Automating the Fourth Estate: Journalistic Norms and Ethics Under the LLM Revolution

Joseph Karaganis

jbk2191@columbia.edu

Faculty Mentor: Prof. Chris Wiggins

Laidlaw Scholars Program Summer 2023

Introduction

In November of 2022, San Francisco-based artificial intelligence (AI) company OpenAI released ChatGPT to the public. ChatGPT was the latest breakthrough in a long line of AI-powered chatbot programs. Trained on a vast "corpus" of human-produced text, it could replicate the patterns it observed in that text to write persuasive and straightforward prose across a wide range of topics—on everything from the aesthetics of Augustan poetry to the workings of its own software architecture. The surprising depth and breadth of ChatGPT's language generation capabilities quickly attracted media attention and controversy. Supporters lauded OpenAI for creating the first viable public-facing product that took advantage of recent developments in large language model (LLM) technology, while critics suggested that the software's frequent "hallucinations"—the industry term for its tendency to invent facts and pursue unpredictable tangents—reflected an incomplete and potentially dangerous product. Nonetheless, users began flocking to OpenAI's website in order to try the chatbot out for themselves. By January of 2023, ChatGPT had over 100 million active users, making it the fastest growing web platform of all time—much faster than Tiktok and Instagram.[1]

AI's now elevated status in the collective consciousness is mostly a product of this ubiquity. Questions that were once only raised in science fiction, Big Tech labs, and academic literature have become accepted by the broader public as important to the future of science and technological innovation. Not all visions of an AI-powered future are optimistic: critics have highlighted the ethical, social, and political risks that may emerge as AI systems become more powerful. These risks exist both in the short-term—academic plagiarism, online misinformation,

[1]Andrew R. Chow, "Why CHATGPT Is the Fastest Growing Web Platform Ever," Time, February 8, 2023, https://time.com/6253615/chatgpt-fastest-growing/.

social isolation, mass surveillance—and the long-term—mass unemployment, national security threats, political repression, targeted assassination, and even human extinction.

While ChatGPT has emerged as a focal point in conversations about the future of AI technology, it represents only a small corner in the expansive field of artificial intelligence and machine learning (ML). ML can be used to develop AI models that are optimized to solve any number of tasks (playing chess, running social media platforms, analyzing difficult engineering problems, etc…). Only a subset of those systems—chatbots, like ChatGPT and Google's Bard—are trained to recognize natural language inputs and respond in turn. Chatbots fall under the umbrella term of "generative AI," which describes systems that can produce content (images, text, audio, etc…).

Many modern chatbot systems are powered by LLMs—complex statistical formulas built from artificial neural networks that approximate human brain function. After a process of "unsupervised learning," during which the model learns to develop associations between different words and phrases using large batches of unlabeled data, the models undergo several more specialized training phases, some of which may be human-assisted. The extent of this human assistance ranges from "semi-supervised" and "supervised" learning paradigms, which use human-labeled data to increase model accuracy, to "reinforcement learning from human feedback" or RLHF, in which humans rate the accuracy of the model's outputs.

ChatGPT is powered by GPT-3.5, an LLM released by OpenAI in 2022 (GPT-3.5 is similar but more powerful than the previous model, GPT-3, which was released in 2020). What made GPT-3 unique—and ChatGPT so powerful—was the much larger quantity of material used to train the model. "Large" language models (as opposed to traditional language models) have a

greater number of parameters (the variables used to symbolize individual data points).[2] GPT-2 had 1.5 billion parameters. GPT-3 had 175 billion. GPT-4, the latest model, has significantly more, although the exact number has remained undisclosed. Unlike GPT-2 and GPT-3, GPT-4 is "multimodal," which means it can interact with other mediums like image and video.[3]

This remarkable new technology has rapidly begun to find applications in fields that depend on standardized writing formats and documentary analysis. While some researchers have claimed that larger parameter counts will reap decreasing (and possibly negative) returns in the future,[4] current adoption trends suggest that LLMs are already powerful enough to be useful across a wide range of basic tasks. Humans will continue to experiment with LLMs—over time, their limits, strengths, and applications will be identified. AI will be further integrated into our social, political, and economic institutions. As systems become more powerful, that integration may have unexpected results.

This fact raises important questions about what integration might look like. How will humans harness LLMs in the workplace? How might LLMs transform, or even eliminate, tasks and functions that are foundational to human work and productivity? This study focuses on an industry that is experiencing a particularly sharp paradigm shift with regards to LLM integration: journalism and the news media. Journalism's reliance on well-defined writing formats, its adherence to strict ethical and normative standards, and its perennial financial struggles have created unique circumstances with regard to experimentation with and the adoption of LLM

---

[2] Kyle Wiggers, "The Emerging Types of Language Models and Why They Matter," TechCrunch, April 28, 2022, https://techcrunch.com/2022/04/28/the-emerging-types-of-language-models-and-why-they-matter/?guccounter=1&guce_referrer=aHR0cHM6Ly93d3cuZ29vZ2xlLmNvbS8&guce_referrer_sig=AQAAAJ2haZlG9HkOb8uxapLpivpm4TeUlQR4FjtsGUJ6O74G38XUot9x_m09nF4OvAsZMjoFZpfISRJnQfNqiw85uuK2Q2Kh9yohCbdNeJW0bh76efU5y5-wlo1mLmSJunWNGlYKEbE1X2MpWqdmH5TQSAWhWY1EjBmLujPQ1V6aW8x0.

[3]Will Douglas Heaven, "GPT-4 Is Bigger and Better than Chatgpt-but Openai Won't Say Why," MIT Technology Review, March 23, 2023. https://www.technologyreview.com/2023/03/14/1069823/gpt-4-is-bigger-and-better-chatgpt-openai/.

[4]Jacob Stern, "GPT-4 Might Just Be a Bloated, Pointless Mess," The Atlantic, April 25, 2023, https://www.theatlantic.com/technology/archive/2023/03/openai-gpt-4-parameters-power-debate/673290/.

systems. As a case study, observing the rise of generative AI-powered journalism may help illustrate a broader tension that will play out across workplaces and institutions in the years to come, where the tendency towards the adoption of new information technology complicates organizational hierarchies and transforms business models. We are living in the “wild west” of AI integration. Among journalists, the feelings of uncertainty brought on by generative AI technologies are particularly acute—and the risks especially concerning.

Background

The Pre-ChatGPT Era

Generative AI language systems have improved and become more widely available over the past decade or so, prompting large numbers of news media publications to experiment with versions of LLM technology and their predecessors. The Associated Press (AP) was a pioneer in the field, integrating an earlier form of NLG (Natural Language Generation) into their sports and corporate earnings stories as early as 2014. These articles were fairly static: an AI model would manipulate structured data (earnings reports, sports results, etc…) and absorb it into a predetermined outline. This application of generative AI was a far cry from the more flexible LLM systems that exist today, but it proved successful within its industry niche—AP's production of corporate earnings stories increased by tenfold after incorporating NLG.[5] While NLG-powered "fill-in-the-blanks" data journalism remained tied to certain simplistic article formats—the only ones where it was viable—the success of AP's experiment corresponded with a growth in generative AI use during the 2010s.[6] Other major news media corporations such as Yahoo News developed and implemented similar use-cases for early generative AI technology.[7]

During the "pre-ChatGPT era," beginning in the early 2010s and culminating with the November 2022 release of ChatGPT, academic research into the journalistic applications of generative AI remained relatively abstract and predictive—in part because the more serious

[5] "Artificial Intelligence: AP," Associated Press, accessed July 20, 2023, https://www.ap.org/discover/artificial-intelligence#:~:text=News%20production&text=This%20ranges%20from%20the%20automatic,both%20sports%20and%20corporate%20earnings.
[6] Nicholas Diakopoulos, *Automating the news: How algorithms are rewriting the media*. Cambridge, MA: Harvard University Press, 2019. (107)
[7] Mathew Ingram, "Is Ai Software a Partner for Journalism, or a Disaster?," Columbia Journalism Review, accessed July 20, 2023, https://www.cjr.org/the_media_today/ai_software_chatgpt_journalism.php.

applications of generative AI were hypotheticals. Publicly-available AI systems were incapable of producing coherent long-form text on their own. While viable in limited circumstances, these systems could not transform the news media production process or substantially reconfigure traditionally human-led roles in journalism.[8]

However, a few major academic studies during this period did produce accounts of how AI integration was being approached and anticipated in the industry. The first was an international survey of 71 news organizations released in 2019 by the JournalismAI initiative at the London School of Economics. The study suggested that news media companies were already concerned with growing applications of AI. 44% of respondents believed that AI already had an impact on their organization—another 19% believed that there would be an impact in the next year.[9]

Not all of this predicted "impact" can be attributed to early experiments with generative AI models. News media organizations—alongside most online content producers—have been using non-generative AI software for years, often in a commercial context.[10] These models wield machine learning to power dynamic paywall systems, advertisement personalization, and SEO (search engine optimization) website-retooling. It is important to make a distinction between these applications of AI technology—which are tangential to the actual newsmaking process—and applications that have a direct impact on how stories are conceived, written, and edited. Unlike debates over the use of generative AI tools, considerations about non-generative AI (when used in a commercial context) have been somewhat less contentious—in part due to their ubiquity and the non-specificity of their implications to the news media. The mass

[8] Diakopoulos, *Automating the news,* (97)
[9] Charlie Beckett, rep., *JournalismAI Report* (Polis: Journalism and Society, November 18, 2019), https://blogs.lse.ac.uk/polis/2019/11/18/new-powers-new-responsibilities/. (46)
[10] This project uses an expansive definition of AI that includes online commercial systems that are built using machine learning software, as well as earlier language models that did not use machine learning. The definition tracks the one used in the surveys and academic papers analyzed.

collection and distribution of data on the internet certainly raises important ethical and social questions—but those cannot be isolated to journalism (rather than, say, social media or online shopping).

The popularity of commercial non-generative AI is reflected in some of the data collected by the LSE survey. 45% of respondents were using AI to "deliver more relevant content to users," and another 20% were using it to "improve business efficiency."[11]

Still, the report does reveal some important—and specific—insights into the use of generative AI tools before ChatGPT. 68% of respondents intended to make the work of journalists "more efficient" by using AI systems. Most attempts to actually use AI as part of the journalism process were relatively small-scale experiments in data journalism—although in some cases, AI models provided crucial research assistance during investigative projects. Machine learning algorithms, when paired with human intelligence and oversight, were being used to extract content from complex datasets and to fact-check statements in real-time.[12] Not all of these algorithms were, strictly speaking, generative AI models that produced consumable content—although some were—but they shared generative AI's ability to identify patterns and irregularities in inputted data. These shared capabilities suggest a continuity in use-cases: while the machine learning tools of the late 2010s may have made the job of wading through large datasets much less time-consuming, multi-purpose LLM models (like GPT-4) can be trained to identify the same patterns *and* then to explain those patterns using clear and concise prose.

The report also suggests that rudimentary generative AI models were used to produce content summaries and to expand data journalism projects that would have been unfeasible without automation. For example, an AI model was used by reporters at *Le Monde* to automate

[11] Beckett, *JournalismAI Report*, (32)
[12] Ibid (24-28)

local election reports across all of France—it would have been impractical, and perhaps impossible, for humans at a single publication to report on the results of elections in all 36,000 French cities.[13]

Still, the LSE study cautions that many of the more ambitious AI-powered journalism projects envisioned by respondents were still in development at the time of the report's release. Most of the applications highlighted in the report are small-scale, experimental, and non-disruptive. The industry is described as broadly cognizant of this (relatively small) impact on data journalism production, but largely unprepared for the bigger normative and ethical questions surrounding the role of much more powerful AI systems in general-purpose news production. 60% of respondents were "concerned" about the impact AI could have on their journalism work[14]—yet 63% did not have a dedicated AI strategy.[15] Many that did report having a strategy also described "confusion around roles and responsibilities." Asked about which newsroom departments were taking the lead on AI policy, respondents were evenly split among five different power-sharing proposals—suggesting that the industry had not settled on an optimal organizational strategy for AI adoption.[16] The takeaway should not be that news organizations were incompetent or lacked proper foresight, but that the generative AI applications that existed at the time were not compelling or concerning enough to spur an industry-wide response (few outside of tech circles were anticipating the rapid improvement in language models that came during the early 2020s).

A similar study published by Oxford University's Reuters Institute for the Study of Journalism in January of 2022 also offered some insight into this "pre-ChatGPT" era of AI

[13] Beckett, *JournalismAI Report*, (25)
[14] Ibid (54)
[15] Ibid (39)
[16] Ibid (40)

integration. The report surveyed 246 senior "media leaders" from a wide range of publications spread across 52 countries. It corroborated the JournalismAI study's suggestion that some AI-adjacent tools were becoming widespread among publishers—according to Oxford, 85% of respondents believed that AI was useful for their content recommendation systems, and 69% believed that it was "critical" to their business operations. But the Oxford study also reflected increased publisher awareness of recently released LLM systems like GPT-3 and Google's Deepmind, and recognition that these systems could have a significant impact on news production. 81% of respondents believed that AI would be used to automate repetitive newsroom tasks like transcription and tagging. A surprising 40% of respondents believed that automated content generation would be an important newsroom use-case. The Reuters report supports these numbers with examples of contemporaneous innovations in generative AI—a BBC experiment in local elections reporting (not unlike *Le Monde*'s earlier experiment), and the Wall Street Journal's Narrativa, which uses generative AI to write stories about financial markets (with human editorial supervision).[17]

A final pre-ChatGPT study was released by the Associated Press in March of 2022. This report targeted a different group of respondents than the studies published at LSE and Oxford: instead of looking at news media leaders from across the world, it looked at local journalists working in the United States. The results suggested that smaller local publishers were less capable of integrating AI into their journalism and experimenting with possible use-cases than the bigger firms that were the focus of the LSE and Oxford studies. Managers at local publishers were described as "unprepared" for AI adoption. Most responded that their AI use was a 0 or 1

[17] Nic Newman, "Journalism, Media, and Technology Trends and Predictions 2023," Reuters Institute for the Study of Journalism, January 10, 2023, https://reutersinstitute.politics.ox.ac.uk/journalism-media-and-technology-trends-and-predictions-2023#header--8.

on a scale of 0 of 6.[18] Despite this result, many local news managers were interested in using AI to automate page layouts and social media feeds. Others were also interested in using an NLG system akin to the AP's "fill-in-the-blank" style content generator.[19]

The ChatGPT Era

Since ChatGPT's release, generative AI has become central to conversations about media innovation and newsroom adaptation to technological change. With generative AI tools freely available to the public, a growing number of newsrooms have begun to feel pressured to adapt to an environment where AI use is not experimental, but expected.

The pace of this attitude shift—from a cautious but optimistic view of all AI to something closer to a frenzy to take advantage of generative AI—has been remarkable. In May of 2023, a survey by WAN-IFRA (World Association of New Publishers) found that 49% of newsrooms were already working with generative AI tools. The same survey found that 70% of respondents expected generative AI tools to be helpful for journalists. Most respondents reported that they were using generative AI to produce summaries—others suggested that it was useful in their research work and made them more efficient.

However, while the pace of AI adoption has been startling, internal policy measures and norm-setting guidelines have failed to keep up. The LSE report suggested that 37% of newsrooms in 2019 had a dedicated AI strategy. According to the WAN-IFRA report, only 20% of newsrooms in 2023 had specific guidelines focused on regulating the use of generative AI. The discrepancy suggests that guidelines regulating past use-cases of AI have not yet translated into regulations for LLM use. This is unsurprising, given that LLMs represent a transformational

[18] Aimee Rinehart and Ernest Kung, publication, *Artificial Intelligence in Local News: A Survey of US Newsrooms' AI Readiness* (Associated Press, 2022). (13-14)
[19] Rinehart and Kung, *Artificial Intelligence in Local News*, 37-38

and unexpected technology innovation with substantial implications for newsroom practice. Developing responsive and appropriate AI policies will be a challenge for newsrooms in the months and years to come. Similarly worrying is that 49% of newsrooms allow journalists to use AI entirely as they wish.[20]

Of the companies that have begun drawing directly on generative AI for newswriting—a use-case that is strictly regulated at most publications with an established AI policy—two stand out in particular: CNET and Gizmodo. Both publications attempted a rapid rollout of generative AI, and both initiatives were met with significant criticism and condemnation for the poor quality of their AI-written articles.

In January, CNET came under fire after it was reported that they had published 77 articles that were at least partially written by an in-house generative AI model engineered by their parent company, Red Ventures. The articles, many of which focused on personal financial advice, were not disclosed as AI-written, and were full of inaccuracies—41 of the 77 were retroactively corrected. Red Ventures had published many of the articles without the knowledge of editorial staff and the resulting debate within the company produced a temporary pause on AI-generated content.[21]

In July, Gizmodo's parent company, G/O Media, experienced similar fallout over its publication of AI-generated content without approval from editorial staff. As with CNET, the articles were filled with inaccuracies. One was supposedly a chronological list of Star Wars films—but the list wasn't in chronological order. Public condemnation followed, including from Gizmodo staff members who felt blindsided by G/O Media's decisions.[22]

---

[20]Teemu Henriksson, "New Survey Finds Half of Newsrooms Use Generative AI Tools; Only 20% Have Guidelines in Place - Wan-IFRA," WAN, May 31, 2023, https://wan-ifra.org/2023/05/new-genai-survey/.
[21]Mia Sato and Emma Roth, "CNET Found Errors in More than Half of Its AI-Written Stories," The Verge, January 25, 2023, https://www.theverge.com/2023/1/25/23571082/cnet-ai-written-stories-errors-corrections-red-ventures.
[22]Wes Davis, "Gizmodo's Staff Isn't Happy about G/O Media's AI-Generated Content," The Verge, July 8, 2023, https://www.theverge.com/2023/7/8/23788162/gizmodo-g-o-media-ai-generated-articles-star-wars.

Current Policies:

Below, I have compiled generative AI editorial policies that have been publicly released by various online publications. The list is not exhaustive but it is representative of the conversations, norms, and policies that are emerging across different parts of the industry.

Financial Times (UK)	● FT's journalism will "continue to be reported and written by humans who are the best in their fields." ● A generative AI working group will consider possible use-cases and applications to journalistic work—mostly with regards to research support. ○ Potential for generative AI summaries to be created with human oversight. ● AI-generated visuals may be used (if clearly attributed to AI), but they cannot be photorealistic. ● All newsroom experimentation with AI will be recorded internally (including any reliance on third-party software) and a series of masterclasses will train journalists on responsible generative AI use.[23]

[23]Roula Khalaf, "Letter from the Editor on Generative AI and the FT," Financial Times, May 26, 2023, https://www.ft.com/content/18337836-7c5f-42bd-a57a-24cdbd06ec51.

VentureBeat (US)	● Will not "blindly" copy-paste generative AI text or allow generative AI to write entire stories. ● AI will be used as a supplementary tool to "inspire and strengthen" work. ● Human copyeditors will continue to review all text, regardless of generative AI influence or assistance. ● Optimistic and hands-on adoption of AI tools—will proactively integrate chatbots into brainstorming and headline/lede writing. ● EIC "fully supports" journalistic use of AI tools throughout each part of the production process. ● "Produce more, better stories at scale."[24]
Insider (US)	● Will promote individual journalists' responsible experimentation with generative AI. ● AI can be used for ideas, but not as a substitute for writing. Stories must still be completely written by reporters. ○ A pilot group will look at the possibility of AI-generated writing.

[24]Michael Nuñez, "Letter from the Editor: How Generative AI Is Shaping the Future of Journalism and Our Newsroom," VentureBeat, May 2, 2023, https://venturebeat.com/ai/letter-from-the-editor-how-generative-ai-is-shaping-the-future-of-journalism-and-our-newsroom/.

	○ Ground rules for the pilot group are to always verify AI-produced facts, verify originality, and ensure that the text satisfies the publication's house style. ● Rules apply exclusively to ChatGPT, because its terms of service have been reviewed by the editorial team. Other generative AI models may not be used. ● Journalists will be encouraged to use generative AI in any of the following ways: producing story outlines/structure, AI-powered editing, SEO headline generation, research summarization, and interview question writing. ○ Advises journalists not to input sensitive information in the system, because OpenAI may have access to it. ● Editors need to take initiative to ensure the work of their journalists is not written by AI. Journalists are still responsible for the originality and truthfulness of anything they submit. ● AI will not replace journalists, but will make

	them "faster and better."[25]
Associated Press (US)	● AP was an early innovator in AI-powered journalism and still uses generative AI systems at multiple points in the editorial process. ● Began automating corporate earnings stories in 2014, now also automates sports news. ○ After using NLG to cover corporate earnings reports, production grew by a factor of ten. ● Uses generative AI for video transcription, video shot-list descriptions, and story summarization. ○ Story summarization done by AI is still subject to human oversight. ● Work with emerging AI startups to experiment with generative AI at low costs.[26]
Buzzfeed (US)	● Generative AI will replace the "majority of static content," becoming embedded into all of Buzzfeed's products. ○ Content will become more interactive, personalized, and "gamified."

[25]Nicholas Carlson, "My Editor's Note to the Newsroom on AI: Let's Think of It like a 'Bicycle of the Mind,'" Business Insider, April 13, 2023, https://www.businessinsider.com/how-insider-newsroom-will-use-ai-2023-4.
[26] "Artificial Intelligence: AP." Associated Press

	● Site has experimented with interactive AI content and found that it leads to higher engagement levels. ● Already using AI for SEO headline generation. ● AI will be used to help journalists meet productivity goals (although how much of this will be "news" is uncertain, since Buzzfeed recently shut down their news website).[27] ● Now also using generative AI to mass-produce SEO-focused travel guides and personalized quizzes. ○ Website discloses that the content was "collaboratively" produced through human-AI interaction.[28]
The Oregonian (US)	● Collaborating with a company called United Robots to produce dozens of real-estate focused stories using generative AI. ○ Underwent a testing phase where the stories were carefully vetted for quality control.

[27]Maggie Harrison, "BuzzFeed Says AI Will 'Replace the Majority of Static Content,'" Futurism, May 18, 2023, https://futurism.com/buzzfeed-ai-replace-content.
[28]Jay Peters, "BuzzFeed Is Using AI to Write Seo-Bait Travel Guides," The Verge, March 30, 2023, https://www.theverge.com/2023/3/30/23663206/buzzfeed-ai-travel-guides-buzzy.

	○ Using the AP model of NLG sports news to produce generic stories within a set structure, instead of using LLM-based generative AI models for a more expansive (and less repetitive) form of content generation. ● AI-produced content is flagged and disclosed to readers.[29]
CNET (US)	● Using an in-house generative AI model called RAMP (Responsible AI Machine Partner). ○ All content will either be produced by a human or by this in-house model—and everything will be fact-checked and edited by a human. ○ RAMP will avoid plagiarization by "prioritizing accurate sourcing" and providing built-in citation standards. ● CNET plans on using RAMP to organize large datasets, reduce administrative work, and contribute to background research.

[29]Therese Bottomly, "Letter from the Editor: How AI and Automated Content Power Information on Oregonlive," oregonlive, July 10, 2023, https://www.oregonlive.com/opinion/2023/07/letter-from-the-editor-how-ai-and-automated-content-power-information-on-oregonlive.html.

	○ Certain passages in a story may be generated by RAMP and then fact-checked by a human (entire stories will not be written by AI). ○ Product reviews, images, and videos will continue to be exclusively produced by humans.[30] ● These new policies come after CNET writers unionized and protested a "lack of transparency and accountability from management" regarding the use of generative AI tools in dozens of articles. The union was not involved in discussions surrounding the new policy, which was crafted internally.[31]
The Guardian (UK)	● Generative AI working group created to assess different policies–looking at other approaches adopted around the industry. ● All AI-produced content will be disclosed to readers and will have been edited by a human (AI content will also require senior editorial

[30]"How We Will Use Artificial Intelligence at CNET," CNET, accessed July 20, 2023, https://www.cnet.com/ai-policy/#:~:text=Writing%20full%20stories%3A%20None%20of,fact%2Dchecked%20by%20our%20editors.
[31]Mia Sato, "CNET Is Overhauling Its AI Policy and Updating Past Stories," The Verge, June 6, 2023, https://www.theverge.com/2023/6/6/23750761/cnet-ai-generated-stories-policy-update.

	oversight). ○ Any inclusion of AI content will also be contextualized with the benefit it provides to the story and reader. ● Generative AI will be used as a tool and research assistant, not as a replacement for human writers. ● There will be certain ethical considerations regarding the transparency and compensation models built into each generative AI system.[32]
WIRED (US)	● Will not publish stories with AI-generated text (unless that is the explicit purpose of the article, in which case it will be properly disclosed and contextualized). ○ This applies to any text, including snippets. Undisclosed generative AI use for text generation will be treated as plagiarism. ● Text will not be edited or condensed using AI. ● Generative AI may be used for headline generation or social media, but with human oversight.

[32] Katharine Viner and Anna Bateson, "The Guardian's Approach to Generative AI," The Guardian, June 16, 2023, https://www.theguardian.com/help/insideguardian/2023/jun/16/the-guardians-approach-to-generative-ai.

	● AI can also be used for idea generation or as a research tool—but it must be scrutinized under the same standards as existing platforms (e.g. Wikipedia). ● Use of AI-produced images will be disclosed, and will only be tolerated if AI is part of a larger process that includes human creativity. ○ AI images will not replace stock photographs.[33]
Bavarian Broadcasting (Germany)	● Will use AI to personalize user experiences, generate content, and make such content more "attractive." ● Commit to continuously studying AI ethics discourse to "avoid a gap between theory and practice." ● Emphasize data security and responsibility. ● New use-cases will undergo multiple prototype and testing phases, including a publicly available beta. ● Will work alongside academia and industry to

[33]How Wired will use generative AI Tools | Wired, May 22, 2023, https://www.wired.com/about/generative-ai-policy/.

	exchange ideas and reinforce standards.[34] ● [This AI policy was released before ChatGPT. It provides an interesting example of a forward-looking strategy, but its relevance to the broader conversation about current AI approaches may be limited]
InformationWeek (US)	● Will not publish anything "primarily authored, edited, or otherwise created" by generative AI systems. ● Allow limited use of generative AI as a supplementary tool, but one to be approached with "extreme caution." ● "Any and all" use of generative AI will be disclosed to readers—appropriate passages will be highlighted, credited to the AI, and clarified. ● Reporters will require editorial permission to use AI in ANY way. ● Ask that sponsors and press releases abide by similar standards. ● Further policy changes or new AI use-cases will

[34]Uli Köppen, Jonas Bedford-Strohm, and Cécile Schneider, "Ethics of Artificial Intelligence: Our Ai Ethics Guidelines," BR24 NACHRICHTEN, April 11, 2022, https://www.br.de/extra/ai-automation-lab-english/ai-ethics100.html.

	be publicized and explained.[35]

Proposed Frameworks:

In addition to official newsroom policies, I have also collected a representative list of guidelines that have been proposed by industry groups and academic organizations in the wake of ChatGPT's release. They serve as frameworks for potential generative AI policies and describe some of the questions that AI strategists should consider. Again, this list is not exhaustive, but it does offer insight into how prominent media leaders are approaching industry-wide norm-creation.

(Note: some of these guidelines have been paraphrased or rephrased for clarity, while others are lifted verbatim from the organization websites. Furthermore, some of the guidelines are broader frameworks that, while created by journalism-industry trade groups, question the appropriate uses of generative AI across all industries.)

RTDNA (Radio Television Digital News Association)	● Newsrooms need a clear AI policy—can be built through a continuous process of skeptical review. ● AI as a tool to assist humans—NOT a replacement for human judgment and thinking. ● 3 pillars of norms: ○ Accuracy, context and clarity ■ Ensure that AI text is reviewed by a human and that text is both accurate

[35]Sara Peters, "Our Editorial Policy on Journalists' Use of Generative AI, Chatgpt," InformationWeek, June 1, 2023, https://www.informationweek.com/big-data/generative-ai-in-the-newsroom-our-policy.

	and original. ○ Transparency and disclosure ■ Consider the risks of not disclosing AI use—potential impacts to journalistic trust and legitimacy. ■ Create a reliable process for the disclosure of AI-produced text across multiple platforms/mediums. ○ Privacy ■ Ensure that AI systems do not infringe on individual rights—make sure they satisfy existing journalistic norms regarding privacy.[36]
Birmingham City University Institute of Media and English	● Focus on diversity and equity: how can AI-generated news uphold existing journalistic standards for DEI? ● 6 principles: ○ Be aware of built-in bias. ○ Be transparent where appropriate. ○ Build diversity into your prompts. ○ Recognize the importance of source material

[36]"Use of Artificial Intelligence (AI) in Journalism," Use of AI in Journalism - Radio Television Digital News Association, accessed July 20, 2023, https://www.rtdna.org/use-of-ai-in-journalism.

	and referencing. ○ Report mistakes and biases. ○ AI-produced text should be viewed with journalistic skepticism. ● Expresses the view that AI should be a "tool" rather than a replacement for journalists. ● LLMs may be trained on biased or inaccurate material—must be vigilant to prevent bias from carrying over into journalistic publications.[37]
Digital Content Next	● Seven principles, split into three sections: ● Intellectual property: focus on the rights that media companies have to protect their content from being scraped by AI companies and turned into training data for LLMs. ○ Developers and deployers of generative AI must respect creator's rights to their content. ○ Publishers are entitled to negotiate for and receive fair compensation for use of their IP. ○ Copyright laws protect content creators from the unlicensed use of their content.

[37]"Generative AI Diversity Guidelines," Birmingham City University, accessed July 20, 2023, https://www.bcu.ac.uk/media/research/sir-lenny-henry-centre-for-media-diversity/blog/six-principles-for-responsible-journalistic-use-of-generative-ai-and-diversity-and-inclusion.

	● Transparency: focus on proper attribution of material and disclosure of AI-produced text. ○ Generative AI systems should be transparent to publishers and readers. ● Accountability: particularly focused on the risk that generative AI might attribute false information to a publisher. ○ Deployers of generative AI technology should be held accountable for system outputs. ● Fairness: focus on regulation of generative AI development ○ Generative AI systems should not create, or risk creating, unfair market or competition outcomes. ● Safety: focus on preventing systems from replicating biases or from being trained on personal information. ○ Generative AI systems should be safe and address privacy risks.[38]
Generative AI in the Newsroom	○ This comparative report of AI policies

[38]Dcn, "DCN's Principles for Development and Governance of Generative AI," Digital Content Next, June 6, 2023, https://digitalcontentnext.org/blog/2023/06/05/dcns-principles-for-development-and-governance-of-generative-ai/.

Project	across different publications identified the most frequently applied regulatory frameworks and policies, and the most widely expressed concerns. ○ **Most prominent concepts:** ● Oversight ○ Human editorial supervision and fact-checking of AI-generated text. ● Transparency ○ Some form of disclosure when an article is using AI-generated text. ● Banned vs. Allowed Uses ○ Which use-cases are definitely prohibited, and which are allowed (as long as they conform with the other guidelines). ● Accountability and Responsibility ○ Publishers are responsible for what they publish, even if it is written by an AI bot. Reinforces qualities of "accuracy, fairness, originality, and transparency." ● Privacy and Confidentiality ○ Caution when providing sensitive information to third-party generative AI

	platforms. ● Cautious Experimentation ○ Balancing risks of AI (particularly misinformation) with potential productivity innovations. ● Strategic Intention of Use ○ Having a particular goal for AI integration—and having that integration stay close to organizational values. ● Training ○ Creating internal education systems that ensure responsible use of AI—includes responses to biases found in datasets. ● Adaptability of Guidelines ○ Recognizing that generative AI is rapidly evolving, and that guidelines will need to accommodate technological change. ○ **Less prominent concepts:** ● Supply Chain ○ Being careful which AI companies a publication works with—potential recourse includes building an in-house LLM that fits publication values.

	● Legal compliance ○ Only using AI in a way that complies with existing law. ● Personalization ○ Implementing AI-powered content personalization in a responsible way. ● User feedback ○ Push for user feedback to guarantee transparency and accountability.[39]
News/Media Alliance	○ 10 principles, split into 5 sections: ● Intellectual Property: suggests that "[Generative AI] developers and deployers should not use publisher IP without permission, and publishers should have the right to negotiate for fair compensation for use of their IP by these developers." ○ Developers and deployers of generative AI must respect creators' rights to their content. ○ Use of publishers' IP requires explicit permission. ○ Compensation agreements must account for

[39]Hannes Cools, "Towards Guidelines for Guidelines on the Use of Generative AI in Newsrooms," Medium, July 10, 2023, https://generative-ai-newsroom.com/towards-guidelines-for-guidelines-on-the-use-of-generative-ai-in-newsrooms-55b0c2c1d960.

	harm generative AI systems may cause publishers and the public. ○ Copyright laws must protect, not harm, content creators. ○ There is an existing market for licensing publishers' news content. ■ Here, the guidelines suggest incorporating existing compensation frameworks like licensing to the debate over LLM training data. ● Transparency ○ Generative AI systems should be transparent to publishers and users. ● Accountability ○ Deployers of Generative AI systems should be held accountable for system outputs. ● Fairness ○ Generative AI systems should not create, or risk creating, unfair market or competition outcomes. ● Safety ○ Generative AI systems should be safe and avoid privacy risks.

	■ Notes that these systems collect significant amounts of user data.[40]

[40] "News/Media Alliance AI Principles," News/Media Alliance, April 20, 2023, https://www.newsmediaalliance.org/ai-principles/.

Research Methods

My study followed two main research pathways:

First, I produced a general literature review of academic sources that explored the relationship between generative AI tools and the journalism industry. This review ultimately incorporated relevant sources from the broader field of communication studies. I tried to focus on recent scholarship that more accurately reflects the current state of generative AI, although this was often a challenge—publicly-available generative AI has grown exponentially more powerful over the past few years and academic research in the social sciences has not caught up. In addition to using academic sources, I constructed a database of internet articles that reported on developments in GPT, commentaries on uses of generative AI in newsrooms, and news media press releases detailing generative AI policies.

Second, I conducted a series of interviews with leading journalists who have experience covering or commenting on issues related to artificial intelligence and large language models. In all, I spoke to 21 journalists and academics. The journalists had written for credible news media publications; the academics were leaders in the field of journalism studies. For the purposes of this study, the names of the interviewees and their employers remain anonymous, and no direct quotes are used to identify the journalists or their comments. Each journalist consented to their inclusion in the study—the interviews ultimately served to frame my conceptual arguments and to highlight the different ways in which practitioners in the field are considering the ethical and social implications of generative AI technology. I did not rely on the interviews for factual information regarding specific publications or newsrooms. Instead, I depended on the information collected during the first half of the research to illustrate particular examples.

Findings

Across the 21 interviews, I identified seven recurring areas of concern. I did my best to compile these into a coherent conceptual outline that describes the major arguments and points of contention that surround these issues. This approach, I hope, can help illustrate current industry discourse: what journalists are concerned about, where the main schools of thought are emerging, and where there is room for further research.

Appropriate generative AI use-cases

Most participants in this study saw generative AI as more of a "tool" than a replacement for the basic work of journalism. This "tool" rhetoric—AI as an enhancement to human productivity but not a replacement for human workers—evoked some of the language used by publications that have pursued aggressive AI implementation policies. But "tool-rather-than-replacement" thinking has also gained traction with the AI-skeptical journalists I spoke to, who believe that generative AI's range of viable use-cases is currently too limited to go beyond moderate research assistance and small efficiency gains.

Many of the journalists who saw AI as a research and efficiency tool claimed to use generative AI in their day-to-day work, variously, to brainstorm article ideas and SEO-ready headlines; to summarize research literatures and scan potential sources; to structure and transcribe interviews. Some were particularly optimistic about generative AI's potential to reduce time spent on busy work—potentially opening up more opportunities for research and real-time reporting. Others suggested that generative AI's ability to create SEO-ready headlines would be helpful for smaller publications that depend on online traffic for revenue. Because of these

productivity benefits, a few participants suggested that personal use of AI tools for work assistance was likely to be more widespread than previously reported—so much so that they thought it was fair to assume that almost every journalist uses AI, at least in some capacity.

This was not a majority view. Many believed that few, if any, journalists are taking advantage of AI tools. These journalists had a narrower assessment of generative AI's usefulness: the technology, in their eyes, is more trouble than it is worth. Even with the potential to generate ideas—and convert them into writing—more quickly than a human, the risk of potentially plagiarizing another journalist's work or of reporting "hallucinated" information is unacceptably high. The amount of time needed to fact-check information produced by an AI during the course of research offsets the efficiency gains, making the tools worthless. A few of these "hyper-skeptical" journalists believed that after some (inevitable) technical improvements, AI tools could eventually have a legitimate place in the newsroom: it was a matter of time and quality rather than form. But others thought that AI processes are so distant from the human reasoning and skepticism foundational to journalism that generative AI is unlikely to take hold in newsrooms anytime soon.

Both those who saw AI as a "tool" and those who saw widespread AI integration as unlikely were generally unconvinced that more diverse use-cases (beyond writing and research) would emerge as AI tools became more powerful. Rather, most believed that existing use-cases would become more reliable as AI writing became more persuasive and less prone to hallucination. A few suggested that increasing model complexity might expand the range of possible use-cases, although it was difficult to specify in advance the ways that it might do so.

There was wide consensus that certain parts of the journalism process—interviews, relationships with sources, and other higher-level work—would be hard or impossible to

automate. A far-future scenario where AI systems can convincingly do these things would be difficult to recognize given the current state of AI technology—and such a future would likely contain a very different set of journalistic institutions.

The other major source of consensus was the belief that AI tools are not ready to write full-length articles—most thought that content seen by a reader should not be produced by AI. This is not a surprising conclusion from a group of journalists: first, because the most prominent early examples of LLM-generated articles were poorly written and filled with mistakes, and second, because AI systems that can plausibly write entire articles would pose a more existential threat to the future of human-led journalism—one that would have both personal (for the journalists) and ethical repercussions. Additionally, some mentioned that AI systems would struggle to replicate the unique voices offered by different publications—this could limit their usefulness, particularly at higher-level firms.

Similarly, some pointed out that for many journalists the process of *writing* articles is secondary to the reporting process—many news-style pieces are already produced in a highly structured house style that emphasizes clarity and efficiency. For those journalists focused on newsgathering and not analysis, writing simply isn't the hard part—the rewards of employing unreliable AI systems in these critical information-heavy sections of the industry would probably not outweigh the risks.

Journalists thinking about potential AI use-cases tended to articulate a distinction between "strong" and "soft" uses of AI. "Soft" uses include AI in the production process—usually as a research or productivity tool. There was debate over how useful or acceptable these "soft" uses were likely to become, but most journalists fell somewhere in the middle: wary of reliance on AI help but open to the possibility that it could be helpful. Strong

uses, on the other hand, anticipated the elimination of humans from the production equation. In these cases, AI is deployed to produce content that might have otherwise been written by a human. The journalists I spoke to universally opposed "strong" uses of AI—at least for now.

Risks of AI adoption

Journalists suggested that the scandals at CNET and G/O Media had encouraged a stigma surrounding AI-written content—although many believed that this stigma would gradually be eroded as generative AI becomes more ubiquitous and proper oversight systems are developed.

The most commonly articulated risk factors for AI integration were the possibilities of releasing plagiarized content and spreading misinformation.

Plagiarism: Journalists noted that integrity and accountability are threatened by the tendency for generative AI systems to incidentally plagiarize, sometimes word for word, the content they are trained on. If any content released by the system could be a direct (and thus, inappropriate) replication of training material, a publication might need to run AI content through a plagiarism detector in order to prevent uncomfortable accidents.

Misinformation: The other major risk that journalists commonly cited was "hallucinations": the tendency of generative AI systems to produce statements that are simply untrue or unfounded. These incidents have been widely publicized, and many believed they pose a substantial roadblock to any serious use of generative AI systems in research and writing contexts. Any AI content would need to be systematically fact checked in order to confirm its veracity. Some were less convinced than others that "hallucinations" limited acceptable AI use-cases. These

participants noted that systems were getting better at reducing the number of falsehoods in their output—and, after all, humans make factual errors too. At some point, they argued, AI technology will be close enough to a human in this regard that the distinction will no longer be a relevant barrier.

The last major concern brought up by participants was also the most divisive: the prospects for automation of newsroom jobs. Many expressed concern that AI integration would not be limited to "tool" use-cases that increase individual efficiency, but instead that AI would be deployed in order to maintain production levels while facilitating substantial layoffs. These journalists pointed out that the news media industry was already in the midst of a significant wave of layoffs in the middle of 2023. While this current wave is not directly attributed to AI, participants believed that companies would become more comfortable laying off employees in the future with the knowledge that generative AI could enable productivity gains. Some referenced earlier waves of automation that eliminated industry roles—although the labor transformation brought on by generative AI, they suggested, could be the most transformational one yet.

Those who believed that some degree of automation was likely in the near-future warned that replacing human journalism jobs could have substantial structural consequences. Premature automation without sufficient consideration about the need for editorial oversight was one scenario. Some were concerned that even partial automation could permanently eliminate opportunities for entry-level journalists who often begin in the industry by writing generic low-impact stories (the "boring" stories that AI systems have already been deployed to write).

Layoffs, some argued, are likely to target junior staff first—cutting off the supply of future editorial leaders.

Others, however, were skeptical of arguments claiming that automation was imminent. They thought of automation as a normal part of technological change and adaptation. The result was likely to be a "skills transition" that places a greater emphasis on technical expertise and familiarity with AI—not mass layoffs.

Some of these "automation skeptics" believed that partial automation was possible at lower-quality outlets, but that certain kinds of high-quality work would be insulated from this transition. This work includes investigative journalism, interview-based work, and longer-form narrative journalism—all of which require the cultivation of relationships and navigation of complex human dynamics, which are difficult tasks for an AI.

"Automaton skeptics" also tended to believe that worries about automation had been exacerbated by alarmist coverage that reinforced fears of an imminent transition. Such fearmongering, some thought, could bring unintended consequences: journalists and editorial staff might be more willing to believe that they have no bargaining power in negotiations with their employers.

Factors impacting integration

Participants outlined a number of factors that they believed would contribute to the adoption of different AI strategies across the industry. Five of these factors were particularly prominent and came up often across different conversations.

1. *Competition*

Participants believed that most publications are playing a "waiting game" in order to see what kinds of policies are being implemented by other companies. No publication wants to pursue an aggressive strategy, only for it to fail. By drawing on lessons from other firms, these publications balance interest in AI adoption with caution. Still, this is not the same as passivity. Many believed that the environment of AI policy was very competitive, with companies rushing to experiment with the most accurate AI systems, discover which use-cases are the most viable, and deploy them in a responsible manner.

Asked about sources of organizational influence, journalists suggested that editorial teams are primarily interested in looking at the decisions made by other news organizations. For editorial staff, the normative structure of journalism makes influence from outside the industry less relevant. Some, however, felt that the real pressure—when it came—would come from the corporate side and draw lessons in downsizing and efficiency from outside the industry.

2. *Business vs. Editorial*

Business-side decisionmaking was widely viewed as the driving factor behind certain AI strategies. Participants believed that the most aggressive implementation policies (Buzzfeed, CNET, G/O Media, etc…) had been driven by business teams rather than by editorial staff. There were no journalistic grounds for any of these policies, since they bypassed the editorial pipeline and did not help journalists produce higher-quality work. The only AI writing use-cases met with significant editorial support were for the "boring" stories that do not require journalistic creativity or skepticism, such as financial earnings updates or sports results tallies that fit a highly determined structure.

Participants believed that the majority of decisions surrounding AI policy would fall on business leadership teams. For many, this was a serious issue, because they believed that business teams are less likely to respect journalistic norms.

A few participants expressed optimism that AI-related business decisions could still be constrained by the influence of editorial staff—for example, that staff writers would not be forced to use certain AI tools in their work if they chose not to. Some believed that the actors and screenwriters guild strikes could provide some framework for union bargaining programs that took AI into consideration. Others were more cynical about these prospects. They believed that journalistic unions are too fragile to mount a sustained defense—news departments could risk being completely shut down by publications. While some unions have effectively lobbied for benefits and severance packages, their ability to push back on industry-wide change may be limited, especially if reporters are seen as replaceable.

3. *Economics*

Journalists tended to believe that larger, more financially stable publications are less likely to adopt aggressive AI policies. Many of these publications are heavily invested in their reputation and credibility, which a rushed AI rollout could put at risk. Some suggested that the presence of an all-human reporting workforce could have prestige value that might be leveraged against rival publications who have implemented more aggressive AI programs. However, credibility also goes in both directions. The AP, which holds a revered position among journalists, introduced responsible AI tools into the mainstream. Other big companies might follow and take advantage of their resources to experiment more extensively with AI before integrating it into their production process. Once an acceptable use-case has been identified, a

publication's reputation for high-quality journalism could deflect backlash. Still, some pointed out that the bigger companies may not have a strong financial interest in generative AI at all, since they tend to operate on subscriptions and are not as reliant on search traffic revenue.

At the other end, journalists mentioned that smaller publications would likely have very strong incentives to pursue AI adoption, especially if they thought that doing so would boost their efficiency. For many small traffic-focused publications, conditions are so dire that any cost-cutting or revenue-boosting measure appears very attractive, even if it carries substantial risk. With production goals steadily rising over the past few years, some believed that AI integration could be a matter of financial survival. Others were less convinced. The risks might be too high, they believed, and the benefits too unproven. A poor rollout of AI could be the end of a struggling newsroom, and many publications would not be willing to take that risk—especially since industry-wide best practices have yet to be established. Without the time or money to internally experiment with AI, small-newsroom adoption could be doomed to fail. This perception was also a source of debate. Some believed that smaller publications could actually produce stronger AI strategies, because they could be more flexible: with smaller teams and less institutional baggage, there might be fewer roadblocks to experimentation and integration.

4. *Publication Genre*

Participants generally believed that publications focused on a broader range of topics may be more likely to pursue aggressive AI integration, because there was more room for coverage and expansion. Especially as production targets increase, tapping the potential for a wider story base could make a lot of sense for such firms. Publications that were more specialized,

analytical, or occupied a different format (magazines, opinions, etc…) were perceived as less likely to pursue AI. For these publications, AI adoption could be harder to reconcile with company values and might not mesh with established story formats.

5. *Corporate Attitudes*

There was also some discussion about more contingent factors that could have an outsized influence on a company's AI policy. The most notable of these factors was the attitudes and interests of corporate executives. An executive who is prepared to take risks with new technology could steer an entire publication towards a more aggressive AI policy—one who is more conservative and concerned about their reputation might hold off to see where industry standards land. These personal factors suggest a great deal of unpredictability: the larger battles of interests could be reduced to the decisions of a small group of individuals, which may or may not align with the beliefs of editorial teams or even business teams.

Different approaches to policy creation

Most participants described their own editorial environments as lacking clear AI strategies. Still, many journalists had developed a strong sense of the kinds of policy approaches that seem to have taken root at major publications. These approaches can be classified into three different "models" of AI adoption.

Model A: The Aggressive Model

This approach reflects the fast-moving policies of publications that have ventured directly into AI content generation, often without consent or collaboration from editorial staff. Often,

these AI strategies are steered by executives who are more interested in efficiency gains than in credibility or norms. At some publications, this has backfired, resulting in significant reputational loss. Justifications for this approach centered on the importance of adapting rapidly to technological innovation despite short-term implementation hurdles—evoking Facebook's famous "move fast and break things" motto. The most cited examples of this model were CNET, G/O Media, and Buzzfeed, although participants believed Model A publications represented a small share of the industry.

Model B: The Cautious Model

Participants believed that the majority of publications roughly fit this approach, which emphasizes due diligence and caution above all else. This reflects uncertainty and unfamiliarity with AI. It is also a product of journalism's tradition of intellectual skepticism and normative standard-bearing (and the small-c conservatism of establishment media). Model B publications are not necessarily hostile towards AI adoption in the long run, but they do hold appropriate use-cases to a higher bar than do Model A firms. For these publications, AI integration should only be realized after a period of experimentation intended to identify use-cases that satisfy a publication's ethical standards and business interests. Because these publications tend to place a greater value on their credibility, the two are intertwined.

Cautious AI experimentation in the newsroom can—and does—take different forms. Responsibility might be assigned to a small "task force" or working group (with AI use discouraged for other staff) or it could be distributed across an entire publication, with staff given broad guidelines about how to personally experiment with the technology. Experimentation can also be broadly applied across different editorial domains or heavily focused on particular

use-cases—depending on a company's interest in AI adoption and familiarity with the technology, breadth of experimentation could be variable. Regardless of the format, an experimental phase would result in the accumulation of feedback that could be used to construct a more durable AI policy that takes editorial needs into account.

While most participants said they preferred this model, some noted that a lot of the policies produced by a "middle ground" strategy seemed relatively diffuse or unorganized, plagued by vague language or gray areas that create more questions than answers.

Model C: The Avoidant Model

The final approach with some currency among journalists is an outright ban on all journalistic uses of AI. Bans were supported by those who were most skeptical of AI's newsroom viability and most concerned about the potential risks of adoption. But this minority position was controversial and viewed by some as shortsighted and unrealistic. At some point, these journalists argued, generative AI will undeniably reap productivity benefits and make pledging not to use it unviable.

The case for rejecting AI seemed strongest for well-established publications that make most of their revenue from paid subscriptions, not ad revenue. For these firms, the risks of public AI rollout were perceived to be large and the benefits small. But this raises the question of the boundary between Models B and C. Many of these big-name publications are still experimenting with AI in-house—they can afford to do this without committing to a particular AI strategy because they are funded by a readership more interested in quality than in quantity. This is different from Model C companies who have made a commitment to ignore AI innovation and have elaborated a policy that formalizes this choice.

Beyond these three models of integration, journalists also brought up issues that were less central to the AI adoption process, although still tangential. One was the question of in-house vs third-party LLM use. CNET claims to have developed a custom in-house model to write their AI stories—while the resulting articles were bland and riddled with errors, the endeavor did show that media companies (or their parent organizations) are capable of building LLM systems strong enough to produce coherent prose. If access to LLM technology continues to be made cheaper and more accessible, other media companies may seek to produce LLM systems fine-tuned for journalistic use-cases. Such efforts would probably be out of financial and organizational reach for the vast majority of publications (for now), but some journalists suggested that industry-wide models (focused on news media content formats) could, in theory, be developed through the collaboration of interested parties.

One driver of in-house solutions is likely to be the need for protection of sensitive information. Chatbots like ChatGPT record all of the content uploaded into their servers—that content is, by default, integrated into the bot's training processes. Because of journalism's reliance on source confidentiality and potentially sensitive, private, or privileged information, participants pointed out that this could be an important roadblock to adoption in certain contexts—similar concerns had impacted usage of AI transcription services. However, on some platforms, content disclosures can be avoided through an opt-out form.

Emerging norms

Many journalists expressed concerns about the potential for generative AI adoption to violate traditional journalistic norms and ethics. The degree of apprehension loosely tracked

opinions on the speed and scope of AI integration: those who did not see a viable use of generative AI in the near future were less concerned. Nonetheless there was broad agreement that strong norms regulating the acceptable use-cases of AI will be essential in the years to come. While most journalists expressed a great deal of uncertainty about the norms-making process, fairly consistent views emerged about the norms themselves. Five came up most frequently in these conversations:

1. *Transparency*

Those who valued transparency believed that publishers have an obligation to disclose to readers any content written by an LLM system. The disclosure needs to be unambiguous—some suggested highlighting certain passages in a different color or displaying a "robot" as the author of AI-produced articles. Transparency was often cited as one of the fundamental conventions of the journalistic process: the explicit and public knowledge of who is behind the news, any conflicts of interest they might have, and their potential biases. Large-scale adoption of AI, participants believed, would complicate this greatly by introducing new questions about the reliability of content and the potential biases found in datasets (racial, gender-related, political, corporate, etc…). Publications outed for using AI systems without disclosure could undermine the public's trust in journalism in general—a trust that has already experienced a precipitous decline over the past few years.

A few participants preferred a higher standard for transparency in which publishers not only have an obligation to disclose content that was *written* by an AI, but any content which may have been influenced or supported by the use of AI tools.

Others thought that the format of disclosure also held ethical implications: publications could have additional obligations to justify, on the basis of journalistic ethics, why any particular use of AI contributed to an article's substantive value. This obligation might operate as a curb on purely commercial or gratuitous uses of generative AI by creating a source of accountability. It would also explicitly reintroduce journalistic norms and ethics into the production of AI-written and AI-supported articles. Many of interviewees felt that the "hype" surrounding AI integration had pushed important ethical concerns out of the spotlight.

Finally, some expressed concern that discourse surrounding AI transparency had become too focused on the relationship between publisher and reader, and had ignored transparency issues that may emerge between employers and employees within publications. Generative AI tools are widely available. Journalists could use tools in ways that violate company policy and are difficult to identify—particularly when writers are freelancers. This dilemma may require greater degrees of trust, not only between news media and its audience, but between news media corporations and their employees.

2. *Human oversight of AI-produced content*

Journalists also identified human oversight of all AI-produced content as a crucial component of any AI strategy. While some journalists preferred a policy that bans AI content generation altogether, all participants agreed that content produced directly by an AI chatbot should be edited and reviewed by humans. The reasoning behind this stance was fairly uniform: oversight from humans can prevent AI systems from releasing misinformation and/or plagiarized content, both of which would seriously threaten a publication's credibility.

Here, again, the CNET and G/O Media scandals served as important examples for several of the respondents. Neither company took human oversight seriously enough to prevent inaccurate information from being released on their platforms. Ultimately, the question returns to cost-benefit analysis: heightened human oversight versus increased publication efficiency. Different publications may come up with different answers. But on a normative level, human oversight is not a controversial issue: it is a necessary standard for any serious AI policy.

3. *Skepticism of AI 'facts'*

This norm is substantially related to the last one, although it involves a different set of generative AI use-cases. Many journalists expressed strong mistrust of information provided by chatbots during AI-assisted research. While generative AI can often be helpful in broad strokes, respondents argued that it should not be treated as a reliable research database. The analogy most often evoked was Wikipedia: a resource that is most useful for general background research and as a directory towards more reliable sources. However, some pointed out that systems like ChatGPT can be even less useful than Wikipedia for research, since they are more likely to produce fake or incorrect citations. In some contexts, this verification process wastes more time and energy than the AI assistance is worth.

The practical "norm" to be taken from this is that publications should encourage their journalists to take AI information with a grain of salt. They should encourage and reinforce skepticism about the quality of information released by generative AI systems—this should be foundational to any systematic workplace interaction with the software.

Skepticism also means holding AI policies and integration methods accountable. Multiple journalists suggested that the work produced by chatbots at CNET was not only error prone, but

unacceptable by broader professional standards. Any journalist who published content on that level would have likely faced professional repercussions. Some participants believed that AI systems should be held to the same standard—if the work they produce is routinely unpublishable, then the policy surrounding their use needs to be changed.

4. *Institutional clarity on acceptable uses*

One of the more frustrating aspects of the CNET and G/O Media scandals was the lack of communication between the business-facing sides of the publications (which developed and released the AI-written stories in both cases) and editorial staff. A deputy editor for io9 (a Gizmodo subsidiary) learned about the impending AI stories just minutes before they were published.[41] CNET, on the other hand, was quietly publishing AI-written stories for months without alerting their editorial team.[42] Writers and editors at both publications were outraged upon learning about what had happened.

The two incidents demonstrate the importance of internal policy clarity and decision-making accountability. AI integration has a tendency to fuel tension within publications and, so far, that sense of distrust has contributed to poor AI strategies that have hurt everyone—both editorial and business.

Journalists also suggested that clarity has a broader value in creating an efficient and well-organized AI-supported production process. If the guidelines for appropriate AI use are well-established, journalists will feel more comfortable using AI within the acceptable parameters. If the guidelines are generally responsible, this should result in positive productivity

[41] Davis, "Gizmodo's Staff Isn't Happy about G/O Media's AI-Generated Content."
[42] Roth and Sato, "CNET Found Errors in More than Half of Its AI-Written Stories."

outcomes for a publication. If guidelines are unclear or too vague, however, there may be a chilling effect on AI use.

5. *Policy Maneuverability*

Perhaps the most important and predictable takeaway from this work is that most people in the industry do not have a good sense of where things are going. In this context, policy flexibility will remain important. New technologies may be released which amplify AI's role in the industry, making existing internal regulations obsolete. Or public backlash against new technologies could encourage companies to eliminate certain AI use-cases. These things can happen fast, like the release of ChatGPT.

Journalists are receptive to this sense of uncertainty—even if it worries many of them. They believe that newsrooms should be receptive, too. The right attitudes and foresight could help make publications more resilient and more ethically responsive to changes in the AI landscape. This attitude might be formalized in policies that are built to be changed if and when circumstances change—but it could also remain an unwritten but important addition to company culture regarding AI use.

Interaction with technology giants

Participants expressed a great deal of distrust towards technology giants like Google, Microsoft, and Meta. Even those who were more sympathetic admitted that big tech's embrace of generative AI has created a whirlwind of ethical and social issues, including some with important consequences for journalism. Some expressed frustration with the web scraping of journalistic work used to build LLM systems—many participants believed that AI companies had infringed

on their copyright. Yet journalists also seemed to acknowledge that tensions between online publications and big tech were now being complicated by new kinds of competition and mutual interdependence.

One of these new complications is the gradual integration of LLM content generation into search engines—an innovation largely led by Google's new SGE (Search Generative Experience) software. SGE uses generative AI to produce paragraph-long responses to search queries by synthesizing information from top search results. While relevant websites will be cited by the SGE system, the entire process could theoretically replace a lot of traffic that is currently directed to online media sites. Journalists worried that this could have substantial financial implications for smaller, content-heavy publications that depend on online advertisement revenue. If traffic is negatively affected, revenue will be too.

Some suggested that widespread adoption of SGE could force a market transition towards the subscription-based revenue models that are more popular among well-established outlets. Yet several noted that news media search traffic has been on the decline for years: SGE did not start this trend. Regardless of SGE, search-based revenue models may become unviable due to market and social factors unrelated to AI adoption.

Others, however, viewed search-based revenue models as more resilient. They pointed out that previous innovations (Google/Apple News, Google's article previews, etc…) had failed to eliminate the viability of search, and they were skeptical of SGE's transformational potential. These "SGE skeptics" made five major arguments:

First, they suggested that the consumers who regularly consume online news are unlikely to be driven away by SGE systems. Websites can provide live updates and compelling

longer-form content—SGE can only offer a rudimentary sketch of current news trends. This makes it appealing in certain fringe cases (limited time, limited interest, etc…), but will not put a dent into consistent readership that values and understands the unique benefits of well-organized news media sites.

Second, the inclusion of website links in SGE might have a positive effect on traffic—driving clicks to the websites that provide newsworthy reporting.

Third, existing disruptions to the search-revenue model have not entirely eliminated its viability. Some of these disruptions, like Google/Apple News, are superficially similar to SGE, but have had little impact on traffic to websites. Even if general search traffic continues a longer-term structural decline, it is unlikely that SGE will substantially speed up the current downwards trajectory.

Fourth, some journalists believed that Google's Bard, which powers that company's version of the SGE platform, is not powerful or sophisticated enough to run a compelling news summarization program. Microsoft's Bing was also described as insufficiently popular for widespread adoption.

Fifth, participants suggested that even if SGE becomes technically viable, it might still be a less attractive option for consumers than highly personalized news websites built using generative AI. SGE (as it exists now) provides a generalized product, while personalized news media will be able to appeal to the interests of particular individuals, creating incentives for interaction with the websites themselves.

Other journalists hoped that SGE might spark larger legal and regulatory challenges to corporate control of generative AI. The size and scope of this AI use-case (all of online search)

could provoke concern about the integration of AI in media information systems. It could also motivate greater skepticism of vertical corporate AI strategies, and, some hoped, support antitrust actions against big tech.

An interesting development in this realm came in July 2023—the AP announced a deal with OpenAI to license their content for AI training.[43] This kind of collaboration could prove attractive for news media companies and may offer a solution to the question of monetization under SGE. It shows that publications have some degree of leverage over AI companies because they possess something of value: large corpuses of high-quality human writing. However, collaborations like these also create a host of questions:

First, if SGE is a financial risk to news companies, why should publications empower the tech corporations developing it by licensing their content? By providing SGE platforms the data they need to train their systems, news media companies might be subverting their own business models.

Second, if news media companies license their content to SGE platforms, they could develop conflicts of interest that affect their reporting on major tech companies. If news companies have a financial stake in tech companies' success, an important public check on the tech giants would be eliminated.

Beyond these issues, journalists also pointed out that SGE adoption creates several contradictions for the tech companies. If SGE models are fully deployed, they could threaten segments of the news media industry—or at least strongly disincentivize news organizations from participating in online search systems. With less high-quality news to draw on, SGE

[43] Matt O'Brien, "ChatGPT-Maker OpenAI Signs Deal with AP to License News Stories," AP News, July 14, 2023, https://apnews.com/article/openai-chatgpt-associated-press-ap-f86f84c5bcc2f3b98074b38521f5f75a.

systems will become less attractive. Some journalists also argued that SGE could have perverse commercial consequences for search platform companies like Google, which also makes money from the ads shown on the websites to which it directs users. SGE and advertising could become competing sources of income.

AI beyond the newsroom + Predictions

Even the journalists who worried less about AI adoption within newsrooms expressed concern over generative AI's potential to complicate established journalistic practices. Some suggested that the proliferation of believable AI-generated images and video content might force journalists to double check and corroborate the media they receive through confidential drops. This would create new barriers for watchdog reporting and whistleblower relationships.

On a broader level, all online information could become increasingly suspect, as the cost of producing and spreading disinformation drops precipitously. While a few thought that AI would not have a substantial impact on misinformation (the most common argument was that it is already fairly easy and inexpensive to propagate online disinformation), many believed that the information ecosystem could, and would, become more polluted with unverifiable images and news accounts. Some expressed the deeper fear that AI-powered misinformation could influence political and financial events (the 2024 election was mentioned), creating additional challenges for journalists seeking to accurately cover world affairs. Most believed that to respond effectively, journalists and news media organizations will need to adapt their norms and practices to address the increased unreliability of information. The positive outlook was that this unreliability would make the role of journalism, and particularly human journalists, more

valuable as arbiters of truth (although some were skeptical of journalism's ability to provide a meaningful check on misinformation).

The pessimistic view of journalism's AI future envisions a growing divide between publications dedicated to human-led reporting and those who have fully committed to AI-produced content. AI-heavy publications more interested in reaping the financial rewards of mass information may resign themselves to circulating misinformation freely. They would become more akin to online content farms that mass produce articles without serious editorial oversight, with the primary purpose of soliciting ad revenue.

On the other hand, companies that market themselves on their human "integrity" may expand their paywalls (as a premium for human labor), making their content less accessible to large audiences. Participants who predicted this future worried about the chaos it could wreck on public institutions and trust in journalism. There was agreement among some respondents that a situation like this was plausible.

Most participants did not expect a strong government response to newsroom AI adoption. In the U.S., regulation has been slow and vague—a belief in free-market innovation has stalled the more serious regulatory proposals. While regulation is being pursued more aggressively in Europe, the implications for journalism are still unclear. Some believed that European newsrooms would emerge relatively untouched from AI legislation. Others were more critical of regulatory approaches in general: they believed that any regulation would either be too vague to be effective or counterproductive. Across the board, there was very little faith in governments' ability to enact strong AI regulation. Some pointed out that the majority of journalistic norms emerge from moral and practical deliberations (e.g. protecting confidential source identity), or legal compliance (e.g. copyright). New regulations or legislative proposals rarely have

substantial impacts for news media. In that vein, some participants felt that legal disputes over copyright and fair use related to LLM content could have more serious implications for newsroom policy. Of these, regulation of SGE technologies was seen as one of the most likely regulatory outcomes—and one that could give news publishers more leverage in negotiations with Big Tech.

Journalist literacy + Coverage of AI

Most of the respondents thought that AI literacy and competency was a serious problem in journalism. AI, they argued, will transform all aspects of social, political, and economic life. Journalists, as guardians of the public trust, need to understand how these systems work and where their limits are. There was a broad consensus that increased training, supported by newsrooms but also by journalism schools and industry groups, could help bridge the gap between AI-literate tech journalists and general-purpose journalists with less experience in the field. This training would have a dual purpose: not only might it make journalistic reporting on AI more accurate and responsible, but it would also help facilitate stronger discourse on AI norms and integration policies. Participants suggested that a journalism workforce that better understands AI will also be better equipped to implement it in their work.

Many of the respondents believed that this "literacy problem" had negatively impacted past coverage of LLM systems and generative AI as a whole. Publications jumped on the "hype" surrounding generative AI and published stories that were often ill-informed or sensational. Some journalists claimed that the publications that had devoted significant coverage to other tech hype cycles, like the ones surrounding cryptocurrency and Web 3.0, had been the first to produce exaggerated accounts of generative AI after ChatGPT's release.

Some of these publications had fueled discourse about AI's "existential risk" potential—the possibility that advancements in AI technologies could eventually lead to human extinction. Participants were split on this reporting trend. Some thought that existential risk stories had been exaggerated by cultural factors (robots turning on mankind are a cliche of science fiction) and researchers who had built a brand on AI alarmism. Others thought that coverage had struck an important balance between criticism of big tech narratives and journalistic skepticism.

Participants had other qualms about the state of AI coverage. A common complaint was that articles had anthropomorphized AI and implied that generative AI systems had gained sentience—or were close to it. This trend was suggested to reflect the current hype cycle surrounding generative AI systems in particular. After ChatGPT's release, coverage of AI skyrocketed, and the cultural conversation surrounding the systems too often reflected sensationalized accounts of AI behavior instead of technical facts about system limitations. Some pointed out that this drive may be encouraged by news media executives who are convinced of AI's transformative potential: broader ideological currents within the industry may account for the current shape of coverage. One proposed solution involved eliminating images that portrayed generative AI systems like ChatGPT as humanoid robots. Robotic images, participants argued, support the unfounded idea that ChatGPT has a degree of autonomy and individuality.

Conclusion

There is a great deal of disagreement among journalists about AI adoption in the news industry. This is unsurprising given the speed of AI innovation and the complexity of crafting AI policy—yet it still speaks to a fundamental need for industry-wide stabilization and effective norm-setting.

Despite this environment of uncertainty, certain patterns are taking hold. Some norms are becoming more entrenched than others, and those norms are often being paired with an integration approach that emphasizes caution and experimentation. Journalists are in agreement that policies should emphasize transparency, skepticism, flexibility, clarity, and human oversight.

External factors will play a substantial role in newsroom AI policy. New technologies like Google's SGE may reconfigure incentives, making existing frameworks of AI integration obsolete. The contingent attitudes of particular executives and business teams could have an outsize influence on the industry as a whole.

Many of the journalists I spoke to were reserved in their descriptions of AI integration—and critical of apocalyptic narratives that warn of widespread automation or the collapse of informational reliability. They had anxieties about the future, but also an optimism that the ethical and normative standards of journalistic inquiry—skepticism, objectivity, and the pursuit of public accountability—would overcome this disruption, as they had in the past.

In the end, the most pronounced divide was between journalism's optimists and pessimists. Both sides had their forms of skepticism: the optimists questioned grand narratives and dystopian prophecies of automation; the pessimists questioned corporate intentions and use-case viability. There was a great deal of overlap, and most found themselves somewhere in the middle. But the divide was there, and it felt like a good symbol of our collective struggle to

understand what AI's place in our social institutions should be, and how we best ought to accommodate it.

Acknowledgements

I would like to express my sincere gratitude towards the Laidlaw Scholars Foundation for supporting and funding my research this summer—the organization offered crucial resources and training sessions that contributed greatly to the shape and structure of this project. In particular, I would like to thank Dean Ariella Lang and Dr. Lisa Del Sol for supporting the work of all the Laidlaw Scholars at Columbia and for running such a well-organized program. I would also like to thank my faculty advisor Professor Chris Wiggins for providing me with indispensable guidance and research support, and my graduate advisor Eun Ji Sally Son for her methodological advice and her thoughtful and supportive mentorship.

Most of all, thank you to all the journalists and academics who took time out of their busy schedules to speak to me. Their insights and perspectives form the core of this project—without them, none of this would have been possible.

Bibliography

"Artificial Intelligence: AP." Associated Press. Accessed July 20, 2023. https://www.ap.org/discover/artificial-intelligence#:~:text=News%20production&text=This%20ranges%20from%20the%20automatic,both%20sports%20and%20corporate%20earnings.

Beckett, Charlie. Rep. *JournalismAI Report*. Polis: Journalism and Society, November 18, 2019. https://blogs.lse.ac.uk/polis/2019/11/18/new-powers-new-responsibilities/.

Bottomly, Therese. "Letter from the Editor: How AI and Automated Content Power Information on Oregonlive." oregonlive, July 10, 2023. https://www.oregonlive.com/opinion/2023/07/letter-from-the-editor-how-ai-and-automated-content-power-information-on-oregonlive.html.

Carlson, Nicholas. "My Editor's Note to the Newsroom on AI: Let's Think of It like a 'Bicycle of the Mind.'" Business Insider, April 13, 2023. https://www.businessinsider.com/how-insider-newsroom-will-use-ai-2023-4.

Chow, Andrew R. "Why CHATGPT Is the Fastest Growing Web Platform Ever." Time, February 8, 2023. https://time.com/6253615/chatgpt-fastest-growing/.

Cools, Hannes. "Towards Guidelines for Guidelines on the Use of Generative AI in Newsrooms." Medium, July 10, 2023. https://generative-ai-newsroom.com/towards-guidelines-for-guidelines-on-the-use-of-generative-ai-in-newsrooms-55b0c2c1d960.

Davis, Wes. "Gizmodo's Staff Isn't Happy about G/O Media's AI-Generated Content." The Verge, July 8, 2023. https://www.theverge.com/2023/7/8/23788162/gizmodo-g-o-media-ai-generated-articles-star-wars.

Dcn. "DCN's Principles for Development and Governance of Generative AI." Digital Content Next, June 6, 2023. https://digitalcontentnext.org/blog/2023/06/05/dcns-principles-for-development-and-governance-of-generative-ai/.

Diakopoulos, Nicholas. *Automating the news: How algorithms are rewriting the media*. Cambridge, MA: Harvard University Press, 2019. (107)

"Generative AI Diversity Guidelines." Birmingham City University. Accessed July 20, 2023. https://www.bcu.ac.uk/media/research/sir-lenny-henry-centre-for-media-diversity/blog/six-principles-for-responsible-journalistic-use-of-generative-ai-and-diversity-and-inclusion.

Harrison, Maggie. "BuzzFeed Says AI Will 'Replace the Majority of Static Content.'" Futurism, May 18, 2023. https://futurism.com/buzzfeed-ai-replace-content.

Heaven, Will Douglas. "GPT-4 Is Bigger and Better than Chatgpt-but Openai Won't Say Why." MIT Technology Review, March 23, 2023. https://www.technologyreview.com/2023/03/14/1069823/gpt-4-is-bigger-and-better-chatgpt-openai/.

Henriksson, Teemu. “New Survey Finds Half of Newsrooms Use Generative AI Tools; Only 20% Have Guidelines in Place - Wan-IFRA.” WAN, May 31, 2023. https://wan-ifra.org/2023/05/new-genai-survey/.

“How We Will Use Artificial Intelligence at CNET.” CNET. Accessed July 20, 2023. https://www.cnet.com/ai-policy/#:~:text=Writing%20full%20stories%3A%20None%20of,fact%2Dchecked%20by%20our%20editors.

How Wired will use generative AI Tools | Wired, May 22, 2023. https://www.wired.com/about/generative-ai-policy/.

Ingram, Mathew. “Is Ai Software a Partner for Journalism, or a Disaster?” Columbia Journalism Review. Accessed July 20, 2023. https://www.cjr.org/the_media_today/ai_software_chatgpt_journalism.php.

Khalaf, Roula. “Letter from the Editor on Generative AI and the FT.” Financial Times, May 26, 2023. https://www.ft.com/content/18337836-7c5f-42bd-a57a-24cdbd06ec51.

Köppen, Uli, Jonas Bedford-Strohm, and Cécile Schneider. “Ethics of Artificial Intelligence: Our Ai Ethics Guidelines.” BR24 NACHRICHTEN, April 11, 2022. https://www.br.de/extra/ai-automation-lab-english/ai-ethics100.html.

Newman, Nic. “Journalism, Media, and Technology Trends and Predictions 2023.” Reuters Institute for the Study of Journalism, January 10, 2023. https://reutersinstitute.politics.ox.ac.uk/journalism-media-and-technology-trends-and-predictions-2023#header--8.

“News/Media Alliance AI Principles.” News/Media Alliance, April 20, 2023. https://www.newsmediaalliance.org/ai-principles/.

Nuñez, Michael. “Letter from the Editor: How Generative AI Is Shaping the Future of Journalism and Our Newsroom.” VentureBeat, May 2, 2023. https://venturebeat.com/ai/letter-from-the-editor-how-generative-ai-is-shaping-the-future-of-journalism-and-our-newsroom/.

O’Brien, Matt. “ChatGPT-Maker OpenAI Signs Deal with AP to License News Stories.” AP News, July 14, 2023. https://apnews.com/article/openai-chatgpt-associated-press-ap-f86f84c5bcc2f3b98074b38521f5f75a.

Peters, Jay. “BuzzFeed Is Using AI to Write Seo-Bait Travel Guides.” The Verge, March 30, 2023. https://www.theverge.com/2023/3/30/23663206/buzzfeed-ai-travel-guides-buzzy.

Peters, Sara. “Our Editorial Policy on Journalists’ Use of Generative AI, Chatgpt.” InformationWeek, June 1, 2023. https://www.informationweek.com/big-data/generative-ai-in-the-newsroom-our-policy.

Rinehart, Aimee, and Ernest Kung. Publication. *Artificial Intelligence in Local News: A Survey of US Newsrooms’ AI Readiness*. Associated Press, 2022.

Sato, Mia. “CNET Is Overhauling Its AI Policy and Updating Past Stories.” The Verge, June 6, 2023. https://www.theverge.com/2023/6/6/23750761/cnet-ai-generated-stories-policy-update.

Sato, Mia, and Emma Roth. “CNET Found Errors in More than Half of Its AI-Written Stories.” The Verge, January 25, 2023. https://www.theverge.com/2023/1/25/23571082/cnet-ai-written-stories-errors-corrections-red-ventures.

Stern, Jacob. “GPT-4 Might Just Be a Bloated, Pointless Mess.” The Atlantic, April 25, 2023. https://www.theatlantic.com/technology/archive/2023/03/openai-gpt-4-parameters-power-debate/673290/.

“Use of Artificial Intelligence (AI) in Journalism.” Use of AI in Journalism - Radio Television Digital News Association. Accessed July 20, 2023. https://www.rtdna.org/use-of-ai-in-journalism.

Viner, Katharine, and Anna Bateson. “The Guardian’s Approach to Generative AI.” The Guardian, June 16, 2023. https://www.theguardian.com/help/insideguardian/2023/jun/16/the-guardians-approach-to-generative-ai.

Wiggers, Kyle. “The Emerging Types of Language Models and Why They Matter.” TechCrunch, April 28, 2022. https://techcrunch.com/2022/04/28/the-emerging-types-of-language-models-and-why-they-matter/?guccounter=1&guce_referrer=aHR0cHM6Ly93d3cuZ29vZ2xlLmNvbS8&guce_referrer_sig=AQAAAJ2haZlG9HkOb8uxapLpivpm4TeUlQR4FjtsGUJ6O74G38XUot9x_m09nF4OvAsZMjoFZpfISRJnQfNqiw85uuK2Q2Kh9yohCbdNeJW0bh76efU5y5-wlo1mLmSJunWNGlYKEbE1X2MpWqdmH5TQSAWhWY1EjBmLujPQ1V6aW8x0.