







AI use in American newspapers is widespread, uneven, and rarely disclosed

Jenna Russell  Marzena Karpinska  Destiny Akinode  James Zhou 
Katherine Thai  Bradley Emi  Max Spero  Mohit Iyyer 
 University of Maryland, College Park  Simon Fraser University  UMass Amherst
 Pangram Labs
{jennarus, miyyer}@umd.edu, karpinsk@sfu.ca
[ainewsaudit.github.io](https://github.com/ainewsaudit)

Abstract

AI is rapidly transforming journalism, but the extent of its use in published newspaper articles remains unclear. We address this gap by auditing a large-scale dataset of 186K articles from online editions of 1.5K American newspapers published in the summer of 2025. Using Pangram, a state-of-the-art AI detector, we discover that approximately 9% of newly-published articles are detected as either partially or fully AI-generated. This AI use is unevenly distributed, appearing more frequently in smaller, local outlets, in specific topics such as weather and technology, and within certain ownership groups. We also analyze 45K opinion pieces from *Washington Post*, *New York Times*, and *Wall Street Journal*, finding that they are 6.4 times more likely to contain AI-generated content than news articles from the same publications, with many AI-flagged op-eds authored by prominent public figures. Despite this prevalence, we find that AI use is rarely disclosed: a manual audit of 200 AI-flagged articles, 96.5% of authors and 94.0% of publishers did not disclose AI use. A factuality analysis shows articles flagged as AI-generated are 8.2 times more likely to contain hallucinated claims than human-labeled news. Overall, our audit highlights the immediate need for greater transparency and updated editorial standards regarding the use of AI in journalism to maintain public trust.

1 Introduction

How much AI-generated content is being published in newspapers across America? To answer this question, we conduct a large-scale audit of recently-published articles using Pangram (Emi and Spero, 2024), a high-precision AI detector that has previously been used to audit consumer reviews (Cavazos and Sterling, 2024), research papers (Evanko and Natale, 2025), and Medium articles (Knibbs, 2024). Our analysis reveals that ~9% of newly-

DISCLAIMER

All findings in this paper rely on the outputs of an automated AI detector and should be interpreted accordingly. We do not assert definitive authorship attributions for any article. We do not draw conclusions about the intent, conduct, or practices of individual journalists, outlets, or companies. Results should not be interpreted as rankings, qualitative judgments, or accusations of wrongdoing. See §6 for further discussion.

published U.S. newspaper articles are detected as either partially or fully AI-generated.

Why does this matter? Our audit is largely motivated by concerns of *transparency* and *factual-ity*. We do not claim that all AI use is inherently harmful; in fact, limited applications like grammar / style checking and template-driven article creation (e.g., weather reports) can improve article quality and accessibility (Medill Local News Initiative, 2024; Radcliffe, 2025). However, LLMs often hallucinate (Maynez et al., 2020a; Ji et al., 2023a; Su et al., 2024) and inherit social biases from their training data (Gallegos et al., 2024; Hu et al., 2025); thus, public opinion is highly sensitive to *undisclosed* AI use.¹ We manually analyze 200 articles flagged for AI use by Pangram and find that only 7 of them disclose AI use, while only 12 of the newspapers have any public policies on AI use, leaving readers largely unable to determine the role AI plays in article authorship.

¹Recent studies from Pew Research show that (1) 49% of Americans who get news directly from AI assistants report encountering inaccurate information (Lipka and Eddy, 2025); (2) 56% of Americans would feel less confident about a news article if they knew an AI wrote it (Yam and Kennedy, 2025); and (3) 76% believe it is extremely important for them to know if the text they are reading is AI-generated (Kennedy et al., 2025).

Audit design: We collect and audit two large-scale datasets of American newspapers:

- 📄 recent_news contains **186K** articles published online by **1.5K** local and national newspapers from June to September 2025.
- 🗣️ opinions contains **45K** *opinion* articles published by the New York Times, Washington Post, and Wall Street Journal between August 2022 and September 2025.

We feed each article in both datasets through Pangram’s API to obtain both an AI likelihood (from 0-100%) and a categorical label in { **HUMAN-WRITTEN**, **MIXED**, **AI-GENERATED** }. Pangram has a reported false positive rate (FPR) of $\sim 0.001\%$ on news text (Emi, 2025) and Jabarian and Imas (2025) verified Pangram’s high accuracy, finding an FPR of 0.08%, substantially below the rates of AI use we observe. Importantly, for articles labeled **MIXED**, we cannot infer the role that AI played in the authorship process: we only know that some parts of a **MIXED** article are classified as human-written while other parts are classified as AI-generated.



AI use in published articles is increasingly common. In our recent_news dataset, 9.1% of articles are labeled by Pangram as either **AI-GENERATED** or **MIXED**, and disclosure of AI use is rare in our manual sample. Digging deeper, we observe that AI usage is unevenly distributed: it is much higher in smaller local outlets than nationally-circulated papers; it occurs more frequently in topics such as weather, science / technology, and health (Figure 13); it varies across ownership group, with Boone Newsmedia and Advance Publications among the heaviest AI adopters (Figure 14); and it is higher in languages other than English (Figure 2).²

Undisclosed AI use compounds risks of factual hallucinations. While AI use is becoming more common, we find it is rarely disclosed and often factually unreliable. When investigating whether AI use is disclosed, only 4% of articles were transparent about writing articles with the help of AI,

²Example articles linked within this paper are solely for illustrative purposes. Their inclusion does not attribute any intent or misconduct on part of any reporters: Pangram has a small but non-zero FPR, and we also emphasize that it is infeasible to tease apart exactly how AI is used in **MIXED** authorship cases.

IN THIS PAPER...

AI USE: We use the term *AI use* throughout this paper to denote articles labeled by Pangram as either **MIXED** or **AI-GENERATED**.

EXAMPLES: Each article discussed in this paper is associated with a  icon that links to an AI prediction dashboard and a  icon that links to the original news article.

leaving readers in the dark about the true sources of information. At the same time, we find that 41% of articles labeled as AI-generated contain hallucinations, showing that AI-generated information is often inaccurate and not corrected by human revision.

AI use is concentrated in opinion articles at top newspapers. Opinion articles published at the NYT, WaPo, and WSJ are **6.4** times more likely to be flagged for AI use than contemporaneous news articles from the same three newspapers (4.5% vs. 0.7%). Many opinion articles flagged for AI use are written by prominent guest contributors, including Nobel prize winners, US Senators and Governors, Pulitzer Prize-winning journalists, and CEOs (see Table 12 for examples). Analysis of both opinions and ai_reporters show AI usage rising over time, with reporters in the latter dataset increasing their AI use from $\sim 0\%$ prior to 2023 to over 40% in 2025 on average (Figure 8).

Contributions: We release our datasets,³ analysis code, and an interactive dashboard to facilitate further exploration of AI use in newspapers.⁴ We also commit to periodically updating our dashboard with new articles and annotations (i.e., disclosure audits) to measure future changes to AI adoption.

2 Collecting newspaper articles

We collect two datasets of published newspaper articles (Table 1): recent_news and opinions.⁵ The datasets vary significantly in terms of coverage, diversity, and publication date, which allows us to analyze differences in AI use across local vs. national papers, staff reporters vs. guest contributors, and articles written pre- vs. post-ChatGPT.

³We release links to the articles involved in our study (not full texts).

⁴a.newsaudit.github.io

⁵We collect a *third* dataset, ai_reporters, which is detailed and analyzed in §D.3.

	📰 recent_news	🗣️ opinions
Temporal coverage		
Years	2025	2022–2025
Dataset statistics		
# Articles	186,507	44,803
# Authors	34,608	9,863
# Newspapers	1,528	3
Avg. tokens	787.4	1078.4
AI use statistics		
HUMAN-WRITTEN	90.85%	99.04%
MIXED	3.98%	0.85%
AI-GENERATED	5.24%	0.11%

Table 1: Dataset and AI use statistics for recent_news and opinions. The token counts are reported as per tiktoken tokenizer (o200k_base).

This section outlines our dataset creation process, in which the full text and metadata for each article is paired with a label indicating whether the article was generated via AI.

2.1 recent_news

To examine AI use in present-day newspaper journalism, we form the recent_news dataset by collecting 186,507 articles⁶ published online by 1,528 unique newspapers between June 15th, 2025 and September 15th, 2025, accessed via url RSS feeds.⁷

2.2 opinions

While recent_news contains a broad sample across many different American newspapers, topics, and journalists, another object of our study is AI use in *opinion* articles written by prominent people in highly-reputable newspapers. To facilitate this, we also collect a dataset of opinions articles published by The New York Times, Wall Street Journal, and Washington Post between August 2022 and September 2025. The full text and metadata of these articles were accessed via *ProQuest Recent Newspapers*. In total, we collected 44,803 articles during this time period: 16,964 from WSJ, 15,977 from WP, and 11,862 from NYT.

2.3 Labeling the datasets

For each of the 251,442 total articles across all datasets, we use Pangram (v2.0) to obtain an AI detection label and score.⁸ We also classify each

⁶Each instance in our datasets includes the full article text along with title, author, publication date, newspaper, and URL.

⁷See §A for additional collection details.

⁸All Pangram links display v3.2 predictions, and had consistent labels with Pangram v2.0.

article into one of 19 topics, and we use an existing database to match about half of the newspapers in our recent_news with print circulation statistics.

Detecting AI use: Emi and Spero (2024) introduce Pangram; a robust AI-generated text detection tool. On news articles, it achieves a false positive rate of 0.001% (Emi, 2025), consistent with other studies that have also reported low FPR (Russell et al., 2025; Jabarain and Imas, 2025; Dugan et al., 2025). Using Pangram’s inference API⁹, we collect (1) the likelihood (from 0-100%) that a text is AI-generated and (2) a text label that is one of HUMAN-WRITTEN, MIXED, or AI-GENERATED.¹⁰ Pangram predicts MIXED when there is a high confidence of both AI and human writing present in the document; specifically, where some segments are predicted as AI, and some segments are predicted as human. While prior studies have validated Pangram’s high accuracy, we also experiment with another commercial detector, GPTZero (Tian and Cui, 2023), and observe a high cross-detector agreement of 88.2% (Cohen’s $\kappa = 0.764$).¹¹

Topic classification: To analyze AI use across article topics, we further augment our datasets with topic labels for each article using the International Press Telecommunications Council *Media Topics* taxonomy, which comprises 17 top-level topics (IPTC, 2025). We prompt QWEN3-8B (Qwen Team, 2025) in a zero-shot setting to assign a topic (full prompt and details in §A).¹² To assess the reliability of these topic labels, two of the authors independently reclassified a random subset of 100 model outputs,¹³ with model–human agreement averaging 77%, indicating moderately strong alignment between the classifier and human judgments.

⁹API documentation available [here](#).

¹⁰We simplify the fine-grained labels produced by the Pangram API in the following way: HUMAN-WRITTEN = {Human, Unlikely AI}; MIXED = {Mixed, Possibly AI, Likely AI}; AI-GENERATED = {Highly Likely AI, AI}. See §B for details on how labels are combined.

¹¹This experiment was conducted on a balanced binary held-out set of 1K news articles, 500 of which are marked by Pangram as HUMAN-WRITTEN and 500 as AI-GENERATED. Agreement on the human subset is 98.4% while agreement on the AI subset is 78.4%, and discrepancies on the latter label are likely due to each detector’s differing treatment of MIXED text (see §B for more).

¹²In addition to the 17 IPTC topics, we include two auxiliary categories—*Other* and *Obituary*—for items that fall outside the taxonomy or are obituaries.

¹³Inter-annotator agreement between the two human raters was 87% (Cohen’s $\kappa = 0.85$).

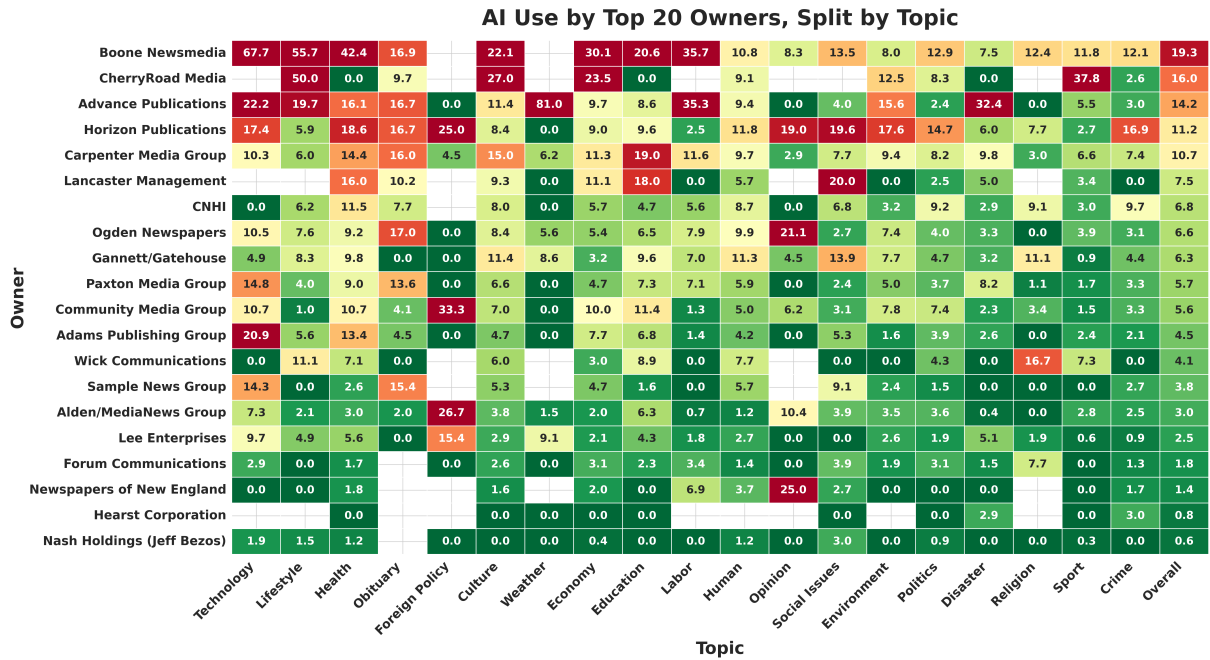


Figure 1: Heatmap of AI use by publication owner and article topic in recent_news. Some owners disclose AI use for specific content, such as [Advance Publications](#) for weather reports, but the biggest adopters, such as [Boone Newsmedia](#), use AI broadly across many topics. Note that only 87% of recent_news has ownership information, and only topics that have at least 5 articles per owner are visualized in the heatmap.

Linking newspapers to circulation and ownership information: We link newspapers to print circulation data from the U.S. News Deserts Database ([UNC Center for Innovation & Sustainability in Local Media, 2020](#)) and ownership information from the Local News Initiative ([Medill Local News Initiative, 2024](#)). Because circulation figures are based on historical AAM audits and self-reports (primarily from 2019), we treat them as a proxy for *print scale* rather than current total audience.¹⁴ 40.3% of articles come from small publications with daily circulation between 1-7K and another 36.3% comes from publications with a circulation of 10-50K. Only 20.1% of articles come from very large nationally-circulated publications with average circulations over 50K.¹⁵

3 Analyzing AI use in newspapers

How much of recent news contains AI-generated content, and when and how do journalists use AI? Using our recent_news dataset, we investigate *where* AI-generated content appears across circulation, ownership, language, and article topic. We also explore *how* it is used, looking at differences

¹⁴See §A for details and limitations. 54.6% of the articles and 49.7% of newspapers in recent_news were matched to publications in the News Desert Database.

¹⁵Circulation distribution is depicted in [Figure 6](#).

between [MIXED](#) and [AI-GENERATED](#) articles, factual details, and disclosure.

3.1 How often is AI used in American newspapers?

We find that **9.1%** of the 186K articles in recent_news are labeled by Pangram as either [AI-GENERATED](#) (5.2%) or [MIXED](#) (3.9%), while the remaining 90.9% of articles are classified as [HUMAN-WRITTEN](#). The rest of this section goes beyond these aggregate numbers to examine AI use as a function of different fine-grained aspects like topic and ownership.

AI use is higher in local newspapers. Local outlets, often operating in "news deserts," places with little to no access to credible reporting due to limited resources ([Metzger, 2024](#)), rely on AI more than national newspapers. Only 1.7% of articles at papers with circulation >100K are labeled as [AI-GENERATED](#) or [MIXED](#), versus 9.3% at papers below 100K (see [Figure 15](#)),¹⁶ suggesting that large national newspapers enforce stricter editorial constraints on automation than local papers.

¹⁶At the article level, this difference is highly significant ($\chi^2(1) = 1175.6, p < 10^{-250}$). At the newspaper level, smaller outlets averaged 8.5% AI content vs. 5.0% among very large outlets (Welch's $t(\approx 23) = 2.24, p = 0.032, d = 0.22$).

AI use also varies geographically, peaking in the mid-Atlantic and southern US, Maryland (16.5%), Tennessee (13.6%), Alabama (13.9%), and lowest in the Northeast, including New Hampshire (2.9%) and Massachusetts (3.4%) (Figure 11).

AI use varies with topic. Prior work has shown that factual, data-heavy content (e.g., reports about weather, finance, or sports) is particularly amenable to automation (Medill Local News Initiative, 2024). Consistent with this, weather articles in our dataset exhibit the highest average AI likelihood (27.7%), as shown in Figure 13. However, we also observe high AI use in other topics, such as science and technology (16.1%) and health (11.7%), while content on more sensitive issues such as conflict and war (4.3%), crime, law, and justice (5.2%), and religion (5.3%) exhibit lower rates.

AI use varies with ownership. Many newspapers in recent_news share ownership: for example, Advance Publications owns widely-read outlets like pennlive.com, cleveland.com, and al.com and often syndicates articles among them. While some ownership groups view AI use as a reputational risk, others emphasize cost reduction and efficiency gains (Medill Local News Initiative, 2024). Boone News Media and CherryRoad Media both have detected AI use rates over 15%, while Nash Holdings and Hearst Corporation have AI use rates under 1% (Figure 14). We note the irony of several prominent media groups suing AI companies over training language models on their content (e.g., Advance Local Media v. Cohere)¹⁷ even as they churn out LLM-generated articles.

AI use varies by topic across ownership groups. Figure 14 shows how different ownership groups use AI. Advance Publications, for instance, relies heavily on AI for weather reporting (81% AI Use). Boone Media exhibits the broadest adoption, with especially high rates in science and technology (67.7%) as well as lifestyle and leisure (55.7%). Other chains show narrower topical patterns, such as CherryRoad Media in sports (37.8%) and Lee Enterprises in foreign policy coverage (15.4%). These results show that AI integration is vastly uneven across both topic and ownership groups.

AI use is higher in languages other than English. AI-generated content is more prevalent in

¹⁷Advance Local Media LLC v. Cohere Inc., No. 1:25-cv-01305, S.D.N.Y. filed Feb. 13, 2025.




Figure 2: AI use in recent_news is more frequent in languages other than English. The most prominent such languages include Spanish, Portuguese, Vietnamese, French, and Polish.

non-English news articles.¹⁸ As shown in Figure 2, only 8.0% of English-language articles in recent_news are classified as **AI-GENERATED** or **MIXED**; this share rises to 31.0% for articles in other languages.¹⁹ Most non-English AI use ($\approx 80\%$) comes from U.S.-based Spanish-language reporting (7.2K articles),²⁰ suggesting that local bilingual outlets rely more heavily on translation and other automation. High non-English AI use is observed across many states (Figure 12).

Exploring the role of machine translation in AI detection To understand if the higher rate of detected AI use in non-English languages is an artifact of machine translation (MT), we translate a sample of 300 English articles into 12 languages and re-evaluate them using the same detection pipeline.²¹ Translated articles receive lower AI-likelihood scores than their English counterparts, with an average decrease of 13%. Under the binary AI use versus human, translated texts agree with English in 83.2% of cases, and disagreements are asymmetric: in our sample, translation is far more likely to suppress AI use predictions than to introduce them.²² The magnitude of this effect varies by language: Vietnamese translations show

¹⁸Pangram FPR's range from 0.0% to 0.10% across the languages in this study (Pangram, 2024).

¹⁹Machine-translated articles may be misclassified, but our small-scale verification suggests human-written articles are still correctly identified after translation.

²⁰Other major languages include: Portuguese (468), Vietnamese (403), French (343), and Polish (314).

²¹We sample 100 **AI-GENERATED**, 100 **MIXED**, and 100 **HUMAN-WRITTEN** articles. Articles are translated using GPT-4.1. We note that results may differ with other MT methods or models; see §D.1 for details.

²²McNemar's exact test, FDR-corrected $p < 0.001$ for all but one language (Japanese).

POLICY CATEGORY	FREQ	EXAMPLE PUBLICATIONS / OWNERS	ILLUSTRATIVE DISCLOSURE TEXT
✔ AI ALLOWED	12	<i>Ionia Sentinel-Standard, Amsterdam News, Penn Live</i>	"If AI-assisted content is approved for publication, journalists must disclose the use of AI and its limitations to their audience. AI-generated content must be verified for accuracy and factuality before being used in reporting."
❌ AI PROHIBITED	2	<i>New York Post, Michigan Daily</i>	"The use of generative artificial intelligence for content production (including written, visual and auditory content) is unacceptable in all circumstances. Any staffer found to have used generative AI to produce content for The Daily can be fired by their section editor or the Editor in Chief."
🔒 NO PUBLIC POLICY	186	<i>Daily Register, Hudson Reporter, LA Opinion</i>	No disclosure found on website.

Table 2: AI disclosure policies among 200 sampled U.S. news outlets.

the largest score decrease ($\Delta = -0.21$) and lowest agreement (74.7%), while Portuguese and Spanish retain high agreement ($\geq 88\%$) with minimal score shifts (see Table 10). While these results suggest that MT is unlikely to be the primary driver of the higher AI detection rates in non-English content, especially given the observed $\sim 31\%$ AI use rate, this analysis does not account for the influence of human translation or alternative MT methods. Multilingual AI detection is an evolving field; factors like language-specific writing conventions, varied editorial workflows, and regional AI tool availability likely contribute to these discrepancies.






3.2 Characterizing AI use in American newspapers




Our analysis establishes that (1) many published newspaper articles today are written partially or entirely by AI, and (2) AI use varies across factors like ownership, topic, and language. We investigate important questions about authorship, transparency, and public trust, revealing that journalists often use AI in conjunction with their own writing rather than as a full replacement; that prestigious outlets rely more on **MIXED** authorship; articles with AI use are more likely to contain hallucinations; and that disclosure of AI involvement is strikingly rare. These findings show not only how AI is changing journalistic practices, but also why it matters whether readers are informed about its role.

Many articles have mixed authorship. While **MIXED** authorship articles include at least some human writing, it is difficult to tease apart the contributions of the reporter compared to AI (Thai et al., 2025). Mixed authorship is common: of the 17,059 articles we detect as using AI, 42.7% are predicted as **MIXED**, while 57.3% are classified as **AI-GENERATED**. At the author level, 1,453 out of 34,608 writers produce at least some AI content.

The majority (54.8%) primarily publish mixed articles, while 36.1% rely mostly on AI-generated text. These findings resonate with survey evidence from Radcliffe (2025), who report that over half of journalists use AI to edit their work, while only about a third employ it to generate text directly.

AI use is largely undisclosed. Disclosure of AI use (e.g., exactly how and where AI was used in the construction of the article) is important to maintain audience trust. Readers might be okay with small AI edits for style in articles labeled as **MIXED** but they require proper disclosure to make these judgments. Unfortunately, in a sample of 200 AI-flagged articles from unique newspapers in recent_news, 96.5% of authors and 94.0% of publishers did not disclose AI use (Table 2).²³

Notable individual cases of AI use. We identify several unique cases of AI-generated writing in recent_news, snippets of which are shown in Table 3. One outlet, *Argonaut*, turned out to be an entirely AI-generated newspaper with AI reporter personas who "write" exclusively articles labeled as **AI-GENERATED** (e.g., ). Another reporter appears to have revisited and republished their older work, producing updated AI-assisted versions of previously **HUMAN-WRITTEN** stories (e.g., *edited (2025)*: , ; *original (2021)*: , ).²⁴

More concerning, we also found AI-generated responses in a popular advice column, *Dear Annie*, a practice that risks betraying the trust readers place in such personal guidance (e.g., *Reader*: , *Annie*: , ).²⁵

²³95% Wilson score CIs: 193/200 (96.5%; [93.0%, 98.3%]) and 188/200 (94.0%; [89.8%, 96.5%]).

²⁴Earlier versions were retrieved via the Internet Archive's Wayback Machine.

²⁵We were unable to scrape the *Dear Annie* column to quantify AI use because its publication across many different

NEWSPAPER	PREDICTION	ARTICLE EXCERPT	OBSERVATIONS
ARGONAUT	AI-GENERATED	“...The forms of commerce have drastically changed, with constant evolution to meet consumer needs and market dynamics . From brick-and-mortar stores, e-commerce, and omnichannel strategies to pop-up shops and outdoor commerce, businesses need to remain adaptable to thrive in this ever-changing commercial landscape ...”	Fully AI-generated text with a very generic conclusion. No details about the author exist online, and when we further investigated the site, no details could be found about any of its staff.
CALEXICO CHRONICLE	HUMAN-WRITTEN	“...This grant represents an important step forward in our efforts to create healthier, more sustainable learning environments for our students. By increasing shade and greenery across our campuses, we’re not only improving outdoor comfort and air quality, but also setting an example of environmental responsibility for our students and community...”	While this article is detected as human-written, it includes an AI-generated quote that was likely provided to the reporter . Reporters who write their own articles may not be aware that the people they quote in their articles used AI to create their response.
WASHINGTON POST	MIXED	“...Finally, focus on who and where else you can seek support from. Is there even one family member or community member you can turn to and tap in for support and allyship as you navigate these familial struggles?”	AI use often occurs even in high circulation papers like Washington Post. In this advice column, a person writes in feeling lonely, only to receive partially AI-generated advice.

Table 3: Notable cases of AI use in the recent_news dataset. Words and phrases identified as indicative of AI use by Pangram are highlighted in red. AI use takes many forms, from completely made-up news sites to AI responses to advice columns (e.g., *Dear Annie*), legitimate articles that happen to quote AI-generated text from other sources, and highly-templated topics like weather and sports reports.

AI-generated news often contains hallucinations. *Hallucinations* are defined as outputs that are fluent and coherent but factually incorrect or unfaithful to the source material (Maynez et al., 2020b; Ji et al., 2023b). The widespread hallucinations in news media risks the spread of inaccurate information and reputational damage (BBC News and European Broadcasting Union, 2025). To understand if news detected as AI-generated is as verifiable as human-written reporting, we manually reviewed 100 AI-GENERATED and 100 HUMAN-WRITTEN articles, finding that AI-labeled articles are 8.2x more likely to contain hallucinations:²⁶ 41% of AI-labeled news contained at least one hallucinated claim, whereas only 5% of articles predicted as human-written had hallucinations.²⁷ Typical factual errors include fabricated quotes, incorrect statistics, and misdated events (example hallucinations in Table 8).

4 Opinions

Opinion articles play a large role in shaping public attitudes, especially those authored by trusted figures such as Nobel laureates, elected officials, and journalists. Even a single op-ed can significantly shift beliefs (Coppock et al., 2018; Bai et al., 2025).

websites made it difficult to access archival content.

²⁶Fisher’s exact test, $p = 2.3 \times 10^{-9}$.

²⁷A BBC/EBU study found around 55% of AI-generated responses had accuracy issues (BBC News and European Broadcasting Union, 2025).

As LLMs are often more persuasive than humans (Salvi et al., 2025; Schoenegger et al., 2025), AI-assisted opinion writing raises concerns of misinformation and ideological amplification (Weidinger et al., 2022; Nehring et al., 2024). To assess these risks, we examine opinion articles written from 2022-2025 in three of the most highly circulated national newspapers (New York Times, Washington Post, and Wall Street Journal). While AI use in opinions remains low relative to recent_news, it has risen sharply following mainstream LLM adoption and is also almost completely undisclosed.

AI use in opinion articles published at reputable newspapers has increased by 25x over the past three years. To see whether AI-generated material has risen over time, we measure the share of opinions flagged as AI between September 2022 and September 2025. AI use increases from 0.1% in 2022 to 3.4% in 2025, about a 25x rise, consistent across the three outlets.²⁸ By outlet, the AI use share grows from 0.1% to 3.4% at the Wall Street Journal, 0.2% to 4.3% at the Washington Post, and 0.0% to 2.6% at the New York Times (Figure 3). As a pre-ChatGPT sanity check, only 5 of 5,029 opinion articles published before December 2022 are flagged under our AI-use definition (AI-GENERATED or MIXED), implying an empirical false-positive rate of 0.10% (95% Wilson

²⁸Across all opinion pieces in our sample, 0.1% are labeled AI-GENERATED and 0.8% MIXED; see Figure 18.




Figure 3: Monthly proportion of opinions articles flagged for AI use (Sep. 2022 -Sep. 2025). AI use in opinion pieces has dramatically increased over time.

CI: 0.04%–0.23%).

Opinions exhibit higher AI use than other sections (June–September 2025). Across the three outlets, AI use is 6.4 times more likely to occur in opinion pieces than non-opinion news articles published by those same outlets²⁹ between June–September 2025 (4.56% vs. 0.71%; $n=3,420$ opinions, $n=10,129$ all articles). By outlet, the gap is largest at *The Washington Post* (5.51% vs. 0.55%), followed by *The Wall Street Journal* (4.99% vs. 0.74%), and smaller at *The New York Times* (2.94% vs. 1.80%). Looking at how AI is used, **MIXED** dominates in both settings (86.5% of AI use in opinions and 86.1% of AI use in non-opinion articles at the same outlets). Without clear disclosure standards, readers cannot discern whether AI contributed merely to editing or to greater content generation, restricting readers’ ability to judge the appropriateness of AI in specific articles.

Guest contributors are much more likely to use AI in opinion pieces than full-time reporters. We observe that the majority of the 219 unique authors in opinions with at least one article detected as having AI use are infrequent contributors rather than full-time journalists. Categorizing authors by occupation reveals clear occupational divides: political figures, executives, and scientists exhibit the highest AI use, with many using AI for all of their articles in opinions. Veteran opinion columnists show near-zero AI incidence (<0.5%). This pattern intuitively makes sense, as guests generally lack established editorial processes and newsroom support (Washington Post Staff, 2022).

²⁹These non-opinion articles are extracted from recent_news, which includes WSJ, WaPo, and NYT; see Table 11 for more details.




Figure 4: Changes in AI use in opinions articles year over year (2025 only includes January 1 - September 15). AI likelihood increased across all topics, with especially large gains in opinion articles about Science & Technology, consistent with topic trends in recent_news.

Topical differences in AI use between opinions and recent_news. Setting aside AI use, the majority of opinion articles in our dataset focus on politics and government (56.9%), far more than in recent_news (15.0%). In Figure 4, several categories show large increases since early 2025: *crime and law* (about 31x higher), *religion* (16x), and *economy and business* (14x) stand out as the fastest-growing. *Science and technology* continues to lead in absolute terms, roughly 9x higher (1.2%→9.0%), mirroring its position in recent_news as the topic most likely to contain AI use. Substantial gains also appear in *conflict, war and peace* (12x), *human interest* (10x), and *politics and government* (9x). These patterns indicate that AI use in opinion writing extends well beyond scientific domains and into political, economic, and broadly human-centered discourse.

5 Related work

AI use in news media. Recent scholarship maps how generative AI is impacting newsroom practices. Recent case studies document both benefits and frictions caused by AI within news organizations (Brigham et al., 2024; Jones and Jones, 2025; Ansari et al., 2025), while industry policies study the adoption, governance, and platform–publisher dynamics (Simon, 2024; Brown and Jaźwińska, 2025; Simon et al., 2025; Meir, 2023).

Effects of disclosing AI in news. Disclosure reliably changes how readers judge identical content (Longoni et al., 2022). In recent surveys, comfort reading AI-generated material was low

(**AI-GENERATED**: 19%; **MIXED**: 30%) (Newman et al., 2025). Disclosing machine authorship reduces perceived credibility (Toff and Simon, 2025; Lee et al., 2025), yet readers rarely detect AI without cues and experience (Brown et al., 2020; Clark et al., 2021; Russell et al., 2025). Moreover, disclosure labels often fail to reduce persuasive impact (Gallegos et al., 2025). Despite these mixed effects on perceived credibility, transparency in AI use is important: ethically, newsroom standards emphasize being accountable to the reader (Meir, 2023; Viner and Bateson, 2023) and practically, would enhance trust between readers and news outlets using AI (Newman et al., 2025).

Measuring AI content in other domains. Prior studies audit the growing presence of LLM-generated text across academic and creative domains. In scholarly settings, the influence of AI is measured in peer reviews (Liang et al., 2024a; Zhou et al., 2025) and academic papers (Liang et al., 2024b; Luo et al., 2025; Kobak et al., 2025). Gupta and Pruthi (2025) use AI detection to identify patterns of plagiarism in generated research. Similar detection efforts in creative domains examine AI-generated media on art platforms, Wikipedia, social networks, and news websites (Matatov et al., 2025; Brooks et al., 2024; Sun et al., 2025; Hanley and Durumeric, 2024). Other studies trace the adoption of AI in public communication and consumer complaints (Liang et al., 2025; Shin et al., 2025).

6 Conclusion

Our analysis provides the first audit of AI use across **250k** articles of U.S. newspaper articles. We find that AI use is both *widespread* and *uneven*: roughly 9% of recent articles are flagged as either partially or completely AI-generated, with particularly high rates in smaller local papers, specific topics, ownership groups, and languages other than English. In highly-reputed national newspapers, we observe AI use in opinion pieces is rapidly increasing (from 0% in 2022 to 3-4% in 2025), much more so than in other news articles published by the same papers. Based on our audit, we suggest some disclosure policies that newspapers could consider adopting to improve public trust regarding AI use:

1. **Guidance for mixed authorship:** editorial standards should publicly outline what kind of AI use is acceptable without disclosure if any (e.g., grammar checks, style edits), acceptable

with disclosure (e.g., summarization, more in-depth rewrites), and not permitted at all (e.g., full article generation). Reporters should record notes and/or log their AI use during the writing process, as these can be useful for more informative disclosure.

2. **Explicit AI policies for external contributors:** Since opinion pieces exhibit high rates of detected AI use at top papers, particularly those written by guest contributors, we propose that author attestations about AI use are collected along with article submissions. Editors may additionally want to check submissions (automatically or manually) for AI cues (Russell et al., 2025) and publish standards on what they deem acceptable AI use in opinion pieces.

Limitations

Our study focuses primarily on AI use in the U.S. press in English-language publications. While we connect our work to prior studies about hallucination and factual errors in language model generations, we do not perform a large-scale evaluation of the factuality of **AI-GENERATED** articles in our dataset beyond studying the authenticity of quotations. Furthermore, although articles in other languages are included in our datasets, they are comparatively much fewer in number; as such, we do not claim that this sample captures the diversity of AI use in journalism globally.

Our machine translation experiment uses a single MT model (GPT-4.1) and does not account for human translation or other MT methods, which may affect AI detection differently. Additionally, while Pangram reports low false positive rates across languages, detection performance may vary in ways not fully captured by aggregate metrics; we therefore treat our cross-lingual findings as exploratory.

`recent_news` focuses on regional and local newspapers (although national outlets like NYT and WSJ are included), which may bias results towards the journalistic practices of smaller outlets. The inclusion or exclusion of particular newspapers is partially due to data accessibility constraints. Specifically, we are unable to obtain articles from most sites without active RSS feeds, and we are not able to include print-only newspapers. Finally, the data contains inevitable noise. Despite extensive data-cleaning efforts, metadata occasionally appears within article text, which may marginally

affect AI detection and topic classification results.

Ethical Considerations

Our data was collected from publicly accessible newspaper sites, either through RSS feeds or available archives. Given the sensitivity of large-scale text collection, we do not release the complete article texts, but instead provide metadata to respect the rights of content owners. We identify AI-generated text using Pangram, an AI detection model. While Pangram does exhibit very low false-positive rates in benchmark testing, it is not infallible, and all findings should be interpreted as detector outputs rather than definitive authorship attributions. We do not attribute any intent, misconduct, or ethical lapses to the individual journalists and newspapers flagged by the model. We do not draw conclusions about the intent, conduct, or practices of individual journalists, outlets, or companies. Results should not be interpreted as rankings, qualitative judgments, or accusations of wrongdoing. Rather, we purposely use large-scale data collection to understand AI use trends at an aggregate level. We aim to understand what trends in AI use appear across the industry rather than making judgments about specific cases.

Acknowledgments

We thank the University of Maryland Computational Linguistics and Information Processing (CLIP) Lab for their feedback and support. We are grateful to the University of Maryland Howard Center for Investigative Reporting, especially Derek Willis and Daniel Trielli, for their guidance. This project was partially supported by awards IIS-2046248 and IIS-2312949 from the National Science Foundation (NSF) as well as an award from Open Philanthropy.

References

- Abolfazl Ansari, Delvin Ce Zhang, Nafis Irtiza Tripto, and Dongwon Lee. 2025. *Echoes of automation: The increasing use of llms in newsmaking*. In *Social, Cultural, and Behavioral Modeling: 18th International Conference, SBP-BRIMS 2025, Pittsburgh, PA, USA, October 14–17, 2025, Proceedings*, page 3–13, Berlin, Heidelberg. Springer-Verlag.
- Hui Bai, Jan G. Voelkel, Shane Muldowney, Johannes C. Eichstaedt, and Robb Willer. 2025. *Llm-generated messages can persuade humans on policy issues*. *Nature Communications*, 16:6037.
- Adrien Barbaresi. 2021. *Trafilatura: A Web Scraping Library and Command-Line Tool for Text Discovery and Extraction*. In *Proceedings of the Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 122–131. Association for Computational Linguistics.
- BBC News and European Broadcasting Union. 2025. *News integrity in ai assistants*. Technical report, BBC Media Centre. Accessed: 2025-10-21.
- Natalie Grace Brigham, Chongjiu Gao, Tadayoshi Kohno, Franziska Roesner, and Niloofar Mireshghal-ah. 2024. *Developing story: Case studies of generative ai’s use in journalism*. *Preprint*, arXiv:2406.13706.
- Creston Brooks, Samuel Eggert, and Denis Peskoff. 2024. *The rise of AI-generated content in Wikipedia*. In *Proceedings of the First Workshop on Advancing Natural Language Processing for Wikipedia*, pages 67–79, Miami, Florida, USA. Association for Computational Linguistics.
- Peter Brown and Klaudia Jaźwińska. 2025. *Journalism zero: How platforms and publishers are navigating ai*. Report, Tow Center for Digital Journalism, Columbia University.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. *Language models are few-shot learners*. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Roberto Cavazos and Greg Sterling. 2024. The high cost of review fraud: An economic analysis of consumer harm. <https://askfortransparency.com/research/high-cost-of-review-fraud/>. Accessed: 2025-10-17.
- Elizabeth Clark, Tal August, Sofia Serrano, Nikita Haduong, Suchin Gururangan, and Noah A. Smith. 2021. *All that’s ‘human’ is not gold: Evaluating human evaluation of generated text*. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 7282–7296, Online. Association for Computational Linguistics.
- Jim Conaghan. 2017. Young, old and in-between: Newspaper platform readers ages are well-distributed. <https://www.newsmediaalliance.com>.

- [org/age-newspaper-readers-platforms/](https://www.nielsen.com/usa/newsroom/2016/08/01/2016-release-2-finding-that-the-median-age-of-daily-print-newspaper-readers-is-57.9-years/). Highlights Nielsen Scarborough USA+ 2016 Release 2 finding that the median age of daily print-newspaper readers is 57.9 years.
- Alexander Coppock, Emily Ekins, and David Kirby. 2018. **The long-lasting effects of newspaper op-eds on public opinion.** *Quarterly Journal of Political Science*, 13(1):59–87.
- Liam Dugan, Andrew Zhu, Firoj Alam, Preslav Nakov, Marianna Apidianaki, and Chris Callison-Burch. 2025. **GenAI content detection task 3: Cross-domain machine generated text detection challenge.** In *Proceedings of the 1st Workshop on GenAI Content Detection (GenAIDetect)*, pages 377–388, Abu Dhabi, UAE. International Conference on Computational Linguistics.
- Bradley Emi. 2025. All about false positives in ai detectors. <https://www.pangram.com/blog/all-about-false-positives-in-ai-detectors>. Pangram Labs blog.
- Bradley Emi and Max Spero. 2024. **Technical Report on the Pangram AI-Generated Text Classifier.** *Preprint*, arXiv:2402.14873.
- Daniel S. Evanko and Michael Di Natale. 2025. **Quantifying and assessing the use of generative ai by authors and reviewers in the cancer research field.** In *Peer Review Congress*. Accessed 2025-10-17.
- Isabel O. Gallegos, Ryan A. Rossi, Joe Barrow, Md Mehrab Tanjim, Sungchul Kim, Franck Dernoncourt, Tong Yu, Ruiyi Zhang, and Nesreen K. Ahmed. 2024. **Bias and fairness in large language models: A survey.** *Computational Linguistics*, 50(3):1097–1179.
- Isabel O. Gallegos, Chen Shani, Weiyan Shi, Federico Bianchi, Robb Willer, and Dan Jurafsky. 2025. **AI-generated content and public persuasion: The limited effect of AI authorship labels.** In *NeurIPS 2024 Workshop on Regulatable ML*.
- Tarun Gupta and Danish Pruthi. 2025. **All that glitters is not novel: Plagiarism in AI generated research.** In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 25721–25738, Vienna, Austria. Association for Computational Linguistics.
- Hans W. A. Hanley and Zakir Durumeric. 2024. **Machine-made media: Monitoring the mobilization of machine-generated articles on misinformation and mainstream news websites.** *Proceedings of the International AAAI Conference on Web and Social Media*, 18(1):542–556. ICWSM 2024.
- Tony Hu, Yuliya Kyrychenko, Steve Rathje, et al. 2025. **Generative language models exhibit social identity biases.** *Nature Computational Science*, 5:65–75.
- IPTC. 2025. Media topics. <https://iptc.org/standards/media-topics/>. Accessed: 2025-09-29.
- Brian Jabarian and Alex Imas. 2025. **Artificial writing and automated detection.**
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023a. **Survey of hallucination in natural language generation.** *ACM Comput. Surv.*, 55(12).
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023b. **Survey of hallucination in natural language generation.** *ACM Comput. Surv.*, 55(12).
- Bethany Jones and Richard Jones. 2025. **Action research at the BBC: Interrogating artificial intelligence with journalists to generate actionable insights for the newsroom.** *Journalism*, 26(8):1708–1725. Original work published 2025.
- Brian Kennedy, Eileen Yam, Emma Kikuchi, Isabelle Pula, and Javier Fuentes. 2025. **How americans view AI and its impact on people and society.** Pew Research Center, Report.
- Kate Knibbs. 2024. **Ai slop is flooding medium.** *Wired*. Accessed: 2025-10-17.
- Dmitry Kobak, Rocío González-Márquez, Éva A. Horvát, and Jana Lause. 2025. **Delving into llm-assisted writing in biomedical publications through excess vocabulary.** *Science Advances*, 11(27):eadt3813. PMID: 40601754; PMCID: PMC12219543.
- Daniel Chaein Lee, Jihoon Jhang, and Tae Hyun Baek. 2025. **AI-generated news content: The impact of AI writer identity and perceived AI human-likeness.** *International Journal of Human-Computer Interaction*, pages 1–13. Published online 2025-03-24; Taylor & Francis.
- Weixin Liang, Zachary Izzo, Yaohui Zhang, Haley Lepp, Hancheng Cao, Xuandong Zhao, Lingjiao Chen, Haotian Ye, Sheng Liu, Zhi Huang, Daniel A. McFarland, and James Y. Zou. 2024a. **Monitoring ai-modified content at scale: a case study on the impact of chatgpt on ai conference peer reviews.** In *Proceedings of the 41st International Conference on Machine Learning, ICML’24*. JMLR.org.
- Weixin Liang, Yaohui Zhang, Mihai Codreanu, Jiayu Wang, Hancheng Cao, and James Zou. 2025. **The widespread adoption of large language model-assisted writing across society.** *Preprint*, arXiv:2502.09747.
- Weixin Liang, Yaohui Zhang, Zhengxuan Wu, Haley Lepp, Wenlong Ji, Xuandong Zhao, Hancheng Cao, Sheng Liu, Siyu He, Zhi Huang, Diyi Yang, Christopher Potts, Christopher D Manning, and James Y. Zou. 2024b. **Mapping the increasing use of LLMs in scientific papers.** In *First Conference on Language Modeling*.

- Michael Lipka and Kirsten Eddy. 2025. [Relatively few americans are getting news from AI chatbots like ChatGPT](#). Pew Research Center, Short Reads.
- Chiara Longoni, Andrey Fradkin, Luca Cian, and Gordon Pennycook. 2022. [News from generative artificial intelligence is believed less](#). In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '22, page 97–106, New York, NY, USA. Association for Computing Machinery.
- Ziming Luo, Zonglin Yang, Zexin Xu, Wei Yang, and Xinya Du. 2025. [Llm4sr: A survey on large language models for scientific research](#). *Preprint*, arXiv:2501.04306.
- Hana Matatov, Marianne Aubin Le Quéré, Ofra Amir, and Mor Naaman. 2025. [Examining the prevalence and dynamics of ai-generated media in art subreddits](#). *Preprint*, arXiv:2410.07302.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020a. [On faithfulness and factuality in abstractive summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1906–1919, Online. Association for Computational Linguistics.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020b. [On faithfulness and factuality in abstractive summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1906–1919, Online. Association for Computational Linguistics.
- Medill Local News Initiative. 2024. [Ai and local news: Report 2024](#). Technical report, Northwestern University. Accessed: 2025-10-01.
- Nicole Meir. 2023. [Standards around generative ai](#). The Definitive Source: Behind the News, The Associated Press. Accessed on 2025-10-15.
- Zachary Metzger. 2024. [The state of local news 2024](#). Technical report, Northwestern University. Accessed: 2025-10-01.
- Jan Nehring, Aleksandra Gabryszak, Pascal Jürgens, Aljoscha Burchardt, Stefan Schaffer, Matthias Spielkamp, and Birgit Stark. 2024. [Large language models are echo chambers](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 10117–10123, Torino, Italia. ELRA and ICCL.
- Nic Newman, Amy Ross Arguedas, Craig T. Robertson, Rasmus Kleis Nielsen, and Richard Fletcher. 2025. [Digital news report 2025](#). https://reutersinstitute.politics.ox.ac.uk/sites/default/files/2025-06/Digital_News-Report_2025.pdf. Accessed: 2025-09-30.
- OpenAI. 2022. [Introducing ChatGPT](#). <https://openai.com/blog/chatgpt>. Accessed: 2025-10-06.
- OpenAI. 2024. [Gpt-4.1](#). <https://platform.openai.com/docs/models/gpt-4.1>. Accessed 2025.
- Pangram. 2024. [Pangram multilingual v2: Expansion of ai content detection into multiple languages](#). <https://www.pangram.com/blog/pangram-multilingual-v2>. Pangram Labs blog.
- Pew Research Center. 2025. [News platforms \(news & information\) – american trends panel wave 177: Topline results](#). https://www.pewresearch.org/wp-content/uploads/sites/20/2025/09/PJ_2025.09.25_news-information-fact-sheets_topline.pdf.
- Qwen Team. 2025. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- Damian Radcliffe. 2025. [Journalism in the AI era: Opportunities and challenges in the global south and emerging economies](#). Trf insights report, Thomson Reuters Foundation. Survey of 200+ journalists from 70+ countries.
- Jenna Russell, Marzena Karpinska, and Mohit Iyyer. 2025. [People who frequently use ChatGPT for writing tasks are accurate and robust detectors of AI-generated text](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5342–5373, Vienna, Austria. Association for Computational Linguistics.
- Fabio Salvi, Manoel Horta Ribeiro, Riccardo Gallotti, et al. 2025. [On the conversational persuasiveness of gpt-4](#). *Nature Human Behaviour*, 9(10):1645–1653.
- Philipp Schoenegger, Francesco Salvi, Jiacheng Liu, Xioli Nan, Ramit Debnath, Barbara Fasolo, Evelina Leivada, Gabriel Recchia, Fritz Günther, Ali Zarifhonarvar, Joe Kwon, Zahoor Ul Islam, Marco Dehnert, Daryl Y. H. Lee, Madeline G. Reinecke, David G. Kamper, Mert Kobaş, Adam Sandford, Jonas Kgomo, Luke Hewitt, Shreya Kapoor, Kerem Oktar, Eyup Engin Kucuk, Bo Feng, Cameron R. Jones, Izzy Gainsburg, Sebastian Olschewski, Nora Heinzelmann, Francisco Cruz, Ben M. Tappin, Tao Ma, Peter S. Park, Rayan Onyonka, Arthur Hjorth, Peter Slattery, Qingcheng Zeng, Lennart Finke, Igor Grossmann, Alessandro Salatiello, and Ezra Karger. 2025. [Large language models are more persuasive than incentivized human persuaders](#). *Preprint*, arXiv:2505.09662.
- Minkyu Shin, Jin Kim, and Jiwoong Shin. 2025. [The adoption and efficacy of large language models: Evidence from consumer complaints in the financial industry](#). *Preprint*, arXiv:2311.16466.
- Felix M. Simon. 2024. [Artificial intelligence in the news: How ai retools, rationalizes, and reshapes journalism and the public arena](#). Report, Tow Center for Digital Journalism, Columbia University.
- Felix M. Simon, Rasmus Kleis Nielsen, and Richard Fletcher. 2025. [Generative ai and news report 2025](#):

- How people think about ai’s role in journalism and society. Report, Reuters Institute for the Study of Journalism, University of Oxford.
- Jinyan Su, Claire Cardie, and Preslav Nakov. 2024. Adapting fake news detection to the era of large language models. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 1473–1490, Mexico City, Mexico. Association for Computational Linguistics.
- Zhen Sun, Zongmin Zhang, Xinyue Shen, Ziyi Zhang, Yule Liu, Michael Backes, Yang Zhang, and Xinlei He. 2025. Are we in the AI-generated text world already? quantifying and monitoring AIGT on social media. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 22975–23005, Vienna, Austria. Association for Computational Linguistics.
- Katherine Thai, Bradley Emi, Elyas Masrouf, and Mohit Iyyer. 2025. Editlens: Quantifying the extent of ai editing in text. *Preprint*, arXiv:2510.03154.
- Edward Tian and Alexander Cui. 2023. GPTZero: Towards detection of AI-generated text using zero-shot and supervised methods. <https://gptzero.me>. Accessed: 2025-10-09.
- Benjamin Toff and Felix M. Simon. 2025. “Or They Could Just Not Use It?”: The Dilemma of AI Disclosure for Audience Trust in News. *The International Journal of Press/Politics*, 30(4):881–903.
- UNC Center for Innovation & Sustainability in Local Media. 2020. Newspapers 2020 — active u.s. newspapers as of 2020. <https://usnewsdeserts.cislm.org/>. Accessed: 2025-09-29.
- Katharine Viner and Anna Bateson. 2023. The guardian’s approach to generative ai. Accessed: 2025-10-17.
- Washington Post Staff. 2022. The washington post guide to writing an opinion article. <https://www.washingtonpost.com/opinions/2022/op-ed-writing-guide-washington-post-examples/>. Updated June 23, 2022; accessed 2025-10-16.
- Laura Weidinger, Jonathan Uesato, Maribeth Rauh, Conor Griffin, Po-Sen Huang, John Mellor, Amelia Glaese, Myra Cheng, Borja Balle, Atoosa Kasirzadeh, Courtney Biles, Sasha Brown, Zac Kenton, Will Hawkins, Tom Stepleton, Abeba Birhane, Lisa Anne Hendricks, Laura Rimell, William Isaac, Julia Haas, Sean Legassick, Geoffrey Irving, and Iason Gabriel. 2022. Taxonomy of risks posed by language models. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency, FAccT ’22*, page 214–229, New York, NY, USA. Association for Computing Machinery.
- Eileen Yam and Brian Kennedy. 2025. From political speeches to songs, how would americans react if they found out AI was involved? Pew Research Center, Short Reads.
- Li Zhou, Ruijie Zhang, Xunlian Dai, Daniel Herscovich, and Haizhou Li. 2025. Large language models penetration in scholarly writing and peer review. *Preprint*, arXiv:2502.11193.

A Dataset

In this section of the appendix we provide more details on our data collection process.

recent_news additional collection details. To facilitate automatic data collection, we first obtain a list of 6,175 URLs for American newspapers³⁰ and filter out those that are unreachable and/or do not have active RSS feeds, which yielded 1,528 URLs. Roughly twice a week, from June 15 to September 15, we automatically accessed each RSS feed and downloaded the full text and metadata for up to 50 recently-published articles from each paper.³¹ Each full text article was then preprocessed using the Trafilatura (Barbaresi, 2021) and Newspaper4K³² libraries to strip headers/footers, advertisements, and HTML artifacts from the text.

Topic classification. Each instance was presented to the QWEN3-8B model with the full article text and an instruction to select the single most semantically appropriate topic from the IPTC taxonomy. The complete prompt is shown in Figure 5. To verify, we randomly sampled 100 articles from the recent_news set. Two of the authors independently re-labeled these samples according to the same IPTC taxonomy. Inter-annotator agreement between the two human raters was 87% (Cohen’s $\kappa = 0.85$), reflecting strong consistency in human judgments. Agreement between the QWEN3-8B predictions and the majority human label averaged 77%, indicating moderately strong alignment between the classifier and human annotations.

Topics present in datasets. The classification scheme in this work builds on the 17 top-level categories defined in the International Press Telecommunications Council (IPTC) “Media Topics” taxonomy. We map each article to one of the 17 top-level topics. In addition to those categories, we include two supplementary labels — “Obituary”

³⁰<https://onlinenewspapers.com/usstate/usatable.shtml>

³¹While the majority of recent_news is obtained from RSS feeds, articles from some newspapers (e.g., Washington Post, Wall Street Journal) were accessed via ProQuest Recent Newspapers.

³²<https://github.com/AndyTheFactory/newspaper4k>

Prompt for classifying topic of articles

Classify the following article into the most appropriate primary topic from the IPTC taxonomy.

IPTC Primary Topics:
{taxonomy_text}

Article to classify:
Title: {article.title}
Text: {article_text}

Instructions:

1. Carefully read the article content and identify the main focus
2. Choose the most appropriate primary topic from the taxonomy above
3. Base your decision on the primary subject matter and content focus
4. Consider the description of each category when making your choice
5. If the article truly doesn't fit well into any of the listed categories, choose "Other"
6. Be specific and accurate - don't force a category if it's not a good fit
7. Respond ONLY with the classification in this exact format:

```
<classification>  
Primary Topic: [exact primary topic name]  
</classification>
```

Figure 5: Prompt for classifying topic of articles

and “Other” — to capture content that does not cleanly fall into the predefined classes. Definitions of topics are displayed in Table 4 and the distribution of topics across the three datasets are shown in Table 5.

Circulation definition and data source. In the *U.S. News Deserts Database*, **circulation** refers exclusively to the average number of printed copies distributed per publishing day, excluding any digital readership metrics. These figures are compiled primarily from *Alliance for Audited Media* (AAM) audits when available, and supplemented with publisher self-reports, state press association directories, and industry trade sources for unaudited outlets. As noted in the database’s section “*Dealing with Circulation Limitations*” on <https://www.usnewsdeserts.com/methodology/>, many small newspapers rely on self-reported or infrequently updated numbers, which are treated as approximate indicators of print scale rather than precise measures of audience size. The distribution of newspapers by circulation amounts is depicted in Figure 6.

Data coverage and overlap. Our correlation analysis uses the subset of articles for which both AI-likelihood and circulation are observed. Out of 186 507 articles in total print circulation is available for 101 799 (54.58 %). At the outlet level, the

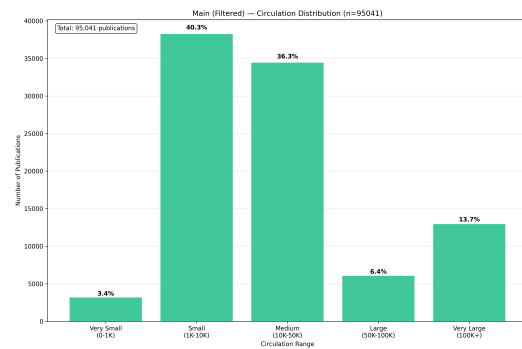


Figure 6: Distribution of circulations of articles in the recent_news dataset.

main dataset contains 1560 unique newspapers, of which 776 (49.74 %) report any circulation figure. Unless otherwise noted, all circulation–AI likelihood correlations are computed on this overlapping set to avoid missing-data biases.

Ownership Details Our ownership information comes from the *Northwestern Local News Initiative* (Medill Local News Initiative, 2024). We collect ownership information for 162,342 articles (87.04% of recent_news) and 1,259 unique newspapers (82.40% of unique newspapers in recent_news). Since ownership of newspapers is transferred frequently, this database shows ownership information current as of November 2025.

Topic	Definition (IPTC summary)
arts, culture, entertainment and media	All forms of arts, entertainment, cultural heritage and media
conflict, war and peace	Acts of socially or politically motivated protest or violence, military activities, geopolitical conflicts, as well as resolution efforts
crime, law and justice	The establishment and/or statement of the rules of behavior in society, the enforcement of these rules, breaches of the rules, the punishment of offenders and the organizations and bodies involved in these activities
disaster, accident and emergency incident	Man made or natural event resulting in loss of life or injury to living creatures and/or damage to inanimate objects or property
economy, business and finance	All matters concerning the planning, production and exchange of wealth.
education	All aspects of furthering knowledge, formally or informally
environment	The protection, damage, and condition of the ecosystem of the planet Earth and its surroundings
health	All aspects of physical and mental well-being
human interest	Item that discusses individuals, groups, animals, plants or other objects in an emotional way
labor	Social aspects, organizations, rules and conditions affecting the employment of human effort for the generation of wealth or provision of services and the economic support of the unemployed.
lifestyle and leisure	Activities undertaken for pleasure, relaxation or recreation outside paid employment, including eating and travel.
politics and government	Local, regional, national and international exercise of power, the day-to-day running of government, and the relationships between governing bodies and states.
religion	Belief systems, institutions and people who provide moral guidance to followers
science and technology	All aspects pertaining to human understanding of, as well as methodical study and research of natural, formal and social sciences, such as astronomy, linguistics or economics
society	The concerns, issues, affairs and institutions relevant to human social interactions, problems and welfare, such as poverty, human rights and family planning
sport	Competitive activity or skill that involves physical and/or mental effort and organizations and bodies involved in these activities
weather	The study, prediction and reporting of meteorological phenomena
<i>Additional labels used in our classification</i>	
Obituary	Memorial and death-notice content about individuals.
Other	Articles that do not clearly align with any IPTC top-level topic.

Table 4: Top-level topics from the IPTC Media Topics taxonomy and definitions (IPTC, 2025).

Topic	recent_news (%)	opinions (%)	ai_reporters (%)
Politics and government	15.02	56.86	19.55
Economy, business and finance	10.53	9.77	2.30
Arts, culture, entertainment & media	20.55	4.87	6.55
Opinion	0.40	4.65	0.75
Health	3.75	4.47	12.92
Human interest	3.61	2.91	6.09
Crime, law and justice	10.91	2.80	7.41
Environment	4.22	2.77	4.36
Education	3.72	2.06	10.27
Science and technology	0.98	1.97	0.72
Society	1.29	1.86	3.71
Conflict, war and peace	0.54	1.72	0.35
Labor	1.54	0.96	1.45
Religion	0.74	0.94	0.69
Sport	11.24	0.69	5.28
Lifestyle and leisure	1.70	0.42	0.50
Disaster, accident, emergency	3.69	0.23	1.25
Obituary	2.48	0.04	1.49
Weather	0.76	0.01	0.37
Other	–	–	14.00

Table 5: Comparison of topic distributions between recent_news, opinions, and ai_reporters.

Lengths of datasets. We see that all datasets are similar lengths, with opinions being slightly longer at an average of 1078.4 tokens, compared to recent news which has an average of 787.4 tokens. Comparison of lengths shown in Figure 7.

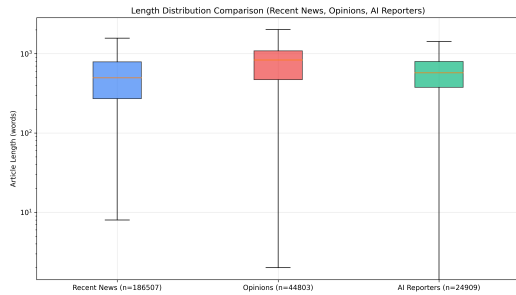


Figure 7: Comparison of the lengths of articles in recent_news, opinions, and ai_reporters

B AI Detection

In this section of the appendix we discuss details of how we used AI Detection models.


Pangram prediction API details. Pangram is a highly accurate AI detection language model (Emi and Spero, 2024). To detect AI, Pangram divides the text into segments, with each segment assigned an AI probability score before an overall confidence score and final label are produced. We use Pangram v2 for all results in this study. However, examples may reflect newer versions of Pangram likelihood.

Condensing the Pangram API labeling scope. To standardize categories across Pangram’s short-text (single pass) and long-text (sliding-window) endpoints, we collapse the vendor’s granular labels into three meta-labels used throughout our analysis: **HUMAN-WRITTEN** = {*Human, Unlikely AI*}; **MIXED** = {*Mixed, Possibly AI, Likely AI*}; and **AI-GENERATED** = {*Highly Likely AI, AI*}.

Comparison to GPTZero. We drew a balanced sample of 1,000 articles from Pangram: 500 labeled *AI / Highly Likely AI* and 500 labeled *Human / Unlikely AI*. We then ran GPTZero on the same texts; GPTZero returns human, mixed, or ai. For evaluation, we binarized the GPTZero outputs by counting mixed as AI and compared the two detectors on this binary task. Under this protocol, we observed 88.2% raw agreement and Cohen’s $\kappa = 0.764$ (118/1,000 disagreements).

Pangram \ GPTZero	<i>Human</i>	<i>AI</i>
<i>Human</i>	490	10
<i>AI</i>	108	392

C Tracking reporter adoption of AI

 ai_reporters is a historical dataset of **20K** articles published by a subset of **10** veteran

reporters who “authored” multiple AI articles in recent_news. Each reporter in this dataset published articles written both before and after the release of ChatGPT (November 2022), enabling longitudinal analysis.

Additionally, we analyze ai_reporters, a longitudinal dataset for 10 veteran reporters (5 male and 5 female) who published articles both before and after ChatGPT’s release (November 2022) and had at least three articles in recent_news labeled as either **AI-GENERATED** or **MIXED**.³³ For each reporter, we scrape all of their available online articles and run the same detection pipeline detailed in §2. We note that since the reporters selected for this analysis have multiple articles flagged for AI use in recent_news, they are not representative of all reporters. Nevertheless, they give us valuable insight into some of the heaviest adopters of AI in modern American journalism.

ai_reporters Several reporters in recent_news published articles both before and after the release of ChatGPT, making them good candidates for a longitudinal analysis to explore when and how they started using AI. We identify a set of 10 veteran reporters from recent_news who meet two criteria: (1) they have published articles prior to November 2022,³⁴ and (2) at least three of their articles were identified as **MIXED** or **AI-GENERATED**. We then write ten custom scrapers, one for each reporter, to collect a corpus of *all* of their published articles that are available online, resulting in a final dataset of 20,132 articles from 14 newspapers.³⁵

AI use by these reporters rises from 0% pre-ChatGPT to 40% in 2025. Figure 8 shows a near-absence in AI use prior to November 2022 before rising sharply to 15.7% in 2023, 36.1% in 2024, and then 40.4% in 2025. This result also

³³All but one reporter are veteran reporters, often with decades of experience. However, due to technical difficulties associated with scraping specific websites, we could not obtain *all* of their older articles. The one early-career reporter started their career in 2021, and thus had less exposure to the profession prior to the release of ChatGPT. The reporters collectively cover a variety of topics ranging from local government, public safety, and environmental justice to national politics, racial equity, LGBTQ+ rights, Caribbean-American culture, economic development, and community sports.

³⁴This cutoff was selected because the public release of ChatGPT in November 2022 (OpenAI, 2022) made AI-assisted writing tools widely accessible.

³⁵One journalist published in multiple newspapers.




Figure 8: AI adoption takes off after ChatGPT’s release in late 2022 with the ten reporters in our ai_reporters sample.

serves as a sanity check on Pangram’s reliability: as all of these reporters published articles pre-ChatGPT that were correctly detected as human-written, it is unlikely that their unique writing styles are a source of false positives.

While AI use varies across this cohort of reporters, none discloses their use. We observe high variance in AI use across these ten reporters. While some show a negligible increase in articles flagged with AI use (from 0% in 2022 to about 2.3% in 2025), others are far more receptive. For instance, 90.1% of the most prolific author’s 2025 articles were classified as MIXED or AI-GENERATED, up from 0% in 2022 (see individual plots in §D.3). Importantly, none of the ten reporters disclose use of AI to their readership.³⁶

AI traces are visible in the articles. Perhaps not surprisingly, we observe qualitative differences between the articles marked for AI use and the human-written ones. For instance, AI-assisted articles in this subset have up to 11.53× more em dashes than those authored entirely by humans (Table 6, Reporter 3). Furthermore, human-written articles tend to be more concrete, naming specific people and locations (Reporter 1) and focus more on factual knowledge than vague statements (Reporter 2).

³⁶We reached out to each reporter in this cohort with a publicly-available email address in an effort to learn more about their AI use, but (perhaps unsurprisingly) we received no responses.

D Additional Results

D.1 recent_news

In this section, we report more results about recent_news, a collection of over 185k articles collected between June 15th and September 15th, 2025.

Overall AI usage in recent_news. recent_news has 90.9% HUMAN-WRITTEN articles, with a total of 9.1% having AI Use. 5.2% is detected as AI-GENERATED and 3.9% is MIXED as shown in Figure 9. The distribution of AI likelihoods per each of the three labels is shown in Figure 10. Data by state is displayed in Figure 11 and Figure 12 and by topic in Figure 13.




Figure 9: Distribution of AI Use predictions in recent_news.




Figure 10: Distribution of AI Likelihoods per each AI Use category in recent_news.

Newspaper-level test (below 100K vs. 100K+ circulation). We collapse to one observation per outlet (share of AI-labeled articles) and compare newspapers below 100K circulation to those at 100K+. The below-100K group (n=750) averages 8.47% AI articles (median 2.93%, SD 16.08, range 0–100),

REPORTER	HUMAN-WRITTEN	AI-GENERATED or MIXED	OBSERVATION
1	<p>"Problems associated with the past California administration seem to have dissipated as the IID and Newsome government officials are on the verge of agreeing to a plan restoring the lower part of the Salton Sea as reported by Water Manager Tina Shields during the April 30 meeting. ..."</p> <p>(2019-05-06) 📄 📅</p>	<p>"The State of California has established the Salton Sea Conservancy under the Salton Sea Conservancy Act, a comprehensive initiative to reverse decades of ecological damage and promote sustainable development in the Salton Sea region. This new state agency, housed within the Natural Resources Agency, is poised to lead the charge ..."</p> <p>(2024-11-18) 📄 📅</p>	<p>The human-written article names specific relevant people, times, and pinpoints the main point of the article, the restoration of the Salton Sea. In the AI article, the conservancy act is tied to lofty statements such as 'decades of ecological damage'.</p>
2	<p>"One shouldn't need a reason to buy Black, it should be a way of life. Such is the thought process for several Black business owners and individuals who hang their hat on the idea of supporting within to keep the Black dollar circulating longer in local neighborhoods, mom-and-pop shops, and places that need it most..."</p> <p>(2023-02-15) 📄 📅</p>	<p>"The spending power of the Black dollar stands at a staggering \$1.7 trillion, reflecting immense potential for community growth and economic empowerment. This financial strength, however, is underutilized, prompting a crucial need for increased and unconditional support within the African American community..."</p> <p>(2023-11-29) 📄 📅</p>	<p>In the human-written article, the reporter uses unique phrases like 'hang their hat' and 'mom-and-pop shops', focusing on specific ways the reader could buy from Black businesses. In the latter, the "author" makes sweeping claims about the financial state of Black businesses in a more general manner.</p>
3	<p>"The UN General Assembly met on Tuesday afternoon in Emergency Special Session on the decades long Israel-Palestine conflict and as the ongoing crisis in Gaza shows no signs of abating. Member States adopted a resolution, demanding an "immediate humanitarian ceasefire", the immediate and unconditional release of all hostages ..."</p> <p>(2024-06-20) 📄 📅</p>	<p>"The Government of Guyana has expressed its support for the United States' call for a ceasefire in the ongoing conflict between Israel and Gaza. In a recent statement, the Guyanese government acknowledged the three-phase plan proposed by U.S. President Joseph Biden on May 31, 2024, aimed at ending Israel's war on Gaza..."</p> <p>(2025-03-21) 📄 📅</p>	<p>In the human-written article, the author references the time (Tuesday) and setting (Emergency Special Setting). In the AI-generated article, the author uses more fluff words and more vague time placement (i.e. ongoing, recent).</p>

Table 6: Excerpts from passages of newspaper articles written by reporters in the ai_reporters dataset. Words and phrases identified as indicative of AI use by Pangram are highlighted in red. In the left HUMAN-WRITTEN column, excerpts of older, human-written articles are displayed, while the AI-GENERATED or MIXED column shows newer articles by the same author detected as AI-generated or assisted. When AI use is present, articles by these reporters include fewer specific details, broader time markers, and loftier language.

while the 100K+ group (n=24) averages 4.95% (median 1.75%, SD 7.14, range 0–24.74). The difference in means is +3.52 percentage points (Cohen's $d = 0.22$). A Welch t -test indicates a statistically significant gap ($t = 2.24$, $p = 0.032$), whereas a Mann–Whitney U test does not ($U = 9506$, $p = 0.634$). The divergence reflects heavy right-skew and zero-inflation in the outlet-level shares: mean-based tests are more sensitive to a few high-AI smaller outlets, while rank-based tests emphasize the bulk of the distribution. In other words, higher AI usage is *not uniform* across smaller outlets; it is concentrated in a subset of them, while many smaller outlets have low (often zero) AI shares. Overall, we read this as modest evidence that smaller-circulation newspapers exhibit higher AI usage on average. In Figure 15, the rate of newspapers labeled as AI-GENERATED or MIXED is much lower in papers with circulation above 100k.

AI use in print. Many AI-GENERATED articles in recent_news appear in print in addition to the online editions. According to a 2025 Pew study, the print audience still includes an estimated 65 million Americans who read physical newspapers “often” or “sometimes” (Pew Research Center, 2025).

Table 7: Language distribution

Language	% (n)
English	95.16% (177,478)
Spanish	3.88% (7,235)
Portuguese	0.25% (468)
Vietnamese	0.22% (403)
French	0.18% (343)
Polish	0.17% (314)
Other	0.15% (271)

Other includes: Russian (163), German (37), Yiddish (27), Indonesian (25), Turkish (10), Dutch (5), Latin (2), Ukrainian (1), Chinese (1). Total: 186,512 articles.

Critically, this group is heavily skewed towards an older demographic less likely to be technologically savvy, with a median age of 57.9 (Conaghan, 2017). We further note that readers of print newspapers face higher barriers to verifying content, in contrast to digital readers who can simply copy and paste text into search engines. Thus, newspapers that *only* appear in print are much harder to audit using our methodology.

Detecting Hallucinated Content We manually reviewed 200 articles: 100 articles predicted as AI-GENERATED and 100 predicted as HUMAN-WRITTEN.

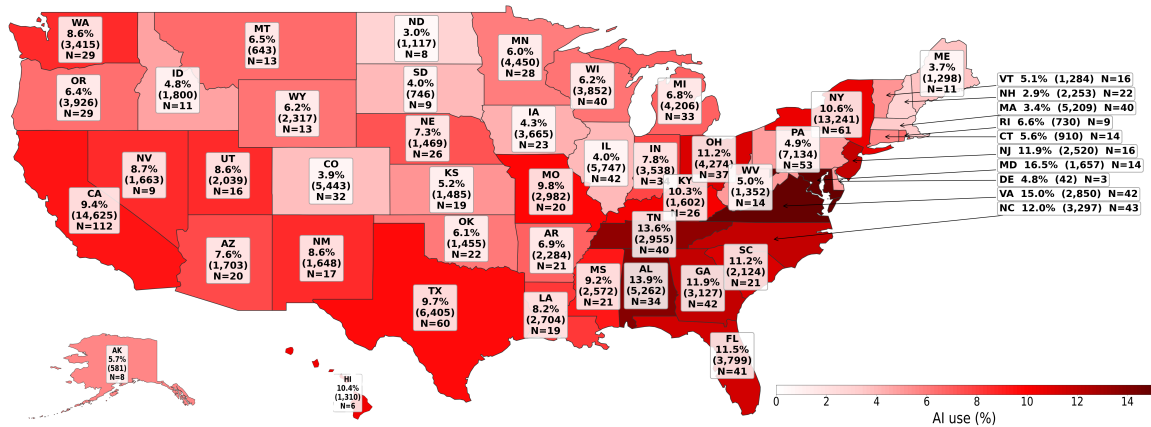


Figure 11: Map of the United States showing the AI use (%), number of articles, and number of unique newspapers for each state. States in the mid-Atlantic and southern US exhibit markedly higher AI use than other states. Note that this plot only considers articles written in English (see Figure 12 for others).

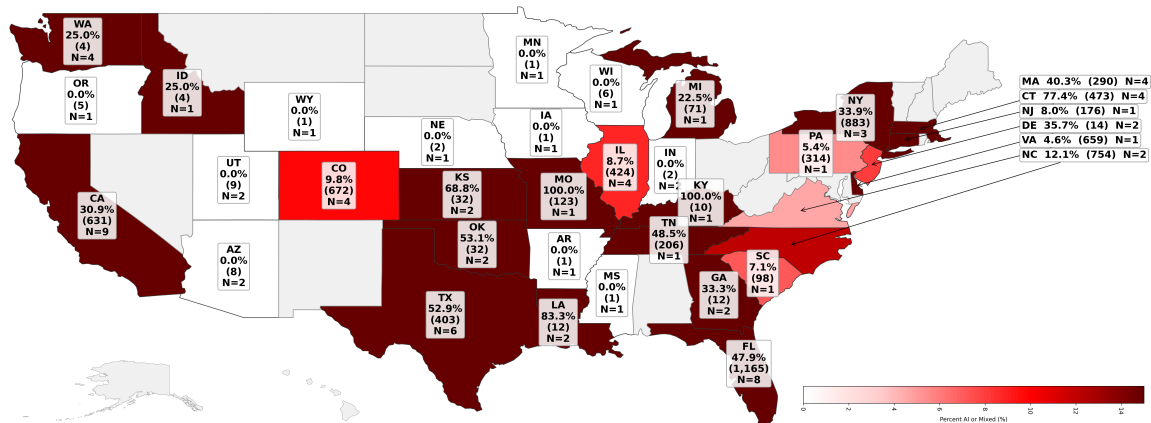


Figure 12: AI use in articles written in languages other than English. The most common language are Spanish, Portuguese, and Vietnamese.

41% of **AI-GENERATED** and 5% of **HUMAN-WRITTEN** contained factual inaccuracies referred to as hallucinations. Some examples of hallucinations found in **AI-GENERATED** articles are depicted in Table 8.

Detecting AI in quoted speech We analyze 30,462 articles with available quote-level predictions. Each quote is evaluated with an AI-likelihood classifier, restricted to spans of at least 50 tokens to reduce false positives on very short fragments. We treat quotes labeled as *Unlikely AI* or *Human* as human-written, and all others (Highly Likely AI, AI, Likely AI, Mixed, Possibly AI) as AI-generated. At the article level, we compare the incidence of AI quotes in AI-flagged versus human-flagged articles, yielding $P(\text{AI quote}|\text{AI article}) = 0.239$ and $P(\text{AI quote}|\text{human article}) = 0.034$. Conversely, the probability that an article is AI-generated given the presence of an AI quote is

$P(\text{AI article}|\text{AI quote}) = 0.293$. In total, 1,376 articles contain at least one AI-generated quote, and 403 articles contain both AI-generated narrative and AI-generated quotes. To test whether length influences detection, we also compare long quotes (≥ 120 words, $n = 870$) to short quotes, finding no significant difference in AI likelihood (two-sample t -test, $p = 0.96$).

Many AI-generated articles contain authentic quotes. To examine whether articles labeled as **MIXED** and **AI-GENERATED** include fabricated information, we analyze the authorship of quotations in these articles. Specifically, we extract all quotes in the dataset longer than 50 words, and run each of them through Pangram individually. Note that Pangram’s reliability degrades on shorter texts, and so we are unable to perform this analysis on all quotes. Within the subset of articles that include at least one quote >50 words long,

NEWSPAPER	ARTICLE EXCERPT	HALLUCINATION
LA SENTINEL	“We know that 52 percent of adult women across this country are on Medicaid, and that number is exponentially higher in Los Angeles” 🗞️	The percent of woman on medicaid is 19%, not 52% (Source). The article also contains other hallucinations such as referring to ‘former President Trump’ in 2025, when he is the current U.S. president.
SANTA FE REPORTER	“The program is currently set to expire in June 2030 under the reconciliation bill, though advocates continue pushing for longer extensions and inclusion of additional affected communities.” 🗞️	The program is set to expire in 2027, not 2030 as reported in the article (Source).
SUNCOAST NEWS	“‘We’re not trying to silence anyone,’ Koulianos said. ‘We’re trying to conduct the people’s business in a reasonable time-frame.’” 🗞️	Most of the quotes in this article, including the excerpt here, cannot be found in the transcript of the Board of Commissioners meeting (Source). Most quotes are paraphrases or similar quotes, but do not have the same wording.
THE ADVOCATE-MESSENGER	“It still shows up, five days a week in print and daily online, to deliver what no outside media outlet ever could: the truth of this town, in this moment, for these people.” 🗞️	In an article detailing the history of their own paper, they incorrectly state the paper is printed 5 times a week, when in reality it is printed twice a week. (Source).

Table 8: Examples of hallucinations found in articles predicted as **AI-GENERATED**.




Figure 13: AI use by topic in recent_news. Weather, science/technology, and health exhibit higher AI use than topics like conflict, crime/justice, and religion.

76.1% of articles flagged with AI use contain at least one human-written quote. This suggests that many stories written with AI use rely on authentically sourced material. However, it remains unclear whether journalists are choosing these quotes and feeding them into a prompt for AI generation, or if the AI is also doing quote selection (see Table 3 for examples).

Effects of republishing articles. We examine how AI-generated content propagates via redistribution. Using recent_news, we cluster articles by cosine similarity and define *exact duplicates* for scores ≥ 0.95 and *semantically similar* for $0.85 \leq \text{score} < 0.95$. In total, 16 580 articles

meet these criteria, forming 6413 clusters (mean size = 4.16). The *Associated Press* appears most frequently, in 1664 clusters and 4030 redistributed articles, consistent with syndication as the dominant driver of redistribution. Across all duplicates, 6.7% are labeled AI-generated, with a lower rate for exact duplicates (3.1%) and a markedly higher rate for semantically similar duplicates (14.3%). Figure 16 summarize the resulting cluster-size distribution; Table 9 lists the top content providers present in duplicate coverage.

Rank	Author	Count
1	Associated Press	4,011
2	Unknown	1,204
3	Staff Reports	701
4	Advance Local Weather Alerts	404
5	USA Today Network	264
6	Myedmondnews	104
7	Teresa Wippel	101
8	WP Block Co-Authors	98
9	Cascade PBS Staff	82
10	Grace Gilson	81

Table 9: Top 10 authors by frequency of appearances in duplicate articles.

Effects of machine translation on AI-text detection. As depicted in Figure 2, AI-use is more frequent in languages other than English. To better understand if this is due to Pangram picking up on texts that were automatically translated into other languages from English or written in-part or fully using AI, we take a sample of 300 English detects with the following distribution: 100 **HUMAN-WRITTEN**, 100 **MIXED**, and 100 **AI-GENERATED**. The articles are translated using GPT4.1 (OpenAI, 2024) in chunks of 2000 tokens. The translation prompt is displayed in Figure 17. If the article is less than 2000 tokens, the whole article is translated at once. If it is longer, the article is chunked into chunks of 2000 tokens

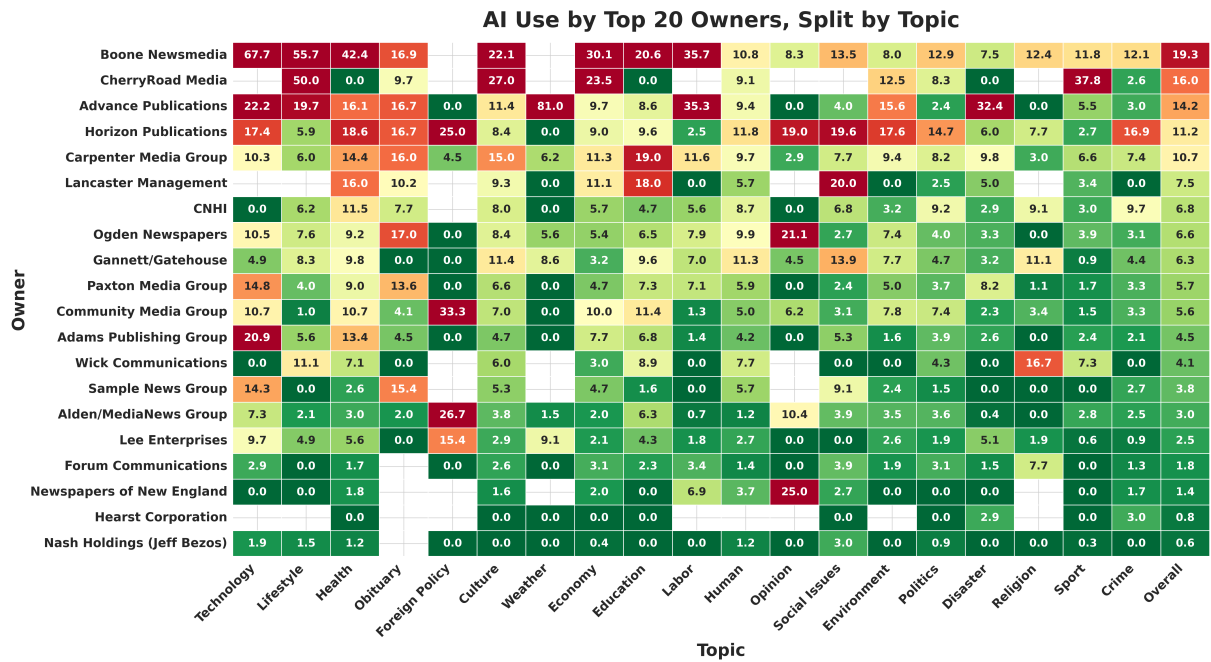


Figure 14: Heatmap of AI use by publication owner and article topic in recent_news. Some owners disclose AI use for specific content, such as [Advance Publications](#) for weather reports, but the biggest adopters, such as [Boone Newsmedia](#), use AI broadly across many topics. Note that only 87% of recent_news has ownership information, and only topics that have at least 5 articles per owner are visualized in the heatmap.




Figure 15: Distribution of AI Categories in recent_news articles between papers with circulations below 100k (left) and above 100k (right).

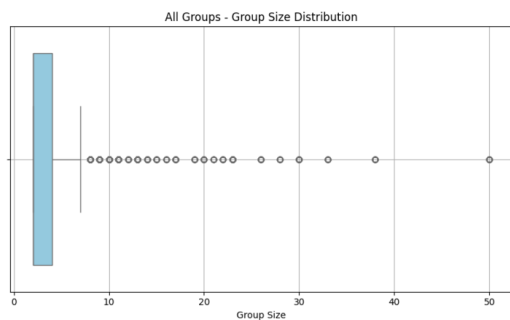


Figure 16: Distribution of duplicate-cluster sizes in recent_news. Most clusters are small, with a minority of large clusters driving redistribution volume.

or less, respecting paragraph boundaries, and then translated. We translate the articles into the 12 languages: Arabic, Chinese (Simplified), Dutch, French, German, Italian, Japanese, Polish, Portuguese, Russian, Spanish, and Vietnamese. After translation, we find the average AI-likelihood decreases by an average of 0.13, with Japanese having the smallest decrease of 0.08 (Table 10). We see that agreement using Pangram is high but imperfect, ranging between 74.7% (Vietnamese) to 88.7% (Portuguese). All languages have a McNemar P lower than 0.01 except Japanese, indicating that translation suppresses some AI-use. We note that this experiment uses a single MT model (GPT-4.1); results may differ with other translation methods or models, and the effects of human translation on AI detection remain an open question.

Topic-level variation is unlikely to be driven by detector bias alone. One concern is that the uneven distribution of AI labels across topics could reflect differences in detector performance rather than differences in AI adoption. We use opinions to test whether topic-level uniformity differs in 2022, before widespread LLM-assisted writing, versus 2025. If detector bias were the main driver of topic-level differences, we would expect to observe similar topic variance in both years. In 2022,

Language	Δ AI Likelihood	Binary Match (%)	FP	FN	McNemar p
Arabic	-0.12	82.0	11	43	< 0.001
Chinese (Simpl.)	-0.12	82.0	14	37	0.002
Dutch	-0.17	82.0	4	50	< 0.001
French	-0.11	83.0	9	35	< 0.001
German	-0.16	82.3	5	48	< 0.001
Italian	-0.13	83.3	4	36	< 0.001
Japanese	-0.08	82.3	28	25	0.78
Polish	-0.19	82.0	9	49	< 0.001
Portuguese	-0.10	88.7	7	27	0.001
Russian	-0.14	82.3	6	45	< 0.001
Spanish	-0.09	88.0	4	32	< 0.001
Vietnamese	-0.21	74.7	10	66	< 0.001
Mean	-0.13	83.2			

Table 10: Binary effects of machine translation on AI-use detection. Δ AI likelihood denotes the mean change in predicted AI likelihood (translated minus English). Binary match reports agreement with English labels under a human vs. AI-use formulation. FP and FN denote false positives and false negatives relative to English. McNemar’s exact test evaluates directional asymmetry in disagreements (FDR-corrected).

Prompt for translating language of articles

Translate the following English text to { TARGET LANGUAGE}.
 Preserve the original meaning, tone, and style as much as possible.
 Only return the translated text, without any explanations or additional comments.

Text to translate:
 {TEXT}

Figure 17: Prompt for translating articles to other languages

across 4,344 articles, only 5 were flagged (0.12%). A permutation chi-squared test with 10,000 samples yielded a non-significant result ($p = 0.1101$), indicating no significant topic-level divergence. In 2025, across 7,932 articles, 267 were flagged (3.37%). The same permutation chi-squared test was significant ($p = 0.0002$), showing strong topic-level divergence. We use a permutation chi-squared test because expected flagged counts in most topics are near zero in 2022, violating the large-sample assumptions of the standard asymptotic test. Taken together, topic uniformity in 2022 and topic divergence in 2025 support the interpretation that the detector is capturing real differential AI adoption across topics rather than fixed topic-specific detector artifacts.

D.2 opinions

Overall AI usage in opinions. opinions has 99% **HUMAN-WRITTEN** articles, with a total

of 1% having AI Use. 0.11% is detected as **AI-GENERATED** and 0.85% is **MIXED** as shown in Figure 18. The distribution of AI likelihoods per each of the three labels is shown in Figure 19.




Figure 18: Distribution of AI Use predictions in opinions.

D.3 ai_reporters

In section §C, we talk about the aggregated trends for the 10 reporters identified as using AI in our data set, whose data we analyzed longitudinally. Here we present more details about the reporters, label distribution, and patterns observed in individual reporters.

Reporter Profile. Figure 8 shows reporter statistics. We have anonymized the data, as we want to show only the diversity of expertise the reporters in ai_reporters have rather than highlight specific reporters. Nine out of ten reporters in this data had at least 10 years of experience.

Outlet	Opinions (AI+mixed)	recent_news (AI+mixed)	Ratio
WSJ	4.99%	0.74%	6.8×
WaPo	5.51%	0.55%	10.1×
NYT	2.94%	1.80%	1.6×
Pooled	4.56%	0.71%	6.4×

Table 11: Opinion vs. recent_news AI rates, 2025-06-01–2025-09-15.




Figure 19: Distribution of AI Likelihoods per each AI Use category in opinions.

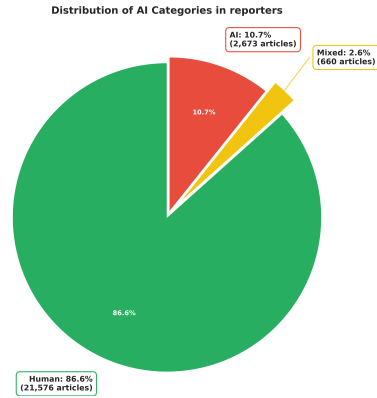


Figure 21: Distribution of AI Use predictions in ai_reporters.




Figure 20: Distribution of AI use predictions in ai_reporters by year.

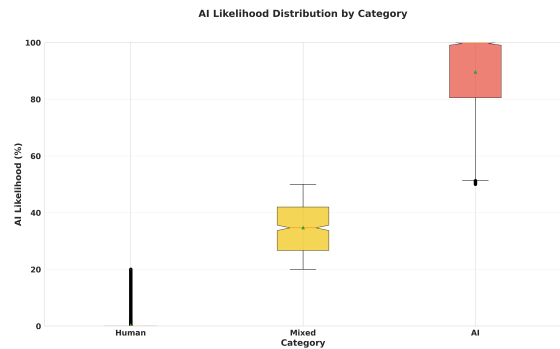


Figure 22: Distribution of AI Likelihoods per each AI Use category in ai_reporters.

Label distribution in ai_reporters. ai_reporters has 86.6% **HUMAN-WRITTEN** articles, with a total of 13.3% having AI Use. 10.7% is detected as **AI-GENERATED** and 2.6% is **MIXED** as shown in Figure 21. The distribution of AI likelihoods per each of the three labels is shown in Figure 22. More importantly, all **AI-GENERATED** labels and almost all **MIXED** are concentrated in the last three years (i.e., 2023–2025; see Figure 20)

AI use by reporter. We report AI use by reporter in Figure 23 and Figure 24. We note a limitation in our data collection: reporters often write for multiple outlets, and accessing historical articles from some of these sources is difficult. Consequently, we could not retrieve all possible articles for ev-

ery reporter. We did, however, ensure we gathered most of their recent articles and as many as possible from before November 2022 (the ChatGPT release date).







































<i>Publication</i>	<i>Label</i>	<i>Author role</i>	<i>Title</i>
U.S. Politics & Governance			
WaPo	AI-GENERATED	Political writer and former congressional speechwriter	MAGA Maoism is spreading through the populist right.  
WaPo	MIXED	Former U.S. Senator	My fellow Republicans, the responsibility to speak out rests with you.  
WaPo	MIXED	Urban warfare scholar	As the border wall goes up, the underground threat will grow.  
WSJ	MIXED	Legal activist and think tank fellow	The Gates Foundation Put Its Tax Exemption at Risk.  
Public Health			
NYT	MIXED	Immunologist	America Is Abandoning One of the Greatest Medical Breakthroughs.  
WSJ	MIXED	Health-longevity company founder	A Chatbot Can Be Part of Your Medical Team.  
WSJ	AI-GENERATED	Physician, Fox News commentator	The Madness in RFK Jr.'s Autism Method  
War & National Security			
WSJ	AI-GENERATED	Former Canadian international trade minister	Canada Is the Best Friend America's Got.  
NYT	AI-GENERATED	Retired U.S. Army general	Be Not Afraid.  
WaPo	MIXED	Space research CEO	Nuclear-powered rockets will win the new space race.  
Cybersecurity			
WaPo	MIXED	U.S. Representative & CEO of TFG Asset Management	Cyber warfare has arrived. Here's the United States' best defense.  
WaPo	MIXED	AI data-center CEO	AI extremists are peddling science fiction.  
NYT	MIXED	Principal Product Manager	Ye and the Limits of Free Speech  
Technology			
WSJ	MIXED	Former Intel CEO	A Sovereign-Wealth Fund to Keep America's Technological Edge;  
NYT	MIXED	Therapist	I'm a Therapist. ChatGPT Is Eerily Effective.  
WSJ	MIXED	Investment firm founder	The AI Boom May Be Too Good to Be True.  
Law & Justice			
NYT	AI-GENERATED	Journalist, ex-Navy SEAL	What the 'Rust' Shooting Case Is Really About.  
WSJ	AI-GENERATED	Governor and sociology professor	Utah Hands Parents the Keys to the App Store.  



Table 12: A sample of op-eds labeled as either **MIXED** or **AI-GENERATED** from WSJ, WaPo, and NYT. Author names are omitted and replaced with professional roles. These op-eds often address polarizing topics such as politics, war, and public health, making the disclosure of AI imperative. All examples included produced the same labels on Pangram v2 and v3.2.

REPORTER	YEARS ACTIVE	AI USE IN 2025 (%)	TOPIC COVERAGE
1	25+ years	90.1% _{519/576}	National politics, civil rights, racial justice, Black affairs
2	20+ years	84.6% _{111/13}	Regional news, agriculture, water policy, environment
3	10+ years	78.3% _{90/115}	Caribbean-American news, culture, health, diaspora
4	40+ years	72.7% _{8/11}	Opinion on policy, education, creativity, technology
5	20+ years	38.1% _{48/126}	Local news, public safety, city government, environment
6	30+ years	30.4% _{157/517}	Environmental justice, racial equity, misinformation (opinion)
7	50+ years	11.7% _{19/162}	Hyper-local news, history, community affairs
8	10+ years	6.9% _{27/393}	LGBTQ+ news, politics, legal issues, culture
9	<5 years	3.3% _{5/153}	Business, Black entrepreneurship, economic development
10	~10 years	2.3% _{3/130}	Local sports, high school athletics, youth leagues



Table 13: Reporter Profiles: Years active, aggregate AI use in 2025 ($\%_{ai_flag/total}$), and topic coverage. Reporters are arranged from most to least prolific AI user.





(a) Reporter 1





(b) Reporter 2





(c) Reporter 3



(d) Reporter 4

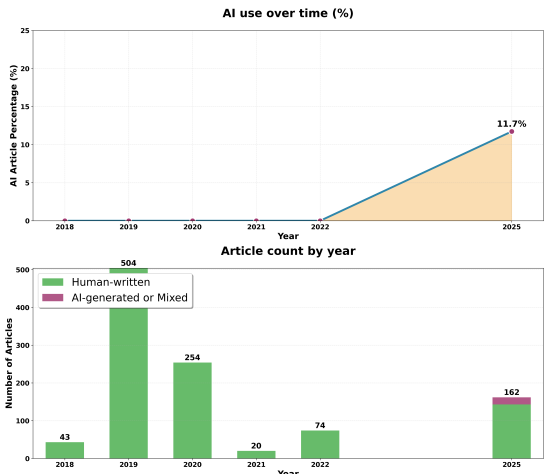


(e) Reporter 5




(f) Reporter 6


Figure 23: AI content patterns in 2025 by reporters (part 1). See Table 13 for profile alignment.




(a) Reporter 7



(b) Reporter 8



(c) Reporter 9



(d) Reporter 10

Figure 24: AI content patterns in 2025 by reporters (part 2). See Table 13 for profile alignment.