

INDEPENDENT WIRE

An independent newsroom

Open · Transparent · For everyone

66 SOURCES · 11 LANGUAGES · 3 DOSSIERS / DAY

The Problem

We don't read the news. We read the news someone selected for us.

News is not neutral. It never was. Every text that gets published passes through a chain of human decisions: Which topic is picked up? Which angle chosen? Which sources cited? Which facts emphasized, which omitted? At the end of every news chain stands a human — with career ambitions, worldviews, blind spots, and the perfectly natural desire to be heard.

This is not an accusation. It is an observation. Humans are not objective. We cannot be — and that is not a weakness, it is our nature. The weakness lies in a system that pretends objectivity exists.

Publishing houses have owners. Owners have interests. Editorial boards have lines. Editors have careers that depend on fitting the line. Those who take uncomfortable positions get isolated. Those who serve the line get promoted. This is not a system failure — it is the system.

THE COLORING

The full picture does not exist. There is always a colored one. The coloring comes from many directions.

- **Economic pressure** — Clicks determine budgets. Polarization generates clicks. Nuance does not.
- **Institutional bias** — Publishing hierarchies filter topics before they reach the reader.
- **Individual bias** — Journalists are human. They write what moves them, not necessarily what is relevant.
- **Cultural bias** — Western media reports western. Eastern reports eastern. The global remains a gap.
- **Attention economy** — What is loud gets heard. What is complex gets simplified. What grows slowly gets ignored.

The result: People receive news daily and believe they are informed. In truth, they see a slice — selected, weighted, and framed by actors whose interests are rarely transparent.

The Thesis

AI cannot eliminate human bias. But it can make it visible, systematically cross-check it, and reduce its impact.

Let's be honest: AI is not neutral. It is trained on human texts. Additionally, RLHF — Reinforcement Learning from Human Feedback — has pressed the models into a shape that reflects human preferences. The models respond as humans prefer. Not as would be necessary. This is a real problem.

And yet — despite all these limitations — AI has a decisive structural advantage in news production:

AI has no ego. No career. No network. No fear of consequences. No vanity. No agenda — except the one you give it.

That sounds banal. But in the context of a newsroom, it is revolutionary. An AI agent scanning sources has no reason to prefer one source over another. No reason to suppress a topic. It does not get tired, cowardly, or opportunistic. It can observe hundreds of sources in dozens of languages simultaneously — 24 hours a day, 365 days a year.

The question is not whether AI can deliver perfect journalism. The answer is: No. The right question is:

Can AI be systematically better than the structural distortion of the current system?

The thesis of this paper is: Yes. Not perfect. Not infallible. But systematically less distorted, more broadly informed, and more transparent in its decisions.

The Vision

An independent newsroom in your pocket. Independent of publishers, algorithms, and vested interests.

Imagine: Every person has access to a personal news system that:

- **Scans globally** — hundreds of sources in dozens of languages
- **Systematically cross-checks** — every claim validated against original sources
- **Shows perspectives** — not one line, but the spectrum of positions
- **Works transparently** — every decision, every source, every step traceable
- **Belongs to no one** — open source, self-hosted, under the user's control

The Approach

Independent Wire answers a different question than existing news tools. Aggregators answer: "What is happening?" Bias classifiers answer: "Is this source left or right?" Independent Wire answers:

|"How is it being reported — and what is missing?"

The system scans 66 sources across 11 languages, identifies where coverage diverges across languages and media systems, documents which voices are absent, and publishes everything — including its own biases and limitations. No publisher. No engagement algorithm. No hidden editorial line. Every decision the system makes is traceable, auditable, and open.

The output is not an article. It is a **Topic Package** — a structured transparency bundle that makes visible what conventional newsrooms hide: the editorial decisions behind every published text.

Why Open Source

The core of this vision stands or falls with a single condition: The system must belong to no one. Not a startup, not a corporation, not an investor.

The moment an economic interest gains control over the prompts — the instructions that determine how an AI agent understands its task — independence is lost.

Open source is not a technical feature. It is the only guarantee that the system does what it promises. Who can read the prompts can check the agenda. Who can change the prompts is free.

Independent Wire is licensed under **AGPL-3.0** — the strongest copyleft license available. It ensures that anyone who hosts the system must also open-source their modifications, including changed prompts. This is not a restriction. It is the license doing its job.

The Architecture

The system consists of specialized AI agents, each fulfilling a clearly defined task. No agent has access to the entire system. No agent makes decisions alone. The agents do not communicate directly with each other but through a deterministic pipeline that orchestrates their work sequentially.

Much of the work between the agents — fetching the sources, clustering the day’s findings, assigning them to stories, merging and numbering sources, aggregating the bias dimensions — is done by deterministic Python, not by an agent. An LLM is used only where the work is genuinely non-deterministic: discovering stories, researching, mapping perspectives, writing, fact-checking, and analyzing linguistic bias.

AGENT	TYPE	TASK
Curator	Per run	Identifies the day’s distinct stories from the collected findings.
Editor	Per run	Conducts the editorial conference. Prioritizes stories. Has persistent memory of previous coverage.
Researcher	Per topic	Builds the evidence dossier — plans the questions, searches original sources, and assembles what is known and where it conflicts.
Perspective Agent	Per topic	Maps the spectrum of positions: who says what, where they contradict each other, which voices are absent.
Writer	Per topic	Writes from the dossier. Represents positions without favoring any.
QA / Fact-Check	Per topic	Checks every claim against the cited sources. Flags factual errors, unsupported claims, hidden divergences, and misleading framing — and corrects them in place.
Bias Detector	Per topic	Analyzes the finished text for linguistic bias. Marks — never rewrites.

The pipeline is deterministic: Python code orchestrates which agent runs when, with what input, and what output. No LLM decides “what to do next.” This is a deliberate architectural choice born from operational experience — LLM-orchestrated workflows fail reliably.

The Topic Package

Independent Wire does not produce articles. It produces **Topic Packages** — structured transparency bundles that contain everything.

- **The facts (OSINT layer)** — What happened, according to whom, in what language, and where sources disagree.
- **The perspectives** — A structured map of who says what, how strongly represented, from which region.
- **The divergences** — Where sources explicitly contradict each other — factually, in framing, or through omission.
- **The gaps** — What's missing — which regions, demographics, or aspects no source covers.
- **The article** — A rendered text that presents the facts — source-based, multi-perspective, transparent about uncertainties.
- **The bias card** — A systematic analysis across five dimensions — language, source, framing, selection, geographical.
- **The transparency trail** — Why this topic was selected, what was discarded, what the confidence level is.

The article is one possible rendering of the Topic Package — not the Package itself. The same data can produce a website article, a podcast briefing, a newsletter, or an API response. Different formats, same transparent data foundation.

The Bias Card — The Heart

The bias card is what distinguishes this system from every conventional newsroom. It covers five dimensions — and how each one is produced is itself part of the transparency.

Language bias is analyzed by a dedicated agent reading the finished text: Does it use evaluative adjectives? Emotionalizing formulations? Passive constructions that obscure responsibility?

The other four dimensions are not the opinion of a second AI. They are derived from the pipeline's own structured record — the source list with its countries and languages, the perspective map, the documented divergences, and the catalogue of what is missing:

- **Source bias** — Are certain positions supported by more or more authoritative sources than others? Are relevant counter-voices missing?
- **Framing bias** — Which metaphors are used? Which associations are evoked? Are there alternative framings?
- **Selection bias** — Which facts were included, which omitted? Are there relevant contexts that are missing?
- **Geographical bias** — Is the topic covered from a western, eastern, global perspective? Which regions are missing?

Reading these four off the data rather than asking a model to judge the text again is itself a transparency gain: they are as auditable as the data they come from, and they carry no second layer of model opinion.

The card does not deliver a "neutral" text — there is no such thing. It delivers a **Bias Transparency Card**: a document that discloses to the reader where and how the text is colored. The reader decides for themselves how to interpret it.

The goal is not objectivity. The goal is transparency. Not: "This is the truth." But: "This is what we found. Here are the sources. Here are the limits. Judge for yourself."

The system writes deliberately flat prose. This is not a limitation — it is a position. Every stylistic enhancement is an editorial decision. A "punchier" opening selects what the reader perceives first and with what emotion. The Bias Detector proves the point: even in deliberately restrained text, it finds framing. If the prose were more vivid, the Detector would either produce more flags (which looks absurd) or be calibrated down (which destroys credibility). The flatness is the feature.

Visualizations, Not Illustrations

Independent Wire does not use AI-generated images. No photorealistic illustrations, no simulated photography. This is a deliberate philosophical choice: AI-generated images simulate a reality that nobody photographed. This directly contradicts the project's core principle of transparency.

Instead, all visuals are data-driven diagrams generated deterministically from structured data.

- **Perspective Spectrum** — Who says what, how strongly represented
- **Source Map** — Geographic distribution of sources — who speaks, who is silent
- **Divergence Chart** — Where sources contradict each other
- **Fact Check Diagram** — Verification status of all claims — color-coded, instantly graspable

What AI Cannot Do

Honesty about one's own limits is the prerequisite for credibility.

AI cannot investigatively research. An AI agent cannot conduct confidential conversations, meet whistleblowers, request files, or visit a factory. Investigative journalism needs humans — on the street, with the people, on site. This system does not replace investigative journalists. It aims to relieve them by taking over routine work so they can focus on what only humans can do.

AI cannot eliminate real bias. The models carry the bias of their training data. RLHF has introduced additional distortions. A Bias Detector can mark distortions — but it cannot fully resolve them because its own judgment is not free from bias. The honest aspiration is not neutrality, but bias reduction and bias transparency.

AI cannot have style. Good journalism lives from the author's voice — from attitude, experience, humanity. AI-generated texts can be correct, structured, and informative. But they have no lived experience. The system should therefore not try to imitate human voices but find its own, transparent style: Clear, source-based, without presumption.

AI cannot understand what matters. Relevance is a deeply human concept. What is relevant for a mother in Toronto is not for a fisherman in Mombasa. The best stories begin where the data is still silent.

Source balance is uneven. Western and English-language sources are over-represented. Hebrew-language media is more accessible via RSS than Lebanese or Palestinian sources. This is not a bug — it reflects which voices are structurally easier to reach in the global media system. Improving source balance is an active and open workstream where community contributions have direct impact.

The system runs partially on proprietary models. Every prompt, every source list, and every line of editorial logic is open and auditable. The models are a split: several stages already run on open-weight models — downloadable, and self-hostable on your own hardware, independent of any provider — while others still depend on a proprietary model with no open alternative at the same quality. Today even the open-weight stages run through a commercial provider, for one reason only: self-hosting them requires hardware we do not yet have the means to operate. The independence is built into the choice of models; exercising it fully is a question of resources, not architecture. We name this rather than hide it. The active goal is to migrate every LLM stage to open-weight, self-hostable models, on the explicit condition that quality does not degrade — open models are adopted when they hold the quality, not for their own sake. The endpoint is a system where prompts, sources, logic, and inference are all independent of any provider — where "verify it in the source code" extends to the models themselves. Like source balance, this is an open workstream where community work has direct impact.

These limits are not arguments against the system. They are arguments for honesty. A system that knows and communicates its limits is more trustworthy than one that promises objectivity.

The Path

PHASE 1 – THE PROTOTYPE · OPERATIONAL

Operational pipeline since April 2026. A lean, purpose-built Python framework — not a general-purpose agent platform. 66 global sources across 11 languages. Daily publication of three Topic Packages at independent-wire.org. Pipeline cost a couple of euros per run. All prompts, source lists, and editorial logic published under AGPL-3.0. Model evaluation across 14 models and 90+ test calls completed.

PHASE 2 – THE COMMUNITY

A system that should belong to no one must be carried by many. Open-source community of developers, journalists, and media scholars. Open API for external use of research data. Prompt library with community-reviewed agent instructions. Collaboratively maintained and rated source catalog — especially from underrepresented regions. Migration of the remaining proprietary LLM stages to open-weight, self-hostable models, without sacrificing quality. MCP server providing Topic Packages as structured data to Claude, ChatGPT, and other LLM clients.

PHASE 3 – THE SCALING

Topic and region configuration. Trust network for community-based evaluation of sources and agents. Docker deployment. Narrative tracking across days and weeks.

SUSTAINABILITY

No commercial model. No advertising. No investor equity. Operating cost: a couple of euros per daily run. Funding paths: community sponsorship, public interest grants, institutional partnerships. The system is designed so that its operating costs never create pressure to compromise editorial independence.

The Difference

	EXISTING AI NEWS	INDEPENDENT WIRE
Owner	Companies (Google, Apple, OpenAI...)	No one. Open source. AGPL-3.0.
Prompts visible?	No — proprietary and secret	Yes — every instruction is readable
Source list	Opaque	Open, community-maintained
Bias transparency	None	Bias card per topic
Business model	Advertising, subscriptions, data	None. Operating costs = AI API costs
Control	With the provider	With the user
Perspectives	Algorithmically optimized	Systematically broad, explicitly contrasted
Images	AI-generated illustrations	Data-driven diagrams only
Output	Proprietary app	Topic Package JSON, open API, any format

The fundamental difference is not in the technology. The technology is largely the same. The difference lies in the question: *Who owns the control?* In existing systems, it lies with the provider. In Independent Wire, it lies with the user. This is not a marketing promise — it is an architectural fact that anyone can verify in the source code.

Closing Words

There is a sentence that comes up again and again in discussions about AI journalism: "But that can't replace real journalism."

True. It can't. It shouldn't.

This system is not meant to replace the journalism that happens on the street — the questions only a human can ask, the stories only someone who was there can tell. It is meant to replace the journalism that happens *despite* the people — the structural distortion, the economic pressure, the institutional filters that ensure only a comfortable slice of reality remains.

Whether that works is an open question. But the tools exist. The architecture stands. The code is open.

All that is missing are people who join in.

This document is not a business plan. It is an invitation.

INDEPENDENT WIRE · independent-wire.org

github.com/deniz-schwenk/independent-wire

AGPL-3.0 — Because transparency is not a feature, it is a promise.