

AI transparency in **European news media**

Transparency and editorial responsibility under Article
50(4) of the EU AI Act

A study of how European news organisations interpret and implement the EU's new transparency rules for AI generated public interest text, what they tell their audiences, and what editorial safeguards they describe.

On 2 August 2026 a new European transparency rule began to apply to AI generated text published to inform the public. It does not prohibit newsrooms from using artificial intelligence, and it does not require a label on everything AI touches. It ties the duty to disclose to how editorial oversight and responsibility are organised inside a news organisation.

That makes it a question about how journalism is run rather than a question about technology. This study examines how European news organisations interpret the new rule, what they tell their audiences about it, what editorial safeguards they say they apply, and how those practices develop across the first year of application.

THE RULE

What Article 50(4) requires

Article 50 of the EU AI Act contains two transparency provisions that are frequently confused. They fall on different parties.

ARTICLE 50(2)

Applies to providers

Companies that build generative AI systems must mark outputs in a machine readable format so they can be detected as artificially generated or manipulated. This is the technical watermarking layer, and news publishers are not primarily subject to it.

ARTICLE 50(4)

Applies to deployers

Organisations using an AI system under their own authority in a professional capacity. Where they use it to generate or manipulate text published to inform the public on matters of public interest, the artificial origin must be disclosed, subject to an exception. This is the provision that reaches newsrooms.

The first limb of Article 50(4) concerns deepfakes, which this study does not examine. Its second limb, on public interest text, is the subject of the research.

Several conditions must be satisfied before the obligation arises. An AI system must have generated or manipulated the text. The organisation must have used that system under its own authority. The text must be published to inform the public. And the subject must be a matter of public interest, a category the European Commission reads broadly to include politics and democratic processes, public services, justice, fundamental rights, public security, health, environmental protection, consumer safety, and economic, financial, scientific and cultural developments relevant to public debate.

Where those conditions are met, disclosure is required unless **both** of the following apply: the text has undergone human review or editorial control, and a natural or legal person holds editorial responsibility for its publication.

The Commission distinguishes these concepts. **Human review** means a genuine examination of the substance of the material by someone with relevant knowledge and professional judgement. **Editorial control** means actual authority over the substance, held by a person able to approve, alter or reject it on

substantive grounds. Correcting spelling or grammar is not sufficient. **Editorial responsibility** concerns ultimate legal responsibility for publication.

The result is a transparency obligation that depends on the strength of editorial arrangements. That is why this rule is of particular interest for journalism, and why it falls within MJRC's field of research.

RESEARCH QUESTION

What the study asks

PRIMARY QUESTION

How are European news organisations interpreting and implementing Article 50(4)'s transparency rules for AI generated public interest text, and how visible are those arrangements to the public?

- 1 What do news organisations publicly disclose about their use of generative AI, human review, editorial control and editorial responsibility?
- 2 How do newsrooms themselves describe their implementation of the rule, including when they label AI generated text and when they consider the review and editorial control exception to apply?
- 3 How do these practices vary across outlet ownership, editorial governance model and national media environment?
- 4 How do practices change during the first year in which the rule applies?

The comparison between the first and second questions, between what an audience can see and what a newsroom describes, is the study's central contribution.

How the study works

The research combines three sources of evidence. A website audit alone can establish what an organisation publishes about itself. It cannot establish what happens before an article is published.

Legal and regulatory mapping

For each country in scope, the study records the environment in which the rule operates: the competent market surveillance authority and its designation status, whether national guidance specific to Article 50 has been issued, the involvement of the national media regulator, the mechanism through which a member of the public could complain, national professional or press council guidance on AI in journalism, and known adherence to the EU Code of Practice on transparency of AI generated content.

Public facing outlet audit

For each outlet, the study records what is publicly observable: whether a public AI policy exists and what it permits, what it states about human review and editorial control, whether editorial responsibility is described and whether a role or body is identified, disclosures observed in sampled material together with their wording and placement, and reader facing AI features such as summaries, chatbots and translation tools. Every observation is recorded with its own evidence, an archived snapshot, and an explicit uncertainty value where the position cannot be determined.

Newsroom survey and interviews

A standardised questionnaire to sampled outlets, followed by semi structured interviews across country and ownership categories. These address matters that cannot be observed from outside: where generative AI is used in editorial production and by whom, what happens before AI generated text is published, whether the reviewer can reject or substantially rewrite, who carries editorial and legal responsibility, when the outlet labels AI involvement, when it considers the exception applicable, and whether anything changed as a result of the rule. Material gathered in this way is reported as self described rather than independently verified.

Participation in the survey and interview component is voluntary. Participants are informed of the purpose and use of the research before taking part, may choose the basis on which they are identified, including full anonymity, and may withdraw. The study is not a compliance investigation, and participants are told so.

LIMITS

What the study cannot establish

The research uses no AI detection tools, and draws no conclusions from the style, cadence or formatting of published text. Detection of AI generated text is unreliable, and findings resting on it would not be defensible.

Several consequences follow, and they govern how every finding is reported.

- The law does not require a news organisation to publish an AI policy. The absence of one is a fact about transparency, not a violation.
- The law does not appear to require an outlet to identify its responsible editor publicly or to provide contact details. The absence of such information is likewise not a violation.
- An article carrying no disclosure may have been written by a person, may fall outside the scope of the provision, or may fall within the exception.
- A visible label is not an admission of weak editorial oversight. Some organisations label generously as a matter of policy.
- The study records that no AI disclosure was observed in the material sampled. It does not record that AI generated material was unlabelled, except where an outlet has itself established that the material was AI generated.

Whether any organisation complies with Article 50(4) depends on facts about internal editorial process that cannot be established from outside. The study makes no such allegation, and makes no claim about whether any particular article was generated by AI.

What the study can establish is how visible, specific and verifiable these arrangements are in public, what newsrooms describe of their internal practice, where those two accounts diverge, and how both vary across ownership types, media systems, and countries that have or have not built the capacity to supervise the rule.

SCOPE

Countries and outlets

The research begins in Hungary, Romania and Germany, with a small group of outlets in each, every one examined independently by two researchers. This first stage establishes which questions researchers can reliably answer, how long each outlet requires, and how far the design can be extended.

The proposed wider frame covers ten EU countries: Hungary, Slovakia, Bulgaria, Romania, Greece, Italy, Poland, Spain, the Netherlands and Germany, selected to span a range of media system conditions and enforcement environments. Within each country the sample follows a fixed composition of public service broadcasters, leading private national outlets, politically owned or state aligned outlets, digital natives, regional and local outlets, a news agency and an aggregator, so that markets remain comparable. Ownership and media system classifications draw on MJRC's Media Influence Matrix and State Media Monitor.

The wider frame is a proposal until the first stage is complete. It is not a representative sample of European media, and results are not generalised beyond the outlets examined.

OUTPUTS

Reporting

NOVEMBER 2026

A first report covering the regulatory picture across the countries in scope, early findings on what news organisations publish about their own arrangements, and an account of what the research instrument is able to observe. The instrument is published alongside it, together with the earlier drafts it replaces, so that the research can be examined, repeated or challenged.

THROUGH 2027

The wider audit, the newsroom survey and interviews, and repeat rounds on the same outlets. A comparative report and an open dataset, in which every observation carries its own archived evidence. A methods paper on measuring public transparency arrangements where the underlying legal concepts, substantive review and editorial responsibility, sit inside organisations and cannot be observed externally. Country notes prepared with the researchers who carried out the fieldwork, in national languages where they prefer.

The study is conducted by the Media and Journalism Research Center, an independent, non profit organisation researching media ownership, media governance and the conditions for independent journalism.

Enquiries, including from news organisations wishing to take part in the survey and interview component:
mjrc@journalismresearch.org