فصلنامه یافته های نوین درعلوم انسانی، علوم اجتماعی و روانشناسی شماره2- دوره 3- صفحات54-42
Journal of New Findings in Social Sciences, Educational Sciences and Psychology
ISSN 3060-611X

Newsroom Generative AI Guidelines Three Years On: A Content Analysis and Implications for Future Journalist Skills and Journalism Education

Pendar Vakilzadeh[1]*
1.PhD Candidate, Atatürk University, Erzurum, Türkiye

Abstract.
Since the release of ChatGPT in 2022, news organizations have published public guidelines to manage the risks of generative AI, but it remains unclear how these documents have changed over time and which skills they implicitly require of journalists. This study examines newsroom AI guidelines across four categories (transparency to audiences, human oversight, accountability for errors, and permitted and prohibited uses) and extracts the skills they imply for journalism education. Using qualitative content analysis, a purposive sample of eight news organizations from the United States, the United Kingdom, Germany and the Philippines was coded with a coding sheet and eight skill tags. Preliminary findings show that transparency, human oversight and accountability appear explicitly in all eight documents. Based on the versions read, one outlet fell into a restrictive level, six into a level of specified uses, and one into a level of generally permitted use under senior approval. The Associated Press (three versions) and the BBC (two) moved from prohibition toward permitted tasks under human review, and AP added a disclosure requirement (2026). Disclosure, verification of machine output, instrumental tool use and human judgment appeared in all eight documents. The guidelines can be read as a map of expected skills; it is proposed that journalism curricula stress verification, disclosure and accountability alongside tool training. The findings reflect stated policy rather than newsroom practice, and the small purposive sample is not generalizable.
Keywords: Generative AI, Newsroom Guidelines, Journalism, Journalism Education, Future Skills

1. Introduction

1.1 Problem statement and significance

The arrival of generative AI in newsrooms after the release of ChatGPT in 2022 has raised new questions about audience trust, professional responsibility and the boundary between human and machine work. News organizations answered these questions before legislators and before educational institutions did, publishing public guidelines for the use of the technology. Beyond regulating staff behavior, these guidelines implicitly sketch the skills expected of a journalist: one who can verify machine output, disclose its use and accept responsibility for content. Yet the potential of these documents for designing journalism curricula has received little attention. The issue matters because universities and news organizations, especially in contexts where institutional frameworks for AI are still forming, can use these documents as a map of the skills the industry expects.

1.2 Literature and background

Becker, Simon and Crum (Becker et al., 2025) compared the guidelines of 52 news organizations from countries including Belgium, Brazil, Canada, Finland, Germany, India, the Netherlands, Norway, Sweden, Switzerland, the United Kingdom and the United States. Drawing on neo-institutional theory, they argue that the policies show signs of homogeneity and that publishers have converged on transparency and human oversight, and they point to gaps concerning technological dependency, sustainability and inequality. Their sample comes mainly from Western Europe and North America and captures the situation in 2023.

At the supra-organizational level, the Council of Europe adopted guidelines on the responsible implementation of AI systems in journalism, covering the decision to use such systems, the acquisition of tools, their incorporation into professional practice and their effect on audiences (Council of Europe, 2023). The Poynter Institute has also published a template for drafting newsroom policies (Poynter Institute, 2024). In Iran, Nik Maleki and Afkhami (1403 SH) surveyed the knowledge and attitudes of editors at the IRIB News Agency toward the use of AI in the newsroom, and Masoudi (1404 SH) used a qualitative meta-analysis to examine the effect of automated news production on the quality and professional standards of journalism. What has been studied less is the analysis of the guideline documents themselves from the standpoint of skills and their link to journalism education.

1.3 Aims and research questions

The study is exploratory and has no testable hypothesis. Its aim is to examine newsroom AI guidelines and to extract the skills they imply for journalism education. The research questions are: (1) What do the published guidelines say about transparency, human oversight, accountability for errors and the limits of use, and how have they changed across

successive versions? (2) Which skills do these documents implicitly require of journalists? (3) What are the implications of this pattern for journalism education?

2. Research Method

2.1 Design and sample

The study used qualitative content analysis of official published documents; no interviews or questionnaires were involved. The unit of analysis was each guideline. Sampling was purposive; the selection criteria were the availability of a public guideline and the simultaneous coverage of public and commercial as well as Western and non-Western outlets. Eight documents were examined, and their characteristics appear in Table 1. The current BBC guidance, AP's July 2026 announcement and the Guardian's editorial code (generative AI section) were read in primary text. Earlier versions (AP's 2023 and interim standards, read in part; the BBC's 2024 guidance), the New York Times' 2025 guidance and some details of the Guardian's update come from secondary reports and are marked as such in Table 1. For Reuters the version read is the May 2023 memo, which is older than the latest published version (see Findings).

Table 1: Characteristics of the guidelines examined

Outlet	Country	Type	Version / date	Source read
AP	United States	News agency	Aug 2023; interim update; Jul 2026	Primary (2023 and interim in part; 2026 announcement)
Reuters	United Kingdom	Commercial news agency	May 2023	Primary (full memo)
BBC	United Kingdom	Public service media	Current guidance; 2024	Primary (current), secondary (2024)
The Guardian	United Kingdom	Newspaper (trust-owned)	Mar 2026	Primary (code) and secondary
The New York Times	United States	Commercial newspaper	2024 and 2025	Primary (2024), secondary (2025)
Tagesspiegel	Germany	Commercial newspaper	Jul 2026	Primary (full text)
WEB.DE / GMX	Germany	Commercial news portal	Aug 2026	Primary (full text)
ABS-CBN	Philippines	Commercial news network	Undated	Primary (full text)

Source: Author's compilation.

2.2 Instrument and coding

Each document was coded on transparency to audiences, human oversight and accountability for errors as "yes", "ambiguous" or "no". "Yes" means the guideline states the point explicitly; "ambiguous" means it appears only for some uses or only implicitly. Permitted and prohibited uses were rated on a three-level scale: level 1, broad prohibition on producing or editing news content; level 2, use permitted in specified, listed tasks; and level 3, general permission conditional on oversight and senior approval. In addition, eight skill tags were defined, and each tag scored 1 if it appeared explicitly in the text read: verification of machine output, detection of fabricated content, awareness of algorithmic bias, protection of sources and confidential data, disclosure of use, instrumental use, continuing training, and human judgment and contextualization. The complete coding sheet is given in the Appendix (Tables A1 to A3).

2.3 Data analysis and reliability

The data were analyzed by counting code frequencies and comparing versions qualitatively. A single coder carried out the coding, and the three-level scale is an interpretive judgment, so inter-coder reliability was not assessed. A zero on a skill tag means only that the skill was not seen in the text read, not necessarily that it is absent from the full document.

3. Findings

3.1 Convergence on human oversight, transparency and accountability

Table 2 shows the codes for each outlet. Human oversight, transparency to audiences and accountability for errors were explicit in all eight documents. In the BBC guidance, accountability takes the form of a senior editorial figure who is responsible and accountable for overseeing each deployment (BBC, n.d.). A representative example is the NYT principle, which holds the newsroom responsible for its reporting regardless of how the report is produced (InPublishing, 2024). Reuters likewise holds editors responsible for published content, whether a human or a machine produced it (Galloni & Freedman, 2023).

Table 2: Coding of the guidelines by category and level of freedom of use

Outlet	Transparency to audiences	Human oversight	Accountability for errors	Freedom of use (1 to 3)
AP	Yes	Yes	Yes	2
Reuters	Yes	Yes	Yes	1
BBC	Yes	Yes	Yes	2
The Guardian	Yes	Yes	Yes	3
The New York Times	Yes	Yes	Yes	2
Tagesspiegel	Yes	Yes	Yes	2

Outlet	Transparency to audiences	Human oversight	Accountability for errors	Freedom of use (1 to 3)
WEB.DE / GMX	Yes	Yes	Yes	2
ABS-CBN	Yes	Yes	Yes	2
Total “Yes”	8	8	8	-

Source: Author's compilation.

3.2 Differences in freedom of use and change over time

On the basis of Table 2, one outlet fell into level 1, six into level 2 and one into level 3 (mean 2.0). This rating refers to the versions read. In a May 2023 memo, Reuters stated that until the tools meet its accuracy standards it would not publish AI-generated stories, video or photographs and would not edit text stories; the same memo described itself as “preliminary thinking” and promised more formal guidelines (Galloni & Freedman, 2023). The newer Reuters version was not read directly, and the rating of this outlet may change with it. The BBC's current guidance says generative AI should not be used to directly create news, current-affairs or factual journalism content unless it is the subject of the content and its use is illustrative; it allows AI to support production and research, for example to analyse data or documents or for ideation, under active human oversight, while warning that AI should not be relied on as a research tool (BBC, n.d.). By contrast, the Guardian's March 2026 code imposes no blanket ban: generative AI may be used carefully and responsibly to enhance journalists' expertise, not replace it; active human oversight and control is essential; and any significant use must be explicitly approved by the relevant senior duty editor and clearly signalled to the audience (Guardian News & Media, 2026). According to a media report, the update also introduces mandatory staff training (Journalism Pakistan, 2026).

Regarding change over time, AP had three official versions. The August 2023 version banned the production of publishable content with generative tools, treated their output as unvetted source material, and did not allow altering elements of photos, video or audio (Barrett, 2023). A later update permitted experimentation in three specific uses, namely translation, news summaries and headline suggestions, with editing and vetting by a journalist (Associated Press, n.d.). The July 2026 update lists approved uses as early-stage research and document summarization, transcription and translation, headline, story-summary and shot-list suggestions, and grammar, spelling and search-optimization assistance, all reviewed and edited by an AP journalist; it also sets disclosure standards where generative AI plays a material role, standards for identifying and contextualizing AI-generated or manipulated content, and a continued prohibition on creating, altering or enhancing news photography (Associated Press, 2026). The BBC also appears to have moved: according to a secondary report, its 2024 guidance prohibited generative AI both for creating news content and for factual research (Broadcast, n.d.), whereas the current

guidance separates direct content creation from production support (BBC, n.d.); the Reuters Institute dates a version of this guidance to 24 June 2025 (Thurman et al., 2025). The NYT likewise moved from its 2024 principles (InPublishing, 2024) to 2025 internal guidance that, according to media reports, permits uses such as headlines, summaries and document analysis (Inquirer.net, 2025). This evidence (AP across three versions, the BBC across two and the NYT from 2024 to 2025) is consistent with a move from prohibition to controlled use. However, comparing different outlets in different years is not the same as observing change within one outlet, and the earlier BBC version and the NYT's 2025 guidance are known only from secondary reports.

3.3 Implicit skills

Table 3 shows the frequency of the skill tags. Disclosure of use, verification of machine output, instrumental use, and human judgment and contextualization appeared in all eight documents, and awareness of algorithmic bias in seven. ABS-CBN was the richest document: it asks journalists to return to the original source even when a machine provides a citation, to check the diversity of perspectives, and not to hand analysis and story angle to the machine (ABS-CBN, n.d.). Tagesspiegel and WEB.DE explicitly make staff training part of the outlet's commitment (Tagesspiegel, 2026; WEB.DE, 2026). The least frequent skill, detection of fabricated or manipulated content, appeared in four documents. This is notable given the importance of deepfakes and misinformation, although, because of the limits of the text read, it should be interpreted cautiously.

Table 3: Frequency of implicit skills in the guidelines (n = 8)

Implicit skill	Number of outlets	Share of sample
Disclosure and labeling of AI use	8	100%
Verification of machine output	8	100%
Instrumental use (prompting, document analysis, transcription)	8	100%
Human judgment and contextualization	8	100%
Awareness of algorithmic bias	7	88%
Continuing training and retraining	6	75%
Protection of sources and confidential data	5	63%
Detection of fabricated or manipulated content	4	50%

Source: Author's compilation.

4. Discussion and Conclusion

The findings are consistent with the conclusion of Becker et al. (2025) that guidelines converge on transparency and human oversight: human oversight, transparency and accountability appeared in all documents examined. This study adds two things: first, limited evidence, especially from AP and the BBC, of guidelines moving from prohibition

to controlled use in the period after 2023; and second, a reading of the documents as a map of skills. The study by Nik Maleki and Afkhami (1403 SH) measures editors' knowledge and attitudes, whereas the present study examines policy documents; the two perspectives are complementary, and directly comparing their results would require a concurrent study.

On the basis of the findings, it is proposed that journalism curricula go beyond tool training and develop three groups of skills: verification skills (checking machine output and tracing the original source), ethical-professional skills (disclosure, accountability and recognizing bias) and application skills (prompt writing, document analysis and transcription). Several outlets stress that judgment, contextualization and relations with audiences remain human work; this is a reminder that education should not reduce the journalist to a tool operator. Outlets that do not yet have public guidelines can draw on existing models, including the Poynter template (Poynter Institute, 2024), to draft a guideline that is regularly updated. Future research can measure the gap between stated policy and actual newsroom practice and include more outlets from the Global South, including Iranian and Persian-language outlets.

From a futures perspective, two exploratory scenarios to 2035 can be imagined. In the “senior approval” scenario, the path of the Guardian and the NYT continues: AI expands in supporting tasks, human approval remains a condition of publication, and universities that build verification, data literacy and accountability into their programs produce more employable graduates. In the “gap” scenario, outlets without guidelines or training use the tools without shared criteria, and the gap in trust between outlets with and without policies widens. These scenarios are based on drivers observed in the documents (convergence of guidelines, frequent revision and mandatory training) and are not predictions.

Limitations: the eight-outlet sample is purposive and mostly Western and cannot be generalized to “news organizations worldwide”; there was a single coder and some sources are secondary; the documents are changing and the paper shows a snapshot; and the study does not measure actual newsroom behavior.

References

ABS-CBN. (n.d.). Artificial intelligence: Policy and guidelines (Integrated News and Current Affairs). https://www.abs-cbn.com/ai-policy

Associated Press. (n.d.). Updates to generative AI standards. https://ds.svcs.associatedpress.com/updates-to-generative-ai-standards

Associated Press. (2026, July). AP updates newsroom standards for artificial intelligence. https://www.ap.org/the-definitive-source/announcements/ap-updates-newsroom-standards-for-artificial-intelligence/

Barrett, A. (2023, August 16). Standards around generative AI. AP Definitive Source. https://blog.ap.org/standards-around-generative-ai

BBC. (n.d.). Guidance: The use of artificial intelligence. BBC Editorial Guidelines. Retrieved September 20, 2026, from https://www.bbc.co.uk/editorialguidelines/guidance/use-of-artificial-intelligence

Becker, K. B., Simon, F. M., & Crum, C. (2025). Policies in parallel? A comparative study of journalistic AI policies in 52 global news organisations. Digital Journalism, 1-21. https://doi.org/10.1080/21670811.2024.2431519

Broadcast. (n.d.). The BBC has set out new rules for applying AI to shows. https://broadcastnow.co.uk/bbc/7-things-you-need-to-know-about-the-bbcs-ai-guidance/5200901.article

Council of Europe. (2023, December 29). New guidelines on responsible implementation of artificial intelligence in journalism. https://www.coe.int/web/portal/-/new-guidelines-on-responsible-implementation-of-artificial-intelligence-in-journalism

Galloni, A., & Freedman, A. (2023, May 14). What Reuters is telling its journalists about using artificial intelligence [Staff memo]. Talking Biz News. https://talkingbiznews.com/media-news/what-reuters-is-telling-its-journalists-about-using-artificial-intelligence/

Guardian News & Media. (2026, March). Editorial code of practice and guidance. https://uploads.guim.co.uk/2026/03/03/Editorial_Code_of_Practice_Guidelines_March2026.pdf

InPublishing. (2024, May 17). NYT issues guidance on its A.I. principles. https://inpublishing.co.uk/articles/nyt-issues-guidance-on-its-ai-principles-24123

Inquirer.net. (2025, February 18). New York Times adopts generative AI tools for journalists. https://technology.inquirer.net/140736/new-york-times-adopts-generative-ai-tools-for-journalists/amp

Journalism Pakistan. (2026, March 5). Guardian updates AI rules to keep humans in charge. https://www.journalismpakistan.com/guardian-updates-ai-rules-to-keep-humans-in-charge

Masoudi, H. (1404 SH). The impact of AI systems in automated news content production on the quality and professional standards of journalism: A qualitative meta-analysis [in Persian]. Journal of News Sciences, 14(2), 102-120.

Nik Maleki, M., & Afkhami, H. A. (1403 SH). Application of artificial intelligence in journalism: Knowledge and attitudes of editors of the IRIB News Agency [in Persian]. New Media Studies, 10(38), 39-74.

Poynter Institute. (2024, March 25). How to create a newsroom artificial intelligence ethics policy. https://www.poynter.org/ethics-trust/2024/how-to-create-newsroom-artificial-intelligence-ethics-policy/

Tagesspiegel. (2026, July 17). Redaktionelle KI-Richtlinien: So arbeiten wir mit künstlicher Intelligenz. https://www.tagesspiegel.de/redaktionelle-ki-richtlinien-so-arbeiten-wir-mit-kunstlicher-intelligenz-15849067.html

Thurman, N., Thäsler-Kordonouri, S., & Fletcher, R. (2025). AI adoption by UK journalists and their newsrooms: Surveying applications, approaches, and attitudes. Reuters Institute for the Study of Journalism. https://doi.org/10.60625/risj-ea11-q402

WEB.DE. (2026, August 26). Leitlinie zum Umgang mit Künstlicher Intelligenz in der Redaktion. https://web.de/magazine/so-arbeitet-die-redaktion/leitlinie-umgang-kuenstlicher-intelligenz-redaktion-38410396

Note: SH = Solar Hijri calendar year (1403 SH = 2024-2025; 1404 SH = 2025-2026).

Appendix. Coding Sheet and Skill Matrix

This appendix reproduces the coding sheet on which Tables 2 and 3 are based. Ratings refer to the versions read (see Table 1); items drawn from secondary reports are marked "(secondary)". Coding scale for freedom of use: 1 = broad prohibition on producing or editing news content; 2 = use permitted in specified, listed tasks; 3 = general permission conditional on oversight and senior approval.

Table A1: Coding of each guideline

Outlet	Transparency	Human oversight	Accountability	Level	Permitted and prohibited uses (summary)
AP	Yes	Yes	Yes	2	Aug 2023: producing publishable content and altering photo, video or audio prohibited. Interim update: experimentation allowed in translation, news summaries and headlines, with journalist editing. Jul 2026: approved uses are early-stage research and document summarization, transcription and translation, headline, story-summary and shot-list suggestions, and grammar, spelling and search-optimization assistance, all reviewed and edited by an AP journalist; disclosure standards where generative AI plays a material role; standards for identifying and contextualizing AI-generated or manipulated content; creating, altering or enhancing news photography remains prohibited.
Reuters	Yes	Yes	Yes	1	May 2023 memo: brainstorming and headline ideas allowed; no publication of AI-generated stories, video or photographs and no editing of text stories until tools meet accuracy standards; unpublished stories must not be uploaded to open AI tools; experiments confined to an internal platform.

Outlet	Transparency	Human oversight	Accountability	Level	Permitted and prohibited uses (summary)
BBC	Yes	Yes	Yes	2	Current guidance: generative AI should not be used to directly create news, current-affairs or factual journalism content unless it is the subject and its use is illustrative (graphics exemptions subject to piloting); AI may support production and research, for example data or document analysis and ideation, but must not be relied on as a research tool and needs active human oversight; use must be transparent to the audience; a senior editorial figure is responsible and accountable; news video and images must not be manipulated beyond a sympathetic crop and minor adjustments; synthetic voices not of a real person must be clearly disclosed, and cloning a real person's voice without consent needs strong editorial justification. The 2024 version (secondary) prohibited generative AI for news content and factual research.
The Guardian	Yes	Yes	Yes	3	No blanket ban. AI may be used carefully and responsibly to enhance, not replace, journalists' expertise; active human oversight and control is essential; any significant use must be explicitly approved by the relevant senior duty editor and clearly signalled to the audience. Tasks such as data analysis, archive search and transcription, and mandatory staff training, are reported in secondary sources.
The New York Times	Yes	Yes	Yes	2	2024 principles: any use must begin with journalist-vetted facts and be reviewed by editors; readers are told how work was created. 2025 guidance (secondary): headlines, summaries, suggested edits, brainstorming questions and document analysis allowed; drafting or significantly altering articles not allowed; AI images or video must be explicitly labelled.
Tagesspiegel	Yes	Yes	Yes	2	No AI-generated text published under a byline; AI revision allowed but does not replace human editing; uses include research, topic planning, data analysis and visualization; final newsroom approval required for original journalistic formats; AI involvement disclosed where substantial.

Outlet	Transparency	Human oversight	Accountability	Level	Permitted and prohibited uses (summary)
WEB.DE / GMX	Yes	Yes	Yes	2	Uses include research, transcription, headline development, fact-check support, drafts for news curation and personalized delivery; all output reviewed and approved by editors; no AI-generated images in news unless the image is itself the subject; disclosure where AI is materially involved in a text.
ABS-CBN	Yes	Yes	Yes	2	AI must not write articles, scripts or stories for publication; no recreation of a person's image or voice except to demonstrate AI, with consent and disclosure; no upload of internal or confidential data to public AI platforms without AI Task Force approval; content using AI needs editorial-gatekeeper approval. The document's statement of intent, which lists writing assistance, is in tension with its ban on AI-written articles.

Source: Author's compilation.

Table A2: Implicit skills identified in each guideline

Outlet	Implicit skills (summary)
AP	Verifying machine output, treated as unvetted source material; reverse image search and tracing the original source; detecting fabricated or manipulated content; not entering confidential information into tools; disclosure.
Reuters	Being sceptical and verifying; checking facts, sense and bias in machine output; protecting unpublished content; logging AI use in an internal channel.
BBC	Monitoring and assessing AI output rather than relying on it for research; awareness of hallucination, bias, plagiarism and mimicry; protecting confidential data entered into tools; attention to the quality of prompts; authenticating user-generated and synthetic media; preserving editorial skills and independence.
The Guardian	Verification with "absolute rigour"; awareness of issues around bias, ownership, plagiarism and intellectual property; human oversight; enhancing rather than replacing expertise; mandatory training (secondary).
The New York Times	Analyzing documents with tools; brainstorming questions; explaining to readers how bias and inaccuracy are mitigated; mandatory training (secondary).
Tagesspiegel	Checking results, sources and methods; sharpening attention to bias; protecting data and sources; learning tools and experimenting.
WEB.DE / GMX	Regular mandatory training on tools, responsibility and error risks; identifying bias; societal contextualization and audience relations remain human work.
ABS-CBN	Going to the original source even when AI cites one; scepticism; checking diversity of perspectives; not delegating analysis and story angle to AI; continuing training.

Source: Author's compilation.

Table A3: Skill matrix (1 = explicit in the text read; 0 = not seen)

Column key: V = verification of machine output; F = detection of fabricated content; B = awareness of algorithmic bias; S = protection of sources and confidential data; D = disclosure of use; T = instrumental use; L = continuing training; J = human judgment and contextualization.

Outlet	V	F	B	S	D	T	L	J
AP	1	1	0	1	1	1	1	1
Reuters	1	1	1	1	1	1	0	1
BBC	1	1	1	1	1	1	0	1
The Guardian	1	0	1	0	1	1	1	1
The New York Times	1	0	1	0	1	1	1	1
Tagesspiegel	1	0	1	1	1	1	1	1
WEB.DE / GMX	1	0	1	0	1	1	1	1
ABS-CBN	1	1	1	1	1	1	1	1
Total	8	4	7	5	8	8	6	8
Share of sample	100%	50%	88%	63%	100%	100%	75%	100%

Source: Author's compilation.

Note: a zero means only that the skill was not seen in the text read, not necessarily that it is absent from the full document. Ratings for the BBC's 2024 version, the New York Times (2025) and part of the Guardian's update rest on secondary reports, and AP's earlier versions were read only in part. Coding was carried out by a single coder.