as growing expectations around personalisation, convenience and interactivity. In this context, AI is not seen as a silver bullet, but as an important aspect of necessary product innovation and commercial agility.

¹ WAN-IFRA 'No longer optional': Why AI is now a strategic priority for journalism 2025-04-10 <https://wan-ifra.org/2025/04/no-longer-optional-why-ai-is-now-a-strategic-priority-for-journalism/>.

² Reuters Digital News Report 2025 page 11 https://reutersinstitute.politics.ox.ac.uk/sites/default/files/2025-06/Digital_News-Report_2025.pdf#page=11



This report brings together insights from the AI Launchpad 2025 programme, delivered by FT Strategies and the Google News Initiative (GNI). The programme supported 24 publishers across EMEA to identify, test and evaluate the impact of AI across editorial and commercial use cases. **This year's AI Launchpad marks a clear transition - from early experimentation to sharper decision-making about where AI can deliver real and sustainable impact.**

Across the AI Launchpad 2025 cohort, publishers focus more on embedding AI into existing tools and workflows rather than building standalone experiments. Internal CMS features, automation tools, and editorial support systems are increasingly AI-enabled. At the same time, many are investing in foundational capability, such as improving tagging, refining data pipelines, and strengthening cross-functional collaboration, to unlock future applications.

INTRODUCTION TO THE AI LAUNCHPAD PROGRAMME

Programme Overview

The AI Launchpad programme, a collaboration between FT Strategies and the Google News Initiative, ran from January to June 2025. It aimed to help news publishers harness the potential of AI by identifying and testing impactful use cases. This was the largest AI Launchpad yet, involving three cohorts of eight publishers each across EMEA and facilitating the exchange of knowledge about AI experimentation and business impact maximisation across a wide range of organisations.

Through a structured three-phase approach - Discovery, Experimentation and Future Planning - publishers were supported in evaluating and implementing AI initiatives that aligned with their strategic goals.

1 Discovery Phase: Identifying and Prioritising AI Use Cases

The programme began by exploring case studies of AI implementation across the media industry, categorised by business area. This year saw particular interest in themes such as investigative reporting, data journalism and video content production. Participants considered whether the solution would enable incremental revenue or cost savings, how quickly the use case could be deployed, and the potential impact on audiences.

We also assessed the feasibility of implementing each use case, considering factors such as the current capabilities of the latest AI models, the data inputs which would be required, and possible ethics considerations. Finally, we looked at the strategic priorities and business goals of every organisation to make sure the selected AI use case was linked to a specific target and area of focus within the 24 organisations.

To further guide decision-making, experts from the Financial Times, Google and other organisations provided additional training on topics such as AI governance, multi-prompt Large Language Model (LLM) workflows and agentic AI.

2 Experimentation Phase: Developing and Testing AI Applications

Each publisher selected one high-impact experiment to focus on in this phase. Publishers were encouraged to leverage advice from the Financial Times' data science team or support from third-party vendors offering software trials, in order to accelerate experiments while laying scalable foundations. The phase concluded with an experiment retrospective, where participants shared insights, reflected on challenges, and discussed strategies to improve the design and execution of future experiments. Inspiration Sessions continued, offering technical insights into machine learning systems and ethical



considerations to refine experimentation.

3 Future Plan Phase: Refining Organisational Governance Structures and Future AI Roadmaps

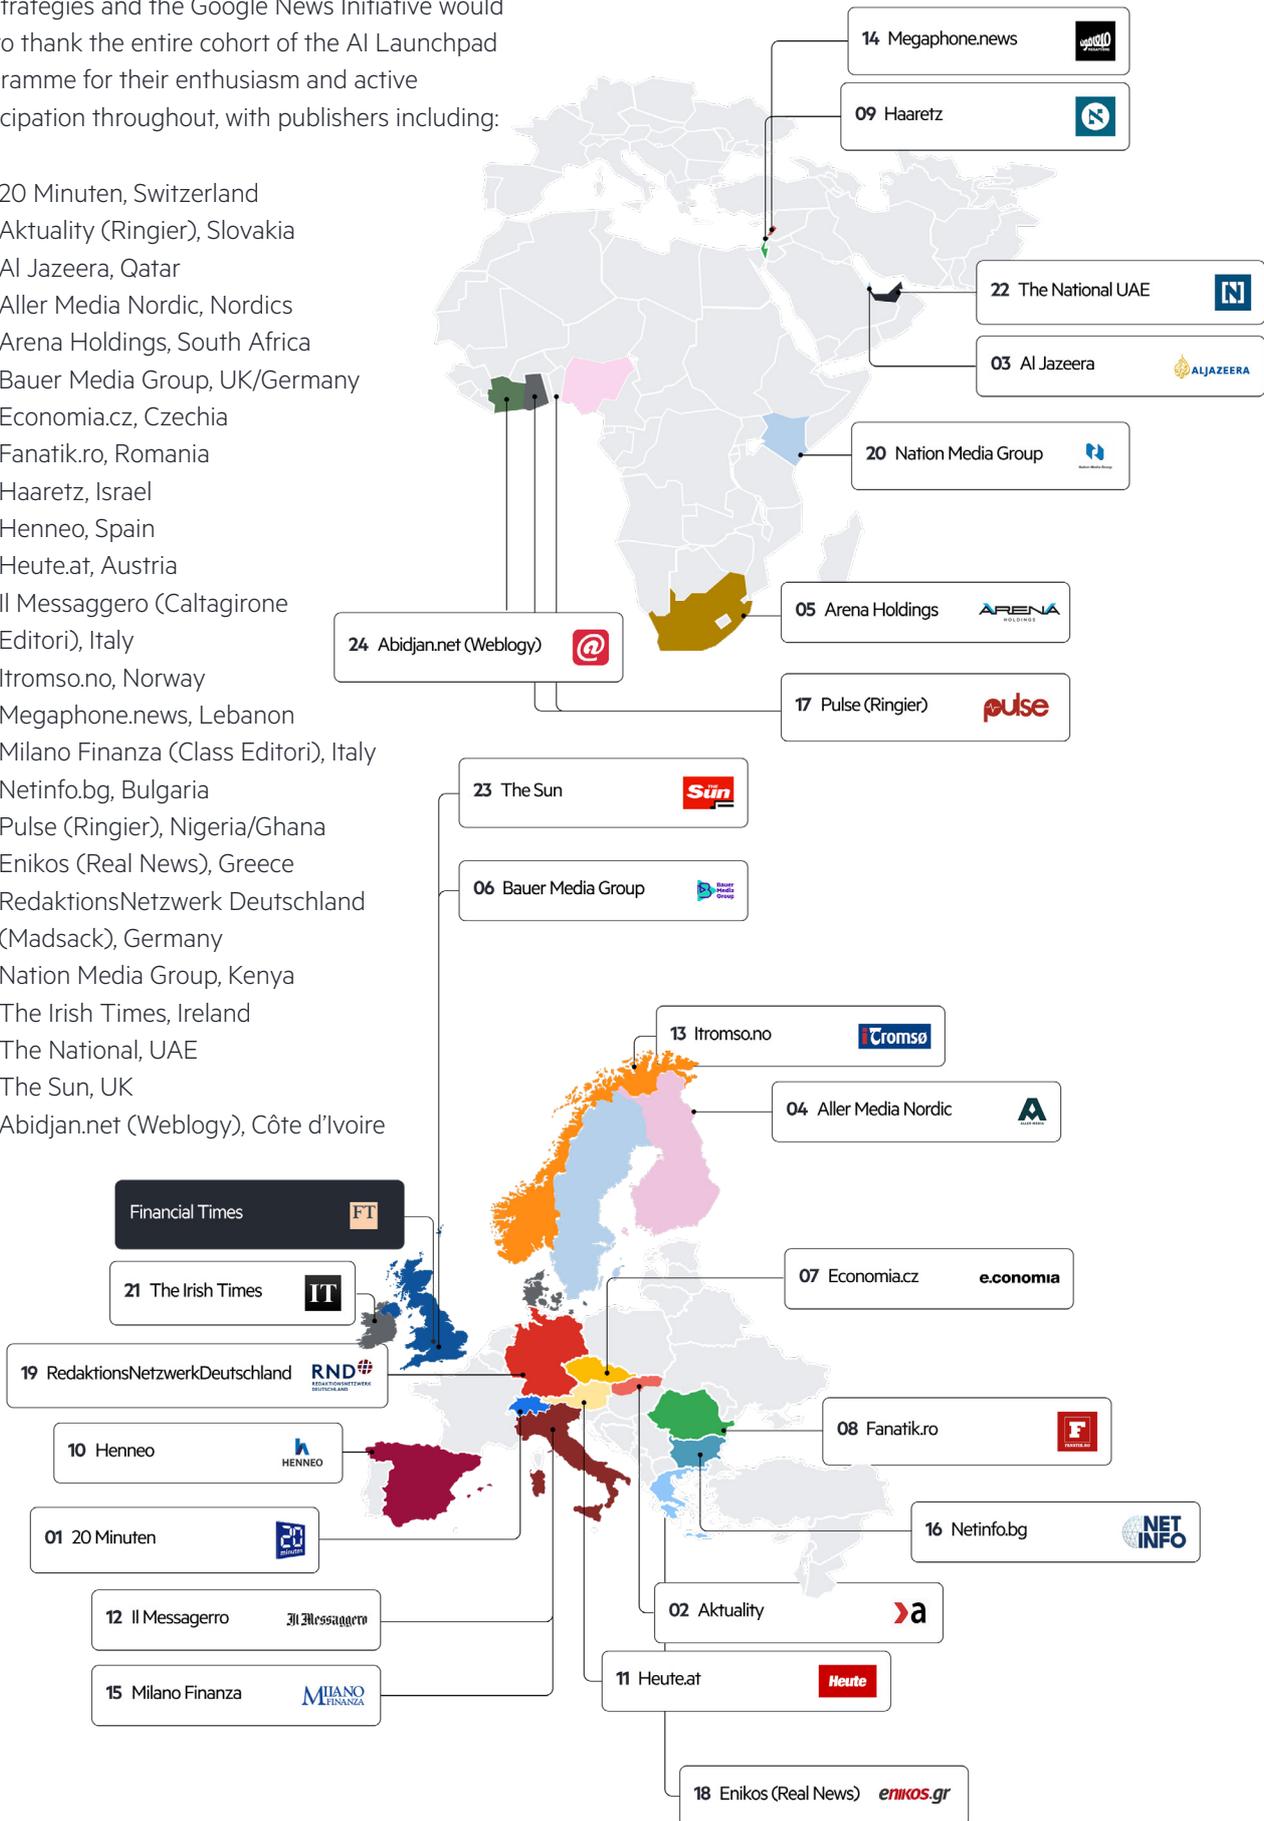
This phase aimed to ensure AI was integrated sustainably within publishers' workflows. FT Strategies consultants delivered in-person workshops that converted strategic discussions into actionable

steps for AI governance and implementation. Publishers received tailored guidance to help sustain progress beyond the programme, with a clear focus on managing change, equipping staff through training, and creating and maintaining clear communication at every level. These recommendations were shaped by the publishers' staff survey insights, which captured organisational attitudes to AI and data practices.

By the programme's conclusion, publishers were equipped with a global range of perspectives, a wealth of examples, and tailored strategies to incorporate AI into their editorial and audience engagement initiatives. The programme's structured approach ensured that a foundation for future success was laid by considering broader technical readiness and organisational alignment around and beyond the prioritised AI project.

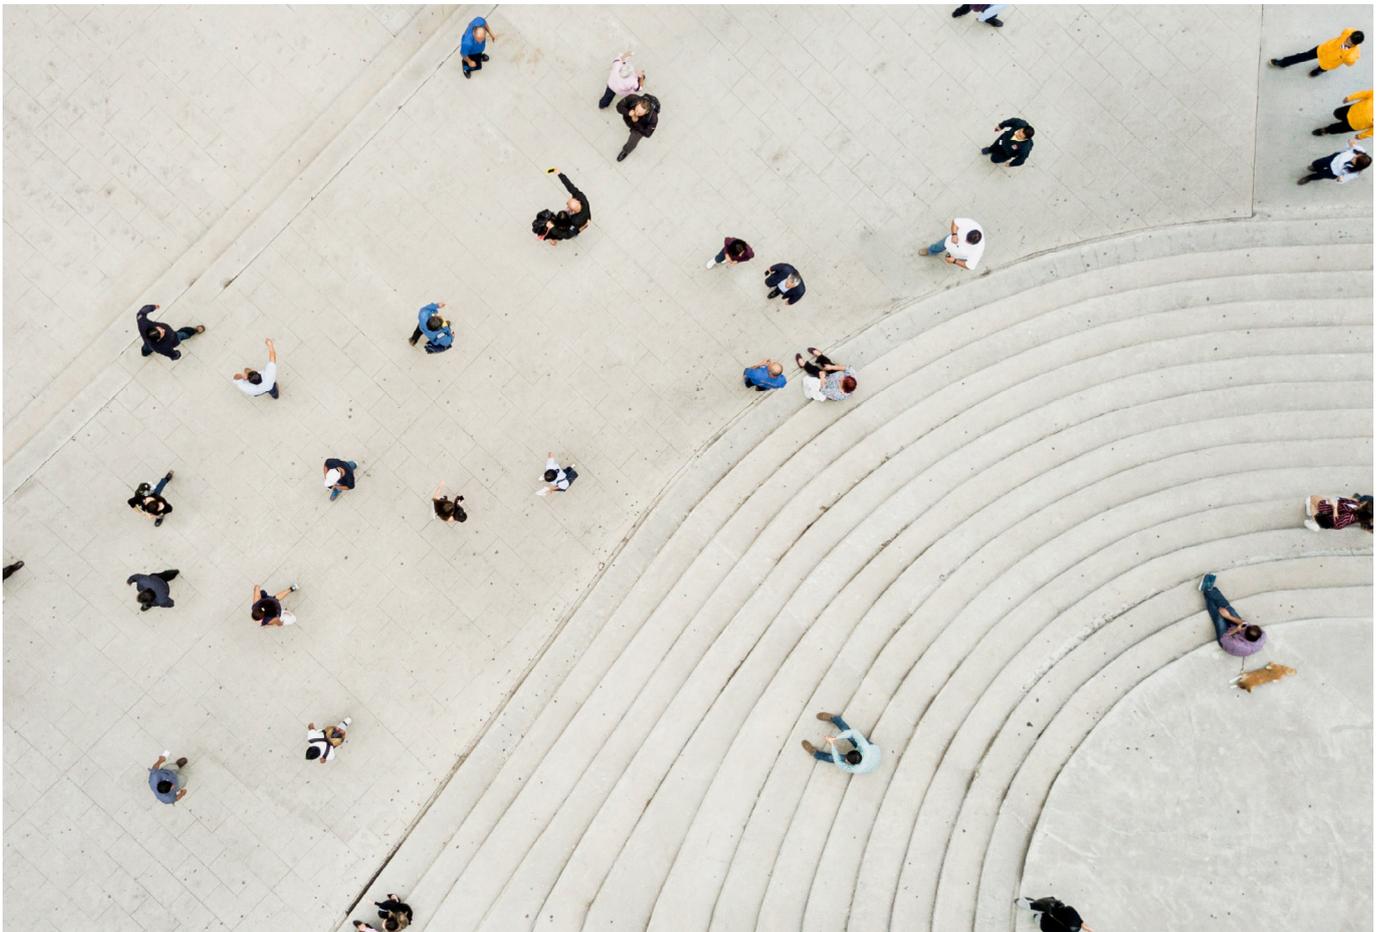
FT Strategies and the Google News Initiative would like to thank the entire cohort of the AI Launchpad programme for their enthusiasm and active participation throughout, with publishers including:

1. 20 Minuten, Switzerland
2. Aktuality (Ringier), Slovakia
3. Al Jazeera, Qatar
4. Aller Media Nordic, Nordics
5. Arena Holdings, South Africa
6. Bauer Media Group, UK/Germany
7. Economia.cz, Czechia
8. Fanatik.ro, Romania
9. Haaretz, Israel
10. Henneo, Spain
11. Heute.at, Austria
12. Il Messaggero (Caltagirone Editori), Italy
13. Itromso.no, Norway
14. Megaphone.news, Lebanon
15. Milano Finanza (Class Editori), Italy
16. Netinfo.bg, Bulgaria
17. Pulse (Ringier), Nigeria/Ghana
18. Enikos (Real News), Greece
19. RedaktionsNetzwerk Deutschland (Madsack), Germany
20. Nation Media Group, Kenya
21. The Irish Times, Ireland
22. The National, UAE
23. The Sun, UK
24. Abidjan.net (Weblogy), Côte d'Ivoire



FROM EXPERIMENTATION TO IMPACT

FROM EXPERIMENTATION TO IMPACT



Publishers in the AI Launchpad 2025 cohort pursued a wide range of experiments, but three strategic themes stood out: foundational capabilities, optimisation of workflows and storytelling enhancements.

Some focused on building foundational capabilities - such as tagging and

prompt design - to enable future AI adoption. Others concentrated on optimising editorial workflows, using AI to generate articles, audio or alerts in ways that support speed and efficiency without compromising quality. A third group pushed into exploring how AI could enhance journalism through visual storytelling, automated data visualisation

or multiformat engagement. While the technologies varied, each project was grounded in a clear business or editorial goal. Full case studies of all participating publishers can be found below or in the appendix.

BUILDING THE FOUNDATION FOR AI

Several publishers focused on strengthening the data that is most crucial to generative AI applications: their content. Tagging and metadata are essential building blocks for future AI use cases, which centre on connecting the right audience to the right story at the right time. Personalisation, recommendations and internal search all rely on consistent, well-structured and well-labelled data. Publishers saw tagging not as a one-off task, but as a reusable capability that could power future AI projects and support long-term editorial strategy.

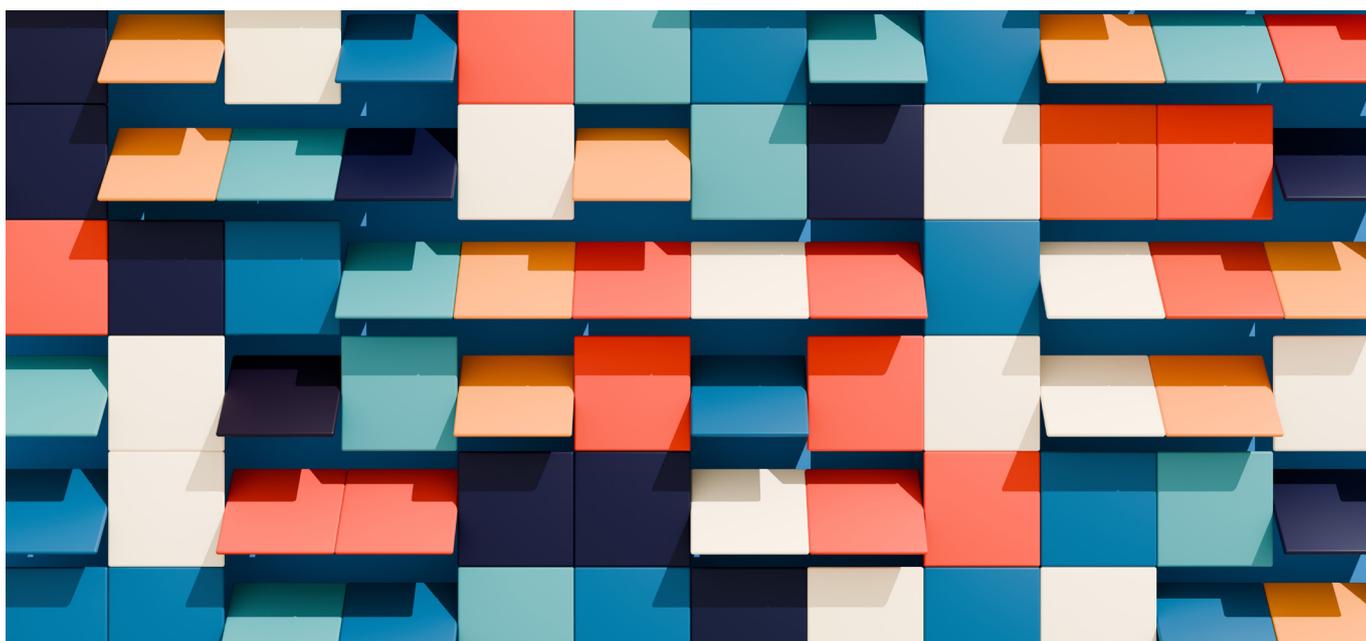
For some, AI-powered tagging became a chance to rethink how content is

categorised. Instead of applying off-the-shelf labels, teams developed custom taxonomies that reflected their editorial values or audience needs. One publisher created a bespoke User Needs framework and used prompts to help an LLM tag articles based on motivation rather than topic. This brought the additional benefit of shifting how the newsroom thought about engaging younger audiences.

Others focused on making AI-generated tags more accurate and trustworthy through editor involvement. Teams used techniques like “gold standard” tag lists to help guide the outputs of AI models. Editors could accept, reject

or amend AI suggestions, and provide feedback to refine LLM prompts and improve future results. These feedback loops show how AI can augment and amplify experienced editorial judgement. In practice, this increased confidence across teams and accelerated adoption.

Some of the publishers in the programme utilised the following framework from FT Strategies to kickstart their tagging efforts:



An important aspect of AI use case ROI is about the capability-building that experimentation can unlock. Developing prompt templates gave non-technical teams a better understanding of how LLMs work, while early tagging pilots helped identify gaps in internal workflows or CMS functionality.

- AI can be used to assist with this process

Assess current content tagging across publications

- List all publications and their CMS/Tagging platforms
- Export and catalogue existing tags from each publication
- Identify current tag types used
- Analyse current tag frequency and distribution
- Identify redundant overly broad and inconsistent tags
- Look for gaps in tagging (e.g. untagged or sparsely tagged content)

- Review current tagging guidelines and processes
- Understand how tags are currently used in business

Identify what tags the content should have

- Identify core tagging dimensions
- Interview stakeholders to establish any tagging gaps

Develop a unified tagging structure

- Design a centralized taxonomy
- Set governance and ownership

Implementation of tagging workflow

- Align technology and workflows
- Re-tag existing content
- Train editorial and associated teams on tagging best practices
- Monitor and refine taxonomy on an ongoing basis

A simple checklist of actions which a publisher can take to begin building a unified tagging structure.

PUBLISHER CASE STUDIES

FANATIK: EXPERIMENTING WITH AI-POWERED HEADLINE TESTING TO DRIVE PAGE VIEWS



The Situation

Fanatik is a Romanian digital publisher best known for its sports coverage, though it also produces general news and lifestyle content. The company operates an advertising-based business model, with no current subscription offering. Given this structure, traffic growth and user engagement are critical drivers of revenue.

To optimise ad revenue, Fanatik has prioritised initiatives that can increase unique users, page views, and repeat visits. At the same time, the publisher is working to streamline its editorial workflows and improve headline performance, which plays a pivotal role in attracting clicks in a crowded digital environment.

Fanatik joined the AI Launchpad programme at a relatively early stage in the company's AI journey. The organisation had not yet implemented a formal AI strategy or governance framework and was eager to use the programme as a springboard into more structured experimentation.



AI Use Case

Fanatik partnered with Smartocto to test AI-powered headline A/B testing, focusing on whether smarter homepage headlines could lift engagement. The tool allowed editors to test different headline variants on live homepage traffic. It did not change SEO or in-article headlines, focusing solely on the homepage version. The tool worked as follows:

- Once connected to Fanatik's website, Smartocto was able to rank existing article headlines based on their performance on the homepage. For example, article headlines placed further down would not be expected to perform as well as the main article.
- Editors then picked an article on which they would like to run an A/B test.
- Smartocto generated alternative headlines using generative AI, or editors could input their own variations.
- The tool then tested headline performance in real time by splitting homepage traffic and tracking click-throughs.

- After a defined testing window, the best-performing headline was then set live permanently.

The experiment targeted three content verticals: sports, news, and lifestyle. Key metrics tracked included page views as well as qualitative feedback from the editorial team on tool usability and functionality.



Results

While headline changes did not produce a significant uplift in click-throughs during the short testing period, the initiative served as a practical introduction to A/B testing methodologies within the newsroom.

Importantly, the tool helped foster a new culture of experimentation. Editors, who were previously unfamiliar with AI workflows, engaged with the tool and contributed feedback on its suggestions and interface. The process also surfaced important technical learnings, such as Smartocto's compatibility with the Romanian language.



Next steps

Fanatik sees this experiment as a first step in a longer-term effort to adopt AI more broadly. The company plans to:

- Continue experimenting with AI solutions that can support audience growth and content optimisation
- Improve internal communication and training to increase confidence in new tools
- Explore complementary use cases such as transcription and news monitoring
- Formalise its AI governance, ensuring future tools are implemented with clarity and accountability

As Fanatik begins to scale its experimentation, the A/B testing initiative has demonstrated the potential of AI to support key business objectives and has helped lay the foundation for a more data-informed, innovation-ready newsroom.



NATION MEDIA GROUP: IMPROVING SEARCH DISCOVERABILITY THROUGH AUTOMATED TAGGING



The Situation

Nation Media Group (NMG) is one of East Africa's largest media organisations, operating across Kenya, Uganda, and beyond through a broad portfolio of print, broadcast and digital brands. With an ambitious North Star goal of reaching 500 million monthly users and generating \$55 million in annual revenue by 2026, NMG joined the AI Launchpad to explore how artificial intelligence could support their growth, particularly by improving content discoverability, streamlining editorial workflows, and laying a foundation for more personalised reader experiences.

At the time, the business had some experience with AI through a few decentralised experiments, such as audio conversion, translation, and automated editing, and had created a basic internal framework for responsible AI use. However, these efforts were ad hoc and not yet embedded in daily operations.



AI Use Case

With improving discoverability, personalisation, and efficiency as key business priorities, the team identified improving metatagging as a foundational enabler. More specifically, during on-site workshops, the team had identified inconsistent and inaccurate tagging as a common bottleneck across the group, with their in-house CMS often surfacing irrelevant or incomplete keyword suggestions. As poor tagging was directly affecting SEO and limiting the performance of evergreen content, the team saw auto-tagging as a high-impact, achievable win. To explore this, NMG chose to experiment with iMatrics, an AI-powered auto-tagging tool.

The experiment was run on Daily Nation and Daily Monitor, two of NMG's highest-traffic titles. Due to time constraints, the tool was tested via iMatrics' standalone web portal rather than being fully integrated into the CMS. Editors manually reviewed the suggested tags and pasted them into the CMS for final publication. Over the course of the experiment, the tool learned from NMG's content, and its tagging suggestions became more accurate and contextually relevant.

While a full picture of the results will become available as the tool is scaled across NMG, significant search traffic gains were observed on already well-performing stories, particularly in sports and entertainment, with some articles doubling or even tripling in reach after enhanced tagging was applied. More moderately performing stories saw smaller, but still encouraging gains in search traffic of about 15-20%.

Although the team has not yet realised efficiency gains due to the lack of full CMS integration, the trial confirmed that better tagging has clear upside for search visibility and laid the foundation for future improvements in personalisation, content recommendations and UX. The test also helped validate iMatrics as a tool well suited to NMG's needs and highlighted the importance of giving AI tools time to adapt to local content nuances.



Fully integrate iMetrics into their CMS to automate tagging and unlock editorial time savings

- Expand usage beyond Daily Nation and Daily Monitor to include more brands, such as radio and broadcasting teams
- Utilise improved tagging to power future UX enhancements like personalised content recommendations and better internal search
- Track SEO impact and article performance more systematically, using dashboards to monitor trends and calibrate tagging quality
- Continue training and feedback loops between editorial teams and iMetrics to ensure tagging logic aligns with NMG's unique voice and taxonomy



Results

While a full picture of the results will become available as the tool is scaled across NMG, significant search traffic gains were observed on already well-performing stories, particularly in sports and entertainment, with some articles doubling or even tripling in reach after enhanced tagging was applied. More moderately performing stories saw smaller, but still encouraging gains in search traffic of about 15-20%.

Although the team has not yet realised efficiency gains due to the lack of full CMS integration, the trial

confirmed that better tagging has clear upside for search visibility and laid the foundation for future improvements in personalisation, content recommendations and UX. The test also helped validate iMetrics as a tool well suited to NMG's needs and highlighted the importance of giving AI tools time to adapt to local content nuances.



Next steps

NMG is now preparing to build on the momentum of this pilot through the following steps:

20 MINUTEN: EXPERIMENTING WITH USER NEEDS TO BETTER ENGAGE YOUNGER AUDIENCES



The Situation

20 Minuten is a print and digital Swiss newspaper written in German (with Swiss grammar rules) and part of the TX Group. Considered a commuter newspaper, it has a wide reach, covering urban Swiss regions as well as more remote mountainous areas. Despite having a loyal reader base, much of that audience is older, and the business had recognised a need to attract and better engage younger readers.

The publication had already taken important steps towards its ambition of attracting younger audiences, including refreshing the website to make the user interface light and intuitive. The team saw the AI Launchpad as a stepping stone to further accelerate this ambition, using AI as a means to get there. With a dedicated AI lead and close collaboration between product and editorial teams, 20 Minuten was well-positioned to make progress on this business opportunity.



AI Use Case

After considering a range of ideas, the team chose to focus on automatically tagging articles with user needs. Most publishers adopting the user needs



model stick to Dmitry Shishkin's six pre-defined categories: Update Me, Educate Me, Inspire Me, Give Me Perspective, Divert Me, and Keep Me On Trend. However, after some editorial reflection, the 20 Minuten team recognised that these categories didn't fully capture the motivations of their audience, who have specific and distinctive reasons for engaging with their content. This led to the decision to build the user needs model in-house with their own user need categories, rather than buying an off-the-shelf tool.

To build it, they created prompts using Gemini that provided the language model with clear guidance on when an article should fall under each of the custom user needs they had defined in their proprietary taxonomy. They tested the model on a sample of 100 articles, supplying the headline, description and full text for each one. To evaluate the model's performance, members of the editorial team manually tagged the same sample, and the model's tagging matched the human labels 75% of the time.



Results

The experiment generated strong cross-functional collaboration, with active engagement from editorial, data and product teams. It helped rebuild confidence in the use of AI, particularly among editorial staff who had previously been cautious. The model demonstrated that LLMs can reliably support editorial tagging, achieving 75% accuracy compared to human inputs. This gave the team confidence that AI could complement - rather than replace - editorial judgement.

The process also reinforced the value of using a custom-built user needs taxonomy, as off-the-shelf frameworks didn't fully reflect 20 Minuten's audience. Editors who typically rely on instinct responded positively to the structure, seeing its potential to guide content planning.



Next steps

20 Minuten is now building on the project's momentum through the following actions:

- Continue testing and improving prompt accuracy, targeting an 85% level of accuracy
- Integrate colour-coded user needs tags into the CMS and Kilkaya dashboard to inform real-time editorial decisions
- Support cultural change within the newsroom by embedding the user needs framework in training and daily workflows
- Experiment with using user needs insights to influence content commissioning and evolve editorial strategy, particularly for younger readers

THE IRISH TIMES: BUILDING A FOUNDATION FOR AI WITH AUTOMATED TAGGING



The Situation

The Irish Times is a leading national media organisation and a pioneer in digital subscriptions among Irish national newspapers. Its strategy is to drive digital growth and revenue to become a sustainable digital media organisation. While The Irish Times has deployed many digital technologies and built strong capabilities, it had relatively limited experience with building AI tools in the newsroom prior to joining the AI Launchpad programme.

Recognising that many promising AI use cases, including personalisation

and content recommendations, depend on well-structured metadata, the team identified automated tagging as a critical first step towards enhancing the utility and future applicability of the organisation's metadata. Their goal was to assess the potential of AI to improve tag consistency and quality across the article archive, and to help position the organisation to adopt more advanced AI capabilities in the future.



AI Use Case

Adopting a multi-phased experiment structure, The Irish Times first focused on the creation of a robust and com-

prehensive sample tag dictionary, generated through the application of an LLM to a raw batch of articles and enriched with a style guide, a historical tag corpus and bespoke Irish seeding to capture Irish place names.

In order to evaluate the effectiveness of this dictionary, a proof-of-concept experiment was conducted to assess the tag quality and relevance of vectorisation-based tagging (using this dictionary) as well as LLM-generated tagging using the Mistral-Large 2 model available within The Irish Times' Snowflake environment. The proof-of-concept experiment worked as follows:

- For the purpose of the POC, a small cohort of Features articles from a pre-publish queue within the CMS were used. These articles had already undergone the subbing process prior to inclusion.
- A bespoke Streamlit application was developed to review these articles and then suggest new and existing relevant tags.
- Editors (human in the loop) reviewed the articles and tags and provided feedback by accepting/rejecting these suggested tags.
- The new approved tag list replaced the original list and was then published as part of the article metadata when the story went live.
- The app recorded editorial feedback via tag acceptance/rejection and this was stored to improve the model's output.

The experiment was led by a cross-functional team from data, technology, strategy, subscriptions and editorial with FT Strategies providing extensive data science support in the initial creation of the tag dictionary.



Results

The proof-of-concept experiment demonstrated AI's potential to enhance the efficiency, effectiveness and consistency of tagging across The Irish Times' content, with editorial participants finding value in many of the tags suggested.

However, the process also revealed several key challenges, particularly around the nascent and limited size of the dictionary and some erroneous thematic linking of entities. Further refinements, including enhanced prompting and strengthened feedback loops, will be necessary before any introduction into the production workflow or retroactive application to The Irish Times' archive



Next steps

The experiment has laid the foundations for a working AI tagging system and for broader AI adoption within the organisation.

Key next steps for The Irish Times include:

- Optimise and fine-tune the tagging tool to address the identified limitations and improve its overall

performance and utility.

- Refine and further develop the base metadata dictionary.
- Formalise internal governance by establishing a new AI structure, rolling out a range of AI policies, appointing AI champions across teams, and establishing a roadmap for further experimentation.
- Broaden internal engagement, including onboarding more editors and product stakeholders in AI initiatives.
- Explore downstream use cases, such as personalisation, semantic search, and content recommendations, enabled by richer metadata.

Although this initial experiment was limited in its scope, it clearly demonstrated the value of AI-assisted tagging (when combined with human oversight) and surfaced a number of highly relevant tags that would likely have been missed otherwise. By starting with a practical, foundational use case, The Irish Times has taken a critical first step in unlocking the value of AI across its newsroom and product operations, setting it up for future experimentation.

OPTIMISING EDITORIAL WORKFLOWS

OPTIMISING EDITORIAL WORKFLOWS

Several publishers in the AI Launchpad 2025 cohort used AI to enhance the speed and output of their content production processes by reducing routine workload and enhancing responsiveness to audience needs. In most cases, the goal was to give editorial teams more time to focus on high-quality, original reporting while delivering faster, more engaging content in formats suited to modern news consumption.



These experiments had a common focus on well-scoped, editorially grounded use cases. Publishers did not attempt to blindly automate whole articles, nor did they aim to remove human oversight. Instead, they built tools that could handle repeatable structures or content types such as travel guides and interview highlights. Journalists remained in control, reviewing and refining the outputs before publication. In some cases, advanced AI techniques like chain-

of-thought prompting or structured inputs were used to improve reliability, especially where factual accuracy was essential; in other cases, this reliability was maximised through the design of user interfaces and feedback loops (such as article drafting agents which could reach out to a human for support during tricky steps).

AI also enabled teams to experiment with new formats and user experi-

ences, including audio and video. One publisher piloted daily AI-generated audio briefings, helping to establish a habit among loyal readers. Another used highlight detection to accelerate video editing, turning long-form content into easily shareable clips across web and social. These tools were aimed at unlocking new ways for journalists to reach audiences, thereby driving business ROI, but without adding significant workload.

PUBLISHER CASE STUDIES

AKTUALITY: USING AI TO GENERATE TRAVEL CONTENT AND EXPLORE NEW MONETISATION PATHS



The Situation

Aktuality is one of Slovakia's leading digital news publishers and a key title in the Ringier media group. Renowned for its investigative journalism, Aktuality launched a digital subscription model in 2022 as part of its North Star strategy to grow reader revenue and decrease dependency on advertising.

As an organisation, Aktuality had already made strong progress on its AI journey - implementing an internal AI ethics policy, regular newsroom training and nominating an AI Implementation Lead alongside the 15 AI tools that they have at their disposal. Therefore, the team was well placed to put AI at the heart of their drive towards increasing reader revenue.

One route to increasing subscribers in the long run is to improve the top-of-the-funnel by increasing reach and engagement among Aktuality's audience. Producing articles on summer holiday destinations was therefore seen as a potential experiment, given its popularity as a topic in Slovakia.



AI Use Case

The team developed an AI-powered article generator, internally referred to as the "AI Travel Service". Built on the Palantir platform and powered by Gemini, the tool allows journalists to produce structured travel articles by filling in a small number of input variables such as destination, language, season and country of origin (from what country

perspective to write). These inputs are automatically fed into pre-defined prompts. Once initiated, the system generates a full article split into thematic sections (e.g., "top things to do", maps, tourist activities), which is then manually transferred into the CMS and undergoes the regular editing process.

The tool was developed in-house by Aktuality's tech team, with FT Strategies providing guidance on prompt design and hallucination mitigation. Techniques such as chain-of-thought prompting were introduced to improve factual reliability. Editors remained involved in reviewing and refining outputs, ensuring the content met editorial standards.



Results

The travel article generator was soft-launched internally, allowing the editorial team to test the tool and assess its effectiveness. As part of the internal proof of concept, 4 articles have been produced in this way.

The project team demonstrated that an article could be produced and transferred into the CMS in 10 minutes, significantly faster than traditional workflows. Although uploading to the CMS for now remains a manual step (by copying and pasting each section), the tool demonstrated clear time-saving potential and strong content consistency. After internal testing, the team gradually started publishing more articles within the travel service. As of the end of July, they had already published more than 50 articles, but the team plans to have launched up to 100 by the end of the 2025 summer holiday season.

Aktuality anticipates that the public launch of these travel articles will help drive increased page views, particularly during Slovakia's peak holiday seasons. Interest from other Ringier group publishers suggests potential for wider adoption of the tool. Furthermore, thanks to the launch of this tool, Aktuality secured a collaboration with a

new business partner, which generated a new source of revenue.



Next steps

Aktuality is preparing to scale the tool across new geographies and seasonal themes, beginning with winter holiday destinations. Additional priorities include:

- Augmenting the articles with additional new formats, such as video and audio
- Expanding content generation to new verticals and countries within the Ringier network
- Integrating the tool more directly into the CMS to reduce manual steps
- Exploring direct monetisation opportunities through partnerships with travel-related advertisers, alongside the anticipated benefit to subscriptions

“The project team demonstrated that an article could be produced and transferred into the CMS in 10 minutes”

NETINFO: AUTOMATING NEWS AGGREGATION TO SUPPORT EDITORIAL EFFICIENCY AND SCALE



The Situation

NetInfo is one of Bulgaria's most prominent digital media organisations and part of the United Media Group. With a market reach of 88%, it holds a leadership position in Bulgaria's publishing landscape. As the company is primarily ad-funded, sustaining scale, quality, and speed of content production is critical to maintaining its business model.

Therefore, when Netinfo joined the AI Launchpad in 2025, they identified improving editorial efficiency as a significant growth lever. By streamlining the process of producing routine news articles, the goal was to give journalists more time to focus on creative and original reporting. As part of its broader digital transformation strategy, NetInfo turned to AI to explore how editorial processes could be augmented.

Through the AI Launchpad programme, NetInfo focused on developing a tool that would help journalists surface, synthesise, and draft articles more efficiently, reducing the manual work and time required by journalists.



AI Use Case

NetInfo developed a newsfeed generator - a dashboard that aggregates updates from multiple news sources, clusters them into topics, and generates article drafts with headlines, suggested images, and reader interest scores.

The system, built entirely in-house, was powered by Gemini and connected to two Bulgarian news agencies and one international English-language source. It automatically vectorised incoming news data, grouped related updates into clusters, and assigned a “newsworthiness” score to each cluster. Clusters above a defined threshold triggered the generation of article drafts, with international content automatically translated into Bulgarian.

The FT Strategies team provided strategic guidance during the project, particularly in regards to prompt design and how to evaluate ‘news-worthiness’.



Results

The dashboard successfully proved its value as a proof of concept. More than 800 articles were generated for internal review during the initial testing phase,

with positive editorial feedback. Editors reported that the generated drafts were largely accurate, with only minor issues such as occasional hallucinations and translation inconsistencies from English sources.

Although the tool had not yet been rolled out publicly at the time of reporting, internal results confirmed the concept’s viability. The team also deepened its technical capabilities, gaining experience in clustering, vectorisation, and multilingual prompt design.



Next steps

NetInfo plans to continue iterating on the tool with further editorial feedback, including improvements to functionality, filtering, and multilingual handling. The next stage of development will focus on:

- UI and CMS integration to streamline the editor workflow and reduce friction
- Scaling the tool across additional newsrooms and regions within United Media Group

- Combining the tool with other internal projects, such as their CMS assistant, to create a more comprehensive editorial assistant
- Rolling out training and internal communications to familiarise editorial teams with the new tools and encourage adoption

The long-term ambition is to embed this AI-powered workflow as part of a complete editorial suite across United Media’s titles - improving efficiency, preserving quality, and enabling journalists to dedicate more time to original, high-impact reporting.

ITROMSØ: BROADENING AI ADOPTION WITH A BREAKING NEWS ALERT BOT



The Situation

iTromsø is a regional news outlet based in Tromsø, Norway, part of the Polaris Media group. Despite being a small newsroom of roughly 20 people, iTromsø has earned national recognition for its investigative journalism and data-driven reporting. Prior to the programme, the organisation had already built several AI-powered tools, including one called DJINN, which scans public documents to assist with investigations and is widely used by journalists.

While tools like Djinn saw high usage among journalists for their simplicity

and reliable responses, iTromsø aimed to take AI a step further, with the newsroom playing an active role in creating the AI tools they needed. To achieve these goals, the team aimed to increase active AI competence and broaden technological knowledge across the newsroom. This would empower journalists to leverage AI for investigations and develop their own tools with a strong underlying foundation from the technical team.

Moreover, the organisation's North Star is to produce the best local journalism in the world and given that the population of the municipality is around 80,000, their focus was more on

deepening engagement with existing readers, rather than increasing reach. Therefore, the team needed to ensure they were providing timely stories on the region's most important events.



AI Use Case

During the Launchpad, iTromsø pursued two parallel initiatives. The "My School" project focused on increasing competence by building a state-of-the-art RAG research system and providing hands-on training to journalists. After this project ran its course (due to the halting school restructuring), the team

pivoted to the breaking news bot to enhance their breaking coverage and competitiveness and to be first to market despite the lean editorial team..

The chosen use case was the development of a breaking news alert bot, designed to help journalists respond more quickly and efficiently to incidents such as avalanches, accidents, or fires. The tool processes police logs from a public API, classifies the severity of each alert using AI, and delivers only high-priority items to a dedicated Slack channel.

Journalists can then ask for further support, such as a list of questions for follow-ups, contact information and the generation of a first-draft article. The system leverages retrieval-augmented generation (RAG) to suggest relevant background context and help structure article content.

Built entirely in-house by two members of iTromsø's product team, the bot was designed with newsroom integration in mind. With support from FT Strategies, the team explored different methods of classification, scoring, and prompt design, and structured the experiment to encourage early adoption.



Results

The tool successfully processed and filtered over 95% of high-priority police alerts and significantly reduced the time journalists spent scanning routine updates. Although the tool is yet to be fully rolled out to the newsroom, the product team considered the prototype a success, with strong internal validation and positive early feedback.

The use case also demonstrated the feasibility of embedding AI in more routine, high-frequency tasks and within tools that the newsroom regularly uses, such as Slack, expanding its relevance beyond data specialists to the newsroom as a whole.

In parallel, the project revealed organisational insights: namely, that having a dedicated project manager, experiment process and regular check-ins (this time aided by FT Strategies) would be vital to any future experimentation that the team undertook.



Next steps

Following the experiment, iTromsø plans to test the tool with journalists in the coming weeks and refine it based on their feedback. The long-term goal

is to integrate the alert bot into the CMS to streamline the publication process end-to-end.

Further priorities include:

- Training and onboarding: Live demos and small-group workshops are planned to encourage adoption and gather feedback
- Internal governance: iTromsø is working on an AI charter to guide ethical and transparent use of AI in the newsroom
- Communication: Email updates and offline channels will be used to raise awareness and build confidence in AI tooling
- Scalability: Success with the alert bot could inform future tools aimed at similar goals - expanding reach, saving time, and enhancing reporting across the group

Through this initiative, iTromsø demonstrated how even a small newsroom can use AI not just to enhance editorial output, but to make innovation a shared, accessible part of everyday journalism.

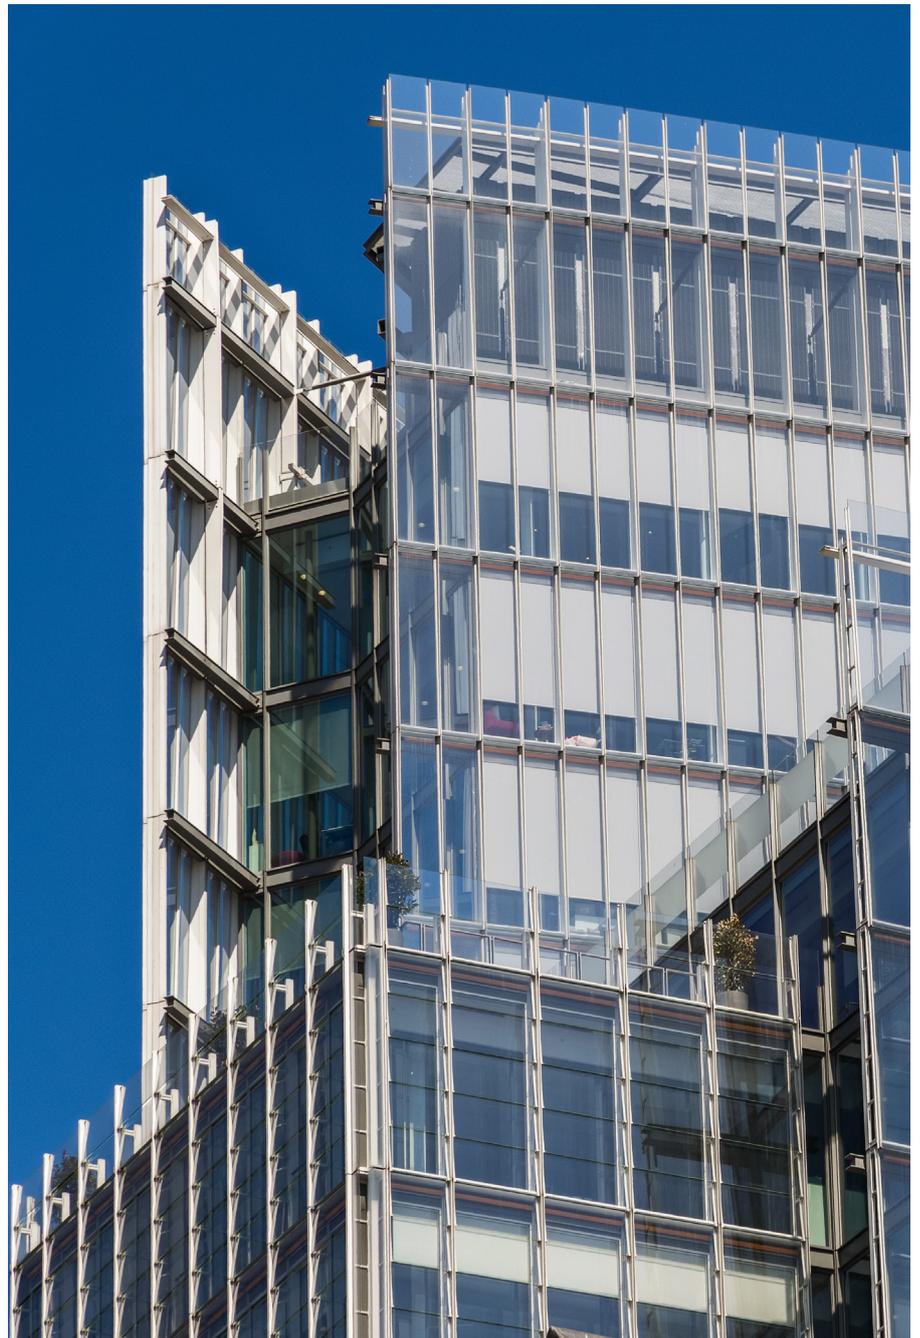
THE SUN: EXPERIMENTING WITH AUDIO SUMMARIES TO BUILD DAILY HABITS



The Situation

The Sun is a leading British tabloid newspaper present both in print and digital formats. It is widely recognised for its mass-market reach and its distinctive editorial style, characterised by short headlines, bold front pages, celebrity scoops, and human-interest stories delivered in a lively, conversational tone. In recent years, The Sun has launched a subscription offering - Sun Club - which focuses less on reaching the masses and more on providing a premium, curated product for its most loyal readers. As part of the strategic shift, the publisher was looking for ways to deepen engagement with its audiences so that they build loyalty and a habit with The Sun, and eventually feel compelled to subscribe.

Some internal AI experimentation had already been carried out, and the team was technically fluent, meaning the foundations for developing AI use cases were strong. The AI Launchpad was a great opportunity to bring a more strategic, problem-led approach to developing AI solutions and creating a reusable and scalable prioritisation framework, supported by FT Strategies, to focus on high-impact, high-feasibility initiatives.





AI Use Case

The Sun partnered with ElevenLabs to test a daily audio summary - a short text-to-speech bulletin including the day's top stories. User research had validated an appetite from readers for these daily briefings, and the aim was to give readers a quick, hands-free way to consume news, helping make The Sun an indispensable part of their daily routines.

The pilot for this use case starts in the morning editorial conference, where editors would choose 5–10 articles to feature in the daily audio summary. These were then summarised by an internal GPT model, which also converted them into a script for the briefing. The editorial team would review and edit the summaries and scripts as needed before passing them to Eleven Labs' text-to-speech tool. This transformed the script into a newsreader-style audio clip with preset speed and pitch. Another member of the editorial team would review the final clip as an extra quality check, focusing on clarity and pronunciation. The final output was a 2-minute audio summary of the day's top news stories.



Results

The internal pilot demonstrated that AI-generated audio summaries could be produced quickly (around 10 minutes end-to-end), sound professional, and provide an adequate listening experience for users. While technology was an enabling factor, clearly spelt out workflows across Product and Editorial teams were ultimately the success factors of the experiment. The team realised the importance of having quality checks throughout the process, to avoid the AI's inaccuracies and focus on key points and the relevant context during the generative process.

The experiment was launched to the public in early June. While results linked to the original goal - increasing engagement and ultimately driving uptake of Sun Club membership - are not yet available, the team is confident that the solution aligns with reader preferences and has been delivered in a scalable, cost-effective way that upholds the group's editorial standards.



Next steps

The Sun is now preparing for a wider audience rollout and further refinement of its audio summary workflow through the following actions:

Approach News UK radio colleagues for advice on audio production and long-term distribution strategy.

Allocate dedicated development time and streamline the workflow to remove extra copy-paste steps, reducing human friction.

HENNEO: ACCELERATING VIDEO WORKFLOWS THROUGH HIGHLIGHT DETECTION



The team had already made progress experimenting with AI in the newsroom - including transcription tools, basic summarisation, and internal automation via their own platform, ThallA - but they wanted to build a more strategic, intentional roadmap for AI adoption. The AI Launchpad gave them the opportunity to address this high-impact bottleneck in their video workflow and further build out their AI maturity.



AI Use Case

Henneo partnered with an external vendor, Video Highlight, to pilot a lightweight video highlight tool that automatically detects and timestamps key moments in long-form video recordings. The goal was to reduce editing time, increase the number of video assets that could be reused across platforms, and allow editorial teams to work more efficiently without adding to their workload.

The process begins when a video editor selects and uploads a raw video to the highlight platform. The system generates a transcript and uses it to identify potential highlight moments. Editors can also input tailored prompts - particularly in Spanish, which proved more effective than English - to refine the analysis. The tool then returns a list of suggested highlights with timestamps,



The Situation

Henneo is a major Spanish media group with a diverse portfolio of publications, including the free national site 20minutos and the subscription-based Heraldo de Aragón. The group is known for its commitment to local journalism and for embracing technology as a way to modernise its newsrooms and grow audience engagement.

At the outset of the AI Launchpad programme, Henneo was facing a clear operational challenge: its video team was producing over 600 videos per month, but identifying and editing highlights from long recordings was a time-consuming and manual process. This slowed content distribution across digital platforms and limited how often videos could be repurposed for articles or social media.

which the editor uses to extract the final clips for web, mobile, and social media.

By choosing a vendor-led approach rather than building the solution in-house, Henneo was able to launch the experiment quickly and focus its efforts on testing usability, integration points, and potential applications across the newsroom. The tool was also explored as part of a wider workflow that could eventually feed into a video-to-text pipeline, converting highlight moments into templated written articles.



Results

The experiment delivered immediate time savings. When tested on a 45-minute interview, the tool produced a highlight version in just 30 minutes, less than a third of the time it would typically take to edit manually. Editors found the suggested clips consistently useful and required minimal review to finalise. The tool was quickly adopted within the team and requested for future use.

Beyond speed, the pilot shifted the team's mindset about AI adoption. Rather than focusing solely on large, complex systems, they saw the value of small, incremental tools that deliver quick wins. It also reinforced the

importance of good inputs, such as clear prompts and relevant context, in driving high-quality AI outputs.



Next steps

- Henneo is now planning to build on the success of the video highlight pilot through the following steps:
- Integrate the tool into ThallA, their existing internal automation platform, to streamline access and reduce tool-switching
- Extend usage beyond the initial test case to more video types and departments
- Explore adjacent use cases, such as using highlight transcripts to generate templated articles or short-form content for social media.
- Continue refining prompt strategies to improve output accuracy, especially for different video formats or editorial needs.
- Assess long-term CMS integration opportunities to fully embed AI into the video production workflow.

“The tool produced a highlight version in just 30 minutes, less than a third of the time it would typically take”

THE NATIONAL: IMPROVE AUDIENCE ENGAGEMENT THROUGH CUSTOMISED USER EXPERIENCES



The Situation

The National is an English-language news outlet based in the UAE, with a clear ambition to grow its digital reach and audience engagement. At the start of the AI Launchpad programme, the business had set a strategic goal of reaching 200,025 registered users by the end of 2025 (up from 150,000 at the time) and was looking for ways to improve operational efficiencies, as well as launching a customer-facing AI use case for its readers.

While the team had begun experimenting with some AI-enabled tools, the capabilities of these tools were underutilised. The team approached the programme with the intent to explore how AI could support content distribution, streamline internal workflows, and drive higher audience engagement, particularly for registered users, whose average onsite engagement time was already strong at over 20 minutes per session.



AI Use Case

During the AI Launchpad, The National team focused on “liquid journalism,” which involves presenting content in various formats to cater to diverse user preferences and behaviours. This approach includes text-to-speech, summarisation, and interactive article engagement (chat- or speak-to-the-article). By offering multiple formats, The National aims to boost user engagement, particularly among registered users, thereby fulfilling a key strategic objective for 2025.

The team approached the experiment with scalability in mind, designing additional modules and features only after having validated user interest and increased engagement. The first module focused on bullet point summarisation for Business articles. This was combined with an existing text-to-speech feature that was already available in both English and Arabic.

The experiment followed the National AI charter and was developed in partnership with the editorial team, reflecting the outlet’s commitment to a human-first approach to content creation. All bullet point summaries are generated by AI and reviewed by editors.



Results

The bullet point summarisation for the Business section launched in July 2025, and while it is too early to understand the impact on user habits and engagement, it laid the foundation to build a full “choose your format” experience for readers of the National.

Part of the results of the experiment include:

- A redesigned toolbar at the top of the article page with the different formats available. At the time of writing, it includes “full article, key points and listen to the article” for the business section. As new modules and sections are added, the toolbar will reflect additional functionalities.
- A new way to approach experimentation and the development of new features that still encourages innovation at the National, but links it to user validation and business priorities.

Overall, this approach to experimentation makes sure resources are allocated effectively, avoiding the underutilisation of available tools.



Next steps

The National is now taking the following steps to build on this momentum:

- Validate user interest and KPI targets of the bullet point summarisation module
- Expand summarisation to other sections of the website
- Explore other features and modules to add to the “choose your format” toolbar

The team has also laid the foundation for tracking success metrics, including user engagement with summaries, time spent on page, and clickthroughs to related content.

MEGAPHONE: TESTING AN AI-POWERED RESEARCH ASSISTANT TO REDUCE NEWSROOM WORKLOAD



The Situation & AI Use Case

Megaphone is a small, independent digital news outlet based in Lebanon. At the outset of the AI Launchpad programme, the team had limited hands-on experience with AI but showed an openness to experimentation, especially after audience research revealed a surprisingly positive public perception of AI in journalism.

Their primary goal was to use AI to streamline editorial workflows, freeing up time for more in-depth reporting and investigative content, while supporting their business model through increased efficiency and stronger audience engagement. One of the most pressing pain points identified was the research process: fact-checking, verifying sources, and reviewing large volumes of background material were time-consuming and resource-intensive for a small team.

To address this, Megaphone chose to test Miso Studio, an AI-powered research assistant designed to help journalists verify claims, retrieve trusted background information, and streamline pre-publication workflows. The tool allows reporters to input questions conversationally (e.g., “What did Trump do during his visit to the UAE?”) and

receive summarised responses based on a curated set of whitelisted sources.

The team selected Miso after evaluating other options based on its usability, source control features, and alignment with Megaphone’s editorial standards. A small team of editors began testing the tool with a dataset focused on political and economic reporting. The aim was to measure whether Miso could reduce the time spent on research without compromising content quality.



Early Insights & Next Steps

As the tool had only recently been rolled out, formal results were not yet available at the end of the AI Launchpad. Initial testing highlighted opportunities to tailor the tool more closely to Megaphone’s needs, and this feedback was shared with the vendor for refinement.

The experiment gave the team valuable exposure to AI workflows and helped build internal confidence in testing new tools. It also surfaced important lessons about onboarding, expectation-setting, and how to evaluate tools against editorial standards.

Megaphone now plans to:

- Collaborate with Miso to address the areas identified during testing
- Conduct a second round of internal testing with an updated version of the tool
- Expand access to a broader set of journalists following successful validation
- Improve the structure of their content archive post-CMS migration to support future AI integration
- Continue embedding experimentation into their workflows with a focus on editorial integrity and practical value

PULSE: IMPROVING EDITORIAL EFFICIENCY WITH AI-POWERED PROOFREADING



The Situation

Pulse is a digital media organisation focused on youth audiences across West and East Africa, with operations in six markets and a strong emphasis on mobile-first, social-led publishing. As part of the Ringier Group, Pulse benefits from access to a shared AI infrastructure and list of approved use cases across the network while also contributing to the wider group with its own, tailored and bespoke AI solutions.

In a fast-paced publishing environ-

ment, Pulse produces a high volume of content each day. Maintaining consistent editorial quality at this scale poses logistical challenges. A small number of editors are tasked with reviewing dozens of articles across markets daily, leading to workload bottlenecks and inconsistent application of editorial standards.

Pulse identified AI as a tool to improve editorial workflows, reduce manual editing time and support the team in upholding quality. Through the AI Launchpad programme, Pulse set out to test whether an AI-powered proof-

reading tool could provide consistent editorial feedback that saves time, without replacing human judgment.



AI Use Case

Pulse partnered with Bridged Media to pilot a proofreading agent: an AI assistant that reviews written articles for grammar, spelling, tone, style, and consistency. The tool was trained using:

- Pulse's internal style guidelines

- A set of high-performing, well-edited articles
- `Feedback from editorial staff to refine both output and user experience

Once the tool suggests edits, editors can approve, reject, or leave comments on AI suggestions, helping the tool to learn and improve over time through machine learning. Early versions of the tool had usability issues - for example, when the agent highlighted issues within the article, it did not take you to the part of the article affected. However, close collaboration with the Bridged Media team resolved most issues.

The system's features include:

- Real-time identification of grammar and spelling issues
- Suggestions for style, tone and clarity adjustments
- Recommendations for stronger headlines or formatting improvements
- Adaptability for different markets or brands within the Pulse network



Results

During the trial period, the AI agent was tested on Pulse's site and integrated into the CMS to find the most efficient way to proofread articles. It was guided by real-time feedback from the editorial teams using it, with a feedback box built directly into the tool. While this process uncovered valuable insights, it became clear that a more tailored solution was needed, which, unfortunately, came at a higher cost.

The experiment also showed the potential value of the use case in reducing editor workload: the tool supports faster article review, helping the QA team prioritise their focus without sacrificing standards.



Next steps

Despite seeing the potential in the tool, Pulse decided to put the tool on hold as they better understand the costs and benefits associated with licensing. As such, the team will continue to:

- Search for suitable tools that can improve their efficiency and integrate easily within their existing workflows

- Evolve existing workflows to ensure they continue providing fast turnaround and serving the needs of the audience

Pulse stands out as a pioneering Ringier African media brand, demonstrating the effective and responsible integration of AI into editorial QA. This case highlights how AI can enhance and support human journalism through careful experimentation, rather than replacing it.

“Pulse stands out as a pioneering Ringier African media brand, demonstrating the effective and responsible integration of AI into editorial QA”

PUSHING THE FRONTIER

PUSHING THE FRONTIER

A number of publishers in the AI Launchpad cohort pursued experiments at the intersection of data, design and automation. Organisations focused on automating visual storytelling, especially through data visualisation. Publishers used AI to generate infographics, maps, videos and charts that would otherwise require manual production or external suppliers.



This allowed them to scale up visual output, reduce turnaround times and integrate graphics more directly into editorial workflows, ultimately helping journalists convey complex information clearly and quickly. Some publishers also explored conversational AI and personalisation.

These experiments also showed that advanced AI projects depend heavily on structured inputs. Teams invested

time and resources in creating clear, reusable data structures, such as blueprints for visual layouts and tagging systems for chart selection. While not all outputs were production-ready (for example, the output quality of video generation tools remained below a publishable threshold), experiments helped teams clarify 'what good looks like' in an emerging medium and built shared understanding around editorial expectation and technical blockers.

What these frontier experiments shared was a clear editorial goal, a willingness to test rapidly and iteratively, and a strong sense of where human oversight adds value. While not all outputs were perfect, each test helped clarify where AI can meaningfully extend a newsroom's reach and capability - and where the technology still has limits.

PUBLISHER CASE STUDIES

BAUER MEDIA GROUP: TESTING AI-POWERED VIDEO GENERATION TO UNLOCK MULTIFORMAT ENGAGEMENT



The Situation

Bauer Media Group is a multinational publisher with operations in the UK, Germany, France and Poland. Known for its diverse portfolio of magazines, radio stations, and puzzle content, the organisation is simultaneously managing the sustainability of its print business while investing in digital transformation. To support this shift, Bauer launched

AI Embrace - a structured internal programme to explore and experiment with AI across markets and brands.

The group were familiar with AI and what it could do. Their capabilities were strong, and they had developed AI tools, including their internal GPT-BauerGPT. However, the complexity of the organisation, spanning multiple countries, brands and formats, made it difficult to identify and prioritise AI use

cases that aligned with the strategic direction set by AI Embrace. The AI Launchpad gave the team the time and space to work on use cases that fed directly into their strategic thinking. The group had a large archive of content across multiple brands and was particularly interested in how to repurpose these assets to grow engagement, especially on social media and video-first platforms.



AI Use Case

Bauer focused on their food brand Lecker.de, exploring whether recipe articles could be turned into visual, social-friendly videos using generative AI. The goal was to better understand if AI-generated video could support multiformat content production in a scalable, cost-efficient way.

Rather than aiming for a polished final product, the team approached the experiment as an exploratory prototype.

The focus was on defining what an ideal output would look like - for example, a realistic video of someone chopping rhubarb - and assessing how close today's AI tools could come to that standard. The editorial team worked on structuring their recipe content to make it more "prompt-friendly," while also designing example prompts to feed into a range of video generation tools. With guidance from FT Strategies, the team tested tools including Kling, VO2 and VO3, assessing their strengths and limitations for Bauer's specific needs.



Results

The experiment revealed that while the concept was promising, the current generation of AI video tools was not yet fit for purpose. Attempts to generate realistic step-by-step cooking visuals often produced unrealistic outputs, such as rhubarb unexpectedly morphing mid-scene. However, the test proved valuable in surfacing what is and isn't possible with today's tooling and in highlighting the importance of structured, consistent source material for future experimentation.

The project also helped instil a stronger experimentation mindset within the team. It created space and time for hands-on testing and reinforced the value of rapid, low-stakes prototyping — even when the output is imperfect. Stakeholders appreciated having dedicated time to explore new tools in a structured way, and the test laid useful groundwork for future video innovation within BauerGPT and beyond.



Next steps

Bauer is continuing to build on the learnings from this early experiment by:

- Continuing to test video generation tools such as VO3 as the technology improves
- Exploring alternative use cases for AI video beyond recipe content, including branded content and simple visual explainers
- Improving the promptability and structure of their evergreen content to better support future experiments
- Sharing learnings from the Lecker.de experiment across other brands in the organisation to inspire and inform new use cases

AL JAZEERA: TRANSFORMING BREAKING NEWS INTO REAL-TIME BROADCAST VISUALISATIONS

This case study represents a work in progress. A comprehensive results assessment will be provided upon completion of the 8-week extension phase currently underway.



The Situation

Al Jazeera is one of the world's leading international news networks, known for its comprehensive coverage of global events and particularly strong reporting from the Middle East and North Africa region. Operating in a fast-paced, 24/7 newsroom environment, the network faces the constant challenge of rapidly converting breaking news content into compelling broadcast-quality graphics that can keep pace with its editorial cycle.

At the start of the AI Launchpad programme, Al Jazeera had already deployed their AJ Now platform for news aggregation and analysis, but identified a critical bottleneck in their workflow: the manual process of transforming unstructured news data into visual content for broadcast was too slow for their real-time reporting needs. This was particularly challenging for regional conflict coverage, where rapidly evolving situations require immediate visual context through maps, timelines, and

location-based graphics.

The network recognised an opportunity to leverage AI to bridge this gap between their rich data sources and their visualisation requirements, building on their existing technical infrastructure while addressing a high-impact operational need.



AI Use Case

Al Jazeera is developing an AI-powered system that automatically extracts key data points from its AJ Now platform and transforms them into broadcast-ready visualisations. The system focuses on regional conflict updates, converting breaking news stories into structured map-based graphics and timelines that can be rapidly deployed on-air.

The workflow begins with the AI system analysing breaking news content to identify locations, event types, and key details. It then maps these elements to appropriate visual components, such as

icons for different event types (drone attacks, explosions, political announcements) and geographic coordinates for accurate placement. The system exports this structured data in formats compatible with Al Jazeera's broadcast systems, particularly their Vizrt graphics platform.

A key technical innovation has been developing dual export pathways: KML files for external mapping integration and XML files for direct import into Vizrt templates. This approach allows editors to either work with external mapping tools or directly import pre-styled graphics into their broadcast workflow, providing flexibility while maintaining speed.



Progress and Early Results

The system achieved its first real-world deployment during breaking news coverage of events in Syria, where it successfully identified and mapped six conflict locations with sufficient accuracy for broadcast use. While editors still

needed to manually adjust labels and icons, the system significantly reduced the time required to create location-accurate maps compared to fully manual processes.

Technical developments have included enhanced location accuracy through improved contextual analysis (considering country and regional information alongside place names), implementation of breaking news filters with intelligent duplicate story detection, and search functionality within the mapping interface. The system now successfully merges similar stories to prevent redundant coverage while maintaining the ability to filter between AJ Now content and breaking news feeds.

The team has made substantial progress on broadcast integration, working closely with Vizrt to optimise data export formats and exploring both KML and XML pathways for graphics integration. A breakthrough came when the team successfully configured XML exports to maintain proper styling and icon mapping, allowing graphics to import directly into Vizrt with correct visual elements rather than requiring manual reconfiguration.

User feedback systems have been implemented and are actively informing model improvements, with structured

testing protocols developed to capture location accuracy, event type classification, and user experience insights. The feedback data has revealed specific geographic areas requiring enhanced accuracy (including Iran, Pakistan, Syria, and Palestine) and informed targeted prompt optimisation strategies.



Next steps

Al Jazeera is now preparing to scale and refine this work through the following actions:

- Complete Vizrt integration optimisation following technical meetings to finalise KML and XML export specifications
- Implement advanced search functionality, including semantic search capabilities and prompt-based querying
- Expand breaking news integration with real-time filtering to prioritise live events over historical content
- Deploy enhanced user testing protocols across broader editorial teams once core functionality is stabilised
- Establish systematic feedback

loops to continuously improve location accuracy and event type classification

- Develop comprehensive training materials and workflow documentation for newsroom-wide deployment
- Create governance frameworks for responsible AI use in live broadcast scenarios, building on industry best practices

IL MESSAGGERO: AUTOMATING DATA VISUALISATION TO SCALE DATA JOURNALISM CAPABILITIES



The Situation

Il Messaggero is one of Italy's leading newspapers, known for its strong national reporting, deep local coverage of Rome and market leadership position across central Italy. At the start of the AI Launchpad programme, the newsroom faced a clear operational constraint: infographics were outsourced to an external agency that could only deliver two visuals per day. This bottleneck limited the number of stories that could be supported with data-rich graphics, raised production costs, and slowed workflows, especially under the

pressure of tight print deadlines and a competitive digital market.

While individual reporters were experimenting with AI tools on an ad-hoc basis, the organisation lacked shared infrastructure, consistent tooling, and coordinated governance around AI. That said, Il Messaggero had already begun to develop meaningful internal capabilities; a small in-house data-science team had built AI-powered functionalities, including a dynamic paywall, or a smart archive with NLP-based tagging. These early investments provided a strong technical foundation for Il Messaggero's team as they ventured into the programme.



AI Use Case

To tackle the visual production bottleneck, the team developed an in-house tool called Data Painter, designed to turn article text and numerical data into publish-ready infographics within minutes.

The tool processes each article using a layered system of AI models and natural language processing. It begins with the tool spaCy, which cleans the text and tags important elements like numbers, names and locations. A first Large Language Model (LLM) identifies the most relevant data points and selects an

appropriate chart type. These decisions are structured in JSON, a format that serves as a blueprint for generating visuals. A second model interprets this blueprint to describe how each chart component should appear, while a third, code-focused model translates those descriptions into fully functional HTML, CSS and JavaScript. A layout builder then assembles the final visual arrangement, which appears live inside the CMS. Reporters can review and edit chart text, regenerate visual elements or reorder layouts before exporting the final infographic for publication in either print or digital formats.

Over the course of the project, the team improved the user interface, added editable elements and expanded the template library to increase the range and flexibility of outputs. FT Strategies supported the design and iteration process, helping the team improve prompt modularity, chart logic and model consistency.



Results

Data Painter delivered immediate benefits to the newsroom. Infographic production time fell from roughly 2 hours to under 15 minutes, allowing journalists

to generate visual content without relying on external suppliers. In one week, the tool produced more graphics than the studio had typically delivered in a month, at no marginal cost. In addition, the intuitive UI made adoption from the digital newsroom easy, with positive feedback on the usability of the tool.

A working prototype was embedded into Il Messaggero's CMS and tested by an initial cohort of three journalists. The results indicated not only time and cost savings, but also stronger integration of visual storytelling into editorial routines.



Next steps

Il Messaggero is now planning to scale and refine the use of Data Painter through several actions:

- Roll out the tool to all 18 web journalists and extend use to print editors
- Expand the template library to include more topics such as economy, culture, sport and society
- Embed a feedback loop so that journalists' edits inform future prompt engineering

- Allocate dedicated developer time to continue iterating on the tool without competing priorities
- Draft internal guidelines to standardise prompt-writing and define editorial review checkpoints

Finally, given the interest in the tool from other news organisations, Il Messaggero's parent company is planning to package the Data Painter as self-standing software with a long-term vision of commercialising the product, potentially creating a new revenue stream for the company.

“Il Messaggero's parent company is planning to package the Data Painter as a self-standing software”

ARENA HOLDINGS: USING AI-POWERED VISUALISATION TO EXPAND EDITORIAL CAPACITY



The Situation

Arena Holdings is one of South Africa's leading media organisations, home to brands such as Business Day, The Sowetan, and The Herald. At the start of the AI Launchpad programme, Arena was focused on advancing its digital transformation by improving the efficiency of editorial workflows, increasing personalisation, and strengthening its value proposition for subscribers.

The team saw AI as a potential enabler of these goals and had begun informal experimentation across areas like data verification and content analysis. However, these efforts were mostly ad hoc and not yet integrated into day-to-day operations. In particular, the team identified data visualisation as a strategic area of interest — one that could help them enhance reader engagement and expand their capacity to deliver compelling, data-rich journalism across brands.



AI Use Case

To address these challenges, Arena partnered with Flare Data - a tool that automates data analysis and visualisation - to test whether AI could help increase the volume and quality of



graphics across their brands. The idea was first discussed during in-person sessions with FT Strategies' team and quickly emerged as the most impactful option.

Editorial teams across Business Day, The Sowetan, and The Herald received training and were given access to test the platform. The Business and Mining desks were selected for the pilot, as these teams regularly work with structured and unstructured data. Flare Data AI ingested datasets on economic topics such as inflation and economic performance, producing over 7,000 insights. The experiment was designed with a focus on internal goals rather than direct audience impact. Success was measured primarily through editorial KPIs, including the volume of visualisations produced, the quality of the outputs, and the tool's usability within newsroom workflows.



Results

The pilot offered a valuable first step in exploring how AI can enhance Arena's visual storytelling capabilities. Although initial uptake was gradual, designated teams at Business Day, The Sowetan, and The Herald reported growing confidence in using the tool as the experiment progressed. Several editors

mentioned that the platform became easier to navigate over time, and its export functionality proved helpful for integrating visuals into existing workflows.

The experiment also surfaced clear opportunities for Arena to deliver more data-rich, engaging journalism, particularly for subscriber products. While the trial version of Flare AI relied on historical data, limiting its immediate use for breaking news, it helped the team understand the broader potential of the tool and how real-time data integration could unlock further impact.

Importantly, the pilot also brought forward key learnings around how to structure and scale future AI experiments, including the need for dedicated onboarding and clear success tracking. These takeaways are already shaping Arena's approach to future experimentation and internal rollout.



Next steps

Arena is now preparing to build on the experiment through the following actions:

- Enhance success tracking using dedicated dashboards and metrics aligned with editorial priorities
- Expand data inputs by onboarding real-time datasets to allow for use in breaking news and live analysis
- Decide on investment in the full version of Flare AI based on refined use cases and user feedback
- Formalise the experiment with a broader set of participating journalists and clearer workflows

ENIKOS: USING AI-GENERATED VIDEO TO ENGAGE YOUNGER AUDIENCES



The Situation

Enikos.gr is a Greek news publisher launched in 2012 by a prominent broadcaster, Nikos Chatzinikolaou and runs a predominantly advertising-led business model. While it has built a loyal readership among older audiences, the organisation has identified a strategic need to better engage younger demographics, particularly through multimedia content. Until recently, the intensity of daily coverage had left little room for exploring more compelling formats like video.

The organisation had previously experimented with AI tools such as image generation and content summarisation, but these attempts had not yielded strong results. Click-through rates on AI-generated visuals were low, and summarised content still required substantial editorial rewriting. As a result, the editorial team remained cautious, placing a strong emphasis on responsible, human-reviewed AI use.

Entering the AI Launchpad programme, Enikos aimed to develop an effective, practical tool that could clearly demonstrate the value of AI to internal stakeholders while helping the newsroom meet a pressing strategic goal: expanding engagement among younger readers.



AI Use Case

The team focused on a **video automation tool** designed to convert articles into short, engaging video summaries suitable for social media. The idea was to generate 10–20 second clips using bullet-point summaries of the article and the linked images. The images are animated through pans, zooms and perspective changes to create an engaging video format for younger readers.

The system, built entirely in-house using Gemini, pulls article text and images into a templated structure governed by pre-defined rules. For example, longer bullet points stay on screen for longer, and all videos include a consistent intro and outro animation to reinforce brand identity. Achieving high visual quality was a key priority for the team: after initial internal feedback, the team revisited font choice, layout, and aesthetic elements to achieve a more premium look and feel.

While the tool is not embedded directly into the CMS, journalists can trigger video creation via a semi-automated workflow and have full editorial oversight. A human-in-the-loop process ensures that all videos meet Enikos' editorial standards before they can be published.



Results

The automated video tool has enabled Enikos to start producing multimedia content consistently - something that was previously out of reach due to limited resources. It has allowed them to:

- Create videos efficiently using existing article content.
- Maintain editorial standards through human review.
- Share visually engaging summaries optimised for younger audiences.
- Integrate video production into daily workflows without overwhelming Enikos' small team.

Although full audience data is still being collected, early feedback from the editorial and product teams has been positive. The collaborative development process fostered cross-team learning, and the tool now serves as a practical example of how AI can meaningfully support newsroom goals.

CONCLUSION

CONCLUSION

In 2025, publishers are no longer asking whether to use AI. They are now deciding where to accelerate, where to pause, and how to ensure AI delivers meaningful value.

Across the AI Launchpad cohort, we see a spectrum of maturity. Some are just beginning to lay the technical and organisational groundwork for future experimentation. Others are integrating AI into daily editorial workflows to increase efficiency in increasingly challenging reporting environments. A smaller group is pushing at the edges of what is technically possible, experimenting with advanced applications aimed at transforming the audience experience.

What unites these efforts is a common objective: to apply AI in service of real business problems. That means making deliberate choices, starting with foundational steps like improving tagging and content structure, refining internal processes to support scalable experimentation, and investing in cross-functional ways of working.

Three consistent patterns emerge across this year's programme:

- Laying strong foundations remains essential. From structured content to effective prompt libraries, the groundwork for AI experimentation often has long-term benefits that extend beyond any single tool or feature.
- Many publishers are using AI to streamline production and reduce time spent on lower-value editorial tasks. This reallocates journalistic effort to areas where human judgment and creativity are most valuable.
- Audience-facing tools are the next frontier. As confidence grows, some publishers are making use of recent advances in generative AI to develop interactive experiences aimed at increasing engagement with existing content.

The most successful projects are those anchored in cross-functional collaboration, built around clear goals, and iterated in response to user feedback. Teams that stay focused on the problem to be solved, rather than the novelty of the technology, are better positioned to build trust, deliver measurable results, and make the case for further investment. Ultimately, when used responsibly, AI can support both the sustainability of quality journalism and the resilience of the organisations that produce it.

FT STRATEGIES AUTHOR CREDITS

This report has been produced by the following FT Strategies contributors, in collaboration with our partners at the Google News Initiative:

Lisa MacLeod, Director EMEA

Brundan Haran, Engagement Manager

Val Imperatore, Engagement Manager

Sam Gould, Manager (AI Lead)

Yuta Nagasaki, Senior Consultant

Shak Ebrahimjee, Consultant

Sofia Louli, Consultant

Azymberdi Taganov, Consultant

Yosef Abebe, Associate Consultant

Emma Copsey, Associate Consultant

Christina Ivey, Associate Consultant

Jack Lee, Associate Consultant

Joseph Miller, Associate Consultant

ABOUT FT STRATEGIES AND THE GOOGLE NEWS INITIATIVE

FT STRATEGIES

FT Strategies is the specialist media consultancy owned by the Financial Times. Staffed by the experts who successfully transformed the FT's business model in the face of disruption, FT Strategies provides the expert insight and consulting excellence required to transform the future of business based on real-world, first-hand experience. FT Strategies has worked with over 800 organisations globally helping them to address strategic challenges, drive sustainable growth and innovate using AI, technology and data.

Google News Initiative

The Google News Initiative represents Google's largest-ever effort to help journalism thrive in the digital age. In collaboration with editors and journalists worldwide, the Google News Initiative (GNI) strives to foster a more sustainable, diverse, and innovative news landscape. Through a comprehensive suite of programs, digital tools, training, and resources, the GNI empowers journalistic endeavors, enabling journalists to effectively discover, verify, and narrate compelling stories while upholding the highest standards of quality journalism.

APPENDIX - ADDITIONAL PUBLISHER CASE STUDIES

ALLER MEDIA (SPIS BEDRE): PERSONALISING NEWSLETTERS TO BOOST ENGAGEMENT



The Situation

Spis Bedre is a Danish food and lifestyle brand published by Aller Media, one of the largest media groups in the Nordic region. While Aller Media operates over 60 brands across Denmark, Sweden, Norway, and Finland, Spis Bedre stands out for its dedicated subscriber base, who receive a printed magazine and a weekly subscriber-only newsletter.

As Aller Media seeks to shift from an ad-led business model toward deeper engagement and digital subscriptions, Spis Bedre offered an opportunity to experiment with AI-driven personalisation, aimed at increasing the value-proportion for subscribers. The existing newsletter had a healthy 10% click-through rate, but the team believed there was room to deliver more relevant, engaging content tailored to individual readers.



AI Use Case

The selected use case focused on personalising the weekly paid meal plan newsletter for subscribed, logged-in users of Spis Bedre. The hypothesis was that if the newsletter included AI-recommended recipes aligned with each user's preferences, click-through rates

and engagement would increase. The experiment involved several stages:

- First-party data selection: One challenge was finding and deciding what first-party data could be used. Spis Bedre collects limited first-party data and therefore does not have the dietary requirements for its readers. However, logged-in users can access the 'saved recipes' function, which was used as a major input for the tool.
- Metadata advantage: however, Spis Bedre did have strong content metadata available, such as ingredients, preparation times, and categories. This was all fed into the model.
- Personalisation testing interface: The team built a dashboard on Streamlit to simulate user behaviours (e.g. "10 saved vegetarian recipes") and observe what the recommendation engine would return. This allowed the team to test what outputs could be generated for hypothetical users rather than having to test initially on real users.
- Editorial integration: Personalised elements were embedded within pre-structured newsletter tem-

plates, with human oversight to override any mismatches.

- Delivery and tracking: Personalised newsletters were sent via the Agillic CRM system to a limited test audience, with open and click rates tracked for feedback.

The solution was developed internally by a small team of two, supported by FT Strategies on technical framing and prompt design.



Results

At the time of reporting, the tool had not yet launched at scale. However, it had been soft-tested internally with promising outcomes, and the team had made substantial progress in preparing for public release. Importantly, the work clarified:

- What user data do they own and could legally use.
- Where user signals were strong enough to support personalisation.
- How to build fallback systems when data was insufficient.



The Streamlit interface proved particularly useful for transparency and debugging, enabling editors to understand and refine how the AI made its recommendations. Although early testing was limited, it helped the team validate the logic and surface edge cases before launch.

Through the process, the Spis Bedre team developed a clearer understanding of how personalisation could work in a low-data environment and laid the foundation for future rollout across other Aller Media titles.



Next steps

The next phase involves a limited live test with a subset of users, followed by iteration based on CTR, session duration, and editorial review. Looking forward, the team plans to:

- Monitor key metrics, including open rate, churn, and recipe engagement
- Refine the model by incorporating more behavioural inputs where available

- Explore enrichment strategies to capture more granular user data (e.g., dietary preferences or cooking skill level) while maintaining compliance

With this experiment, Spis Bedre has positioned itself at the forefront of audience-first innovation within Aller Media. The project illustrates how even modest, data-light environments can begin implementing intelligent personalisation through thoughtful design and editorial partnership.

HAARETZ: BOOSTING USER ENGAGEMENT THROUGH A CULTURE-FOCUSED RECOMMENDATION CHATBOT



The Situation

Haaretz is one of Israel's leading liberal publications, known for its high-quality journalism, in-depth investigations, and influential cultural commentary. As a subscription-first business, Haaretz has cultivated a loyal and discerning readership, both in Hebrew and through its English-language edition.

However, the size of the Israeli media market inherently limits the total addressable audience. As a result, Haaretz's strategic focus is not on traffic growth but on deepening engagement among existing readers to increase subscriber retention and overall value per user.

One of Haaretz's core strengths is its cultural coverage, particularly in film, television, and columns by well-known writers. Recognising an opportunity to leverage this content to deliver more value to existing readers, Haaretz chose to develop an AI-powered chatbot that could offer personalised movie and TV recommendations based on its editorial archive.



AI Use Case

Designed in-house with support from FT Strategies, the tool enables users to receive tailored movie and TV recommendations based on genre and platform, from Haaretz's archive of 2,400 metadata-tagged reviews. The chatbot was built using Gemini, selected for its Hebrew-language capabilities after other LLMs had proved inconsistent in tests.

How the tool works:

- Users access the chatbot via a dedicated mini-site on Haaretz. The chatbot then prompts users to select the genre(s) they are interested in, as well as the platforms they have available - this is so the tool does not recommend movies or content that the reader does not have access to.
- The tool uses function calling to extract the parameters from the user query and match them with the tags in Haaretz's review article database using RAG (Retrieval-Augmented Generation). This ensures that the chatbot will only produce recommendations from Haaretz's own archive.

- The chatbot then generates a response, presenting the recommended show or movie, a link to the Haaretz review and the author's image or name.
- Users can provide free-text feedback to refine the suggestions or use the 'heart' button to indicate a successful recommendation. These interactions are logged and used to refine the future outputs of the model.
- Finally, the team were also concerned that users might ask questions unrelated to movies and TV, prompting unusual or inappropriate responses from the chatbot. To mitigate this, they used prompt engineering to ensure that the bot would give a standard response in these instances, removing the risk of these occurrences.



Results

While the tool is still in early testing, the process has already delivered meaningful outcomes:

- Cross-functional collaboration: The experiment was the first time Haaretz's data science and product teams had worked closely together

on a shared project, setting a precedent for future experimentation.

- Metadata maturity: The process of preparing content for the chatbot significantly improved understanding of the role of data structuring in producing effective AI tools.
- Risk management: Thoughtful prompt design and topic restriction ensured the chatbot aligned with the publication's values and avoided reputational risks.
- Monitor key engagement metrics, including return usage, click-throughs to review pages, and time spent per session.
- Invest in metadata structure - the team realised that only 2,400 of their review articles were correctly tagged. Improving the metadata structure for all articles would not only enhance the chatbot but also be essential for the development of further AI tools.

Haaretz's hypothesis is that this tool will drive stronger reader interaction with its cultural archive, increase session duration, and create a more engaging digital experience for subscribers.

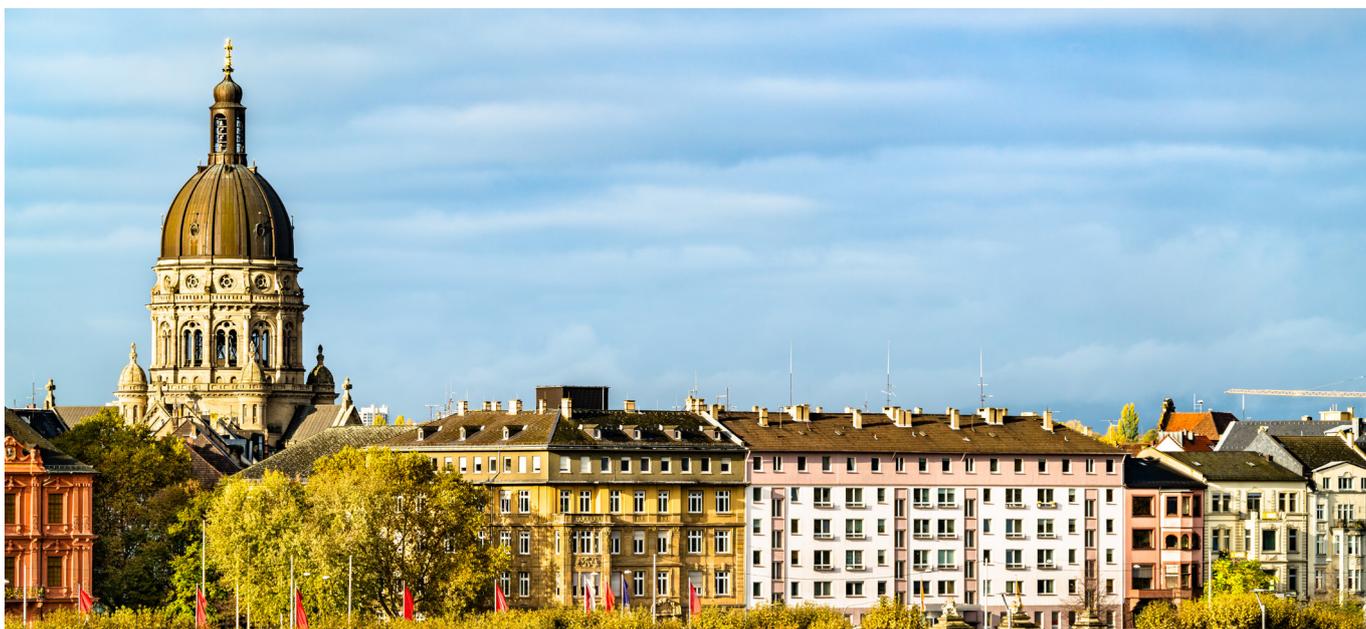


Next steps

At the time of writing, Haaretz launched the chatbot for a small number of users. The next steps will be as follows:

- Launch a marketing campaign promoting the chatbot for subscribers, with feedback collection mechanisms such as a ratings function (replacing the 'heart' function) will inform future refinement.

HEUTE: BUILDING AN IN-HOUSE AI COMMENT MODERATION SYSTEM TO ENSURE ETHICAL USER-GENERATED CONTENT



The Situation

Heute is a widely read Austrian news outlet part of DJ Digitale Medien GmbH, known for its breaking news coverage and broad national reach. The outlet attracts around 4 million unique visitors, with particularly strong performance in direct traffic. They are currently committed to growing the engagement of their reader community, as well as increasing content discoverability, particularly via Search.

The publisher had previously participated in FT Strategies' North Star Founda-

tion programme in 2022 and entered AI Launchpad with nascent but growing AI capabilities. While Heute had tested AI tools in summarisation and podcast production, it had not yet embedded AI systematically across its operations. The team saw the programme as a chance to explore AI use cases that could improve efficiency while supporting their core editorial and commercial goals.



AI Use Case

Initially, the team had planned to focus on SEO optimisation. However, a shift in business priorities presented a more

pressing challenge: the contract for their existing comment moderation and proofreading tool was coming to an end. This left Heute without a solution to manage user-generated content, a critical gap for a publication that regularly receives large volumes of comments.

In response, the team pivoted and used the AI Launchpad to build their own in-house comment moderation tool. Developed with support from the CTO and technical team, the tool was built using Gemini as the underlying LLM. It works by evaluating comments against a set of editorial guidelines listed in

the prompt. The model assigns each comment a score from 0 to 100, where 0 indicates an acceptable comment, and scores above 50 are flagged for review. Comments are then either published, held back, or rejected based on these scores.

The team initially tested the tool on Newsflix, at one of their smaller titles, before planning to expand its use to the main Heute platform. Editors check the flagged comments manually twice a week, forming a feedback loop to help iterate and refine the prompts over time.



Results

The experiment demonstrated how a smaller organisation can be agile in the face of shifting priorities. Despite limited resourcing, the Heute team successfully built a working version of a moderation tool that is already in active use. The choice of Gemini as the LLM proved cost efficient: the system moderated over 500 comments with an estimated cost of just 5 pence — a significant saving compared to the previous outsourcing model or other vendor tools.

The tool proved satisfactory in filtering at-risk comments, and the integrated feedback loop improves its efficiency

over time, underlying the importance of continuous improvement to successfully deploy AI tools and use cases.

Just as importantly, the project showed the value of rapid prototyping and internal alignment. Editorial and technical stakeholders were highly engaged, and the project increased organisational confidence in the use of LLMs to support day-to-day newsroom functions.



Next steps

Heute is now preparing to expand and improve the tool through the following actions:

- Roll out the comment moderation system to the main Heute platform
- Monitor for anomalies as comment volume increases to assess the model's robustness
- Refine the Gemini prompt to improve accuracy and reduce false positives
- Continue weekly manual checks to support iterative prompt tuning and feedback loops

MILANO FINANZA: PREDICTING SUBSCRIBER CHURN TO IMPROVE RETENTION



The Situation

Heute is a widely read Austrian news outlet part of DJ Digitale Medien GmbH, known for its breaking news coverage and broad national reach. The outlet attracts around 4 million unique visitors, with particularly strong performance in direct traffic. They are currently committed to growing the engagement of their reader community, as well as increasing content discoverability, particularly via Search.

The publisher had previously participated in FT Strategies' North Star Foundation programme in 2022 and entered AI Launchpad with nascent but growing AI capabilities. While Heute had tested AI tools in summarisation and podcast production, it had not yet embedded AI systematically across its operations. The team saw the programme as a chance to explore AI use cases that could improve efficiency while supporting their core editorial and commercial goals.



AI Use Case

To address retention, Milano Finanza developed a churn prediction model to help the business identify subscribers at risk of cancelling. The model was designed in collaboration with FT Strat-

egies' Head of Analytics and trained on a "snapshot" dataset aligned to each user's renewal date. Key features included activity levels, subscription plan type, recency of visits, past payment patterns and engagement depth.

Milano Finanza developed the model in

phases. First, they cleaned and brought together the data they needed. Then, they explored the data to understand patterns, defined how to label users depending on their likelihood to churn, and chose the right inputs for the model. Once that was done, they trained the model and began designing a dash-

board to show the results in a clear, understandable way. At the same time, the team planned how the marketing team could use the model's predictions to automatically trigger actions, such as sending tailored offers, making sure the tool would be useful in real business settings.

Once deployed, the model would output weekly risk scores for each user. These scores would feed into a dashboard accessible by different teams (such as the marketing team), enabling them to tailor retention strategies like targeted discounts, value-based messaging or re-engagement campaigns. The use case was prioritised not just for its business impact but for its ability to build long-term AI capabilities inside the organisation.



Results

Model training has just ended, and deployment plans are underway. Although live results are not yet available, the process has already delivered meaningful internal value. Building the model forced early conversations around data quality and ownership, revealing gaps in how subscription data was captured and labelled, including the discovery that duplicate “€0 promo” accounts

were skewing churn labels. Fixing these issues helped resolve broader data hygiene problems and accelerated the creation of a more reliable data infrastructure.

By the end of the experiment, Milano Finanza had not only trained its first churn prediction model but also developed the processes and tools to support future AI experimentation.



Next steps

Milano Finanza is now preparing to scale and refine this work through the following actions:

- Complete deployment of the churn model and begin A/B testing of retention tactics informed by the scores
- Reuse the model's feature store for future use cases such as Next Best Action and dynamic paywall models
- Allocate engineering team time to continue refining the model, including structured sprints for model iterations
- Establish clearer data ownership across the business to maintain the data hygiene achieved throughout the programme
- Expand cross-functional collaboration by appointing AI champions across departments and formalising feedback loops

RND: EXPLORING AI-POWERED PERSONALISATION TO INCREASE CONTENT RECIRCULATION



The Situation & AI Use Case

RedaktionsNetzwerk Deutschland (RND) is a Hanover-based central newsroom operated by the Madsack Media Group, which supports over 60 daily newspaper titles and reaches a combined print circulation of 2.3 million. With 230,000 digital subscribers, RND's focus during the AI Launchpad was on growing reader revenue through stronger audience engagement.

Technically, the organisation had a solid foundation, with early experimentation underway and a roadmap beginning to take shape. However, a formal experimentation culture had yet to take hold. The programme provided an opportunity to shift from localised efforts to more structured, organisation-wide approaches to testing and learning.

As part of this effort, RND partnered with the vendor Froomle to pilot a personalisation tool designed to improve content recirculation. The goal was to increase pages per visit by surfacing contextually relevant article recommendations. The pilot launched on one of their titles, Reise Reporter, with click-through rates and pages per visit tracked as key success metrics.

The setup involved two core members of the technology team, with editorial



staff brought into the process over time to ensure alignment on quality, relevance, and workflow fit.



Early Insights & Next Steps

As the experiment was still underway at the time of reporting, formal results were not yet available. However, the pilot marked an important first step in operationalising AI within RND's newsroom workflows.

The team used the experience to explore new governance approaches, involving both technical and editorial stakeholders. It also prompted internal discussions about how to scale exper-

imentation responsibly, particularly in the context of RND's complex organisational structure and diverse publication portfolio.

An AI lead has since been appointed to support future initiatives, and RND is already considering additional use cases such as AI-generated newsletters, push notifications, and paywall propensity modelling.

Next steps include evaluating the impact of the current pilot, assessing whether to continue the Froomle partnership, and formalising internal testing structures to support a growing pipeline of AI experiments.

ABIDJAN.NET: USING AN AI-POWERED CONVERSATIONAL SEARCH TOOL TO SURFACE ARCHIVAL AND RELEVANT CONTENT



The Situation

Abidjan, operated by Weblogy, is one of Côte d'Ivoire's most established digital media brands. Widely known as a trusted source of news, the platform is also a repository for official communications, legal notices, death announcements, and public records, making it a crucial information hub for Ivorian society.

However, the site's vast archive can be difficult to navigate, and users can struggle to find specific articles or documents. With national elections approaching and heightened interest in public information, the need for a more user-friendly discovery experience became a strategic priority.

With strong technical capabilities and a team accustomed to experimentation, Abidjan.net identified AI-powered conversational search as a way to surface their deep archive in a more accessible, intuitive manner



AI Use Case

The chosen experiment focused on building a chatbot to enhance site search, using political content as a testbed. The chatbot was designed to allow users to ask questions in natural

language and receive relevant, summarised results from Abidjan's content archive.

To deliver this, Abidjan partnered with Bridged Media. The chatbot was trained on a set of political articles and built using retrieval-augmented generation (RAG) to provide contextual responses. The project team included representatives from the editorial department, ensuring the tool aligned with journalistic expectations and practical newsroom needs.

By enabling users to ask open-ended questions like "What did the opposition leader say about the budget?" the chatbot is intended to transform the rigid search experience into a more conversational one, especially valuable in a context where public trust and information accuracy are paramount.



Next steps

Currently, the chatbot is not yet publicly available. A working version has been demoed by Bridged Media and reviewed internally by the project team. Plans are now in place to:

- Test the chatbot internally to gather editorial feedback

- Prepare for implementation on the website ahead of the 2025 Ivorian elections
- Explore how the chatbot could be positioned as part of a future premium offering
- Extend the chatbot's coverage beyond politics to include legal notices, governance documents, and other high-value archive content.

Through this experiment, Abidjan.net is taking a first step toward aligning its technical capabilities with the needs of its audience, bridging the gap between trusted information and user-friendly discovery.

