



How Amedia Got Half Its Journalists Using AI – By Keeping It Simple

Mathias Høibakk Bergquist
Head of Editorial AI, Amedia

17.12.2024, Mjøndalen



This is Amedia

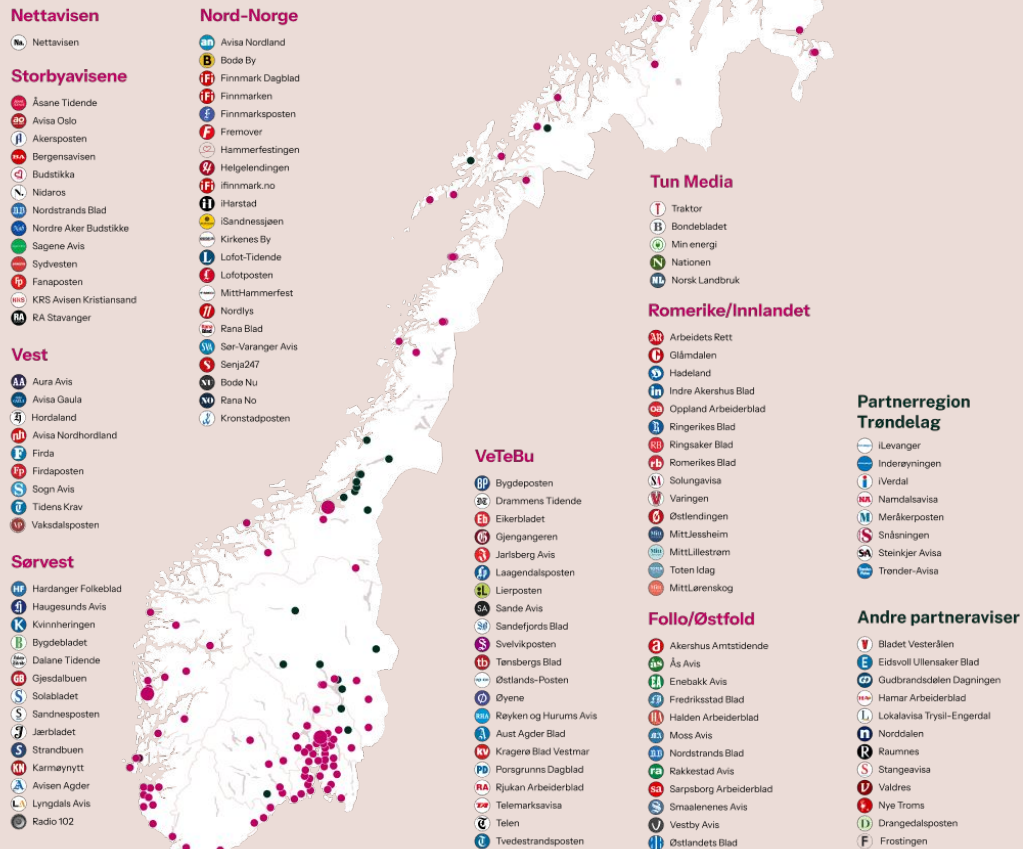
Norway's Largest Publisher: Amedia is the largest publisher of editorial media in Norway, owning 107 newspaper titles and reaching over two million readers daily.

Digital Transformation Leader: Recognized as one of Europe's most innovative media groups, Amedia has successfully increased demand for digital journalism, with nearly 90% of subscribers reading news digitally.






Digital Subscription Growth: By the end of 2023, Amedia's owned newspapers had 684,000 subscribers, with over 68% being fully digital—a testament to the company's successful digital transformation efforts.

Editorial Employee Base: Amedia employs over 1,300 editorial staff who are dedicated to delivering news, debates, and entertainment to readers across Norway.

Centralized Expertise: Amedia's centralized teams in Oslo handle IT, product development, technology, and editorial support, developing and maintaining tools, solutions, and services for readers, subscribers, advertisers, and newsrooms across its newspapers.



For years, Amedia explored the many possibilities within AI along five main tracks

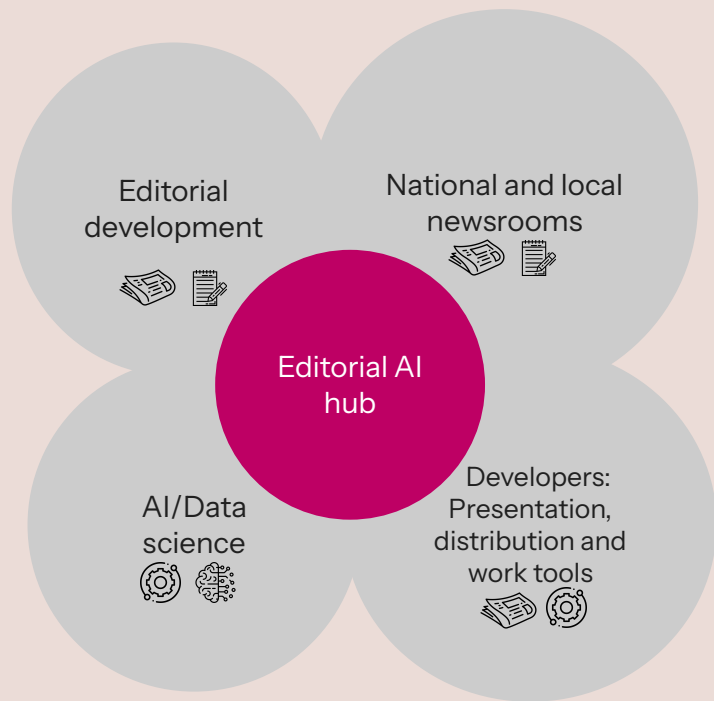
 Language Technology	Extraction of content information - categorisation, keywords, people, geographies, organisations, etc., translation, and recently, advanced language generation capabilities.	Cat20, sports, reader needs, "stories that sell" Entities, sources, gender balance Summaries Title suggestions
 Personalisation	Utilises data to create personalised user experiences/recommendations. Applied to front pages, newsletters, and article pages.	Top lists User preferences Reader recommendations (more of what you like, collaborative filtering) Contextual recommendations
 Insights	Insights into readers' usage patterns and preferences, including forecasts of future developments and demand.	Loyalty and engagement (churn/customer value) Reader profiles Measuring impact and value Fundamental drivers
 Format	Implementation of speech-to-text based on OpenAI's Whisper model.	Speech-to-text Images
 AI Platform	Adopt, adapt, or develop tools to deliver robust solutions for data preparation, accessibility, processing, model training, generation, and result dissemination.	GCP Vertex AI Codebase for creating testable SQL components Adopting DBT



And then this product was launched...



In February '24 we established an editorial AI hub with representatives from eight newspapers



Process



Mathias Høibakk Bergquist
Hub leader



Maren Hornes Martinsen
Business Developer



Oda Oldertrøen
Editorial Developer

Editorial



Pål Nisja-Wilhelmsen
Nettavisen.



Natalie Kristiansen
Moss Avis



Tonje E. Skjørtvedt
Østlands-Posten



Marie Olaussen
Teisberg's Blad



Rune Endresen
Nordlys



Martin Fosslund
NA



Halvor Farsund Storvik
Sogn Avis



Martin Berge
Dadeland

Engaged as needed

Expertise/ Development



Filip Steinhovden
AI magician



Tor Skjøtskift
Frontend



Jørgen Veum-Wahlberg
Backend



Helge Valvik
Presentation and distribution



Bjørn Tore Øren
Work tools



Sigve Harang
ML Engineer



Emiliano Guevara
NLP / Senior Data Scientist



We swiftly grasped that generative AI could transform every aspect of our editorial workflow



Brainstorming and concrete ideas

Sparring partner

Suggestion for story angles



Research

Analyze documents

Video analysis

Audio analysis

Query the archive

Generate timeline from documents



Interview and photo/video/audio

Transcription tools



Production and quality control

Shorten articles

Language enhancement

Headline suggestions

Article generation



Publishing and distribution

Summaries

Automatic fact boxes

AI voices

Automated social media content



Analysis and learning

Categorization

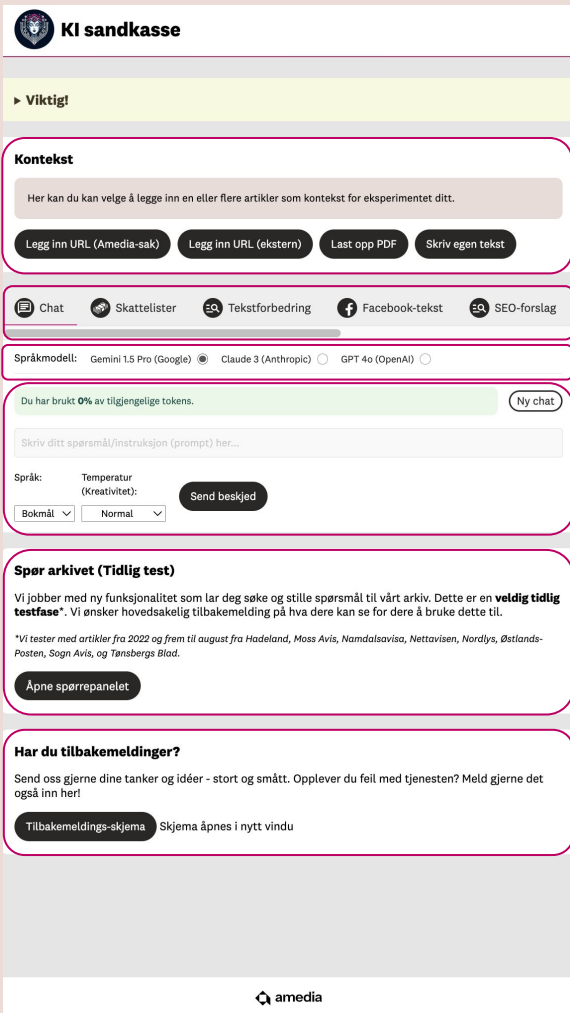
"Which stories work"



We needed a secure environment to experiment and test new functionality — so we built our own AI Sandbox.

Amedia's AI Sandbox was launched for all of our more than 1,000 journalists at the end of February 2024.





Context: The Sandbox can incorporate stories—whether published or drafts—from Amedia’s news sites. It can also scrape external URLs, process PDFs, or allow users to input their own text.

Main functions: We have implemented nine core functions in the Sandbox: Plain chat, talk with the tax records, improve your text, Facebook-post-helper, SEO optimization, summary, title suggestions, image generation and text-to-speech

Language models: Through APIs, we have integrated the Sandbox with the three leading providers of large language models (LLMs). Whenever a new version is released, we promptly update the Sandbox to ensure it stays current.

Chat field: This is where the magic happens. By pressing a button, users receive responses based on predefined prompts. Users can also customize the experience by selecting the desired language (Norway has two official written languages) and adjusting the LLM’s temperature/creativity level.

Ask the archive: We are testing a Retrieval-Augmented Generation (RAG)-based query system, which draws on stories from the past two and a half years across eight of our newsrooms.

Feedback: While we can develop anything we want from Oslo, we wouldn’t have come this far without feedback from reporters and editors across Norway. Whenever they have input on editorial AI, they can use the feedback button, which immediately notifies the entire team on Slack. Some of our best ideas have emerged through this button.



Why build our own sandbox, instead of using off-the-shelf products?



Ensure data security by owning service-generated data and using secure login



Access to multiple providers and increased flexibility



Cost control and efficiency



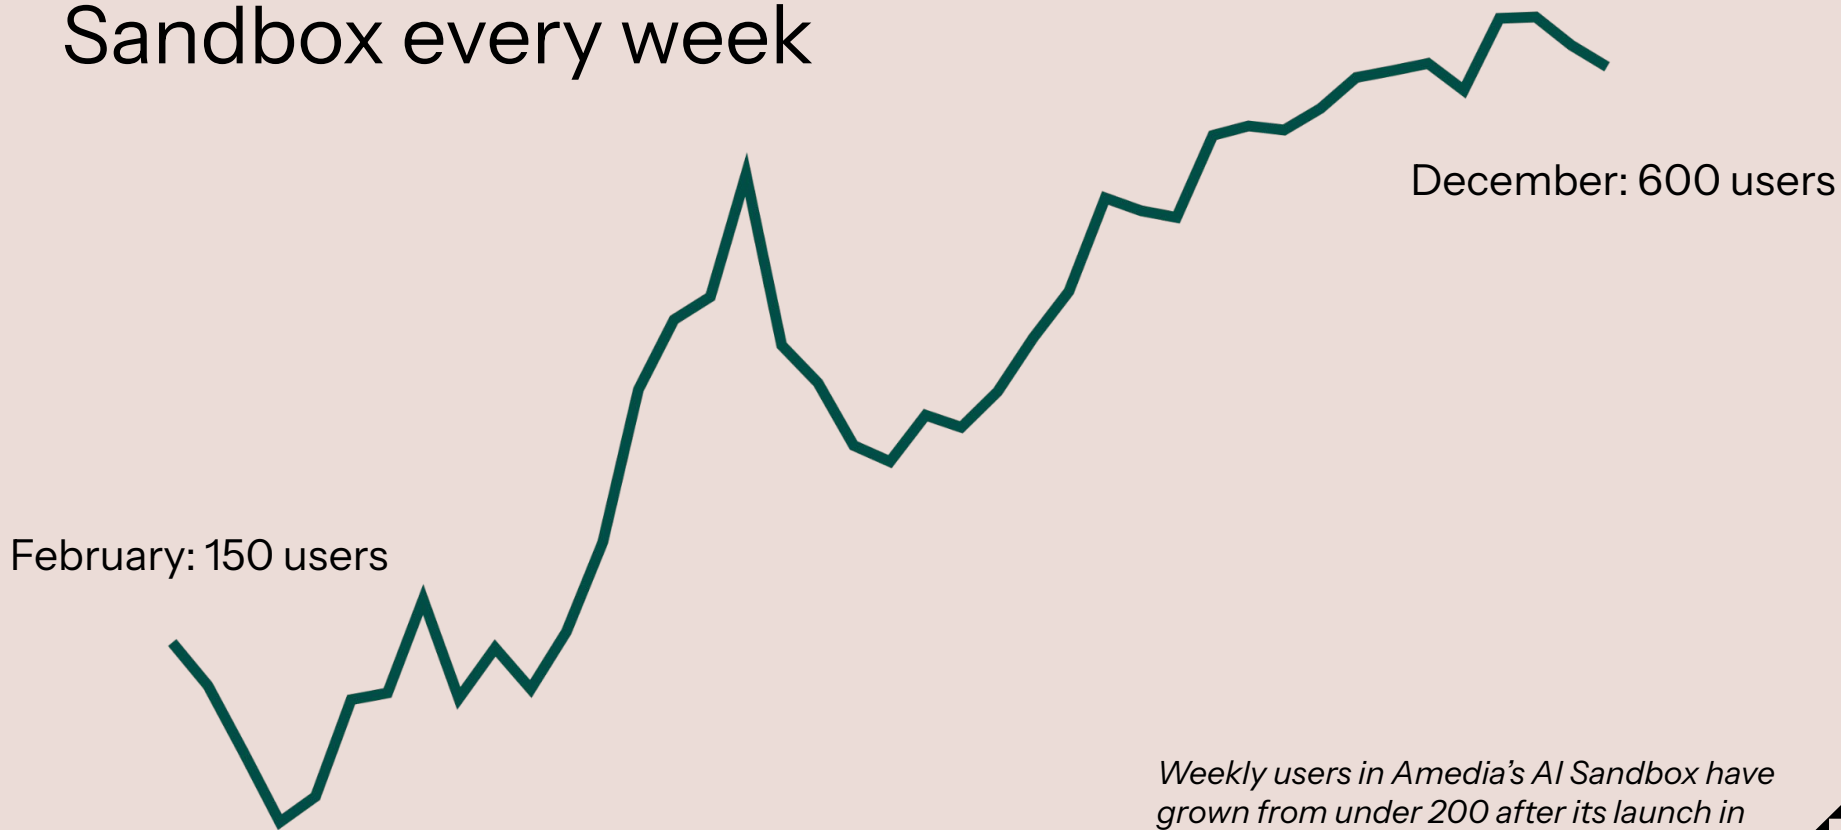
In-house knowledge of business needs and potential use cases, allowing tailored solutions



Enhances cross-departmental collaboration and ensures skill development



Half of Amedia's journalists now use the AI Sandbox every week



Weekly users in Amedia's AI Sandbox have grown from under 200 after its launch in February to 600 by the end of 2024.



This is not because the tools or technology are extremely powerful (though most of them are quite helpful).



This is because we approached AI adoption through a distinctly Norwegian lens - practical, collaborative, and focused on solving real problems rather than chasing technological possibilities.



Let's look at a few examples of what
we did in practice



Reach out. To everyone!

This is a snapshot from our spreadsheet tracking AI talks and courses conducted in October and November alone. Through these 13 sessions, ranging from 3 to 250 participants, we engaged with over 500 editorial staff at Amedia. They gained valuable insights—and we learned even more.

Where	Who	Number of participants
Vestby	Vestby Avis, Akershus Amtstidende, Østlandets Blad, Enebakk Avis og Ås Avis	15
Oslo	Nettavisen	20
Oslo	Avisa Oslo	20
Oscarsborg	Tun Media	35
Oslo	D.C Thomson from Scotland	10
Kongsberg	Chief Editors from the region of Buskerud, Vestfold and Telemark	17
Digital	Editors and leaders from the region of South West	8
Digital	Introduction to Amedia's AI Sandbox	70
Trondheim	Avisa Gaula	3
Drammen	Innholdsforum (Editorial Conference)	250
Oslo	Amedia's partner papers	9
Kongsvinger	Glåmdalen	15
Digital	Amedia's AI Sandbox for advanced users	60
Lillestrøm	Romerikes Blad	20



It's all about show, don't tell

We met journalists wherever they were – literally. From one-on-one sessions helping reporters with specific needs, to newsroom workshops with 3 to 50 participants, to advanced sessions for power users. We trained all summer interns, ensuring new talent started with AI tools in their toolkit.

The momentum built naturally: individual success stories led to newsroom curiosity, which led to broader adoption. By the time we presented at Innholdsforum, our annual editorial conference with 250 participants, we weren't selling a vision – we were sharing hundreds of real success stories from newsrooms across Norway.

No fancy training program – just journalists teaching journalists, at whatever scale made sense.




Oda Oldertrøen from Amedia's Editorial Development Department showcased the possibilities of using AI in the editorial workflow at a conference attended by 250 reporters and editors from Amedia.



The Parade Schedule

A veteran journalist at Østlands-Posten needed to write a preview for the local Constitution Day celebrations. Usually, he would spend hours manually transcribing PDF documents containing parade schedules and event details. With the AI Sandbox, he simply uploaded the documents and had the text ready in seconds.

A small win? Yes. But it showed journalists exactly how AI could help them spend less time on tedious tasks and more time on actual journalism.



Bjørn-Tore Sandbrekkene, Østlands-Posten




Breaking Language Barriers

At Tønsbergs Blad, an experienced Iraqi journalist was doing language practice in the newsroom. She knew journalism - she'd been doing it for years. But writing in Norwegian was still a challenge.

Using the AI Sandbox's text improvement tools, she could focus on the journalism while getting help with language structure. This meant her conversations with colleagues could center around storytelling and angles rather than grammar and syntax.

This wasn't just about improving text - it was about empowering a skilled journalist to fully contribute despite language barriers.



Aryan Arif, Tønsbergs Blad



When a journalist ask - we listen

When Natalie Kristiansen, Front Page Editor at Moss Avis, asked in October if we could develop something cool with AI for Tax List Day* in December, we could have easily said 'too late - priorities are set.'

After all, building a natural language interface for a complex database with 55 million data points isn't exactly a small task.

But Tax List Day is one of the biggest events of the year for our newspapers. And more importantly: when journalists come to us with ideas, we listen. Even if it means throwing out our carefully laid plans.

So we dropped everything and built something completely experimental. Because sometimes the best innovations come from a simple question from a local newsroom.



Natalie Kristensen from Moss Avis

**Every December, Norwegian tax authorities publicly release income and wealth data for all taxpayers. This unique tradition of fiscal transparency makes Tax List Day one of the biggest annual events for Norwegian media, driving massive reader engagement and subscriptions.*



And when we listen - magic things happen

We built a tool that lets journalists explore 55 million tax records through natural language queries. The agent works in three steps:

1. AI translates Norwegian questions into SQL queries
2. These queries fetch data from our tax database
3. Results are returned both as raw data and as AI-generated story summaries

While it can handle complex analyses (like comparing wealth trends across municipalities over time), we designed it for accessibility over sophistication. When Nettavisen wanted to track the growth in millionaire incomes over the past decade, they got their answer in seconds rather than hours.

It's experimental and sometimes rigid in its interpretations - but that's exactly the kind of learning we need as we develop more tools combining structured data with AI for journalism.

The screenshot shows a web interface with a search bar at the top containing the query: "Fredrikstadkvinner med inntekt mellom 368.000 og 372.000 kroner i 2023". Below the search bar, there is a list of results for women in Fredrikstad with income between 368,000 and 372,000 kroner in 2023. The results are listed in a table format with columns for name, birth year, municipality, and income. The first few results are: LINE CECILIE HARVEY (f. 1977), BERIT KRISTI STRAND JOHANSEN (f. 1965), INGER JOHANNE LIE-HOLT (f. 1957), and ASTRID MARIE TØNNESEN (f. 1939). The interface also includes a chat button, a text-to-speech button, and a search button.

When Fredriksstad Blad was planning to write a story about women earning around the median wage, they were able to query the tax records through Amedia's AI Sandbox to find specific cases for their article.

This was the response they received when they entered the query: "List all persons with gender F in Fredrikstad with income between 368,000 and 372,000 kroner in 2023. Use all available information."

Through this data, they could identify and reach out to potential interview subjects who represented their target demographic.



Solving simple problems is far more impactful than showcasing advanced tools that journalists don't fully understand.



AI won't replace the work we do,
but it can help us get **more** out of
what we do



Don't just take our word for it - here's what a couple of the users say

We use the sandbox for everything from proofreading to Google titles to research assistance. As a journalistic tool, it's like having a team of editorial assistants on standby at all times. Simply revolutionary.

The AI Sandbox has removed fear and increased understanding of AI in the newsroom. We use it to improve language, sort and interpret large datasets, plan interviews, transcribe recordings, and extract text from PDFs. When we used it to analyze tax records by postal code, we discovered numerous new sources in our coverage area that we didn't know about - stories we'll be telling for the rest of the year! We love the sandbox and use it every day.




Pål Nisja-Wilhelmsen,
Development Editor,
Nettavisen



Marie Olaussen,
Development Editor,
Tønsbergs Blad





amedia

mathias.bergquist@amedia.no

Head of Editorial AI

Amedia Produkt og Teknologi