

— SECTOR REPORT · AGE OF INTELLIGENCE 2026

The State of AI in *Media*.

What artificial intelligence is actually doing to the industry in 2026: where it saves time, where it breaks the model, and which asset it makes more valuable.

Editorial

Two Things Happening *at Once*.

The industry celebrates every new tool while the channel that carried its work to readers is dismantled. The two stories are almost never read together.

Every week another newsroom announces an artificial intelligence tool. Automatic tagging, summaries, press-conference transcription, translation, synthetic voices, subscriber assistants. The list grows fast and the industry celebrates it with good reason: these are real improvements on real work.

Over the same period, trust in news has fallen to its lowest level on record, search traffic has collapsed, and almost half of Spanish news sites have shut the door on the crawlers that feed the large models. These facts are contemporaneous, not consecutive.

This paper reads them together, and that reading produces an uncomfortable thesis: artificial intelligence is solving cost problems for publishers while dismantling their distribution model. A newsroom can automate most of its mechanical work and still lose its relationship with its audience, because that work was never the problem.

What follows is not a tool catalogue or an inventory of use cases. The cases are here, concrete and verifiable, but they serve a position: where intelligence belongs in a media organisation, which asset has to be protected before any other, and how a programme proves itself when the committee that asks about return finally arrives.

Thinkia publishes this paper in its *Age of Intelligence* series, which defends one simple and demanding idea: an organisation does not have artificial intelligence because it bought models, but because one of its processes would stop if it switched them off. Applied to media, that idea cuts through a good deal of noise.

One note on tone. Nothing here suggests that newsrooms are making fools of themselves. They are doing the reasonable thing with the information they have, in an industry with less room than ever. The problem is not the effort. It is where the effort is being spent.

EXECUTIVE SUMMARY

Artificial intelligence entered the newsroom through the door marked “tasks”: tagging, summarising, transcribing, translating, verifying. Twenty news organisations documented in 2026 confirm it, from El País to the Associated Press. Almost all of those applications are assistance; very few touch a process that, switched off, would show up in the accounts.

Meanwhile the same technology is dismantling distribution. Eight in ten searches that return a generated summary end without a click to the open web, and publishers have started buying the traffic they used to get for free. Trust is falling in twenty-nine of the forty-eight markets surveyed.

The asset that survives that shift is verifiable provenance: being able to demonstrate, item by item, where a piece of information came from, who reviewed it and which tool produced it. Everything published without that capability is **provenance debt**, a liability that appears on no balance sheet until it appears all at once. This report maps the cases by function, examines sport as the industry’s test bed, and closes with four scenarios, each with the condition that decides it.

45%

TRUST IN NEWS BRANDS

against the 19% Spanish readers grant to news from social platforms and AI chatbots.

83%

ZERO-CLICK SEARCHES

when the search engine shows an AI-generated summary above the results.

46%

SITES THAT BLOCK

of Spanish news sites restrict access to the crawlers of the large models.

CONTENTS

01	The Ground Moved	p. 05
02	Editorial Theatre	p. 07
03	Provenance Debt	p. 10
04	Governance Is No Longer Optional	p. 12
05	What Is Already Being Done	p. 13
06	Where the Publisher Recovers Price	p. 19
07	Sport, the Test Bed	p. 22
08	What Comes Next	p. 26
09	What to Measure	p. 28
10	Conclusion	p. 29

01

The Ground Moved While We Watched *the Tool*.

For two years the conversation about artificial intelligence in media has circled the newsroom: which tasks get automated, which job is at risk, which byline is still human. It is a legitimate conversation and it is pointed at the wrong place. The violent change did not happen inside the newsroom. It happened on the road that connected the newsroom to its reader.

83% of searches that return an AI-generated summary end without a single click to the open web; in the search engine's conversational mode the figure rises to 93%.¹ For informational searches, which is exactly what someone runs when they want to find out about something, 74% produce no visit. The Spanish industry's response was immediate and revealing: publishers raised paid-search spending by 274%.² They have started buying the traffic they received free for fifteen years.

“

The change did not happen inside the newsroom. It happened on the *road* that connected it to its reader.

At the same time the floor of credibility has given way. Trust in news is falling in twenty-nine of the forty-eight markets tracked by the Reuters Institute and sits at its lowest level since measurement began in 2015.³ In Spain it stands at 33%.⁴

One figure changes the sign of everything above. Asked whether they can tell truth from falsehood online, 74% of Spaniards say they are worried: the highest reading in the series.⁴ That is not the industry's problem. It is unmet demand.


01 · THE GROUND MOVED

Spanish audiences still tell one source apart from another, and they do it with a clarity the industry tends to forget when it measures itself by traffic alone.

EXHIBIT 01

News brands hold more than twice the credibility of social platforms and chatbots.

Trust declared by Spanish audiences, %



Source: Digital News Report Spain 2026, University of Navarra / YouGov · n > 2,000 adults.

Twenty-six points separate the masthead from the chatbot. That gap is the one asset synthesis cannot manufacture, because it does not rest on the quality of the text but on who answers for it. The operational question is whether a publisher can prove it on demand.

The first bar also deserves to be read backwards. That only 19% of the public grants credibility to what reaches them through platforms and assistants does not mean they consume little there: 60% reach news through social networks and aggregators, and 40% consumed information from creators in the past week. They consume where they do not trust. That mismatch is the commercial space left open.

The worrying part of the data sits at the base of the age pyramid. 53% of under-25s combine low interest in news with low trust, against 35% of the population as a whole. A brand does not inherit the credibility it enjoys today: it renews or loses it with every cohort.

02

Editorial Theatre.

There is a quick test for how much artificial intelligence a publisher really has, and it is not run in the innovation committee. It is run in the workflow: if AI has not touched the place where the newsroom decides what gets published, how it gets verified and who gets billed, it has not touched the business. The rest is material for a deck.

The diagnosis is not an accusation, and that is worth saying plainly because the opposite temptation is strong. Nobody chooses to build a stage set. It happens through inertia: decisions about artificial intelligence get taken in places that can open a project but cannot change a process, and the result is a collection of pilots that never scale and that, instead of being cancelled, get renamed.

Observer bias completes the picture. People see artificial intelligence where it is most talked about, not where it is most used. The lab produces material to show; the newsroom produces editions. If the only evidence of a programme is a demo, the programme does not yet exist.

“

In a mature newsroom, a pilot that does not scale is cancelled. In one that has entered theatre mode, it is *renamed*.

It helps to separate symptom from cause. A pilot rarely fails because of the model: it fails because whoever sponsored it could open an initiative but could not touch the edition deadline or the verification flow, which is where the business works.

One concession, so the diagnosis is not unfair: part of what looks like theatre is legitimate learning. Testing before committing the publication flow is prudent in an industry where mistakes are paid in credibility. The problem is not the pilot, but the pilot with no decision date.

02 · EDITORIAL THEATRE

Five signs that a publisher's AI is stage scenery.

They travel together. Three or more at once are enough to raise a hand in a board meeting.

01

The perpetual pilot.

The tool has been in trial for eighteen months with the same small team. Every quarter the scope widens and the decision slips. Nobody has set the condition that would move it into production, or the one that would kill it.

02

The policy that exists only as a PDF.

The publisher has a responsible-AI document, approved and published. At the moment the model is invoked, that document does not exist: no control enforces it and no record proves it was enforced.

03

The declarative indicator.

What gets measured is how many journalists have access to the tool, how many licences were consumed and how many training sessions were delivered. None of those figures says whether a single editorial decision changed.

04

The lab without authority.

An innovation team has its own budget and no power to change the publication flow. It produces prototypes the newsroom admires and does not adopt, because adopting them would require changing a process the lab does not govern.

05

The byline without traceability.

The publisher states that every AI-assisted piece goes through human review. Asked about a specific piece published three months ago, it cannot say who reviewed it, which model generated it or what was changed.

02 · EDITORIAL THEATRE

Four results the industry rarely cites.

A paper that argues for measuring the process rather than the tool is obliged to look at the evidence that runs against its own industry. The four cases below are published, solid, and largely absent from the slide decks.

01

Klarna, and above all its epilogue.

The support assistant absorbed 2.3 million conversations in its first month, two thirds of all chats, across 23 markets and 35 languages, and matched the work of 700 agents. Fourteen months later the company hired people back for complex calls, because customers preferred talking to someone.³⁹ The case is almost always cited without the second half. With it, it stops being a promise of replacement and becomes an architectural instruction.

02

Calling someone who did not call you can trigger the cancellation.

A peer-reviewed field experiment with 64,147 customers proactively contacted some of them to recommend a better-fitting plan, incentive included. Churn rose from 6% in the control group to 10% in the treated one: the call breaks the customer's inertia and makes their consumption pattern salient.⁴⁰ Inbound contact is an opportunity; outbound is a decision with a cost.

03

Auto-renewal wins a year and loses two.

A two-year experiment across more than two million prospective subscribers of a large European newspaper compared auto-renewal with auto-cancellation contracts. The first produced 20% more sign-ups in the first four months; at two years the balance was 10% fewer new subscribers, and close to half of those who signed up ended up paying for something they did not want and barely read.⁴¹ A retained subscriber who does not consume is not lifetime value: it is deferred attention cost.

04

Perceived productivity does not match measured productivity.

A randomised trial with experienced developers across 246 real tasks measured that they took longer with AI assistance, although they had expected to go faster and, afterwards, still believed they had. The authors themselves later qualified the result after finding severe selection bias, and redesigned the methodology in February 2026.⁴² The citable part is not the sign: it is the distance between what a team believes it gained and what the stopwatch says.

Provenance Debt.

Once automated synthesis takes the visit and the reader can no longer tell where a text came from, value moves from the text to its origin. Information that cannot account for its provenance becomes indistinguishable from its imitation, and competes in a market where imitation has zero marginal cost.

We call that gap **provenance debt**: everything a publisher puts out without being able to show, item by item, where it came from, who reviewed it and which tool produced it. It is a silent liability. No quality audit detects it, it blocks no publication and it sits on no balance sheet, until the day someone asks about one specific piece and there is no answer.

The debt is not repaid by publishing an editorial policy. It is repaid by installing the record at the point where the piece is produced, the only moment when information about its origin exists without effort. Capturing it afterwards means reconstructing it, and reconstructing provenance is exactly what a publisher cannot afford to do under pressure.

The corollary is awkward for anyone who has ordered their priorities the other way round: the provenance layer has to be built before the generation layer. A publisher that generates before it can trace accumulates debt at machine speed.

The archive enters here as an asset, not as a storage cost. Chicago Public Media opened forty years of radio and Eater connected twenty years of recommendations to conversational search. In both cases the value was not in the model: it was in a proprietary body of work nobody else could offer, and that can now be cited.

KEY TERM

Provenance debt

The liability accumulated by every published item that carries no verifiable record of its origin, its review and the tools used to make it.

- Editorial quality control does not detect it.
- It does not block publication: it surfaces much later.
- It grows with every model-assisted piece.
- It is only amortised by capturing origin at the moment of production, never afterwards.

A memory that cannot cite its sources is not memory. It is rumour with an interface.


03 · PROVENANCE DEBT

The Spanish industry's defensive reaction shows where it believes the risk sits, and the belief does not match the damage. Almost half of all news sites block the crawlers of the large models.⁹

EXHIBIT 02

Publishers block whoever trains on them, not whoever answers in their place.

Spanish news sites restricting each crawler, % · n ≈ 150



Source: Vipnet360 audit with AIBrandPulse360 across almost 150 Spanish news sites, June 2026. The last bar is the average restriction applied to agent crawlers, against 37.2% for training crawlers.

The asymmetry is the story: 37.2% blocking for crawlers that train models and 10.6% for those that read the page in real time to answer in the publisher's place. The industry has taken a position on its past and not yet on its present.

The internal split confirms the reading. The large groups average more than five blocks per site, while independent digital natives sit at 1.1. The first are defending an archive built over decades; the second need somebody to cite them. Both positions are coherent, and neither answers the underlying question: what happens when the answer is delivered without anyone visiting the source.

Blocking is a decision about who comes in. Provenance is a decision about what the publisher can prove when somebody doubts one of its pieces. Only the second still has effect once the content circulates beyond its own site.

04

Governance Is No Longer *Optional*.

Since 2 August 2026, Article 50 of the EU Artificial Intelligence Act has applied.¹⁰ It requires synthetic content to be marked in machine-readable form and the artificial nature of manipulated material to be disclosed. For a publisher, the decisive part is not the obligation: it is the exemption.

Content generated or manipulated with AI may omit disclosure where there is effective human editorial review and an entity that assumes responsibility. Read quickly it sounds like permission; read carefully, it is a burden of proof. It turns “we have human oversight” into something that must be demonstrated item by item: the provenance debt of the previous chapter, now with a penalty attached.

Article 50, in force since August 2026

The transparency obligations that bear directly on a newsroom:

- Inform users that they are interacting with an AI system, unless it is obvious. This covers subscriber assistants and archive chat.
- Mark generated audio, images, video and text in a machine-detectable way.
- Disclose the artificial nature of *deepfakes* and of manipulated news content.
- Editorial exemption: disclosure may be omitted where there is real human review, effective editorial control and an entity that assumes responsibility.
- Non-compliance: up to €15 million or 3% of annual worldwide turnover.

05

What Is Already Being Done, and What *Still Isn't*.

The 2026 inventory is solid and documented. The Associated Press reviews overnight output, verifies images through geolocation and has turned thousands of Supreme Court documents into structured, searchable information. Le Monde built its translation style guide into the model and added a voice track to every article. The Philadelphia Inquirer uses *Scribe* to monitor public meetings across dozens of municipalities, summarise them and score their news value.⁷

In Spain, El País has launched **GulA**, an editorial assistant built into the content management system that tags, links related stories, generates summaries and improves contextualisation between pieces, with the declared ambition of creating a specialist AI capability inside PRISA Media and extending it to other titles in Spain and Latin America.⁶ The group also runs a trend tracker, country audio briefings and Vera, a conversational assistant for subscribers.

Unidad Editorial runs **Calíope** across *El Mundo*, *Marca* and *Expansión*: editorial assistance plus automatic generation of recurring items, from weather to Euribor rates. It is a €995,724 project with €390,289 in public funding.⁴⁶ It is not an isolated case: Spain's digital transformation ministry distributed €6.49 million across publishing groups, with El País leading at €1.07 million.

Seen together, the list has a clear pattern: almost all of it is task assistance. Very few of these applications touch a process that, switched off tomorrow, would leave a visible hole in the accounts. The ones that do can be counted on one hand.

250M

reader questions answered
by Hey_ at BILD

40

years of WBEZ audio
archive, now searchable

20

newsrooms with
documented uses in 2026

Real uses, sorted by layer of intelligence.

LAYER	WHAT IT SOLVES	DOCUMENTED EXAMPLES IN 2026
Perception	<i>What is happening?</i>	Image and video verification through geolocation at the Associated Press; monitoring of council meetings with <i>Scribe</i> at the Philadelphia Inquirer; breaking-news detection at WELT.
Comprehension	<i>What does it mean?</i>	Tagging and contextualisation between pieces with <i>GulA</i> at El País; thousands of court rulings turned into queryable data at AP; forty years of audio archive opened at Chicago Public Media.
Action	<i>What do we produce?</i>	Style-guided translation and per-article audio at Le Monde; country audio briefings and translation at PRISA Media; headlines and photo captions at Axios.
Relationship	<i>Who are we talking to?</i>	<i>Hey_</i> at BILD; summaries personalised by interest and location at The San Francisco Standard; conversational search across twenty years of archive at Eater.
Business	<i>What do we live on?</i>	Dynamic paywall deciding the offer in real time at the <i>Financial Times</i> ; propensity scoring that reorders the telemarketing list at Schibsted; retention segmented by risk at Mediahuis; a downgraded product as a cancellation net at <i>The Australian</i> ; ad prospecting at The Seattle Times.

The fifth layer is the least cited and the only one that decides whether the programme survives the next budget.

05 · WHAT IS ALREADY BEING DONE

The business layer is not empty: it is badly told.

The cases exist, they have published figures and they have been in production for years, but they circulate in the subscription and monetisation circuit and almost never in the editorial innovation one. They are also the only ones that answer the board's question without hedging, because they were born from it.

01

The Financial Times puts the paywall in charge.

A probabilistic model picks in real time which wall and which offer each reader sees, across some 250 user states and between 30% and 40% of the audience. The published result is a 290% increase in conversion rate.¹⁹ What makes it the best case in the inventory is not the figure: it is that the publisher itself acknowledges the sample may be biased by purchase intent and is building control groups to isolate the real uplift. It publishes its result and its doubt about its result at the same time.

02

Schibsted does not change the product: it changes who gets called.

A weekly model scores individual propensity to subscribe and decides who telemarketing calls. At *Aftenposten*, conversion in the top decile moves from 1% under random selection to 6%; at *Fædrelandsvennen* it reaches 8.8%; in the *Bergens Tidende* app, the impression rate on target audience is 28.2% against 6.3% random.²⁰ The product was untouched, the sales team is the same and so is the budget. All that changed was the order of the call list.

03

Mediahuis spends the effort where the risk is.

The same kind of model applied to retention, with the channel differentiated by risk level: a personal call for high risk, an email with a video from the editor for medium risk, and nothing for low risk. Retention rises 14.17% in the first group and 8.99% in the second.²¹ The decision not to act on a segment matters as much here as the model.

04

Subscriber service has a known ceiling.

Der Spiegel automates close to 80% of the queries on its subscription pages with a declared 75% satisfaction; the *Frankfurter Allgemeine Zeitung* resolves 71% and converts 18% into a click inside the same chat that handles cancellations.²² These are the only two serious published publisher cases, and both come with vendor-supplied figures: read them with that caution.

05 · WHAT IS ALREADY BEING DONE

And three cases where the bottleneck was not the model.

05

The Australian turns cancellation into a step down.

In the cancellation flow, instead of a discount, they offer a genuinely downgraded product: a registration wall with ten free articles in four weeks and no payment details. Whoever leaves does not leave entirely; they stay as an active registered user. The save rate is 40%, against 6–16% for the classic discount and 0% for the flow that only asks why. 7% of those downgraded return to subscription, and 85% of returning subscribers had passed through it.²² It is the case with the least artificial intelligence and the most product design, which is exactly why it belongs here.

06

NZZ: 80% of the value was not in the model.

The *Neue Zürcher Zeitung* replaced the fixed wall with between one hundred and one hundred and fifty rules, plus a propensity model that decides the moment, the copy, the placement and even the colour of the prompt. During the commute they suppress it on mobile; the heavy registered reader is offered the annual plan. Conversion multiplied fivefold in three years, and the personalised greeting alone raised it by 25%.²⁴

07

Registration is the multiplier nobody disputes.

The New York Times turned registration and some fifty free newsletters into the mandatory step between anonymous audience and subscription: 135 million accumulated registrations, fifteen million weekly newsletter readers, and registered users who convert to subscription at **forty times** the rate of anonymous ones.²⁵ In an industry where the median share of identifiable users does not reach 2%, the margin is in the intermediate step, not in the height of the wall.


05 · WHAT IS ALREADY BEING DONE

The cancellation flow is the most heavily measured moment in the whole subscription business, and also the one that best shows where the real constraint sits.

EXHIBIT 03

What saves a cancellation is not the discount: it is having a step below.

Save rate in the cancellation flow by what is offered inside it, %



Source: aggregated Piano data across its global publisher base, and INMA analysis of The Australian, January 2019 – March 2025.

The jump from 16% to 40% is not produced by a better model but by the existence of an intermediate product to downgrade the leaver into. It is the cheapest reminder that in the business layer the constraint is usually catalogue, not technology.

05 · WHAT IS ALREADY BEING DONE

Two places where AI already operates, and nobody looks.

The sector's inventory walks through the newsroom, the archive, the reader relationship and the commercial department, and stops just short of two places where artificial intelligence has been producing measured effects for years. Neither appears in a conversation about the future of journalism, because neither is remotely futuristic.

The printing plant. *L'Équipe* replaced manual distribution planning with a model that predicts daily demand across some 25,000 points of sale in France, using history by period, public holidays, weather and the sporting calendar. The reduction in the unsold rate exceeds 5%: roughly two million copies that stopped being printed.²⁶ The same supplier's aggregate across France and Germany sets the realistic expectation and is worth citing alongside it, because it guards against promising too much: in national dailies, between 4% and 8% fewer unsold copies with between 0.8% and 1.5% more sales.²⁷

There is one design nuance that decides the outcome. *The Telegraph* optimised not the accuracy of the forecast but the profit, building in the cost of printing and transporting a copy that comes back: it is not a prediction problem, it is a decision under asymmetric cost.

“

The best-documented return in the sector does not decide what gets published. It decides how many *copies come back*.

The ad server. While committees debate which tool to buy, part of the artificial intelligence operating inside a publisher is already switched on by default within the systems it has under contract. Google announced in June 2026 a conversational agent inside its ad manager that diagnoses delivery issues in real time and generates reports from a natural-language instruction.³⁰ It publishes no impact figures, and that is precisely the point.

The question is not which tool to buy for the commercial department, but what the one that comes included is already doing, and with what effect on the price the publisher gets for its inventory.

06

Where the Publisher Recovers *Price*.

The conversation about artificial intelligence and advertising has been fought almost entirely on the advertiser's side, which has spent a decade automating its buying. On the publisher's side there is less noise and three levers with published figures, and all three share the same pattern: the value is not in the data, it is in the point of the chain where it is applied.

AI-driven contextual targeting is a publisher capability. News UK was the first publisher to apply Illuma's AI categorisation across *The Sun* and *The Times*, with an average 10% improvement in segment accuracy, up to 25% in beauty and 20% in sport, and 37% in brand-safety categories.³⁷ News Corp took the same engine to the expansion of live campaigns within its own inventory, with average performance three times higher than its direct activity without it.³² Both figures are reported by the publisher and the supplier, without external audit.

First-party data pays when it is applied on the publisher's side, not when it is exported. The same information is worth more inside the publisher's platform than handed to the buyer: differences of 37% in the price obtained have been measured depending on which end of the chain applies it.³⁴ Direct collaboration with the advertiser in a secure environment points the same way: matching Renault's CRM against Axel Springer's audience, without sharing personal data and with a control group, produced 18% more conversion, 38% more delivery against the target profile and 15% lower cost per acquisition.³⁵

And the strongest price lever is not the algorithm: it is curation. Packaging enriched inventory into curated deals, rather than selling it loose, is the only mechanism that gives pricing power back to the publisher. Whoever does not curate their own inventory does not avoid its being curated: a third party curates it, and that third party keeps the margin.


06 · WHERE THE PUBLISHER RECOVERS PRICE

There is one warning worth giving on the same page as the figures above, because it is the most useful number on the whole advertising front.

EXHIBIT 04

Contextual targeting accuracy varies by almost three times depending on who sells it.

Contextual relevance verified by an independent annotator, % · 1 million impressions, same inventory



Source: 2020 study of live campaigns from four brands, with third-party cost measurement and relevance annotation by an external specialist. The data is years old; what has not been refuted is the spread.

The conclusion is not which supplier wins, which is in any case a figure with some years on it. It is that the difference between best and worst is almost three times over exactly the same inventory,³³ and that no buyer asks for the number. Buying contextual targeting without auditing its accuracy is buying noise at the price of data, and auditing it is the cheapest purchase on the entire front.

06 · WHERE THE PUBLISHER RECOVERS PRICE

The frontier that has not yet moved a price.

A survey of more than seventy-five publisher branded-content studios worldwide, published in July 2026, gives the sector's most honest self-portrait: 65% use artificial intelligence routinely, but 73% of those only with general-purpose assistants, and just 2.7% believe it is changing their work. 81% have not touched their pricing model and 74% have not touched their headcount.³⁶ The technology has entered the branded studio and has moved neither the price nor the structure.

The cases with metrics all point the same way, and it is not the way usually assumed: artificial intelligence wins in interactive tools and visual assets, not in writing the article. A menu-inspiration assistant built for a soup brand beat its publisher's benchmark dwell time by more than 570%.³⁷

And the only case of a publisher awarded for innovative use of artificial intelligence says everything worth knowing about where the advantage lies: it was built by analysing more than 4,500 corporate earnings calls to contrast what executives say against what analysts ask.³⁸ What made it impossible to replicate was not the model, which anyone can rent, but the proprietary corpus it ran on.

“

In AI-assisted branded content, the advantage is *the archive*.

07

Sport, the Sector's *Test Bed*.

If there is one vertical where this discussion has stopped being theoretical, it is sport. Modern sport produces structured data in real time, at volume, with demand that does not wait: positions, events, minutes, performance. It is exactly the kind of raw material a machine handles well, which is why automation has arrived earlier and gone further here than in any other section of a publication.

It is also where the trust problem bites hardest. No other kind of coverage generates as much synthetic content in circulation, or as many incentives to manufacture it: a fabricated quote from a player spreads in minutes, and denying it costs more than publishing it.

Both things happen in the same newsroom and on the same budget. That is why sport works as a test bed for everything this document argues: here, before anywhere else, you can see which part of the craft the machine absorbs, which part gains value, and what happens when a publisher cannot certify what it publishes.

The sector's own conclusion is less apocalyptic than usually told, and considerably more demanding. "AI will not kill sports journalism, but it will punish lazy journalism."¹⁴ Anyone using the machine to produce more of the same will compete against a marginal cost of zero. Anyone using it to reach first what nobody else has will not.

There is also an economic reason to look here first. Sport was the section that contributed the most search traffic and has lost the most, because a match result is exactly the kind of question an automatic summary answers without anyone visiting anything. The pressure on the model arrives earlier and is sharper.

And it is the section with the most direct relationship with its audience. Nobody follows a league by accident: there is a team, a routine and a community. That bond is precisely what remains when the intermediary disappears, and it explains why the sector's most interesting bets are coming out of sports desks rather than innovation labs.

What a sports desk already automates, and where it stops.

TASK	WHAT THE MACHINE DOES	WHERE THE LIMIT IS
Data and statistics	<i>Sorts, cross-references and detects performance patterns</i>	No relevant limit: this is the tool's natural ground and it frees hours of mechanical work.
Match report	<i>Gathers and orders the facts of the game</i>	A study of 28 LaLiga match reports, half human and half generated, concludes that the machine is effective with the data and lacks analytical and interpretive character. ¹²
Video and highlights	<i>Detects the play and cuts the clip</i>	Opta Pulse, launched in May 2026, generates broadcast-quality clips in under a minute, up to 80% faster than the traditional workflow. ¹¹ The judgement of what deserves a clip remains human.
Audio and briefings	<i>Turns coverage into new formats</i>	PRISA produced country-by-country audio briefings for the World Cup at Diario AS. The limit is one of brand: the synthetic voice works on data, not on signed opinion.
Alerts and monitoring	<i>Flags injuries, transfers and trends</i>	It accelerates detection, not verification. Every alert still needs a source before it is published.

The pattern repeats across all five rows: the machine keeps the gathering and the assembly; judgement, interpretation and the byline stay on the other side.

The 2026 World Cup was the first stress test.

The first World Cup of the generative AI era left an inventory no sports desk should file away. Fake statements attributed to FIFA circulated, along with a manipulated video of its president proposing that matches be suspended; in Mexico, a deepfake spread of the head of government announcing the tournament's cancellation.¹³

Quotes that never existed were manufactured, with Lamine Yamal and Erling Haaland among the worst affected. A generated video of Messi dancing with an English defender racked up hundreds of thousands of views before fact-checkers identified it. And synthetic *fancams* proliferated, driving traffic to pornographic platforms, classified by fact-checkers as gendered disinformation.

What matters to a publisher is not the catalogue of hoaxes but what it reveals about its own position. When the volume of synthetic material exceeds the capacity to deny it, the denial stops being the main tool. What remains is being able to demonstrate, on your own material, what is authentic and where it came from.

That is provenance debt again, now with a full stadium. A sports publisher that can certify the origin of its material competes on a field where forgery is free; one that cannot shares the field with it.



When denial costs more than forgery, the defence stops being the denial. It is *provenance*.

The practical reading for an editorial board is short: in the next global event, the cost of not being able to certify your own material will exceed the cost of verifying everyone else's.

The first sports AI product that charges for itself.

La Gazzetta dello Sport launched in March 2026, together with FairPlay Sports Media, a standalone subscription sports-analysis platform that identifies value opportunities in betting markets. It calculates expected value by comparing the probability implied in the odds against the model's estimate, and adds historical analysis of team and player behaviour.²⁸ It is not a feature inside the newspaper or an assistant for subscribers: it is a product with its own subscription, built on an extension of the odds comparator the title already had.

It belongs here for what it demonstrates about the business layer. Sport produces the only kind of data a machine handles well without supervision — structured, real-time and verifiable after the fact — and betting is the only market where that analysis has an established price. The combination turns the section that has lost the most search traffic into the first one that can sustain a new revenue line.

With one warning the case itself makes unavoidable: the model does not export as is. In Spain, Royal Decree 958/2020 obliges gambling tipsters to publish the entirety of their results and severely restricts promotion and affiliation, which changes both the product design and the revenue model.²⁹ A product like this is not copied from one market to another: it is redesigned.

08

Before the Scenarios, What Is Already *Signed*.

It is worth separating what might happen from what has already been announced with a date attached. Five moves in the sector are declared and under way, and they set the starting point for any European publisher.

The *Financial Times* has announced the second phase of its dynamic payroll, aimed at retention rather than conversion: the same decision machinery pointed at the part of the business where published evidence is scarcest and most modest.¹⁹

RCS MediaGroup has held a global technology partnership with OpenAI since July 2024, explicitly bounded to the reader relationship and not to content. Its first implementation was a *Corriere della Sera* assistant on the business vertical, with more than 30,000 articles in the corpus and free for subscribers.⁴³ This is the pattern consolidating across Europe: a group-level framework agreement, an assistant over the publisher's own archive, and a launch through a single thematic vertical.

The same group has publicly presented its automation agenda outside the editorial perimeter: the transition from classic robotic process automation to autonomous agents operating continuously, with a declared 75% reduction in operating times and a 50% improvement in return.⁴⁴ It is one of the few return figures a European publishing group has put in writing.

The *Corriere* newsroom answered the OpenAI agreement through the union. The assembly voted unanimously to mandate its committee to negotiate a professional protocol setting out what may and may not be done with artificial intelligence, invoking the clause of the national collective agreement that requires consultation on technological innovation. The stated grievance is that the agreement was signed without consulting the journalists.⁴⁵ Any European editorial board should assume that conversation is coming, and that it is better to arrive before it does.

08 · WHAT COMES NEXT

Four scenarios, each with its condition.

A document that spends a whole chapter denouncing artificial intelligence theatre cannot close with a list of predictions. What follows are scenarios, and each carries its condition printed on it. The starting point is a sector that has lost confidence in itself before losing it in the technology: only 38% of executives describe themselves as confident about the future of journalism, against 60% four years ago, while search traffic falls 33% worldwide.¹⁷ None of the four began in 2030: all four have already begun.

01

The reader stops reading: the agent reads instead.

Assistants such as ChatGPT Pulse already compose personalised briefings from the user's calendar and interests. More than 75% of media executives expect a large or very large impact from these applications.¹⁵ **Condition:** the publisher's primary audience becomes a bot, unless it first publishes machine-readable terms on how its work must be cited and represented.

02

Provenance stops being a principle and becomes a format.

C2PA content credentials already travel in consumer devices such as the Pixel 10 and in professional cameras, with more than 6,000 members in the authenticity initiative.¹⁶ **Condition:** it becomes general if platforms display the credential by default. The publisher that arrives with its archive already signed will have nothing to rebuild.

03

Content becomes liquid.

The piece stops being a fixed object and adapts in real time to the reader's context, location and moment.¹⁷ **Condition:** it works as long as personalisation does not break traceability. If every reader receives a different version and none is recorded, the publisher loses the ability to answer for what it published.

04

The licensing market splits in two.

The large groups sign deals and the rest are left out. Against that, publisher coalitions have appeared in the United Kingdom, Germany and Indonesia, the last with more than 500 outlets claiming collective compensation.¹⁵ **Condition:** the gap is avoided if those coalitions secure common terms before the habit of use without permission is taken as settled.

09

What to Measure When *the Question* Arrives.

None of these five metrics measures technology. All of them measure process, which is where it is decided whether an artificial intelligence programme survives the next budget.

METRIC	WHAT IT ANSWERS	HOW TO READ IT
Dependent process	<i>What would stop tomorrow?</i>	A valid answer is a noun: verification, archive, headline writing, advertising prospecting. If it begins with “we are working on”, the answer is none.
Traceability per piece	<i>Can we demonstrate the origin?</i>	Share of published pieces with a complete record of model, review and changes. Below 100% there is open provenance debt.
Time to verify	<i>How long do we take to confirm doubtful material?</i>	Measured in minutes on real cases, not in editorial policy. It is the metric that connects with the 74% of the audience who are worried.
Direct relationship	<i>How much audience arrives without an intermediary?</i>	Weight of direct visits, subscription and newsletter in the total. It is the indicator that degrades while search traffic evaporates.
Attributable return	<i>Does it show up in the accounts?</i>	A new revenue line or avoided cost attributable to the system, not licence spend. Without this figure, everything above is technology without a business.

What Remains When *the Intermediary* Falls.

Artificial intelligence has not come to write the news. It has come to occupy the place from which the news was distributed: the search engine, somebody else's front page, the feed. That place was another company's business and now it is even more so. Arguing about whether the machine writes better or worse than a journalist is arguing about the part that was never at stake.

What remains when the intermediary falls is the direct relationship and the ability to answer for what is published. Both are built with technical work, not with statements: an open, searchable archive, a provenance record installed in the workflow, autonomy that advances at the pace of the audit, and one concrete process that would stop if the system were switched off.

The calendar helps set the urgency. Article 50 already applies, trust remains at its lows and search traffic is not coming back. In the coming months the question will ask itself in any audit committee, and it will arrive phrased the way questions that have been put off too long always arrive: out loud and in front of everyone.

It is worth ending on the optimistic part, which is also true. The public grants journalistic brands more than twice the credibility it grants any automatic intermediary, and three out of four people fear being unable to tell the true from the false. That describes a market looking for what a publisher knows how to do. The question is not whether there is demand, but whether the publisher can prove it deserves it.

Three questions for the board

- 01 **Which process would stop tomorrow if we switched off our AI?** If the answer is not a noun, the programme has not yet touched the publication.
- 02 **Can we demonstrate the origin of a piece published three months ago?** Without that, the editorial exception in Article 50 offers no protection.
- 03 **Is our autonomy growing faster than our ability to audit it?** If it is running ahead, the publisher is not being fast: it is losing governance through success.

“

What is scarce is no longer the information.
It is the *provenance* somebody signs.

REFERENCES

Sources

Data, quotations and regulation referenced throughout the document.

Accessed September 2026. How to cite this document: Thinkia (2026). The

State of AI in Media. Thinkia, Age of Intelligence series 2026. thinkia.com

- 01 Omnibound, *Zero-Click Search Statistics 2026*, compiling Bain-Dynata (2024), SparkToro/Datos (2024) and Semrush (2025).
- 02 Comunicación Marketing, *Publishers raise paid search investment by 274%*, 14 September 2026.
- 03 Reuters Institute, *Digital News Report 2026*, 48 markets.
- 04 University of Navarra, *Digital News Report Spain 2026*, YouGov survey, more than 2,000 adults.
- 05 Asociación de la Prensa de Madrid, on the *Digital News Report Spain 2026*, 16 June 2026.
- 06 Laboratorio de Periodismo Luca de Tena, *El País launches Gula*, 12 May 2026.
- 07 OpenAI, *How news organizations are using AI*, 2026.
- 08 Laboratorio de Periodismo Luca de Tena, *Twenty examples of AI use in newsrooms*, 2026.
- 09 Sierra, M., *The Spanish press blocks AI access*, The Objective, 30 June 2026. Vipnet360 audit.
- 10 Economist & Jurist, *AI Act: Article 50 activates mandatory transparency*, 2026. Regulation (EU) 2024/1689.
- 11 Stats Perform, *Opta Pulse*, Opta Forum London, 12 May 2026.
- 12 Murcia Verdú, F. J., Ramos Antón, R. and Calvo Rubio, L. M., *AI sports reports versus journalists*, Revista Latina. University of Castilla-La Mancha.
- 13 Ecuador Chequea (Lupa) and Maldita.es, *fact-checks on AI disinformation during the 2026 World Cup*.
- 14 Columna Deportiva, *AI enters sports journalism*, 31 May 2026.
- 15 Columbia Journalism Review, Tow Center, *AI agents are coming for news*, 2026.
- 16 Content Authenticity Initiative, *The State of Content Authenticity in 2026*. C2PA standard.
- 17 Newman, N., *Trends and Predictions 2026*, Reuters Institute, via Nieman Journalism Lab.
- 18 El Español · Disruptores, *More than 20 publishing groups address the key year for AI in journalism*, 9 September 2026.
- 19 Digiday, *The Financial Times' AI paywall drove conversions up 290%*, 29 October 2025; and FT Strategies, *How dynamic paywalls are transforming reader revenue*.
- 20 Cody-Kenny, C. and Hjertholm Fiskerud, E. (Schibsted), *Schibsted grows digital subscriptions with predictive analytics*, INMA.
- 21 Bulthé, J. (Mediahuis), *WAN-IFRA Digital Media Europe*, Vienna, June 2023.
- 22 Piechota, G., *Turn your churn into a growth driver*, INMA, 5 June 2025. Aggregated Piano data, January 2019 – March 2025.
- 23 moin.ai (Der Spiegel) and the equivalent *Frankfurter Allgemeine Zeitung* case. **Supplier figures**, no independent audit.
- 24 Digiday, *Swiss news publisher NZZ built flexible paywall using machine learning*, June 2018, and subsequent INMA presentations.
- 25 FT Strategies, *Implementing a first-party data strategy*. Publisher figures.
- 26 BearingPoint, *L'Équipe optimizes newspaper sales with AI*, 12 February 2021; coverage in Consultancy.eu.
- 27 Better Retailing, *BearingPoint's DemandSens / Nitro*, 13 December 2021. **Supplier claim**.
- 28 Mediakey and Imgpress, *Gazzetta AI Predictor*, 26 March 2026. Platform with FairPlay Sports Media.
- 29 Royal Decree 958/2020, on commercial communications for gambling activities, in force since 5 November 2020.
- 30 Google, *Introducing Ask Ad Manager*, official blog, June 2026. In beta; no impact figures published.
- 31 New Digital Age, *News UK and Illuma partner to set a new gold standard for contextual targeting*, 8 August 2024. **Early validation**; methodology not published.
- 32 News Corp and Illuma, *Smart Deal*, 9 April 2026. **Figure reported by publisher and supplier**.
- 33 ExchangeWire, *Study: contextual targeting more cost efficient than behavioural*, 10 September 2020. One million impressions.
- 34 PubMatic, *Connect platform for curation delivers for publishers*, 20 March 2025. **Supplier claim**, no independent verification.
- 35 InfoSum, case study *Renault and Axel Springer improve advertising performance with first-party data*. **Supplier case** with a declared control design.
- 36 Native Advertising Institute and FIPP, *State of Native Advertising & Branded Content 2026*, July 2026. Survey of more than 75 branded studios.
- 37 Dotdash Meredith and The Campbell's Company, *Dinner Inspirations*, two-month campaign in 2025.
- 38 Bloomberg Media Studios and BCG, *CEO Radar*, Digiday AI Awards, 28 October 2025. Outcome metrics not published.
- 39 OpenAI, *Klarna case*, February–March 2024; and Forbes, *Klarna reverses on AI*, 18 May 2025.
- 40 Ascarza, E., Iyengar, R. and Schleicher, M., *The perils of proactive churn prevention using plan recommendations*, Journal of Marketing Research, 2016. n = 64,147.
- 41 Sahni, N., Miller, K. and Strulov-Shlain, A., *Stanford Graduate School of Business*, 2023. Two-year experiment on more than two million potential subscribers.
- 42 METR, *Measuring the impact of early-2025 AI on developer productivity*, July 2025; and *We are changing our experiment design*, 24 February 2026, with the redesign following the detection of selection bias.
- 43 Engage.it, *RCS MediaGroup in partnership with OpenAI*, 23 July 2024.
- 44 iKN Italy, *Dall'RPA all'Agentic Process AI*, case presented by RCS MediaGroup, 12 March 2025. **Figures from the case itself**.
- 45 *Corriere della Sera* newsroom assembly, February 2026; working group open since August 2024.
- 46 Laboratorio de Periodismo Luca de Tena and Dircomfidencial on Caliope, by Unidad Editorial, February 2026. Grant from Spain's Ministry for Digital Transformation.

thinkia

thinkia.com

© 2026