

SpinozAI

AUGMENTING JOURNALISM
WITH INNOVATION AND ETHICS

TABLE OF CONTENTS

FOREWORD:	3
Yes, AI can help journalism and produce reliable information!	
BACKGROUND: timeline and key players	4
1 BUILDING AN AI TOOL FOR CLIMATE CHANGE JOURNALISM: CHALLENGES AND CONSTRAINTS	6
A/ An opportunity to be seized – under certain conditions	6
B/ The four pillars of designing a trusted AI tool for journalists	8
C/ Developing the first prototype: creating a competent research assistant	9
2 SPINOZA: A FIRST CHALLENGE MET – ON TO FURTHER DEVELOPMENTS	12
A/ Enhancing databases: giving added value to a functional tool	12
B/ The collective rewriting of Spinoza's prompts	13
C/ From technical experimentation to collective intelligence: the shift to open-source	13
3 RSF RECOMMENDATIONS FOR DEVELOPING AN AI TOOL THAT RESPECTS JOURNALISM AND THE RIGHT TO RELIABLE INFORMATION	15

Director general: **Thibaut Bruttin** • Editorial director: **Anne Bocandé** • Head of the journalism and technology desk: **Vincent Berthier** • Editors: **Anne-Laure Chanteloup** and **Charlotte Hervot** • English translator: **Anne Duthoit** • Graphic Designer: **Sandra Héno** • Research: **Alice Lobu** and **Julia Pilod** for Econovia

The Spinoza project is developed in collaboration with l'Alliance pour la presse d'information générale (l'Alliance) and with the support of French Ministry of Culture and the French Copyright Licensing Center (CFC).



Support of



AVANT-PROPOS

YES, AI CAN HELP JOURNALISM!

Thibaut Bruttin

RSF Director General

For more than two centuries, technological advancements have played a role in the development of journalism. The rise of social media and search engines has created a digital space largely devoid of democratic safeguards, which has weakened news outlets – eroding both their central role in public debate and their economic model – while also offering unprecedented opportunities for content production and distribution.

The shock of ChatGPT's emergence just a few years ago and the recent launch of DeepSeek should spur the journalism community to confront this new technological shift head on, as it is perhaps one of the most significant in journalism's history. For the first time, large language models are not merely transforming the distribution of information but directly affecting content production, the very core of journalism.

Reporters Without Borders (RSF) firmly believes that defending at-risk journalists worldwide is inseparable from the promotion of honest, independent, and pluralistic journalism. In the face of artificial intelligence (AI), two prevailing approaches have emerged. On one hand, there is individual and collective negotiation with tech companies and an emphasis on the importance of experimentation. On the other, there is a reticence – sometimes worn as a badge of pride – to testing the potential of these tools.

Recent history has demonstrated the consequences of this divide within the sector and the edge it provides to tech companies: news media are often treated as mere data sources and their full value goes unrecognised; media outlets are excluded from the governance of technological solutions essential to their work; and reliable media are marginalised in the digital space, drowned out by content that does not adhere to journalistic ethics or professional standards.

RSF is committed to forging a responsible middle path, where ethics go hand in hand with innovation. This vision prompted the Paris Charter on Artificial Intelligence and Journalism, which was developed by RSF and 16 partner organisations and presented at the Paris Peace Forum in November 2023. Centered around ten fundamental principles, this international reference framework sets the conditions for using AI tools in a way that preserves the integrity of information – a world first.

This conviction is also what drives the Spinoza project, an AI tool developed by RSF and l'Alliance de la Presse d'Information Générale (l'Alliance) – the first tool of its kind, designed in collaboration with French journalists and news publishers. We thank l'Alliance, its leadership, and its members who expressed interest in this initiative and a willingness to engage with generative AI, for their commitment to the project.

Spinoza is not designed to replace journalists. On the contrary, it reinforces the central role of newsrooms in selecting, prioritising, and producing high-quality content. It aims to enhance journalism by using reliable databases that include scientific reports, legislative texts, and over 28,000 articles published by the French press since 2022. This first experiment, which helps journalists cover climate change, demonstrates that creating an ethical AI aligned with journalistic values is not only possible, but essential.

As 2025 kicks off and industry leaders gather in Paris for the Artificial Intelligence Action Summit, the findings of this experiment serve as a call to action: we must not allow the development of generative AI to be controlled solely by tech companies, whose economic priorities and ideological biases go against the needs of the news media. Spinoza proves that media innovation is possible – but only if journalists actively participate in the process and technological advances are shared across newsrooms.

Journalists and media outlets have the power to reinvent journalism – provided they reclaim technological sovereignty of the sector.

I THE SPINOZA PROJECT I

The Spinoza project aims to develop an open-source AI tool that ensures the integrity, traceability, and ownership of its data. It is co-designed with journalists and publishers and its first prototype focuses on climate change issues. The initial concept was launched by Reporters Without Borders (RSF), and its execution was made possible by its partner, the French press alliance l'Alliance pour la presse d'information générale (l'Alliance), which represents nearly 300 political and general news outlets in France – 120 of which

participated in this experiment, sharing their reflections on how media outlets can engage with AI.

Spinoza was developed through frequent professional workshops, bringing together diverse profiles from both the news media and data science sectors. Publishers, digital media managers, editorial directors, editors-in-chief, and journalists from l'Alliance-affiliated media collaborated with project managers, UX designers, and data scientists from the project's technical partner, Ekimetrics.



THE CO-DEVELOPMENT PROCESS HAPPENED IN SEVERAL PHASES:

OCTOBER
2023

Project launch with **120 volunteer media** outlets from **12 publishing companies** (*Actu.fr, EBRA, L'Équipe, La Nouvelle République, La Provence, Le Télégramme, Libération, l'Union, Nice-Matin, PMSO, Sogemedia, and Sud Ouest*).

JANUARY
2024

Delivery of the first prototype, integrating data from leading scientific and institutional sources on **climate change**, including the Intergovernmental Panel on Climate Change (IPCC), the Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services (IPBES), the French Environment and Energy Management Agency (ADEME), and legislative texts.

MARCH
2024

The first prototype is enhanced with **12,000 articles** from national and regional French press.

MAY
2024

Testing is expanded to partner newsrooms.

JUNE
2024

Introduction of a new collaborative prompt, enabling a more refined exploration of articles (broad questions vs. precise facts).

AUGUST
2024

The prototype is in a closed group on the platform **Hugging Face**.

AUGUST
SEPTEMBER
2024

Integration of regional data and content from *Agence France-Presse (AFP)*, which joins the project.

OCTOBER
2024

First anniversary of the project. The prototype now provides access to **six types of data** available to all partner newsrooms: scientific data, legislative texts, public institution reports, resources from ADEME, journalistic content and *AFP* data.

NOVEMBER
2024


A hackathon is organised to experiment with **new features** and deploy Spinoza across all participating outlets.

I METHODOLOGICAL NOTE I

THE RSF STUDY ON JOURNALISTS' PERCEPTION OF GENERATIVE AI

A panel of **281 journalists** working in France responded to a **questionnaire designed in collaboration with the agency Econovia to assess journalists' relationship with AI**. The survey was sent to media professionals, including journalists, editors-in-chief, presenters and commentators. The questionnaire, distributed via a secure form, allowed recipients to respond easily and anonymously.

A PANEL OF 281 JOURNALISTS



1 | BUILDING AN AI TOOL FOR CLIMATE CHANGE JOURNALISM: CHALLENGES AND CONSTRAINTS

AI opens up new possibilities for journalists, provided that ethical and technical requirements are met. The first step is to explore the benefits and challenges integrating AI into the journalistic process, as well as the fundamental principles for designing a reliable tool.

A/ AN OPPORTUNITY TO BE SEIZED — UNDER CERTAIN CONDITIONS

Developing an AI tool for and with journalists **requires an understanding of their professional practices and relationship with the technology.** However, perceptions of AI within the profession fluctuate between fear of losing editorial control and hope for increased productivity. While certain specific tasks are already entrusted to AI tools and could be further automated, others generate strong distrust due to the risk of misinformation and the potential impact AI could have on journalists' working conditions, job security, and the sustainability of media organisations.

The Spinoza project — which is backed by outlets and journalists convinced that generative AI represents “the next revolution in [their] profession” — reflects these contradictions well. As one editor-in-chief, who has been involved with Spinoza since the project's inception, explained:

“*For me, AI is not just a source of concern; it is also an opportunity. But we still lack a clear framework to use it effectively.*”

The RSF study on media professionals confirms the prevalence of these sentiments, as participants expressed both high expectations and significant concerns about generative AI. The study also found that journalists recognise they have a role to play in AI's development, as **86%** believe they are best positioned to ensure the structured and responsible use of these tools in journalism.

French media companies and newsrooms are already taking action. Some have created internal working groups dedicated to exploring the use of generative AI in journalism. Far from being mere innovation incubators, these spaces serve primarily to examine both the potential and limitations of these technologies. While some media organisations promote cross-disciplinary discussions among different professional roles, others adopt a more hierarchical approach. However, two sentiments dominated every case: curiosity and caution.

45% of the French journalists surveyed by RSF already use generative AI in their professional work and 93% plan to use it in the future.



AI IN NEWSROOMS: ESTABLISHED USES AND DEEP-SEATED FEARS

Journalists see AI tools as **a way to save time on specific tasks**, such as:

- translating documents;
- generating interview transcriptions;
- summarising information;
- rewording texts;
- generating article summaries.

However, journalists are **extremely reluctant to use generative AI** for certain specific purposes:

- writing newsletters;
- writing articles;
- synthesising voices;
- fact-checking.

“ AI tools are now essential, just like Word, Excel, and search engines, but the sources need to be verified every time. They provide a base to work from. I always use multiple generative AIs to compare their outputs or to get even more ideas!”

Distrust towards AI spans multiple aspects:

> Lack of reliable information:

69% of the journalists surveyed do not trust chatbots like ChatGPT to provide reliable information.

> Amplification of filter bubbles:

77% fear that algorithmic systems will reinforce filter bubbles, isolating audiences in their own personalised information space and limiting their exposure to new perspectives.

> Replacement of journalists and impact on their status and reputation:

83% of the media professionals surveyed worry that AI will lead to job losses or reduce reliance on freelancers. Others fear that AI will generate competing content at lower costs.

“ We’re already seeing drops in traffic on our news websites: people search for information on ChatGPT or Gemini, and that’s enough for them. This question has come up before with social media. What lessons have we learned from it? Should the profession align itself with companies like OpenAI and Microsoft?”

“You don’t need to be a fortune teller to foresee that AI will, at first, replace jobs [...] Those who master the tool will likely fare better. But for how long?”

“AI is not journalism, nor is it even intelligence. It collects and assembles whatever it finds on the internet. It stifles thought and research. It dehumanises society and pushes us toward an Orwellian world.”

“I see three issues at stake with AI: the influence of AI on elections when it is used en masse, the use of AI as a weapon for destabilisation – and, a contrario, AI as a means of detecting such attacks – and, finally, the need to develop algorithms that limit ‘noise,’ since what spreads fastest on the internet is false but attention-grabbing information.”

B/ THE FOUR PILLARS OF DESIGNING A TRUSTED AI TOOL FOR JOURNALISTS

The initial workshops of the Spinoza project, which brought together journalists, publishers, and digital managers, helped identify four key requirements that shaped the technological choices for the prototype and its future developments, ensuring **speed, factual accuracy, and the reliability of sources**.

I ASSIST JOURNALISTS RATHER THAN REPLACE THEM

“It seems important to me to be clear from the outset, so there is no confusion: this is not about developing a robot that produces news articles,” stated an editor-in-chief from a regional newspaper during Spinoza’s inaugural workshop. For him and other experienced journalists, this is a line that must not be crossed, as they consider writing to be the core of their professional identity.

The fear that AI tools could be used to “replace journalists” was repeatedly expressed during the workshops. Instead, the tool should primarily serve to sift through large volumes of reliable data. Therefore, the foremost function expected of Spinoza concerned **document research, improved by the rewording capabilities of generative AI**.

I PRODUCE CONTENT THAT SERVES THE PUBLIC INTEREST


The **climate change theme** was one of the key factors that motivated partner media outlets to join the Spinoza project. **The subject has become a priority** for many newsrooms and journalists. **Spinoza facilitates access** to previously scattered and hard-to-analyze data, opening up new opportunities for deeper and more diverse media coverage of these issues. The journalists emphasised that they do not want to be mere reporters of scientifically accurate bad news. Instead, **they seek to evolve towards solutions journalism**, exploring concrete, grassroots initiatives led by local actors – insights they could access more quickly through Spinoza.

I DEVELOP THE TOOL BASED ON REAL-WORLD USE CASES

“Suppose a company sends me a press release announcing that it has invested in infrastructure to reduce its carbon footprint. What I lack today,” explained a journalist from a regional daily newspaper, “is a tool that allows me to quickly verify whether this is a genuine initiative or simply a move to comply with existing regulations while trying to pass it off as a virtuous action.” This example illustrates a clear need that AI could address: **analysing institutional, regulatory, and legislative texts on the subject** more quickly without losing any accuracy.

I CONTAIN THE GENERATIVE ALGORITHM

The AI tool must mitigate language models’ tendency to “hallucinate,” or generate false information. To achieve this, it is essential to build **databases composed of reliable sources**, which are indispensable for creating well-documented and detailed reports. These sources must be **properly evaluated** so they are not all accepted as having the same level of credibility. The prototype must also allow for the **traceability of sources**, facilitating fact-checking and verification.



C/ DEVELOPING THE FIRST PROTOTYPE: CREATING A COMPETENT RESEARCH ASSISTANT

The construction of the first prototype began in November 2023. Initially, it contained four databases and used the GPT-3.5 language model combined with an embedding algorithm – a search engine that scans texts for terms related to the user's query. The first version of the prototype did not include databases of news articles, which were added in March 2024.

TECHNICAL CHOICES

Why GPT-3.5?

When the prototyping phase began, GPT-3.5 offered several advantages over its open-source competitors and GPT-4, its more advanced version. Execution speed, response quality, and cost efficiency for large volumes of queries were key factors. A comparative analysis conducted by Ekimetrics identified GPT-3.5 as the ideal algorithm to rapidly develop a prototype that could be tested by journalists.

The embedding algorithm

The algorithm needed to be fast enough to ensure a smooth user experience while also functioning effectively in both French and English. Ekimetrics conducted a comparative analysis of various open-source solutions and recommended using the **intfloat/multilingual-e5-base** algorithm.

Interface and databases

The Ekimetrics team recommended building the prototype with technologies that are efficient on both the front-end and back-end, ensuring the tool could operate on relatively lightweight infrastructures. The Gradio web interface was selected for its easy use, and the Qdrant library was chosen to create the vector database.

THE SYNERGY BETWEEN SCIENTIFIC CONTENT AND PRESS DATA

Building a corpus of press articles requires time to determine the right selection methodology, appropriate usage, and a secure way to process the data. The first press data corpus was developed based on the following criteria:

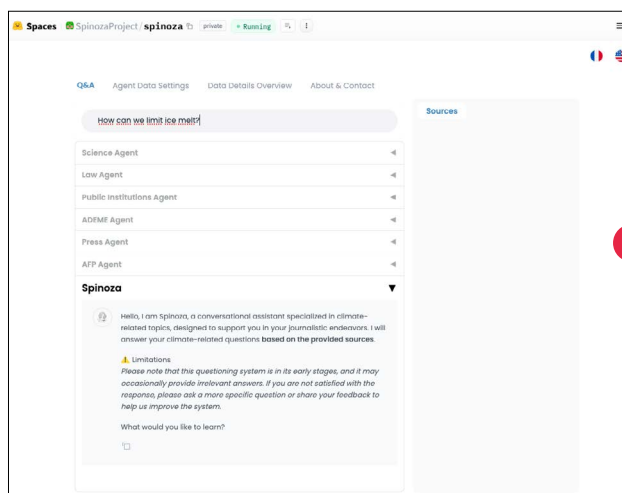
- Target volume: **12,000 articles** – approximately **1,000 articles** per media company.
- Content selection: information was chosen based on **keyword associations** related to climate change.
- Recency: the content must be **less than five years old**.
- Usage limitations: The database is not intended to retrain or refine the algorithm, but rather to provide **extracts for summarisation**. This guarantees publishers that their content remains intact and will not be processed on OpenAI's servers, the algorithm will not be optimised to generate articles, and that media companies can easily opt out of the experiment if they choose to.
- Access restrictions: Full press articles must not be available in their entirety to users and must not be hosted on any server other than that of the project's technical partner, Ekimetrics.

Summary chart of the technical choices based on the needs and concerns of surveyed journalists

Ethical requirements for an AI tool based on press data	Technical proposals developed for the Spinoza tool
The reliability of information that serves the public interest	Development of a specialised, reliable database, and use of the RAG ¹ technique, which enables generative AI models to enrich their responses by recovering information from a collection of documents.
Control over the AI's performance	A user interface that provides visibility on the processing and treatment of information, transparent processes, search prompts that can be consulted and that limit the algorithm's capacity to hallucinate, and the ability to verify the reformulation of questions.
Protecting editorial independence and autonomy	Access to original sources, specialised consultable syntheses, a tool that identifies the relevant documents but does not suggest angles for how to treat subjects.
Compliance with the intellectual property rights of media outlets	A controlled selection of press articles, the algorithm is not trained to copy press data (no fine-tuning press data), and data is hosted on secured Ekimetrics servers.
Avoiding the danger of replacing journalists	An AI assistant that synthesises documents and does not produce content destined for publication.
Transparency and traceability	The responses generated by Spinoza systematically refer back to their source documents and allow the journalist to verify allegations.

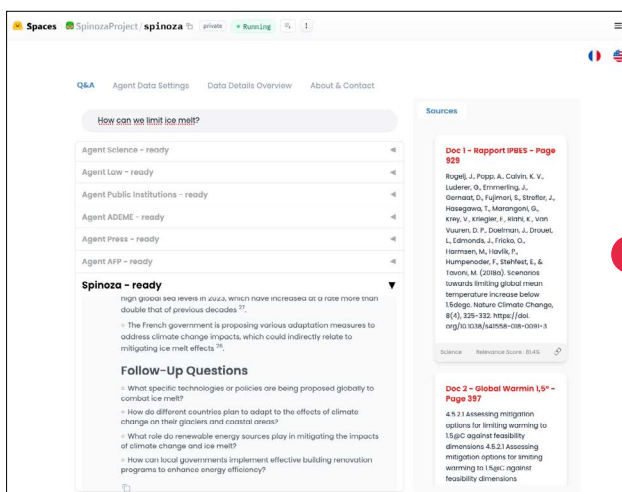
HOW THE SPINOZA TOOL WORKS

The Spinoza assistant is built on a search and synthesis system designed to address journalists' concerns regarding environmental topics. The tool is used by following these steps:



1

The user enters a **question** or **keywords** into the search bar.



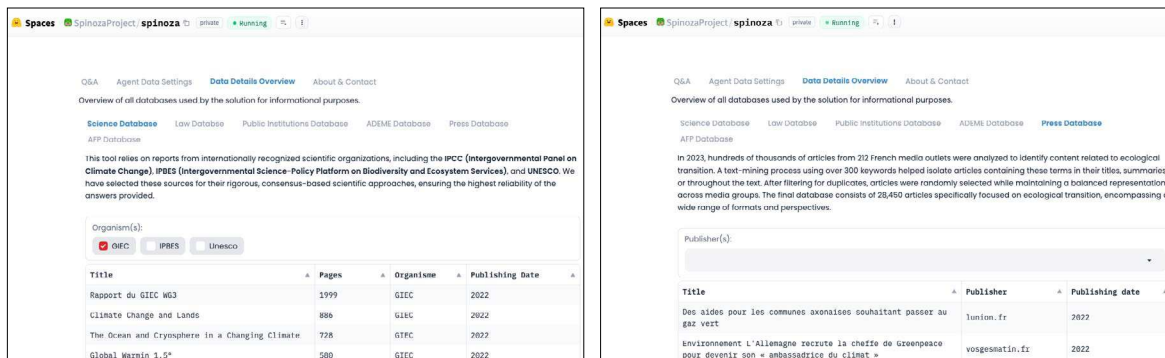
2

The tool reformulates the question according to the different databases to be explored. For example, if the user asks, "How can we limit ice melt?", the tool may rephrase it as, "What regulatory measures have been implemented to limit ice melt?" to obtain relevant results from a database of legal texts. Alternatively, it may be reworded as, "How can ice melt be limited? Which articles discuss this?" for a more precise search in a database of press articles.

¹ Retrieval augmented generation

3 An embedding algorithm explores four distinct databases, each specialised in a specific domain (e.g. scientific reports or press articles). This engine selects relevant excerpts from each database, forming four separate corpora. A relevance score shows the similarity between the terms used in the question and the words found in the database.

4 Each corpus is processed by a specialised agent, using a language model fine-tuned for its domain through specific instructions (prompts). For example, the “science agent” synthesises data from IPCC reports, while the “press agent” extracts and processes journalistic articles. Each of these agents produces a summary based on its respective corpus.



5 The prompts are designed to prevent the agents from generating content when relevant data is missing. If no information in the database answers the user's query, the system will display a message indicating that no relevant information is available. The instructions also prioritise numerical data and include clickable links, directing users to the original source document.

6 The main agent, Spinoza, then consolidates all the summaries generated by the specialised mini-agents. It aggregates the information, highlighting differences and discrepancies between sources.

THE TRANSPARENCY OF SOURCES: A FOUNDING PRINCIPLE OF THE SPOINOZA PROJECT

The user has access not only to the final summary but also to the intermediate summaries generated by each mini-assistant, as well as the sources used (via clickable links), and some original documents. This transparency allows journalists to verify and deepen their understanding of the information at every stage of the process. The tool is not a chatbot. It does not retain previous interactions – it functions as a research and synthesis tool rather than a conversational text generator. This design choice addresses journalists' and publishers' concerns regarding reliability, transparency, and editorial control, while still delivering a powerful tool that can navigate the complexity of dense scientific documents.



2 | SPINOZA: A FIRST CHALLENGE MET — ON TO FURTHER DEVELOPMENTS

The first phase of Spinoza's development was successful: the tool was operational. The generative algorithm, initially a source of concern, was under control. Hallucinations were rare and the measures in place to ensure transparency proved effective. The singular press database provided figures, experts, and local initiatives – areas where other databases fall short or offer limited results. It was time for the second phase: the project team turned its focus to enriching the databases with more diverse and precise information while optimising the tool's prompts to better aid investigative journalism.

A/ ENHANCING DATABASES: GIVING ADDED VALUE TO A FUNCTIONAL TOOL

The tool worked but early testers reported that it did not provide enough added value or new and sufficiently localised information. The challenge was now figuring out **how to expand the range of data and improve the summaries' quality** without compromising the integrity of the information presented.

| STRENGTHENING THE PRESS DATABASES

The first database of press articles demonstrated its potential to participants, but they noted that the quality of the database's articles varied significantly. **A new database** was built with refined selection criteria.

- Articles must contain **at least 300 characters** and have been published **within the last two years**. Environmental journalism, largely based on scientific advancements, requires recent data. What's more, coverage of environmental issues has been steadily increasing in French media: **+17% from 2021 to 2023**, according to an Aday-ObSoCo study published in April 2024.
- The 1,000-article limit per participating media outlet was raised, **enriching the database** with more content that meets the selection criteria. The database now contains **28,450 articles**.
- Articles were processed to **eliminate duplicates** – due to the presence of media groups among the project's partners, the same article could appear in different regional newspapers. The initial filtering missed some duplicate articles, which were removed in the second phase.

| ADDING DATA FROM AGENCE FRANCE-PRESSE (AFP)

AFP joined the project in July 2024 as a contributor to the database. In exchange for access to the prototype, the press agency provided a dataset composed of news bulletins and other highly factual, explanatory news formats on climate change. This dataset includes **964 pieces of content**.

| INTEGRATING LOCAL PUBLIC DATA

At the launch of the Spinoza project, **the team established a shared file accessible to all journalists** from participating media outlets. This document, an Excel spreadsheet, allowed them to suggest public information sources they wanted to see integrated into the database. What initially served as a practical tool quickly became essential.

In May 2024, journalists specialising in climate issues from both regional and national press joined the Spinoza test group. Their main contribution was **expanding the database** by adding reports and sources they already used in their work, such as reports from the French Biodiversity Agency (OFB) and the High Council on Climate, to verify the tool's reliability.

This process had a collective impact. As the number of integrated sources grew, the tool's analyses became more precise and relevant. **The knowledge base grew** thanks to journalists who, by pooling their resources, enhanced the quality of the information that was available to everyone.

TESTIMONIALS FROM SPINOZA USERS

“*As journalists, we have a real responsibility to address climate issues effectively. Sharing public sources, collaborating with a network of colleagues covering major events, and pooling data... If this allows us to produce more reliable and insightful content for our audiences, everyone benefits.*”

“It’s truly a meta-editing process. We know these documents have been validated by newsrooms, which allows us to work with confidence.”

B/ THE COLLECTIVE REWRITING OF SPINOZA'S PROMPTS

The Spinoza project highlights the importance of human expertise in designing AI tools tailored to the specific needs of journalism and the right to reliable information. A key example is the collective rewriting of prompts to synthesise information from Spinoza's press article database.

| THE NEED FOR VISIBILITY AND CONTROL OF GENERATED DATA

The journalists expressed several key expectations in cases when their initial prompt was too restrictive. They needed the tool to:

- Display relevant documents even if they did not fully match the initial query;
- Provide clear summaries enriched with indirectly related information;
- Differentiate the processing of local and national data;
- Have limited autonomy when selecting information and encourage a proactive exploration of sources.

| A USER-CENTRED COLLABORATION

To meet these expectations, the technical team brought the journalists into a collaborative process to create concrete solutions for their needs while also balancing technical constraints:

- The journalists were given a testing tool that allowed them to experiment, modify, and visualise the impact of new prompts in real time.
- The workshop brought together the journalists, developers, and editorial managers, which helped prioritise needs and adapt the tool to real-world journalistic practices.

| HYBRID TECHNICAL SOLUTIONS

The final approach used structured prompts designed like algorithms with “IF” conditions incorporating journalistic logic. This method allowed for:

- An improved differentiation between local and general information;
- Certain limitations of language models, such as geographical or semantic confusion, to be bypassed.

C/ FROM TECHNICAL EXPERIMENTATION TO COLLECTIVE INTELLIGENCE: THE SHIFT TO OPEN SOURCE

| JOURNALISTIC EXPERTISE AS A DRIVER OF INNOVATION

The Spinoza project **stands out for actively involving journalists**, who transformed an initially rigid, cautious tool into a **relevant, efficient prototype**. Thanks to a rich database of press articles, the tool filled gaps left by other sources, such as scientific reports, which are often limited to specific regions or, conversely, describe phenomena on a very broad scale. This **collaborative effort** demonstrated how journalistic content complements and brings value to other data sources — proving its crucial role in developing generative AI tools.

By using the tool as a data exploration assistant, the journalists reinforced its **role as a complementary resource** rather than a replacement for their expertise. Its **transparency and specialised abilities** make it a valuable ally, especially for expert journalists who can efficiently navigate trusted sources.

By leveraging journalists' skills, the tool provides powerful editorial support. It is designed to assist and enhance journalists' work — not replace it.

TESTIMONIALS FROM SPINOZA USERS

“*What’s great about Spinoza is that it’s editorially driven and doesn’t do the work for us — it helps us do journalism.*”

| *“Spinoza is a game changer. It helps us improve content.”*

| *“With Spinoza, we can access vast amounts of information that inspire for articles.”*

| *“I use Spinoza very often to verify numbers.”*

| *“For finding figures, Spinoza is better than Google.”*

| *“I save time on research. And time, in terms of mental workload, is invaluable. Spinoza isn’t a poor replacement for journalists — it’s a great assistant for journalists.”*

| *“The fact that Spinoza’s databases are built by journalists makes us more inclined to trust the tool.”*

MORE POTENTIAL TO EXPLORE

The Spinoza project does not claim to meet all journalistic needs but **lays the foundations for a suitable method and infrastructure**.

Areas for Spinoza’s future development could include:

- **enriching existing databases**, especially on themes like climate change, by integrating new journalistic sources;
- **exploring new topics of public interest** with dedicated databases: geopolitics, international trade and organised crime, for example, are complicated themes that would benefit from a tool that helps journalists master their complexity;
- **developing new functionalities** to meet specific uses: vocal input, the development of a version that can run on a personal computer, the creation of a library of prompts for open-source algorithms, etc;
- **benefitting from future improvements** in language and reasoning models, thanks to the technology-agnostic approach adopted at the start of the project.

“*When our team’s journalists tested Spinoza, their first question was, ‘But why doesn’t this exist for other topics?’*”

| *“My question is: where can we find other databases to further improve the Spinoza tool?”*

TOWARDS AN OPEN-SOURCE COMMUNITY

To ensure broad, collaborative adoption, the project code will be released as an open-source resource in spring 2025 on the platform Hugging Face, where the project is currently accessible to a community of at least 200 participants. This open access will enable journalists and developers **to customise and adapt the tool to their needs**, and allow it to be used as a **discussion tool within editorial departments**, fuelling conversations between journalists and editors on the use of AI in news production.

To make this transition a success, we will need to **optimise prompts and create a dedicated catalogue**. The project’s final workshop, a hackathon held in November 2024, confirmed that the prototype’s duplication and customisation process was simple and easy to use. The hackathon highlighted three types of potential users:

- **non-technical journalists** who are able to copy and use the tool with open-source algorithms;
- **data journalists and developers** who are able to enrich databases and adjust prompts;
- **advanced developers** who are able to create customised functionalities and deploy the tool on specific infrastructures.

3

RSF REPORTERS
WITHOUT BORDERS

RECOMMENDATIONS

FOR DEVELOPING AN AI TOOL THAT RESPECTS JOURNALISM AND THE RIGHT TO RELIABLE INFORMATION

RSF presents ten recommendations for media outlets and newsrooms that wish to adopt generative AI tools for news production.

> **Regulate AI usage and practices in newsrooms based on ethical standards**

Drawing inspiration from the Paris Charter on AI and Journalism, led by RSF and 16 partner organisations, under the presidency of journalist and Nobel Peace Prize laureate Maria Ressa.

> **Develop extensions of the Spinoza project**

Share them within the journalism community to allow for third-party audits.

> **Foster collaborative projects between media organisations to develop AI tools**

Work with media outlets and groups to foster innovation and co-develop ethical tools.

> **Respect copyright and related rights for press publishers and journalists.**

Do not use publishers' data without their consent.

On the design of RAG-based tools:

> **Integrate journalistic methods when designing generative AI systems and appoint editorial supervisors**

Clearly define who is responsible for overseeing and supervising AI training and database creation.

> **Keep databases updated and ensure their reliability**

Ensure constant updates so that the information is as current as possible.

> **Track the origin of all sources used in databases**

List documents in a precise manner so journalists can clearly identify accessible information.

On writing AI prompts tailored to journalism in RAG systems:

> **Adapt prompts to each database**

If a database is updated, ensure the prompt is adjusted accordingly.

> **Define prompts collectively**

Keep prompts updated and available to users.

> **Maintain independence in the use of language models**

Experiment with various large language models (LLMs) and select the best-performing ones based on criteria established with users. Adapt prompts accordingly.



REPORTERS WITHOUT BORDERS (RSF) works for journalistic freedom, independence and pluralism all over the world. Headquartered in Paris, with 13 bureaus and sections around the world and correspondents in 130 countries, it has consultative status with the United Nations and UNESCO.