

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/379668724>

Generative AI in Journalism: The Evolution of Newswork and Ethics in a Generative Information Ecosystem

Technical Report · April 2024

DOI: 10.13140/RG.2.2.31540.05765

CITATIONS

0

READS

1,159

7 authors, including:



Hannes Cools

University of Amsterdam

10 PUBLICATIONS 36 CITATIONS

SEE PROFILE



Charlotte Li

Northwestern University

2 PUBLICATIONS 15 CITATIONS

SEE PROFILE



Ernest Kung

The Associated Press

3 PUBLICATIONS 12 CITATIONS

SEE PROFILE



Aimee Rinehart

Associated Press

3 PUBLICATIONS 12 CITATIONS

SEE PROFILE

Generative AI in Journalism: The Evolution of Newswork and Ethics in a Generative Information Ecosystem

Authors:

Nicholas Diakopoulos, Hannes Cools, Charlotte Li,
Natali Helberger, Ernest Kung, Aimee Rinehart

Edited by:

Lisa Gibbs




Table of Contents

Introduction	03
The Future of Newswork	09
Ethical Considerations	22
Looking Ahead	34
Appendix	37



Introduction



A large section of the page consisting of 27 horizontal red lines, providing space for writing the introduction.

The introduction of ChatGPT by OpenAI in late 2022 captured the imagination of the public—and the news industry—with the potential of generative AI to upend how people create and consume media. Generative AI is a type of artificial intelligence technology that can create new content, such as text, images, audio, video, or other media, based on the data it has been trained on and according to written prompts provided by users. ChatGPT is the chat-based user interface that made the power and potential of generative AI salient to a wide audience, reaching 100 million users within two months of its launch¹.

Other big tech companies quickly flocked to compete with their own AI models: Bard and then Gemini from Google, Claude from Anthropic, Copilot from Microsoft, and open source offerings like LLaMA from Meta, not to mention new search products like Perplexity, browser experiences like Arc, or the creation of interfaces like Adobe's Firefly and Photoshop that integrate the technology and transform how end-users create or interface with information. Although versions of the tech have been around since 2018, by late 2022 it was suddenly working (sort of), spurring its integration into various products and presenting not only a host of opportunities for productivity and new experiences but also some serious concerns about accuracy, provenance and attribution of source information, and the increased potential for creating misinformation.

¹ ChatGPT is growing faster than TikTok. CBS News. Feb 1, 2023.
<https://www.cbsnews.com/news/chatgpt-chatbot-tiktok-ai-artificial-intelligence>

Throughout 2023 the news industry scrambled to figure out what all this new technology would mean for news gathering, production and distribution practices, products and user experiences, for their already precarious business models, and the value of their intellectual property. Understanding how audiences might interact with and consume information in the future is again being challenged, after the disruption that social media wrought. Initiatives like the Generative AI in the Newsroom (GAIN) project², AI, Media and Democracy Lab, the Open Society Foundation AI in Journalism Challenge³, The Reuters Journalism Institute roundtables on generative AI⁴, and the LSE's JournalismAI survey in mid-2023⁵, as well as others, including an early survey from WAN-IFRA⁶ and the AP's own convenings⁷ have all contributed to advancing the industry's understanding of the technology and what it might mean for journalism.

This report serves as a snapshot of how the industry has grappled with the initial promises and challenges of generative AI toward the end of 2023. In this report, we present a survey of 292 individuals in the news industry about how they use and want to use generative AI and what they see as the main ethical and practical issues around developing responsible usage. We collected survey responses for three weeks, from December 4 to December 22, 2023, with the AP circulating the survey among its email lists of news organization practitioners and through various social media and Slack-group postings.

Fully 81.4% of respondents indicated that they were knowledgeable about generative AI (See Figure 1), and 73.8% indicated that they or their organization had already used generative AI in some capacity (See Figure 2). The average number of years worked in the news industry by respondents was 18 years (See Appendix A). In other words, *our sample reflects how some of the more savvy and seasoned members of the profession are reacting to the technology.*

2 <https://generative-ai-newsroom.com>

3 David Caswell. Rising to the Challenge: Applying Generative AI in Newsrooms. Generative AI in the Newsroom (October, 2023).

<https://generative-ai-newsroom.com/rising-to-the-challenge-applying-generative-ai-in-newsrooms-283d5bb3de53>

4 Jessica Cecil. 2023 round tables on AI and the global news industry. Reuters Institute.

<https://reutersinstitute.politics.ox.ac.uk/news/2023-round-tables-ai-and-global-news-industry>

5 Charlie Beckett and Mira Yaseen. Generating Change. JournalismAI.

<https://www.journalismai.info/research/2023-generating-change>

6 Teemu Henriksson. New survey finds half of newsrooms use Generative AI tools; only 20% have guide lines in place. WAN-IFRA. <https://wan-ifra.org/2023/05/new-genai-survey>

7 Local News AI: Building tools and training newsrooms. <https://ai.ap.org>

Figure 1.

More than three quarters of respondents state that they are knowledgeable about generative AI.





Figure 2.

When asked “Have you or your organization used generative AI in some capacity?” almost three quarters of respondents answer affirmatively.



The sample skews heavily toward people working in North America (61.7%) and Europe (24.8%) with a smattering of respondents working in Asia (7.9%), Africa (2.8%), Oceania (1.7%) or South America (1.0%). There was an over-representation of men responding (58.3%) though this appears roughly consistent with expected base rates in the media industry⁸. And while respondents in Editor roles dominated (34.5%), we also captured responses from Executives (20%), Reporters (18.3%), Technologists (9.3%) and people in other roles such as Product or in roles wearing multiple hats (17.9%). For more details on the sample we collected, see Appendix A.

⁸ See: Women and leadership in the news media 2023: evidence from 12 markets. Reuters Institute. March 8, 2023. <https://reutersinstitute.politics.ox.ac.uk/women-and-leadership-news-media-2023-evidence-12-markets>

Based on participants' responses and our analysis, we find that generative AI is already changing work structure and organization, even as it triggers ethical concerns around use. Here are some of our key takeaways:


- **Applications in News Production.** The most predominant current use cases for generative AI include various forms of textual content production, information gathering and sensemaking, multimedia content production, and business uses. Respondents expressed interest in expanding the use of generative AI for information gathering and sensemaking, working with data, business uses, and metadata creation, with some interest in exploring new user experiences with chatbots and through personalization. Overall, attention is focused on improving and making existing workflows more efficient with considerably less attention to exploring and innovating new experiences.
- **Changing Work Structure and Organization.** There are a host of new roles emerging to grapple with the changes introduced by generative AI including for leadership, editorial, product, legal, and engineering positions. Almost half of respondents indicated that tasks or workflows have already changed because of generative AI. New work is created in devising effective prompts and in editing outputs. Perceived efficiency gains are variable and additional research is needed to evaluate any real performance gains across a range of common tasks. Overall, these findings underscore the need for training initiatives and for more fine-grained evaluations to measure actual shifts in productivity.
- **Work Redesign.** There is an unmet opportunity to design new interfaces to support journalistic work with generative AI, in particular to enable the human oversight needed for the efficient and confident checking and verification of outputs. Journalists will need well-designed editing interfaces in order to effectively use generative AI for various tasks. Respondents are also open to getting help from generative AI for tasks related to analyzing, getting, or processing data and information, which are perhaps not coincidentally also the kinds of work activities that respondents rated as boring, repetitive, or tedious.
- **Ethical Concerns and Responsibility.** Ethical considerations are paramount, with concerns about human oversight, accuracy, and bias most prominent. The industry is grappling with how to balance the benefits of generative AI with the need for ethical journalism practices, including the banning or limiting of use for particular use cases such as for the generation of entire pieces of published content. Overall, editors, managers, and executives (rather than technologists) were the roles that respondents thought should be more responsible for ensuring effective and ethical uses of generative AI.
- **Strategies for Responsible Use.** While many organizations are developing or following guidelines for the ethical use of generative AI, there is a call for clearer, more concrete guidelines, training, and enforcement to navigate the ethical landscape effectively. On top of guidelines, there is recognition that additional training is needed to support responsible use. Other strategies that might also improve responsible use of generative AI, like more robust procurement of tools that include AI and automation as well as internal testing and auditing, are rarely mentioned.

- **Ambivalence in Content Rights.** Respondents expressed a degree of uncertainty about whether tech companies should be allowed to train models on news organizations' content, with some emphasizing the negative commercial impacts and others advocating to advance the accuracy and reliability of models which could benefit society.

In the next section we examine what the future of newswork could look like in the era of generative AI. Then we turn to the ethical considerations and approaches needed if generative AI is going to be incorporated into responsible journalism practice. We finish the report with a conclusions section where we argue that the industry will require investments in policy, practices, research, design, and training to further advance and best capture the value of this technology while aligning it with journalism's norms and practices of responsibility to society.



The Future of Newswork



Automation often inspires anxiety about how new technologies like AI might threaten the status quo of work and undermine a person's livelihood. If generative AI systems can do basic news gathering and writing, could they *replace* reporters and editors? Or will these tools be more complementary and help to augment the work? How is all of this going to change the jobs people in the news industry are asked to do, particularly as user experiences and expectations also evolve?

Survey respondents are keen to explore a wide range of tasks to *augment* their workflow and increase their efficiency, but there's also wide variance with more research needed to establish any actual productivity gains. However, it's already clear that generative AI is *changing the structure and organization* of work and putting pressure on individuals to learn new skills to keep up, while also creating new roles and opportunities within organizations.

Current Usage

If respondents indicated that they or their organization had used generative AI in some capacity, they were then asked about what tasks they or their organization had used it for. Responses are coded into broader categories of tasks shown in Figure 2.


The dominant category of use is unsurprisingly related to *content production*. This category included responses about using generative AI tools in processes of producing public-facing or newsroom internal content, including creating, editing, and transforming media formats. Specifically, the responses were further divided into six different categories: *text* (69.6%, 126 of 181 responses), *multimedia* (20.4%), *translation* (8.8%), *transcription* (7.2%), *user experience* (2.8%), and *metadata* (0.6%).

In the *text* category, respondents state they have used generative AI for generating content such as news headlines, social media posts, newsletters, quizzes, text from data, taglines, and story drafts. As one respondent noted, “We use AI to help us create outlines, briefs, and first drafts.” They have also used generative AI for copy editing and summarizing articles, rewriting for a different medium (e.g. script production) or to reduce jargon or produce a press release, and fact checking. Respondents also mentioned using generative AI for generating *multimedia* content, such as illustrations (e.g. for social media posts), videos, audio (e.g. text to speech), or for editing images. AI-assisted *translation* and *transcription* also came up as part of the content production process, using tools such as Otter or Whisper. A handful of respondents mentioned using generative AI to support the *user experience* to create consumer-facing chatbots and for assistance with creating *metadata* such as the creation of alt-text for images or metadata for audio files.

Another somewhat common usage of generative AI is for *information gathering and sensemaking* (21.5%). This category encompasses ways AI is used for news discovery, research, ideation / brainstorming, and curation, or as one respondent put it, “automation of research steps, newsgathering, and notification systems.” We also identified some more technical tasks supported by generative AI: *coding* (5.0%) encompasses responses that use generative AI for software development tasks such as code review or “writing and refining HTML code,” and *working with data* (7.7%) spans tasks that involve manipulation and analysis of data or its extraction from documents such as “manipulating spreadsheets” or “small data cleanup tasks.” Lastly, *business* (16.6%) is a category used to capture responses that mention using generative AI for internal business operation purposes, like creating presentations, drafting emails (e.g. “to sales prospects”), outputting ads or marketing content, or outputting material for search engine optimization.

Figure 3.

The distribution of tasks mentioned in response to the question “What tasks have you or your organization used generative AI for on an experimental or regular basis?” (N=181)



Aspirational Usage

Next, to understand better how some journalists want to use generative AI, respondents were asked to “List at least three tasks that you would ideally like to use generative AI for in your work, if it were capable of producing quality results.” Results are shown in Figure 4 below.


Beyond the tasks already shown in Figure 3, respondents also listed *planning* (3.9%), *distribution analytics* (2.6%), *personalization* (1.7%), *layout* (1.3%), and *fake news detection* (0.9%) as ways they would like to use generative AI. In *planning*, responses mainly request the use of generative AI to improve daily workflow and news cycle plans. For *personalization*, responses point to content customized, suggested, or curated based on reader/user information (e.g. “personalization of newsletters and homepage”). *Fake news detection* reflects a need for identifying and debunking untruthful news content. *Layout* points to the need of respondents to have generative AI help with paper news layout during production. *Distribution analytics* differs from the need of analyzing data for production processes and captures the need of collecting and analyzing user engagement data (e.g. “answer questions about website analytics”). In comparison to current usage, we observe that some aspirational usage categories gain in interest such as *user experiences*, *working with data*, *information gathering and sensemaking*, *metadata*, and *business* use-cases. The largest absolute gain in interest was for *information gathering and sensemaking*, which includes a variety of reporting-relevant activities such as news discovery, research, ideation, and curation.

For a few of the categories in Figure 4, we also observe new tasks. In responses that were identified as *text content production*, respondents mention the need for having generative AI help with the production of event calendars, which can be a highly structured form of text production and lends itself to automation. For *information gathering and sensemaking*, we observe more needs for assistance with monitoring and scanning different media (social media, news media, and local government) and alerting when newsworthy information is identified (e.g. “news-scanning in the local market” and “identifying possible sources”). A few respondents mentioned the need for AI to help with news aggregation and curation. For *multimedia content production*, there are increased needs for AI assistance with video and audio modes of production (e.g. “short video news explainers”). Additionally, there were more frequent responses mentioning the creation of AI chatbots for *user experiences* (e.g. “Chatbot to act as an interface to all of our content”).

Respondents were also asked directly about the opportunities they perceived for the use of generative AI in journalism and these responses largely concurred with what they talked about in aspirational tasks. There was considerable interest in using generative AI to support data analysis and research, to reduce repetitive tasks and save time, allow for efficient editing, and also to enable creativity through brainstorming and to explore new possibilities in personalization. Responses among editors, technologists and executives were similar.

Figure 4.

The distribution of tasks mentioned in response to the question “List at least three tasks that you would ideally like to use generative AI for in your work, if it were capable of producing quality results.” (N=229)



What's Working and Not Working

Respondents talk about saving time and enhancing *efficiency*, but also augmenting and *supporting creativity* with story discovery, idea generation or brainstorming, and mentioning specific activities where AI might offer some gains such as *content production* with headlines and illustrations, for *research* in gathering background content, *data work* for scraping and extraction from documents, and *expanding audiences*.

When respondents found AI to be ineffective, they often refer to *quality* issues relating to *accuracy*, *trustworthiness*, and *content quality* such as the relevance of headlines generated or the blandness of the text that commercial large language models produce. They also point out that sometimes use of the models costs more time than they save or *creates more work* like editing, can output *biased* text, and that also there are issues with prompting and *controlling* the models effectively. Some found that it simply took too much time to edit and “craft prompts that are responsive to needs” to yield much efficiency gain.

New Work, New Roles

Respondents mention new roles being established as a result of generative AI. These include *leadership* roles like “innovation officer,” “AI Expert,” “Head of AI,” *product* positions like “AI Product Owner,” *editorial* jobs like “prompt designer/editor/specialist,” “fact checker,” “AI Video editor,” *legal* roles, and *engineering* positions including “Software Engineer,” “AI + Automation Engineer” and “quality assurance.” Roles were mentioned at different levels, from managers to individual contributors and even an internship.

Some of the new positions involve watching for new innovations and keeping up with the pace of change with a role described as: “a person who monitors turbulent developments in the field of AI.” New editorial roles include fact checking, prompting, and technical roles like engineering and user-interface and -experience design. Overall, there is a shift in human work toward management, product, some new positions, and a lot of technical work to incorporate and maintain systems that include AI and automation.

Changing Work

Almost half of respondents (49%, 92 of 181 respondents) indicated that tasks or workflows have already changed because of generative AI. Among the respondents who said tasks and workflows have changed, they indicate that generative AI models have shaped the *structure and organization of work*. In some cases, the structural and organizational changes to their work are already being reflected in Content Management Systems (CMSs), Slack channels, or via commonly used office software.

AI models take on roles as collaborators and are used as a sounding board “to bounce off ideas” or as an editor that catches “things that may have been missed.” AI can shape relationships between people as well, with one respondent noting: “instead of asking a colleague for help with a heading, I always ask ChatGPT first.” The idea of AI as a collaborator also means that people need to think about how to formulate a task for delegation which might involve determining “if there is a prompt that will increase my efficiency and productivity.”

Much like for a self-service checkout system in a supermarket, respondents indicate that new work is created for them when they use these systems, primarily in terms of having to edit or proofread the outputs of the AI to ensure it is acceptable. In at least one case, a respondent indicated that this self-service mentality also shaped the relationship with freelancers: “We [have] stop[ped] hiring freelancers for certain tasks, like basic translations or copywriting.” As generative AI is increasingly used outside of the newsroom, it can also create new editorial work to evaluate sources of information, as one respondent described: “We had to define protocols to detect AI-generated content. We had to put guardrails in place because we receive a lot of text from external authors... .”

Efficiency Hopes and Realizations

Tasks where efficiency gains from using generative AI seemed to be supported include image editing, monitoring sources, producing alt-text, SEO text generation, press releases, emails, social media posts, and reducing time to produce drafts of text. As one respondent put it: “All the genAI tools in Photoshop save our graphics team hours of time each day.” In terms of creativity support, another respondent wrote: “being able to quickly see numerous visual iterations of a concept makes it easy to explore options and ideas I’d otherwise not pursue due to time or resource restrictions.”

Despite the creation of more editing, prompting, or evaluative work in some cases, many respondents continue to apply an efficiency frame to how they talk about task and workflow changes: “We can scale some tasks (such as finding topic ideas) much faster.” Hopes for and actual evidence of efficiency gains often goes hand-in-hand with talking about how the time saved by using generative AI can be reinvested into other activities: “The graphics team is able to redirect resources into working on other aspects of their job with the time saved by using genAI in Photoshop.” Efficiency also shapes the pace and iterative nature of the work: “because it’s so fast to adjust a prompt and generate more iterations ... [AI] has provided more opportunities to dabble in different styles and methods.”

Delegating Work

Respondents were asked about how they would evaluate the work they got back if they delegated a chosen task to a colleague. We asked this to better understand the criteria for success for the chosen task and with the idea that knowing these criteria could enable ways to better evaluate task performance if delegated to an AI system.

There was wide variance across different tasks among the responses. A story discovery task might rely on criteria of “originality” and “relevance” whereas for a content generation task “clarity” and “concision” might be important. Respondents at times referred to both subjective personal and formal organizational (including legal) rubrics for help in evaluating tasks. A few respondents mentioned criteria related to efficiency, productivity (e.g. volume of output) or general utility, but many more discussed criteria related to *content and information quality*, various key *news values*, *audience-oriented* factors and whether the output was *checkable or verifiable*.

Content and information quality includes many factors that might be applicable in different contexts including clarity, concision, specificity, timeliness, readability, context, completeness/thoroughness, publishability, or even just common sense. Also often mentioned were news values of accuracy, validity (e.g. aware of limitations), relevance, and originality (e.g. including something exclusive or surprising). Respondents also looked to their audience to define what it means to have done a task at an acceptable level, mentioning factors like audience engagement and user feedback. Finally, respondents talked about whether the output from AI was verifiable or could be checked, which included aspects of replicability, provenance, explanation, and ease of fact checking.

AI Help


Respondents were also asked if they would want help from AI in their chosen task, which was then categorized according to a set of generic occupational activities that are relevant to journalism⁹. In Figure 5 it's clear that there is strong interest in delegating activities related to *analyzing data or information* (16 “yes,” 4 “maybe”), and some interest for *getting information* (11 “yes,” 6 “maybe,” and only 3 “no”), *processing information* (9 “yes,” 5 “maybe,” 1 “no”), and *communicating with people outside the organization* (4 “yes” and 3 “maybe”)¹⁰. But there's more resistance and uncertainty around getting help from AI for activities such as *thinking creatively* (9 “yes,” 11 “maybe,” 8 “no”), *making decisions and solving problems* (6 “yes,” 6 “maybe,” 6 “no”), or internal *communicating with supervisors, peers, or subordinates* (2 “yes,” 4 “maybe,” 3 “no”).

⁹ We base these on the occupational information available from O*Net online: <https://www.onetonline.org>

¹⁰ See Appendix B Q13 for definitions of these work activities.

Figure 5.

The distribution over categories of responses “No,” “Maybe” and “Yes” in response to the question “Would you be interested in having AI help with this task?”



For respondents who indicated that they would want AI to help with their identified task (50%, 93 of 186) they mentioned several specific aspects of the tasks. Some indicated that they would want automation to help with efficiency, reduce repetition, or increase precision, augmentation to help them do their job better, or to transform the task to a review task so that they could more quickly check and complete it. These suggestions are informative because they help frame how people want to leverage AI along various levels of automation, with a few calling for actual automation but with more looking for augmentation or task transformation to keep the human in control.

Among respondents who indicated that they might want AI to help with their identified task (32.2%, 60 of 186), they expressed uncertainty about a range of factors including around the capability and accuracy of the models, and whether their use of generative AI would make them overly reliant on the technology. Respondents raised issues of trustworthiness, humanness, and whether there was enough of a payoff for using generative AI. In addition,

respondents raised issues related to intellectual property and ethics such as confidentiality. For respondents who indicated that they would not want AI to help with their identified task (17.7%, 33 of 186), they focused on shortcomings of the technology including of knowledge, capability, performance, accountability, trustworthiness, or humanness (e.g. requiring human judgment or human-to-human contact or relationship) and not knowing if there was sufficient payoff for the investment needed.

Another way to think about whether people might want help from AI is to look at how tedious, boring or repetitive they find a task. For their chosen task we asked respondents to rate “To what extent do you find this task (or parts of it) boring, repetitive, or tedious?” on a scale from 1 (“not at all”) to 4 (“a lot”). The results tabulated by aggregate work activity are shown in Figure 6. We find that respondents find tasks such as *processing information*, *communicating with people outside the organization*, and *analyzing data or information* had a high average rating, whereas activities like *thinking creatively*, *interpreting the meaning of information*, and *making decisions and solving problems* were rated as much less tedious. In terms of specific task categories we found that transcription and metadata production as well as distribution analytics and working with data were rated toward the higher end of the tediousness scale. These findings reinforce the findings in Figure 5 about what work activities might benefit workers if AI could help, both in terms of their own satisfaction as well as alleviating a sense of tedium or boredom induced by the activity.

Figure 6.

Average ratings for selected task in response to the question “To what extent do you find this task (or parts of it) boring, repetitive, or tedious?” aggregated according to work activities. (N=186)



Key Learnings and Opportunities

News organizations are already extensively exploring the use of generative AI for content production, from text and multimedia to transcription and translation. Aspirational use reflects interests in exploring more utility for information gathering and sensemaking, working with data, business uses, and metadata creation. And while both current and aspirational uses reflect an emphasis on enhancing existing workflows, there is some indication of increased interest in creating new user experiences through chatbots and personalization. At the same time, not all aspirational use cases are well-suited to the technology, reflecting some potential misunderstandings about what the technology can do. For instance, tasks like layout of content or the analysis of user analytics are likely best served by non-generative AI systems such as optimization algorithms or rule-based natural language generation systems. In assessing the work activities where respondents would like AI to help, it's clear that getting information, processing information, and analyzing data or information are key areas where there is demand and a recognition of moderate-to-high levels of tedium, indicating potential opportunities to develop better prompts, interfaces, and workflows.

However, even where there is considerable interest and activity around content production, questions linger about just how much productivity generative AI can yield. Workflows and roles are already changing, and in many cases new work is created in prompting models effectively and in editing outputs. It seems that some tasks may benefit overall, but others may not. Additional research might study specific tasks over time to evaluate performance and delve into how new roles are evolving and relate to each other in the overall organization. Another opportunity is to invest more in training journalists how to effectively control models through prompting. As one respondent noted: "I get way better results, if I invest more time and thinking in writing longer and more elaborate prompts."

Our findings also suggest there is an unmet opportunity to design new interfaces to support journalistic work with generative AI. In articulating the criteria used to assess work quality when delegating tasks, respondents indicated that information quality and verifiability were key dimensions. Journalists still see themselves as editors and checkers of the outputs of generative AI, which suggests opportunities to create tools and interfaces that encourage those checks and enable efficient editing and verification. This might, for instance, include details to support replicability of an analysis, provenance for a number, fact, or source, or a general explanation that could be used to help assess an output. Perhaps what journalists need in order to effectively use generative AI are well-designed *editing interfaces* to support human oversight.


Looking at respondents' hesitations toward getting help with a task from generative AI, it's helpful to see that there may be some dimensions, such as capability, performance, efficiency gain or trustworthiness that might be addressed through technical testing and evaluation, or additional interface design work. In addition, articulating the criteria for success for a task, whether delegated to a human colleague or an AI system, can inform how to evaluate whether a task is performed at a high-enough level of quality. Considering both the hesitations for help and criteria for successful delegation on a task-by-task basis could inform engineering design approaches that overcome issues and benchmark acceptable performance levels. For example, in use cases where there is a concern for confidentiality, running models on local or organizationally owned infrastructure could alleviate that hesitation. At the same time, some dimensions, such as issues of model accountability or "humanness" are intrinsic and immutable limitations of the technology.

Novel technologies have often influenced journalism as a process and as a product. Generative AI is no exception, presenting journalists and media professionals with challenges and ethical considerations. These include but are not limited to challenges around source material, intellectual property concerns, and the bias that is ingrained in these technologies. In the following subsections we explore challenges and concerns focusing on the ethical dimensions of developing responsible practices.

Concerns and Challenges for Ethical Use

We asked respondents if they had ethical concerns about the use of generative AI in journalism. The most prominent concerns were *lack of human supervision* (21.8%, 48 of 174), *inaccuracy* (16.4%, 36 of 174), and *bias* (9.5%, 21 of 174). See Figure 7. *Lack of human supervision* feeds into the belief that generative AI would be used without human oversight. Respondents mention that they are not that concerned as long as the output is reviewed by an editor, while one noted that “I think reporters must independently verify information or pay the consequences.” There were also concerns about *inaccurate information*, including mis- and disinformation, as generative AI might produce a lot of incorrect output: “I have large, gaping voids of concern about AI in journalism. Incorrect information, fake images, bad stories, terrible grammar, job losses, all of it.” *Bias* is also a prominent concern, as respondents state that they are aware that the input and the output of generative AI might contain (hidden) bias. “Hidden biases and inaccuracy are the primary concern. Writing articles should be kept to humans, and gathering materials should also be done by humans, even though it’d be much more difficult to verify” noted one respondent.

Figure 7.
Responses reflecting concerns around the use of generative AI. (N=174)



Other concerns highlight the potential risks associated with using generative AI to produce journalistic content, particularly the erosion and *reduced quality* of news (7.7%, 17 of 174). Respondents worry that too much AI content will devalue journalism at a time when monetizing content directly through reader revenue is proving increasingly crucial. The *lack of transparency* of the use of generative AI is mentioned (6.8%, 12 of 174), as respondents fear that people in the news industry will not disclose whether they have used models like ChatGPT or Bard (now Gemini). Still other concerns include the threat of *job displacement* (6.8%, 12 of 174), the risk of *plagiarism* (3.2%, 7 of 174), and the *lack of originality* (2.7%, 6 of 174). Less mentioned concerns were *copyright* issues (2.7%, 6 of 174) and *privacy and data protection* (1.8%, 4 of 174).

Respondents were also asked to formulate some challenges they experience in addressing ethical concerns. One of the main challenges, respondents mention, is the *lack of training* (18.2%, 36 of 196). Training needs include teaching staff about the best practices and risks of generative AI. Other respondents state that smaller organizations might not have sufficient resources to invest in training: “Training is lovely, but time spent on training is time not spent on journalism – and a small organization can’t afford to do that.” In other words, there is not only a feeling that respondents are insufficiently prepared for the generative AI transformation, but also and maybe even more worrisome that there is simply insufficient time for an investment in training.

Another challenge concerns not having *regulation and guidelines* in place (11.1%, 22 of 196) or as one respondent put it: “We should have basic guidelines on what kind of things we check when taking on a tool.” The last prominent concern deals with the *lack of quality control* (8.1%, 16 of 196), as respondents worry that outputs from generative AI will not be verified sufficiently. A respondent states: “I worry that we do not hav[e] ‘standards staff’ in place to fact check AI. News organizations could be viewed as more trustworthy if we can show that real people enforce the news standards.” A potential gateway into dealing responsibly with these concerns and challenges, respondents mention, is by deciding what uses should be *banned* (15.5%, 34 of 196), which we elaborate further in the next section.

Banned Uses


We asked respondents if there were any uses of generative AI that should be discouraged in journalism. Among respondents that mentioned bans as part of the response to ethical concerns, a majority agreed that the generation of entire pieces of content by generative AI should be banned, as models are not yet reliable for this task (55.8%, 19 of 34). One respondent stated: “Any generative AI used to create content is concerning. We view generative AI like a police scanner. We use them to gather information, but still confirm and decide to report on our own.” The specific bans are also rooted in the general belief that journalism requires skills that cannot adequately be performed by a machine, and that outputs of generative AI could contain hallucinations. Other potential uses where bans were suggested include the generation of interview questions (17.6%, 6 of 34) and replicating

artists' styles using generative AI due to concerns regarding accuracy and authenticity (17.6%, 6 of 34). A respondent wonders: "Interviews often switch gears midway because of a reporter's instinct. How will AI match that?" Several respondents proposed to ban the use of generative AI to create content to mislead or deceive, as doing so would conflict with journalism's commitment to trust and integrity in journalistic practices (8.8%, 3 of 34). A respondent states: "We should not generate text, images, or any other reader-facing information that violates the trust they put in our editors and reporters." Additionally, some respondents suggested not to use generative AI for local news coverage and investigative reporting, underscoring the recognition that AI does not possess the nuanced understanding or ethical judgment required for these journalistic endeavors.

These suggested bans point to an emerging belief that there are some forms of using generative AI in journalism that are simply unacceptable. In other words, apart from concerns about actual productivity gains, ethical considerations and public expectations toward the role of journalism can be another important reason to refrain from using generative AI for certain tasks. Of course, this survey is only a snapshot, and it may be worth revisiting the topic further, once generative AI practices have been more firmly integrated into journalistic routines and roles. The responses highlight a collective effort within newsrooms to uphold journalistic standards, safeguard against misinformation, and prioritize the role of human judgment and ethical considerations in news production. Responses underscore that having guidelines in place could contribute to upholding these journalistic standards.

Strategies for Ethical Use

Figure 8.
Responses on various
strategies for ethical
use of generative AI
(N=145)



As shown in Figure 8, the most frequently mentioned strategy for use in overcoming ethical concerns and challenges was *not using generative AI* (20.0%, 29 of 145). In other words, 1 out of 5 respondents stated that a strategy to ethically use generative AI is to avoid its use altogether. This also means that ethical concerns can be an important obstacle to the deployment of generative AI in newsrooms. One respondent mentions: “I think the use of generative AI in my work is unethical, full stop.” Another strategy that plays a role is adhering to existing *guidelines and legal frameworks* (14.5%, 21 of 145). As one respondent states: “We apply the same standards to AI-generated content/information that we would to anything else that we publish or rely on. We have to be able to understand it and stand by our decision to use it.” The strategy of consulting guidelines is closely followed by relying on personal moral compasses and gut feeling (8.3%, 12 of 145).

Apart from a total ban of using generative AI, respondents emphasize that they should *limit its use* (8.3%, 12 of 145) as well as *verify the output* (15.2%, 22 of 145). Respondents underscore that they only use it on a test basis and “compare with other known materials to gauge whether it is accurate.” Additionally, respondents emphasize the need for human oversight and thorough fact-checking. As one respondent put it: “We strongly rely on our editorial core values such as facticity, transparency, impartiality, and accountability. These values have been the foundation of our journalism for almost 80 years. They are ideally suited to creating an ethical framework.”

Some organizations are awaiting further advancements in generative AI that address copyright and intellectual property concerns before considering implementation. There is a mix of readiness, caution, and proactive measures being taken to navigate the challenges associated with generative AI in newsroom settings. At the same time, responses show that not all news organizations have strategies in place to overcome these ethical concerns and challenges. Guidelines play a role, but also an internal gut feeling and moral compass. Other strategies that might improve responsible use of generative AI, like responsible procurement of tools that include AI and automation and internal testing, auditing, and verifying the input are rarely mentioned.

Guidelines

Most respondents (61.2%, 104 of 170 respondents) are aware of various guidelines surrounding the use of generative AI in journalism, though specific knowledge and adoption vary among organizations. Some respondents express familiarity with guidelines from news outlets like *The Guardian*, *NPR*, *BBC* and *AP*¹¹, and regulatory frameworks like the EU AI Act¹² or the UK’s AI white paper¹³. Others mention that their organization has its own set of guidelines (42.3%, 72 of 170 respondents). Common themes in existing guidelines echo the ethical concerns and challenges including transparency, human oversight, and avoiding and banning the use of generative AI for producing content entirely.

There are variations in approaches to crafting guidelines; some organizations adopt a more bottom-up approach by forming working groups to establish guidelines for generative AI, while others have a more top-down approach by relying on existing industry standards or guidelines. Additionally, a few respondents indicated that there should be industry-wide standards in place for the use of (generative) AI, either in combination with self-regulation or in the form of guideline documents.

11 <https://blog.ap.org/standards-around-generative-ai>

12 See: AI Act, Shaping Europe’s digital future.
<https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>

13 See: A pro-innovation approach to AI regulation.
<https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach/white-paper>

Results from the survey emphasize that guidelines should be regularly reviewed, and adaptable to the latest developments of generative AI. Respondents find many guidelines high level. There is a need for concretization and operationalization to make them meaningful for practitioners. Responses include requests for clear delineations of use cases or specific generative AI tools that should be allowed or banned. Some respondents emphasize the importance of mentioning which uses of generative AI should be disclosed to the audience, as well as addressing the specific biases in AI-generated content. Enforcement challenges and the need for a balance between experimentation and regulation are also noted as essential by respondents. As one respondent put it: “As AI evolves, it is not a black and white issue. There has to be room for testing, experimenting.” Another respondent mentions: “It’s more like a judgment call than a clear set of rules. Also, we can only enforce them to an extent. How do I know for sure whether specific paragraphs were AI generated?”


Responses mention that a potential solution for adding more specificity in the guidelines could be to include an external and internal version of the guidelines. Internal guideline documents tend to be more detailed, providing information about banned processes, and what specific software applications to use. The external guidelines are often presented at a higher level, focusing on broader principles and are more targeted toward transparency with the audience. Among the organizations that indicated they had their own guidelines, 22.8% (16 of 70 respondents) noted that they have a separate internal version. Having guidelines in place is only one requirement for potential responsible use of generative AI, or as a respondent states: “I think we are paving the road as we are driving – it is a new technology that seems to explode out of a box and now we are trying to navigate a world where new ‘amazing’ AI tools are dropping left right and center.”

Who Is Responsible?

While a focus on guidelines would seem to assign responsibility to the users of generative AI for ensuring responsible use, we also asked respondents to rank various other stakeholders who might be tasked with ensuring the responsible use of generative AI in journalism, including reporters, editors / managers, technology vendors, executives, staff technologists, the legal department, and unions. Each respondent ranked each of those stakeholder roles from 1 to 8, where 1 presents greater responsibility for ensuring the responsible use of generative AI in journalism (See Figure 9). Overall, respondents report that editors and managers should have the greatest responsibility for ensuring the ethical use of AI (average rank: 2.5), followed by executives (average rank: 3.3) and reporters (average rank: 3.8). At the bottom of the ranking, respondents tend to put unions (average rank: 6.2) and technology vendors (average rank: 5.6) as having less responsibility for ensuring the responsible use of AI.

Figure 9.

Average ranking across 8 stakeholder roles when asked “Who do you think bears greater responsibility for ensuring the responsible use of generative AI in journalism?” Lower numbers indicate a higher ranking of responsibility.



News Content as Training Data

Respondents were asked if they thought other companies should be allowed to train their AI models on news organizations’ digital reporting and information. The biggest group of respondents seems torn, as they responded “maybe” to this question (53.6%, 87 of 166). One respondent states: “In theory this sounds like a good idea but it’s scary to think of not having control over how content is used. (...). Is it fair to let tech giants profit on the shoulders of the reporters grinding out the hard work?” While recognizing the potential for revenue generation and advancements in AI tools, these respondents emphasize the need for careful consideration of copyright issues, transparency, and accountability to protect intellectual property and journalistic integrity. Additionally, there is a need for transparency and accountability in how the data is used and whether proper attribution is given to the original creators. Another respondent mentions: “Good inputs means good outputs. News is verified and high standard content. The ‘maybe’ is about proper compensation and tech companies taking their safety remit seriously.”

Those opposing the idea of allowing other companies to train models on their digitized information (32.5%, 54 of 166) express concerns about copyright infringement, unauthorized use of proprietary content, and the potential negative impacts on the competitiveness and sustainability of news organizations. Skepticism surrounds the fairness of contributing valuable data without compensation or control over its use, with worries about bias, misinformation, and loss of public trust in journalism. A respondent states: “Why would you? If a company that wants to make profit needs our content to do so, they can either pay for it or share the profits. Taking someone’s work and using it for your own benefit is simple theft.”

Those advocating for allowing companies to train on their digitized archives (13.9%, 23 of 166 responses) argue that such collaboration could significantly advance the field by improving AI model accuracy and reliability, benefiting the news industry and society. Collaboration is seen as vital for producing accurate, fact-checked content while adhering to professional and ethical standards. Additionally, respondents highlight the potential for revenue generation and practical benefits for reporters, emphasizing the importance of transparency and copyright adherence in collaborative efforts. Overall, allowing access to news data for AI training is viewed by some respondents as a mutually beneficial endeavor that can enhance the quality of AI-driven journalism while respecting journalistic integrity and legal considerations.

Key Learnings and Opportunities

Our results reveal that respondents have a range of ethical concerns about the use of generative AI. About a quarter of our respondents even indicate that these ethical concerns are a reason for not using generative AI, or only in a limited way. Addressing these ethical considerations and challenges are vital for the responsible implementation of generative AI.

The most pressing concern is linked to losing control, or having a lack of human oversight. Other prominent concerns address the quality of the output (accuracy, bias, originality, transparency). Less prominent concerns include copyright issues or job displacement which suggest few of our respondents are worried about job losses due to the technology. Having said so, these results can also have been influenced by the composition of our respondents and their roles in the organization. Other less prominent concerns were the lack of transparency or disclosure when generative AI is used, both internally (inside the news organization) and externally (toward the audience). Among the respondents who stated that uses of generative AI should be banned, the majority mentioned the generation of entire pieces of content. Other suggested bans include generating interview questions and replicating artists' styles using generative AI due to concerns regarding accuracy and authenticity.

When asked about overcoming these concerns, 1 in 5 respondents mention that they require training to use generative AI more responsibly. In other words, there is currently not only a feeling that respondents are insufficiently prepared for the generative AI transformation, but also and maybe even more worrisome that there currently is simply not sufficient room for and investment in training. At least for Europe, in the forthcoming AI Act there will be a requirement for providers and professional users, such as media organizations, to take measures to ensure the AI literacy of their staff, taking into account also the context in which the technology is being used. One in 10 respondents also emphasized the importance of having guidelines in place.

We observe that the concerns that were mentioned are closely related to the use of generative AI as a tool. It is also valuable to report issues that we know to be of concern but which were not mentioned at all by respondents. In the public discourse around generative AI, concerns about the environment and the ecological footprint of generative AI, extractive labor practices and the working conditions of AI workers, the growing power imbalance and dependency of the media on large tech companies, the danger of further reinforcing social injustice and disparate treatment or even more alarmistic calls about existential threats to humanity figure prominently. None of these concerns were reflected in the responses, which remained focused on daily journalistic practices. Considering that journalism does have an important role in informing the public discourse, there is a need to explore whether this lack of concern for broader ethical issues is the result of a mental disconnect, lack of awareness or the way the questions were framed.

Strategies for using generative AI responsibly were focused on monitoring the output and far less on strategies to monitor the input and the actual models. Some of the prominent concerns that were mentioned like bias, lack of transparency, and lack of accuracy can already manifest themselves earlier in the generative AI development process and can also be addressed (potentially more efficiently) by the model provider. Put differently, throughout the survey responses, we observed very limited critical engagement with ethical and legal concerns at the level of the input (training data) and the model, and by extension the trustworthiness of the technology itself. In part, this could be explained by the fact that most commercial proprietary large language models are not particularly transparent about training data and the model, and partly this could also be a consequence of the need for more training and AI literacy that many survey respondents flagged. It remains to be seen to what extent forthcoming legal mandates will be more transparent about the way generative AI models have been trained (for example under the European AI Act) might also result in actual scrutiny and more critical assessments of the tools used.

Guidelines are an important instrument for using generative AI responsibly, but respondents emphasize the need for a more dynamic approach. Some results reveal that the responsible use of generative AI should be seen as a living document rather than a static set of rules. These guidelines should also be more concrete, with more specific examples of which tools should and should not be used. However, guidelines that are too specific could undermine experimentation with generative AI by overspecifying behaviors. When evaluating the responsible use of generative AI, we have observed that gut feeling and personal moral compasses play an important role for some respondents, although we could question if this “subjective” feeling is sufficient for deciding what responsible use of generative AI is. One of the challenges is how to align, enforce and translate often vague principles and guidelines into practices on the work floor.

When asked about the use of news content as training data for generative AI, the biggest group of respondents, namely half of them, are torn. Results reveal that respondents see potential for revenue generation, but at the same time they emphasize the need for careful consideration of copyright issues, transparency, and accountability to protect intellectual property and journalistic integrity. The ones that are in favor of allowing companies to train on their news reporting, fewer than 14% of respondents, argue that such collaboration could significantly advance the field by improving AI model accuracy and reliability, benefiting the news industry and society.

Our results reveal that the meaning of responsible use of generative AI depends on the outlet involved, and a few respondents mentioned interest in industry-wide guidelines that could be adapted. Respondents did mention some common guidelines at a more abstract level that include transparency, human oversight and specific banned uses. In short, responsible use and implementation of generative AI takes time and resources. Respondents state that we are in the early stages of finding out what responsibility in relation to generative AI means. The news industry needs time to learn new skills, respondents say, and they need to actively experiment in line with already existing guidance and guidelines.

The news industry has rapidly reacted to the wave of generative AI technology that is working its way through society. By adapting workflows, adding new roles, and developing approaches to use generative AI technology in its practices we see signs of the evolution of newswork and of responsible practice in light of the capabilities and limitations of the technology. Yet, there are a whole host of areas where additional investment and action is needed:

- Usage **policies** such as guidelines could be made more concrete to better steer practitioners toward responsible use around specific tasks and use cases. And tools themselves could be evaluated more rigorously and systematically to ensure alignment with journalistic expectations and norms for accuracy, bias, privacy, and so on so that use is more responsible by default.
- Guidelines alone are not enough though, and need to be effectively implemented into working processes and routines to establish **practices** of responsible use, including practices of human oversight, responsible experimentation and the creation of dedicated support and learning structures.
- Additional **research** is needed to establish an evidence base around which tasks and use cases actually benefit in terms of efficiency and performance gains, as well as to elaborate the criteria to evaluate success and quality output for a range of tasks.
- **Design** and prototyping might be used to explore more powerful interfaces to support human oversight and editing of generative outputs, while also exploring genuinely new experiences rather than just the optimization of existing workflows.
- And new **training** programs are needed, not only in prompt writing but also in responsible use and adherence to usage guidelines (or other policies) as well as in thinking systematically about how to evaluate and refine workflows or strategically about how to develop entirely new ones.

In short, news organizations are still in the early phases of the proliferation of this powerful new technology, and much work remains to realize its full potential for journalism by advancing on policy, practices, research, design, and training.

Disclosures and Acknowledgements

The Associated Press has licensed select text archive content to OpenAI for training. OpenAI provides access to select technology as part of the agreement¹⁴. Effort for Nicholas Diakopoulos and Charlotte Li on this report is supported by a grant from the John S. and James L. Knight Foundation. Hannes Cools and Natali Helberger are members of the AI, Media & Democracy Lab¹⁵, Amsterdam, which supported their efforts for this research.

14 <https://www.ap.org/media-center/press-releases/2023/ap-open-ai-agree-to-share-select-news-content-and-technology-in-new-collaboration>

<https://apnews.com/article/openai-chatgpt-associated-press-ap-f86f84c5bcc2f3b98074b38521f5f75a>

15 <https://www.aim4dem.nl>

Appendix A. Participant Sample

Figure A1.

Respondents occupy a variety of roles in their respective news organizations. The majority of the survey respondents were either an editor or an executive. Of people who chose the “Other” category, we observe roles such as visual journalists, product managers, multiple role positions, and consultants. (N=290)

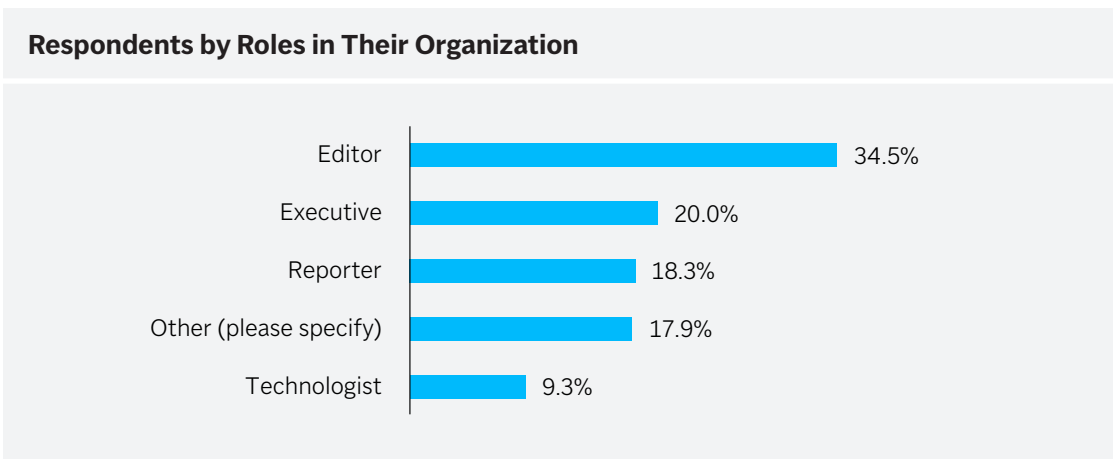


Figure A2.

The survey reached an expansive range of respondents in terms of their length of time worked in the news industry. Among the respondents, the most senior indicated 54 years, and the newest to the industry has had less than a year of experience in the industry. To present the data, responses were binned into 5-year intervals. (N=290)




Figure A3.

A majority of survey respondents conduct work in the West. Survey responses of countries were aggregated into continental regions based on Our World in Data¹⁶ classifications. (N=290)





Figure A4.

Respondents represent a diverse range of news organizations, ranging from digital native media to broadcasters (public and private). (N=290)



16 <https://ourworldindata.org/grapher/continents-according-to-our-world-in-data?overlay=data>

Figure A5.

Respondents represent a variety of newsroom sizes in terms of full-time editorial employees, with people from both very large (100+ editorial employees) and small (1-10 editorial employees) news organizations. (N=290)

Respondents by News Organization Size




Figure A6.

A majority of respondents report the technical team size of their news organizations are smaller than 10 employees, with 17.6% of these respondents reporting they do not have a technical team at their organizations and 14.5% of respondents reporting having technical teams larger than 100 employees. (N=290)

Respondents by Technical Team Size





Figure A7.

A majority of respondents identify as men. Among respondents who self-described their gender identities, one described their identity as non-binary femme. (N=290)



Appendix B. Survey Questions

Q1. Please indicate your current job title/role (space provided)

Q2. How would you classify your role?

- Executive
- Reporter
- Editor
- Technologist
- Other (space provided)

Q3. How many years have you worked in the news industry? (space provided)

Q4. What country do you work in? (space provided)

Q5. What kind of news organization do you work for?

- Digital Native Media
- Legacy Newspaper
- Magazine
- Media Group
- News Agency
- Public broadcaster
- Commercial broadcaster

Q6. What is the size of your news organization in terms of full-time editorial employees (i.e. reporters, editors, etc.)?

- 1-10
- 10-25
- 25-50
- 50-100
- 100+
- Not sure

Q7. What is the size of your news organization in terms of full-time technical employees (i.e. data science, software developer, etc.)?

- 0
- 1-10
- 10-25
- 25-50
- 50-100
- 100+
- Not sure

Q8. Indicate your level of agreement with the following statement:

“I am knowledgeable about generative AI.”

- Strongly disagree
- Disagree
- Neither agree nor disagree
- Agree
- Strongly agree

Q9. We are interested in issues of gender diversity relating to AI.

Please indicate your gender: How do you identify?

- Man
- Woman
- Non-Binary
- Prefer not to answer
- Prefer to self describe (space provided)

Q10. Have you or your organization used generative AI in some capacity? (Yes/No)

- You responded that you or your organization use generative AI in some capacity.
 - Q10a. What tasks have you or your organization used generative AI for on an experimental or regular basis? (space provided)
 - Q10b. Based on the tasks where you or your organization have regularly or experimentally used generative AI, please explain how it has or hasn't been effective in meeting your needs and expectations. (space provided)
 - Q10c. Have any of your tasks or workflows changed as a result of generative AI? (Yes/No)
 - Q10d. You responded that tasks or workflows changed as a result of generative AI. How so? (space provided)

Q11. List at least three tasks that you would ideally like to use generative AI for in your work, if it were capable of producing quality results. (Five spaces provided)

Q12. Has your organization created any new positions that are specifically geared towards using generative AI? (Yes/No)

- Q12a. You responded that your organization has created new positions that are specifically geared towards using generative AI. What new job titles/roles were created

and what does that person do? (space provided)

Q13. From the following options, please indicate all areas of work activity that you find important in your daily work:

- Getting Information — Observing, receiving, and otherwise obtaining information from all relevant sources.
- Communicating with People Outside the Organization — Communicating with people outside the organization, representing the organization to customers, the public, government, and other external sources.
- Interpreting the Meaning of Information for Others — Translating or explaining what information means and how it can be used.
- Identifying Objects, Actions, and Events — Identifying information by categorizing, estimating, recognizing differences or similarities, and detecting changes in circumstances or events.
- Communicating with Supervisors, Peers, or Subordinates — Providing information to supervisors, co-workers, and subordinates in various modalities.
- Establishing and Maintaining Interpersonal Relationships — Developing and maintaining constructive and cooperative working relationships with others.
- Performing for or Working Directly with the Public — Performing for people or dealing directly with the public.
- Updating and Using Relevant Knowledge — Keeping up-to-date technically and applying new knowledge to your job.
- Thinking Creatively — Developing, designing, or creating new applications, ideas, relationships, systems, or products, including artistic contributions.
- Documenting/Recording Information — Entering, transcribing, recording, storing, or maintaining information.
- Organizing, Planning, and Prioritizing Work — Developing specific goals and plans to prioritize, organize, and accomplish your work.

- Working with Computers — Using computers and computer systems to program, write software, set up functions, enter data, or process information.
- Analyzing Data or Information — Identifying the underlying principles, reasons, or facts of information by breaking down information or data into separate parts.
- Making Decisions and Solving Problems — Analyzing information and evaluating results to choose the best solution and solve problems.
- Processing Information — Compiling, coding, categorizing, calculating, tabulating, auditing, or verifying information or data.
- Monitoring Processes, Materials, or Surroundings — Monitoring and reviewing information from materials, events, or the environment, to detect or assess problems.
- Evaluating Information to Determine Compliance with Standards — Using relevant information and individual judgment to determine whether events or processes comply with laws, regulations, or standards.
- Scheduling Work and Activities — Scheduling events, programs, and activities, as well as the work of others.
- Judging the Qualities of Objects, Services, or People — Assessing the value, importance, or quality of things or people.

Q14. Now select one broad category of work activity that you indicated was important, for which you will answer some more specific questions: (select one category from prior question)

Q15. Within the one broader category of activity you selected, please describe a related specific task that you do in your work? (space provided)

- Q15a. How often do you do this task in the course of your job?
 - Never
 - Rarely
 - Sometimes
 - Often
 - All the time
- Q15b. To what extent do you find this task (or parts of it) boring, repetitive, or tedious?
 - Not at all
 - A little
 - Somewhat
 - A lot

- Q15c. If you were to delegate this task to a colleague you managed, and were responsible for the output, what criteria would you use to evaluate whether the task was done to an acceptable level of quality? (space provided)
- Q15d. Would you be interested in having AI help with this task? (Yes/Maybe/No)
 - You replied 'yes' to the question about having AI helping with the task. What aspects of this task specifically would you want AI to help with? (space provided)
 - You replied 'maybe' to the question about having AI helping with the task. What are you unsure about in terms of having AI help with this task? (space provided)
 - You replied 'no' to the question about having AI helping with the task. Why do you not want AI to help with this task? (space provided)

Q16. What do you see as the opportunities for the use of generative AI in journalism? (space provided)

Q17. Do you have ethical concerns about the use of generative AI in journalism? Are there any specific uses that should be discouraged? (space provided)

Q18. What are the greatest challenges for responsibly using generative AI within your organization? How might your organization overcome those challenges to help you use generative AI more ethically? (space provided)

Q19. Are you aware of any guidelines around the use of generative AI in journalism? (Yes/No)

- Q19a. You answered that you are aware of guidelines around the use of generative AI in journalism. Which ones are you aware of? (space provided)

Q20. Does your news organization have its own set of guidelines for the use of generative AI? (Yes/No)

- You responded that your news organization does have its own set of guidelines for the use of generative AI.
 - Q20a. To what extent do you find them helpful in deciding what is ethical use?
 - Not at all
 - A little
 - Somewhat
 - To a large extent
 - To a great extent
 - Q20b. What do you think might be missing? (space provided)
 - Q20c. Are the guidelines enforced? (Yes/No)
 - Q20d. Are there separate externally and internally facing versions of the guidelines? (Yes/No)
 - You responded that there are separate external and internal versions of guidelines. Please elaborate any differences between the two. (space provided)

Q21. What strategies do you use to decide what is the ethical use of generative AI in your work? (space provided)

Q22. Who do you think bears greater responsibility for ensuring the responsible use of generative AI in journalism? [Rank order the following]

- Reporters
- Editors/Managers
- Technology vendors
- Executives
- Staff technologists
- Legal department
- Unions
- Legislators

Q23. Do you think news organizations should allow other companies to train generative AI models on their published data/content? (Yes/Maybe/No)

- You responded that news organizations should allow other companies to train generative AI models on their data. Why? (space provided)
- You responded that news organizations should NOT allow other companies to train generative AI models on their data. Why? (space provided)
- You responded that news organizations should maybe allow other companies to train generative AI models on their data. Why? (space provided)

Q24. What do you think labor unions should be requesting when it comes to the use of generative AI in news production? (space provided)

Appendix C. Additional Resources

Below are links that can provide guidance for your use of AI in the newsroom:

Generative AI in the Newsroom, a collaborative effort led by Nick Diakopoulos to figure out how and when (or when not) to use generative AI in news production.
<https://generative-ai-newsroom.com>

AI, Media & Democracy Lab, an ethical, legal, and societal laboratory focused on the implications of AI for media and democracy led by Natali Helberger.
<https://www.aim4dem.nl>

AI @ AP, the Associated Press' work on AI including its first report published in 2022, free online courses and its five AI projects for local newsrooms.
<https://ai.ap.org>

AI Transparency initiative led by Nordic AI Journalism.
<https://www.nordicaijournalism.com/ai-transparency>

Council of Europe Guidelines on the responsible implementation of artificial intelligence systems in journalism.
<https://rm.coe.int/cdmsi-2023-014-guidelines-on-the-responsible-implementation-of-artific/1680adb4c6>

Partnership on AI offers a procurement guide on AI tool adoption for newsrooms.
<https://partnershiponai.org/ai-for-newsrooms>