MASTERARBEIT | MASTER'S THESIS

Titel | Title

Emotional Tone and Foresight in Media Discourse on Artificial Intelligence (2018–2024): A Longitudinal Analysis of GPT Release Events

verfasst von | submitted by

Daniel Paluch BSc

angestrebter akademischer Grad | in partial fulfilment of the requirements for the degree of

Master of Science (MSc)

Wien | Vienna, 2026

Studienkennzahl lt. Studienblatt | Degree programme code as it appears on the student record sheet:	UA 066 840
Studienrichtung lt. Studienblatt | Degree programme as it appears on the student record sheet:	Masterstudium Psychologie
Betreut von | Supervisor:	Univ.-Prof. Mag. Dr. Claus Lamm
Mitbetreut von | Co-Supervisor:	Mauricio Martins PhD

Abstrakt

Keywords: Künstliche Intelligenz, Medienstimmung, Zukunftsblick, NLP, GPT, politische Orientierung

Die vorliegende Studie untersucht, wie sich emotionale Tonalität und zukunftsbezogene Sprache in der US-amerikanischen Berichterstattung über Künstliche Intelligenz (KI) zwischen 2018 und 2024 verändert haben und ob sich diese Verläufe nach politischer Ausrichtung der Medien unterscheiden. In einem longitudinalen Beobachtungsdesign wurden 14.752 US-amerikanische Nachrichten- und Magazinartikel aus 14 Medien über ProQuest erhoben und in linke, zentristische und rechte Outlets gruppiert. Als sprachliche Variablen dienten LIWC-22 Tone sowie zwei Indizes zukunftsbezogener Sprache, wobei sich die konfirmatorischen Analysen auf positive versus negative *Foresight* konzentrierten. Polarisierungssignale zeigten sich am deutlichsten in der langfristigen Entwicklung zukunftsbezogener Sprache, wobei insbesondere zentristische Medien den Anstieg positiver relativer *Foresight* trugen, während kurzfristige Ereigniskontraste eher in der affektiven Tonalität als in der *Foresight* sichtbar wurden. Lineare gemischte Modelle schätzten langfristige Zeittrends, Ereignisfenster um die Releases von GPT-1, GPT-3 und ChatGPT sowie Moderation durch politische Orientierung. Die emotionale Tonalität nahm über die Zeit leicht, aber verlässlich ab, während positive relative *Foresight* zunahm. Der Rückgang der Tonalität unterschied sich nicht nach politischer Orientierung, der Anstieg der *Foresight* war jedoch in zentristischen Outlets signifikant stärker als in linken oder rechten. Ereignisfensteranalysen ergaben nur begrenzte Hinweise auf release-spezifische Effekte: Nur GPT-3 war mit einem signifikanten post-release Rückgang der Tonalität verbunden, während sich für Foresight keine verlässlichen Prä-Post-Unterschiede zeigten. Insgesamt deuten die Befunde darauf hin, dass die KI-Berichterstattung affektiv etwas negativer, zugleich aber in ihrer Zukunftsrahmung chancenorientierter wurde.

Abstract

Keywords: artificial intelligence, media sentiment, foresight, NLP, GPT, political orientation

This study examined how emotional tone and future-oriented language in U.S. news coverage of artificial intelligence (AI) changed between 2018 and 2024, and whether these trajectories differed by outlet political orientation. Using a longitudinal observational design, 14,752 U.S. news and magazine articles from 14 outlets were retrieved via ProQuest and grouped into left-, center-, and right-leaning categories. Linguistic outcomes were LIWC-22 Tone and two future-oriented indices, with confirmatory analyses focusing on positive versus negative foresight. Polarization signals were strongest in how future-oriented language evolved over time, with centrist outlets driving the increase in positive relative foresight, whereas short-window event contrasts emerged primarily in affective tone rather than foresight. Linear mixed-effects models estimated long-run time trends, event-window changes around GPT-1, GPT-3, and ChatGPT, and moderation by political orientation. Emotional tone showed a small but reliable decline over time, whereas positive relative foresight increased. The tone decline did not differ by political orientation, but the foresight increase was significantly stronger in centrist outlets. Event-window analyses yielded limited support for release-specific effects: only GPT-3 showed a significant post-release decline in tone, whereas foresight showed no reliable pre-post changes across releases. However, short-window tone effects varied by political orientation, especially with ChatGPT, with more positive post-release changes in centrist and left-leaning outlets than in right-leaning outlets. Overall, the findings suggest that AI news discourse became modestly more negative in affective tone while simultaneously more opportunity-oriented in future framing. Interestingly, centrist outlets became particularly optimistic over the long term and positive after ChatGPT was released.

Introduction

Advances in technology have repeatedly reshaped human societies, transforming how people live, work, and make sense of the world. From steam engine and mechanized production to digital computing and the internet, major technological waves have altered social organization, economic structures, and collective imaginaries (Gregersen, 2025; Roser & Ritchie, 2023). For example, the printing press transformed literacy, access to knowledge, and the distribution of political power. Building on this, later phases of mechanization and automation reconfigured everyday relations between humans and machines (Gregersen, 2025; Teixeira & Jones, 2021). Such transformations have often been accompanied by ambivalent public responses that combine optimism about progress with concerns about control, dehumanization, and social displacement (Roser & Ritchie, 2023; Teixeira & Jones, 2021). Technological change is therefore not only a technical process but also a cultural and political one, shaped by how societies communicate, evaluate, and anticipate technological futures (Nelkin, 1995).

Against this backdrop, the contemporary rise of artificial intelligence (AI) can be understood as the latest iteration of this recurring pattern. Earlier technologies transformed physical labor, communication, and access to information. Whereas AI increasingly affects cognitive work by influencing how information is produced, distributed, and interpreted. Historically, AI has developed from early symbolic and rule-based approaches to contemporary machine learning systems and large language models capable of generating human-like text (Roser, 2022; West & Allen, 2018). Importantly, this trajectory is not only technical. It also marks a shift in public visibility. Systems once confined largely to expert and research contexts have become widely discussed technologies with immediate relevance to education, work, governance, and everyday media use. As AI becomes more publicly salient and politically contested, questions of how it is narrated in public discourse become increasingly consequential.

News media occupy a central position in this process. They do not simply document technological change but actively shape how new technologies are interpreted by selecting angles and foregrounding consequences. Simultaneously, these developments have been influencing broader moral and political narratives. Framing research has long shown that such selective emphasis influences how audiences understand causes, consequences, and appropriate responses to complex issues (Entman, 1993). Media coverage is therefore not only descriptive but also constitutive. It helps stabilize which futures appear promising, threatening, plausible, or desirable. In addition, media reporting is one arena in which societies articulate and contest “imagined futures”. These shared expectations consist of what technology will bring, for whom, and at what cost (Jasanoff & Kim, 2009). This is especially relevant for AI, where public understanding often depends less on direct experience than on mediated accounts of what AI is, what it may do, and which societal trajectories it is expected to produce.

A growing body of research has examined how AI is represented in news and public discourse. Across different national contexts and time periods, content analyses generally do not support the idea of uniformly alarmist AI reporting. Instead, opportunity-oriented narratives such as innovation, productivity, and competitiveness often coexist with recurring concerns about ethics, surveillance, labour displacement, and social inequality (Garvey & Maskal, 2020; Ouchchy et al., 2020; Sanguinetti, 2025). This suggests that AI coverage is better understood as a mixed and evolving discourse than as a simple positive-versus-negative media narrative. Research also shows that the thematic composition of AI coverage changes as AI diffuses into everyday institutions, governance, ethics, and public-interest concerns. With a high interest in these prominent sectors, political attention increases (Nguyen, 2024; Brantner & Saurwein, 2021).

More recent work has strengthened this temporal perspective. Longitudinal research on technology-focused news media suggests that AI coverage has become more emotionally polarized over time. Both positive and negative sentiment increased in the post-2022 period as topics such as governance, regulation, and investment gained prominence (Jain & Ranganathan, 2025). Related large-scale analyses of generative-AI coverage likewise point to shifting sentiment and topic structures across time and place, especially after the public diffusion of ChatGPT (Xian et al., 2024). At the level of elite and quality news, newer studies show that AI is often framed less through simple hype than through a more cautious negotiation of present-day uses, societal impacts, and responsibility claims (Vicsek et al., 2026). Generative AI is also increasingly discussed in terms of its cultural, social, economic, and political consequences, including who should be held responsible for associated ethical concerns (Lee & Park, 2026). Research on quality news coverage further suggests that the launch of ChatGPT coincided with a notable shift in both the quantity of reporting and the prominence of frames extending beyond technology itself toward regulation, public consequences, and journalistic implications (Hendrickx & Van Coppenolle, 2026). Taken together, these studies indicate that changes in AI discourse are not only about whether coverage becomes more positive or more negative, but also about how AI-related futures are narrated.

In ideologically structured media environments, news outlets do not interpret issues in the same way. How issues are interpreted is shaped by incentives, audience segmentation, and political alignment. These can produce durable differences in issue emphasis and emotional tone across the media spectrum (Budak et al., 2016; Gentzkow & Shapiro, 2010; Iyengar & Hahn, 2009; Stroud, 2011). Empirically, more targeted studies indicate that portrayals can differ by political orientation and partisan media environment, including differences in sentiment and the specific concerns foregrounded (Yi et al., 2024; Ittefaq et al., 2025). More directly, Chang (2025) found that media partisanship is associated with systematic differences in AI issue coverage. This supports the expectation that aggregate temporal patterns may conceal ideological divergence in both the evaluation of AI and the framing of AI-related futures.

Beyond general media framing, the present study focuses on two related but analytically distinct dimensions of discourse: emotional tone and future-oriented language. Emotional tone concerns the overall affective valence with which a topic is discussed. Meaning whether coverage is expressed in more positive, neutral, or negative terms. Future-oriented language captures a different aspect of discourse by addressing how coverage projects AI into anticipated consequences, possibilities, and risks. This distinction is especially relevant in the context of emerging technologies, because news media not only assess present developments. They also construct expectations about what these technologies are likely to mean for society, politics, and everyday life. A text may therefore adopt a cautious or critical tone while still framing AI's future in terms of opportunity. Distinguishing between tone and future-oriented language thus enables more precise examination of whether changes in AI discourse concern present evaluation, future expectations, or both.

Seen in this light, the existing literature still leaves an important gap. Some studies focus on technology-focused media, others on elite news outlets, and still others on specific dimensions of AI coverage such as framing, sentiment, or responsibility attribution (Chang, 2025; Jain & Ranganathan, 2025; Vicsek et al., 2026). What remains comparatively rare are studies that examine long-run change in emotional tone and future-oriented language simultaneously within a multi-year corpus design, while also accounting for outlet-level clustering and testing whether such trajectories differ by political orientation. This matters because changes in tone need not parallel changes in future framing. Coverage could become more cautious in affective terms while still becoming more opportunity-oriented in how it narrates AI's future. Likewise, patterns that appear uniform at the aggregate level may in fact be driven by ideological segments of the media system. A design that combines long-run trend estimation with explicit tests of ideological moderation is therefore needed to determine whether AI news discourse changed in a common direction across outlets or along politically differentiated trajectories.

The current study

The present study addresses this gap by analyzing a large corpus of AI-related U.S. news and magazine articles published between 2018 and 2024. Using a longitudinal observational approach, it examines whether emotional tone and the valence of future-oriented language changed systematically over time and whether these developments differed across outlet political-orientation groups. In addition, because major technological releases can create short-run shifts in news attention and framing, the study tests whether such changes occurred around key GPT release events (GPT-1, GPT-3, and ChatGPT) and whether these event-proximate shifts varied across ideological contexts. By combining multi-year trend analysis with event-based comparisons and explicit tests of political moderation, the study aims to clarify how AI was evaluated in mainstream news as large language models became increasingly visible in public life, and whether this emotional trajectory unfolded differently across the political media spectrum.

Building on the preceding literature, the present study focuses on whether AI news discourse changed over time and whether these changes were shaped by outlet ideology. Furthermore, major GPT release events are of interest as they may have produced short-run shifts in coverage. Prior research suggests that AI reporting is not uniform but instead combines opportunity-oriented narratives with recurring concerns about ethics, surveillance, labor displacement, and governance (Garvey & Maskal, 2020; Ouchchy et al., 2020; Sanguinetti, 2025). At the same time, more recent work indicates that AI coverage has become more politically salient and increasingly oriented toward questions of regulation, responsibility, and societal consequences, while sentiment and framing may vary across partisan media environments (Brantner & Saurwein, 2021; Ittefaq et al., 2025; Nguyen, 2024; Yi et al., 2024; Chang, 2025). Taking together, this literature suggests that long-run changes in emotional tone and future-oriented language may occur overall, but not necessarily in the same way across ideological groups.

A second expectation follows from the event-like character of major technological releases. News coverage often intensifies around highly salient moments, and such spikes can temporarily shift framing priorities and emotional emphasis (Owen, 2019; Schema Design & Google News Initiative, 2018). In the present context, the releases of GPT-1, GPT-3, and ChatGPT constitute comparatively clear milestones in the public visibility of generative AI. If such moments function as focusing events, they should be reflected not only in increased attention but also in short-run shifts in emotional tone and future-oriented language after GPT version releases. Because ideological environments shape which implications of AI are foregrounded, these shifts pre vs post release may also differ across outlet political-orientation groups (Entman, 1993; Gentzkow & Shapiro, 2010; Iyengar & Hahn, 2009; Stroud, 2011).

Finally, the GPT milestones differed in how visible and societally legible they were. Compared with GPT-1 and GPT-3, which initially circulated more strongly within technical and specialist contexts, ChatGPT was launched as a public-facing conversational system and rapidly entered mainstream attention (Hu, 2023; Similarweb, 2023; UBS Evidence Lab, 2023). Recent research also suggests that generative AI, and ChatGPT in particular, intensified debates about journalism, authorship, automation, and public consequences

(Breazu, 2024; Lewis et al., 2024; Munoriyarwa, 2025; Hendrickx & Van Coppenolle, 2026; Lee & Park, 2026). This suggests that ChatGPT functioned as a stronger trigger of short-run change in public discourse than earlier GPT releases, and that this relative effect may itself vary by political orientation.

Based on this theoretical and empirical background, the present study tested the following hypotheses:

H1. Between 2018 and 2024, AI coverage will show an overall upward trend in positive emotional tone and positive foresight.

H1a. These temporal trends will be moderated by outlet political orientation, such that left-, center-, and right-leaning outlets will exhibit distinct trajectories over time.

H2. The release of GPT-1, GPT-3, and ChatGPT will impact emotional tone and foresight, measured as the growth rates one month pre vs post release.

H2a. These pre- vs. post–GPT release effects will be moderated by political orientation, indicating differential event responses across ideological groups (left, center, and right-wing news outlets).

H3. The release of ChatGPT will have a greater impact on emotional tone and positive foresight than the releases of GPT-1 and GPT-3.

H3a. The difference in event impact between ChatGPT and the earlier GPT releases will be moderated by political orientation, similar to H2a.

Method

All data and scripts are freely available on https://osf.io/3haus/files/osfstorage.

Pre-registration

The present study was preregistered on the Open Science Framework (OSF) prior to conducting the confirmatory analyses and before inspecting model-fit results for the final dataset (https://osf.io/r9pmj). The preregistration followed the logic of a replication-and-extension of a pilot analysis (https://osf.io/3haus/files/osfstorage, Appendix A) on AI-related news coverage (2018–2024), done with the explicit aim of (a) formalizing the three main hypothesis families (H1/H1a, H2/H2a, H3/H3a), (b) specifying the measurement of the linguistic outcomes, and (c) fixing the statistical models, exclusion rules, and multiple-testing procedure in advance. The complete scope of the preregistered parameters can be found in Appendix B.

Data sampling

The present study used a query-defined corpus design rather than a probabilistic sampling strategy. Articles were retrieved from ProQuest TDM Studio using the exact phrase *"artificial intelligence"* and limited to the period from January 1, 2018, to December 31, 2024. Records were further restricted to English-language news and magazine articles with complete core metadata (publication date, full text, publisher/outlet).

Following the pilot study, opinion/editorial pieces, non-English items, duplicates, and records missing date, text, or outlet information were excluded. To enable confirmatory tests by outlet ideology, only outlets with an available AllSides political-orientation label were retained (AllSides, 2022).

The exact-phrase query was chosen to prioritize conceptual specificity and reproducibility. Using *"artificial intelligence"* as the search anchor increases the likelihood that retrieved articles explicitly frame AI as a central topic, rather than only mentioning product names, adjacent technical concepts, or downstream applications in passing. This is important for the present study because the confirmatory analyses focus on the emotional tone and future-oriented framing of AI as an object of media discourse, not simply on all reporting that may indirectly involve AI-related technologies. In that sense, the search strategy was designed to favor a more coherent analytical corpus, even at the cost of narrower coverage.

At the same time, this strategy narrows the inference population. The confirmatory analyses, therefore, generalize most directly to articles in the ProQuest news and magazine database that explicitly use the phrase "artificial intelligence," rather than to all media discourse about AI-related systems. This distinction matters because terminology in public and journalistic discourse has become more differentiated over time, especially during the generative-AI period, when coverage may focus primarily on terms such as *generative AI*, *large language models*, or *ChatGPT* rather than on *artificial intelligence* itself. If later coverage increasingly relied on such alternative labels, the exact-phrase query may under-capture parts of the post-2022 discourse. In turn, this could attenuate estimated time trends or weaken event-window effects by excluding some of the most product-specific or generative-AI-specific reporting. A formal coverage-sensitivity analysis comparing the exact-phrase query with alternative search terms (e.g., *"ChatGPT," "generative AI,"* or *"large language model"*) was not part of the preregistered confirmatory workflow and is therefore not reported here. Accordingly, the present results should be interpreted as characterizing a clearly defined and reproducible AI-labeled corpus, while leaving open the possibility that a broader multi-term query would yield somewhat different temporal patterns, especially in the later years of the observation window.

Building on the pilot, which assembled a multi-country corpus (U.S., U.K., and Canada) for descriptive comparison, the preregistration designated the U.S. subset as the confirmatory sample because the non-U.S. strata were too sparse and imbalanced by political orientation to support stable moderation tests. Accordingly, the present confirmatory dataset comprises U.S. outlets only (N = 14,752; left = 4,094, center = 3,518, right = 7,140). For transparency, the broader pilot collection totaled N = 16,619 articles across countries (U.S. = 14,752; U.K. = 799; Canada = 1,068), with the non-U.S. articles retained only for preregistered exploratory robustness analyses.

Variables and Measures

The first outcome is Emotional tone, operationalized with the LIWC-22 (Boyd et al., 2022) Tone dimension (continuous valence index; higher = more positive, lower = more negative). The second outcome set is based on a custom "foresight" dictionary obtained from the LIWC user repository (Hogenraad, 2019, 2020). Compiled in the standard LIWC dictionary format, it defines four categories targeting future-oriented evaluation and risk emphasis: Positive Foresight, Negative Foresight, Anticipation, and Fear. In keeping with LIWC conventions, entries may be listed as stems (e.g., *anticipat**, *threat**) to capture morphological variants. Scoring follows LIWC's method: for each article and category, we compute the proportion of tokens matched (percentage of words in that category relative to total words), then z-standardize these proportions for analysis.

As preregistered, we form an index to capture relative and future-oriented valence: positive vs. negative foresight = z(Positive Foresight) − z(Negative Foresight). Together with Tone, this index yields the two confirmatory dependent variables analyzed across all hypothesis families.

Time is modeled as date (days), defined as the number of days since the start of the observation window (2018-01-01 to 2024-12-31).

Political orientation was coded at the outlet level using AllSides ratings (AllSides, 2022) and harmonized into a three-level factor (Left, Center, Right), with *Lean Left* and *Left* combined into Left and *Lean Right* and *Right* combined into Right to maintain comparable group sizes and stable estimation.

To capture short-run responses to salient milestones, the *period* was a binary variable distinguishing the window from one month before (*pre*) to each release (GPT-1, GPT-3, ChatGPT), and from the period up to a one-month window after (post). For comparative event analyses, *version* contrasted GPT-1/GPT-3 with ChatGPT.

In all mixed-effects models, outlet (publisher) was included as a random intercept to account for repeated articles within outlets.

Tone and foresight measurements were first computed at the article level and then z-standardized prior to modeling. In accordance with the preregistration, outlier exclusion was applied after standardization, removing observations with $|z| > 3$ on the respective dependent variable before model estimation. Confirmatory analyses were conducted on the final filtered U.S. sample with complete data on text, date, publisher, and political orientation (Tone: $N = 14{,}677$; Foresight: $N = 14{,}588$).

Statistical Analysis

All confirmatory analyses were preregistered (https://osf.io/r9pmj). For the long-run hypotheses (H1/H1a), we fit linear mixed-effects regression models with a random intercept for outlet to account for repeated observations from the same publisher. The fixed-effects structure included time (days since 1 January 2018), political orientation (Left/Center/Right), and their interaction. The H1 model estimated the main effect of time, whereas H1a tested moderation by political orientation via the time × orientation interaction. Time was entered as a continuous predictor and interpreted as change per day. For model reporting, we convert effect sizes as variations per year.

For event-window hypotheses (H2/H2a), our windows spanned ±30 days around each preregistered release date (June 11, 2018, for GPT-1; May 28, 2020, for GPT-3; and November 30, 2022, for ChatGPT). Within each event window, time was centered on the release date (*days_centered* = 0 at release), and we fit models with *days_centered*, period (pre vs. post; pre as the reference), and their interaction. The interaction term (*days_centered* × period) captures whether the slope (growth rate) differs between the pre- and post-release periods. For H2a, political orientation was added as a moderator using the full three-way structure *days_centered* × period × orientation, with follow-up contrasts comparing the pre–post slope change across ideological groups.

For comparative release hypotheses (H3/H3a), we pooled event windows for GPT-1 and GPT-3 as "earlier releases" and contrasted them with the ChatGPT. Models included *days centered*, period (pre vs. post), version (GPT1_3 vs. ChatGPT), and their interactions, while controlling for political orientation as an additional fixed effect. The key H3 test was the difference in the pre–post slope change between ChatGPT and earlier releases (i.e., the period × days-centered effect differed by version). H3a tested whether this difference varied by political orientation via the corresponding interaction terms.

All dependent variables were analyzed as z-standardized outcomes. Statistical inference was based on two-sided tests with 95% confidence intervals. To control for multiple comparisons in preregistered post hoc contrasts (e.g., pairwise group comparisons within moderation tests and the key version contrast), p-values were adjusted using the Benjamini–Hochberg false discovery rate procedure (Benjamini & Hochberg, 1995). Unless otherwise noted, FDR adjustment was applied within each hypothesis family's set of planned contrasts (e.g., the set of pairwise slope-difference contrasts for H1a; the set of between-group ΔΔ slope contrasts for H2a; and the key ChatGPT-versus-earlier-release contrast for H3), while primary omnibus tests were reported unadjusted.

Results

Long-term trends (H1)

Contrary to the hypothesized positive trend, emotional tone declined slightly but reliably between 2018 and 2024. The effect of time (days since January 1, 2018) was significantly negative, $\beta = -8.51 \times 10^{-5}$, $SE = 1.14 \times 10^{-5}$, *95% CI* [-1.08×10^{-4}, -6.27×10^{-5}], $z = -7.45$, $p < .001$, corresponding to an average change of −0.031 SD per year, *95% CI* [−0.039, −0.023]. In contrast, the positive-versus-negative foresight index increased over time, consistent with H1. The time effect was significantly positive, $\beta = 7.19 \times 10^{-5}$, $SE = 1.45 \times 10^{-5}$, *95% CI* [4.35×10^{-5}, 1.00×10^{-4}], $z = 4.96$, $p < .001$, corresponding to an increase of 0.026 SD per year, *95% CI* [0.016, 0.037]. These trends are shown in **Figure 1**.

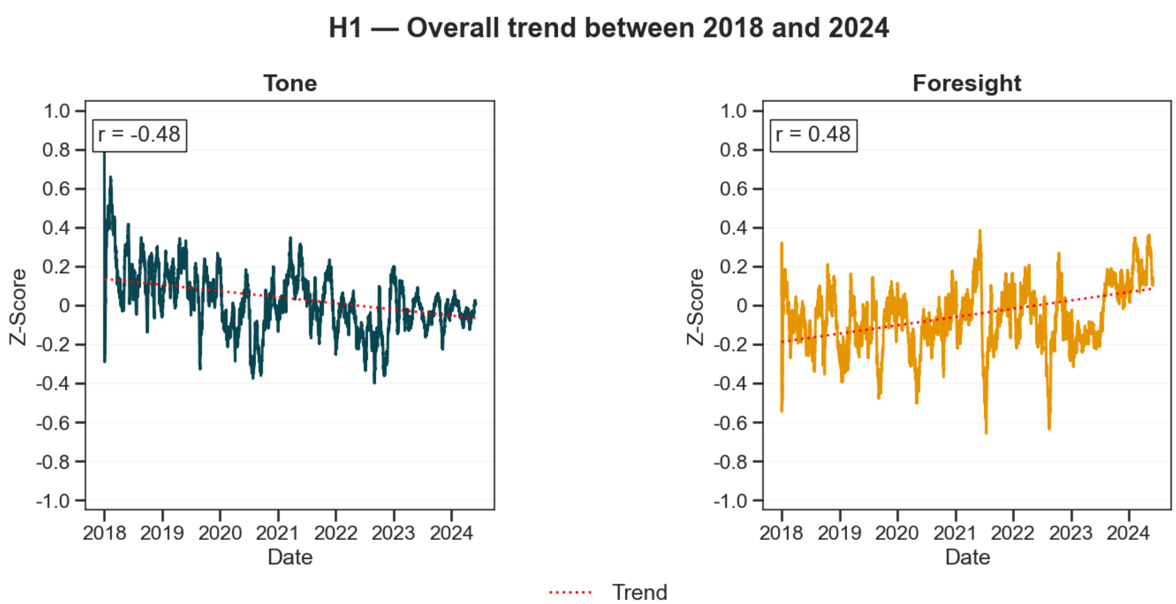

Figure 1. **The 30-day rolling averages of tone and foresight of US news articles between 2018 and 2024**. Tone significantly decreased, whereas foresight significantly increased, as indicated by the dotted red trendline.

For H1a, we tested whether these time trends differed by political orientation. For emotional tone, the time-by-orientation interaction was not supported: all three ideological groups exhibited similarly small negative time trends. Pairwise comparisons of the slopes revealed no significant differences (Left vs. Center: $\Delta\beta = -5.57 \times 10^{-6}$, $p = .874$, pFDR = .874; Right vs. Center: $\Delta\beta = 1.52 \times 10^{-5}$, $p = .633$, pFDR = .874; Left vs. Right: $\Delta\beta = -2.08 \times 10^{-5}$, $p = .429$, pFDR = .874), and the omnibus Wald test likewise indicated no reliable differences in time slopes, $\chi^2(2) = 0.69$, $p = .950$.

For the positive-versus-negative foresight index, H1a was supported. Time trends differed significantly across political orientation, $\chi^2(2) = 21.66$, $p < .001$. Relative to centrist outlets, both left-leaning and right-leaning outlets showed significantly less positive time slopes (Left vs. Center: $\Delta\beta = -1.90 \times 10^{-4}$ per day, *95% CI* [-2.77×10^{-4}, -1.03×10^{-4}], $p < .001$, pFDR < .001; Right vs. Center: $\Delta\beta = -1.73 \times 10^{-4}$ per day, *95% CI* [-2.53×10^{-4}, -9.39×10^{-5}], $p < .001$, pFDR < .001). The difference between left- and right-leaning outlets was small and not significant ($\Delta\beta = -1.71 \times 10^{-5}$, *95% CI* [$-8.25 \times 10^{-5}$, 4.82×10^{-5}], $p = .608$, pFDR = .608).

Overall, while the decline in emotional tone over time was comparable across left-, center-, and right-leaning outlets, the increase in positive relative to negative foresight was strongest in centrist outlets, whereas left- and right-leaning outlets showed weaker trends over time (Figure 2).

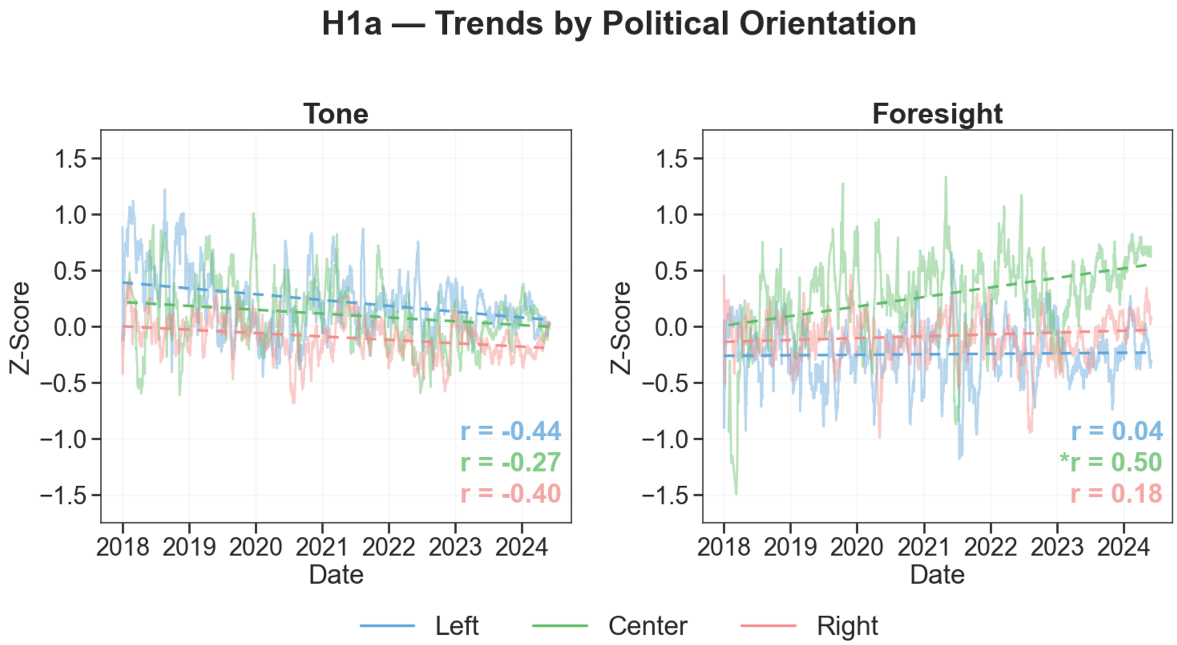

Figure 2. 30-day rolling averages of tone and foresight of US news articles are divided by the political orientation of the news outlet between 2018 and 2024. While tone shows no differences between orientations, the rise of positive foresight is driven by centrist media. ** significant difference in slope*

Effects of each GPT version release (H2)

For H2, we tested whether GPT releases affected emotional tone and positive foresight by comparing trends in the 30 days before (pre) vs. the 30 days after (post) each release (GPT-1, GPT-3, and ChatGPT).

For emotional tone, the pre–post change in time slope was not significant for GPT-1, $\Delta\beta = -0.020$ per day, *SE* = 0.016, *95% CI* [−0.052, 0.012], *p* = .219, pFDR = .307. In contrast, GPT-3 showed a significant negative shift ($\Delta\beta = -0.044$ per day, *SE* = 0.017, *95% CI* [−0.076, −0.011], *p* = .008, pFDR = .024): tone changed from a non-significant positive pre-release slope ($\beta = 0.015$, *95% CI* [−0.006, 0.037], *p* = .163) to a significant negative post-release slope ($\beta = -0.028$, *95% CI* [−0.052, −0.004], *p* = .020. For ChatGPT, the pre–post slope did not significantly change ($\Delta\beta = -0.015$ per day, *SE* = 0.014, *95% CI* [−0.043, 0.014], *p* = .307, pFDR = .307).

For the positive-versus-negative foresight index, none of the release events yielded a significant difference between pre- and post-release slopes (Δβ-GPT-1 = −0.023, *SE* = 0.018, *95% CI* [−0.060, 0.013], p = .204, pFDR = .611; Δβ-GPT-3 = −0.011, *SE* = 0.016, *95% CI* [−0.042, 0.021], p = .504, pFDR = .755; Δβ-ChatGPT = −0.004, *SE* = 0.018, *95% CI* [−0.039, 0.031], p = .840, pFDR = .840).

Overall, H2 received limited support: a significant decline in positive emotional tone was observed after the release of GPT-3, but no other significant effects on emotional tone were observed, and no effects of version release were found for foresight (Figure 3).

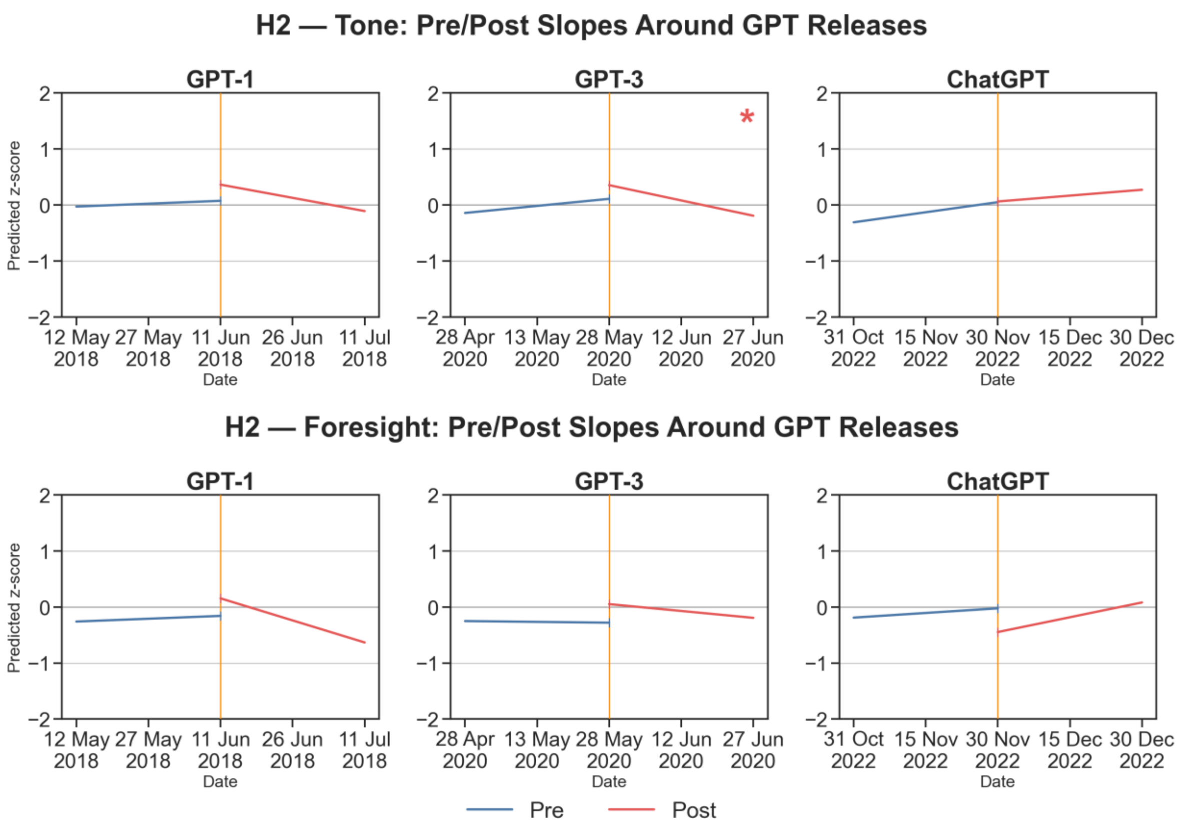

Figure 3. The fitted pre- vs. post-release (±30 days) trends of emotional tone and foresight in US news articles for GPT-1, GPT-3, and ChatGPT. While emotional tone remained largely stable across most events, the release of GPT-3 was associated with a

significant post-release decline. In contrast, foresight shows no clear shifts in trajectory across any of the model releases. * *significant pre vs post release differences (period × time).*

Interestingly, we found significant effects of political orientation (H2a) (Figure 4, Table 1). Emotional tone became more negative after GPT-1 release in centrist outlets ($\Delta\beta = -0.258$ per day, *SE* = 0.015, *95% CI* [−0.288, −0.229], $p < .001$), while right-leaning outlets showed a smaller but significant decrease ($\Delta\beta = -0.012$ per day, *SE* = 0.003, *95% CI* [−0.019, −0.005], $p < .001$); the change for left-leaning outlets was not statistically significant ($\Delta\beta = -0.033$ per day, *SE* = 0.019, *95% CI* [−0.070, 0.003], $p = .074$). Comparing these pre–post changes between groups, both left-leaning and right-leaning outlets differed from centrist outlets (Left vs. Center: $\Delta\Delta\beta = 0.225$, *SE* = 0.024, *95% CI* [0.178, 0.272], $p < .001$, pFDR < .001; Right vs. Center: $\Delta\Delta\beta = 0.246$, *SE* = 0.015, *95% CI* [0.216, 0.276], $p < .001$, pFDR < .001), whereas left and right did not differ ($\Delta\Delta\beta = -0.021$, $p = .262$, pFDR = .394).

A similar moderation pattern emerged with GPT-3. Left-leaning outlets showed a significant decline in slope post release ($\Delta\beta = -0.042$ per day, *SE* = 0.013, *95% CI* [−0.068, −0.016], $p = .002$), as did right-leaning outlets ($\Delta\beta = -0.046$ per day, *SE* = 0.007, *95% CI* [−0.060, −0.032], $p < .001$), whereas the centrist change was not significant ($\Delta\beta = 0.027$ per day, *SE* = 0.018, *95% CI* [−0.009, 0.063], $p = .135$). Accordingly, the pre–post change differed between left and center ($\Delta\Delta\beta = -0.069$, *SE* = 0.023, *95% CI* [−0.113, −0.025], $p = .002$, pFDR = .003) and between right and center ($\Delta\Delta\beta = -0.073$, *SE* = 0.020, *95% CI* [−0.112, −0.035], $p < .001$, pFDR < .001), but not between left and right ($\Delta\Delta\beta = 0.004$, $p = .769$, pFDR = .769).

Around the release of ChatGPT, tone changes again varied by orientation. Left-leaning outlets showed an increase in slope post release ($\Delta\beta = 0.018$ per day, *SE* = 0.007, *95% CI* [0.004, 0.033], $p = .015$), and centrist outlets showed a smaller increase ($\Delta\beta = 0.010$ per day, *SE* = 0.005, *95% CI* [0.000, 0.020], $p = .048$), while right-leaning outlets showed a decrease ($\Delta\beta = -0.034$ per day, *SE* = 0.006, *95% CI* [−0.045, −0.022], $p < .001$). The difference in pre–post change was significant for right versus center ($\Delta\Delta\beta = -0.044$, *SE* = 0.008, *95% CI* [−0.059, −0.028], $p < .001$, pFDR < .001) and for left versus right ($\Delta\Delta\beta = 0.052$, *SE* = 0.010, *95% CI* [0.033, 0.071], $p < .001$, pFDR < .001), but not for left versus center (ΔΔβ = 0.008, $p = .377$, pFDR = .377).

For the positive-versus-negative foresight index, evidence for H2a was limited (Table 1). Right-leaning outlets showed a modest slope decrease after GPT-1 release relative to centrist outlets (Right vs. Center: ΔΔβ = −0.024, *SE* = 0.011, *95% CI* [−0.046, −0.002], $p = .034$), but this contrast did not remain statistically significant after correction for multiple tests (pFDR = .101). All other between-group differences in pre–post changes for foresight were non-significant for GPT-1, GPT-3, and ChatGPT (all pFDR ≥ .380). Thus, unlike emotional tone, GPT release effects on foresight were not consistently moderated by political orientation in the preregistered event windows.

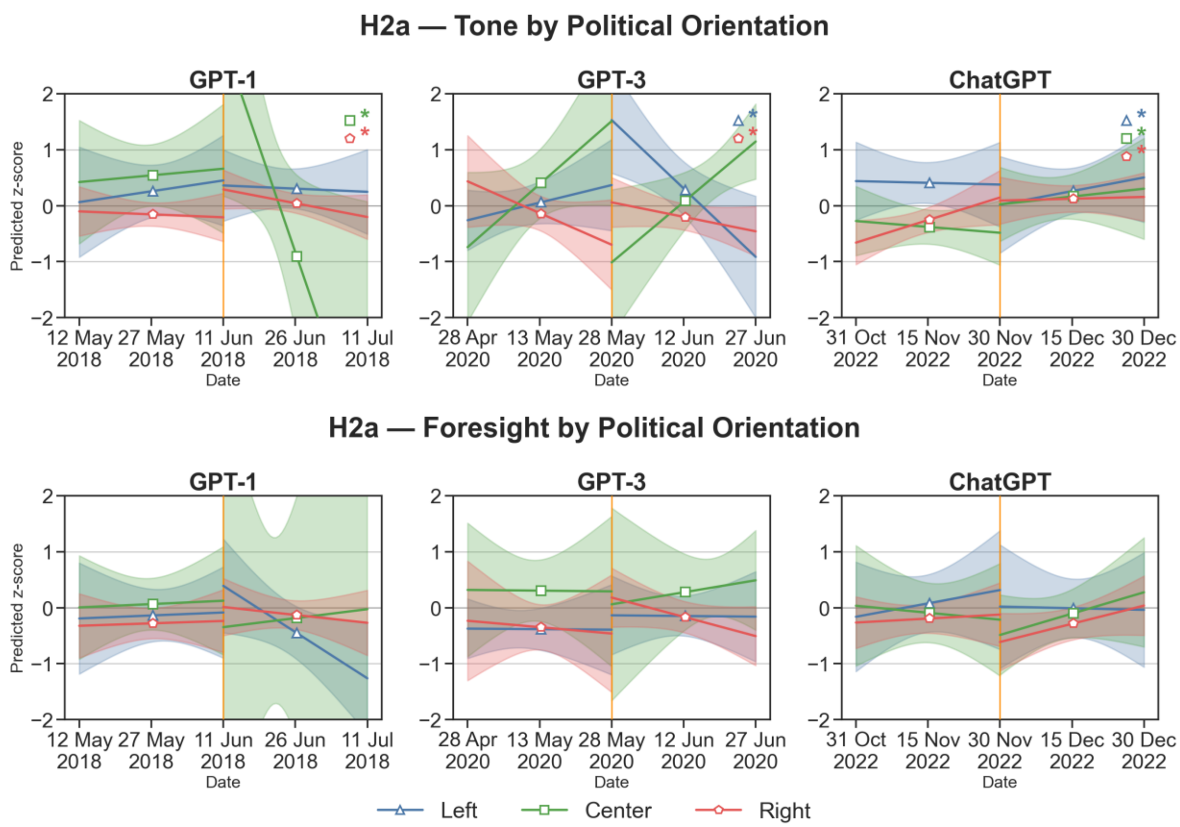

Figure 4. The fitted pre- vs. post-release (±30 days) trends of emotional tone and foresight in US news articles for GPT-1, GPT-3, and ChatGPT, moderated by political orientation (Left, Center, Right). Solid lines represent fitted trends, and shaded bands indicate 95% confidence intervals around the fitted values. * *significant pre vs post release differences across political orientations (period × time × political orientation)*

Release	Contrast	ΔΔβ	*SE*	*95% CI*	*p*	pFDR	Sig.
GPT-1	L vs. C	0.225	0.024	[0.178, 0.272]	< .001	< .001	***
GPT-1	R vs. C	0.246	0.015	[0.216, 0.276]	< .001	< .001	***
GPT-1	L vs. R	-0.021	NR	NR	.262	.394	
GPT-3	L vs. C	-0.069	0.023	[-0.113, -0.025]	.0023	.00347	**
GPT-3	R vs. C	-0.073	0.020	[-0.112, -0.035]	.00020	.00020	***
GPT-3	L vs. R	0.004	NR	NR	.769	.769	
ChatGPT	R vs. C	-0.044	0.008	[-0.059, -0.028]	< .001	< .001	***
ChatGPT	L vs. R	0.052	0.010	[0.033, 0.071]	< .001	< .001	***
ChatGPT	L vs. C	0.008	NR	NR	.377	.377	

Table 1. Slope differences (ΔΔβ) in *Tone* pre- vs. post-release (±30 days) between political orientations ([L]eft, [C]enter, [R]ight). Significance: * p < .05, ** p < .01, *** p < .001. NR = not reported. pFDR = adjusted p-values in accordance with the false discovery rate (Benjamin-Hochberg).

Specific effects of Chat-GPT release (H3)

To assess the specific effects of Chat-GPT release, we contrasted it with the pooled effects of earlier versions (GPT-1 and GPT-3). We did not find specific pre-post slope differences, neither for tone (Δδβ = 0.022 per day, *SE* = 0.015, *95% CI* [−0.007, 0.050], *p* = .136, pFDR = .272) nor for foresight (ΔΔβ = 0.013 per day, *SE* = 0.014, *95% CI* [−0.015, 0.041], *p* = .354, pFDR = .354). Overall, H3 was not supported for either outcome.

However, we found significant effects of political orientation. Compared to earlier versions, the release of ChatGPT was associated with a stronger increase of positive tone in centrist (ΔΔβ = 0.039 per day, 95% CI [0.013, 0.066], p = .003) and left-leaning outlets (ΔΔβ = 0.061 per day, 95% CI [0.045, 0.076], p < .001), while it decreased slightly (non-significantly) in right-leaning outlets (ΔΔβ = −0.002 per day, 95% CI [−0.005, 0.000], p = .098).

Comparing these ChatGPT-versus-earlier differences across orientations, the version contrast differed significantly between right and center (Right − Center: ΔΔΔβ = −0.042 per day, 95% CI [−0.068, −0.015], p = .002, pFDR = .003) and between left and right (Left − Right: ΔΔΔβ = 0.063 per day, 95% CI [0.047, 0.079], $p < .001$, pFDR < .001), whereas the difference between left and center was not significant (Left − Center: ΔΔΔβ = 0.021 per day, 95% CI [−0.009, 0.052], p = .175, pFDR = .175).

For the foresight index, H3a was not supported. The ChatGPT-versus-earlier differences in pre–post slope change were not significant within any orientation (Center: ΔΔβ = −0.009, *p* = .537; Left: ΔΔβ = 0.016, *p* = .441; Right: ΔΔβ = 0.010, *p* = .585), and the corresponding orientation contrasts were also non-significant after correction (Left − Center: pFDR = .625; Right − Center: pFDR = .625; Left − Right: pFDR = .835).

Overall, H3a received partial support. The emotional tone increased more positively after the release of ChatGPT among centrist and left-leaning outlets, but not among right-leaning outlets, for which no significant change was observed. No effects of version release were found for foresight.

Discussion

This study investigated how the general tone and future-oriented language used to discuss artificial intelligence in US news media changed over time (H1) and how they were affected by the release of specific versions of GPT (H2/H3) and by the political orientation of outlets (H2a/H3a).

We found that tone became more negative over time across the political spectrum, while foresight became more positive, particularly in centrist media. We also found that political orientation moderated the effect of GPT version releases on the tone of news articles. In particular, the tone of centrist and left-leaning articles became more positive after the release of ChatGPT relative to previous GPT versions, while it tended to become more negative after the releases of GPT1 and GPT3. Finally, the trends of foresight were not modulated by GPT version or political orientation. We now discuss these findings in more detail.

First, the direction of long-run foresight change aligned with predictions, whereas tone did not. The gradual decline in tone suggests that AI-related coverage did not become more positive throughout the study period. Instead, it became modestly more negative. A plausible interpretation is that AI has increasingly shifted from "novel technology" to infrastructure with governance-relevant consequences, prompting greater emphasis on responsibility, regulation, and social risk. This shows, that journalism about AI and automation often foregrounds questions of technological risk and responsibility alongside innovation (Brantner & Saurwein, 2021), and it is compatible with findings that AI coverage is not uniformly alarmist but can move toward more socially and ethically oriented discussion as public salience rises (Garvey & Maskal, 2020; Moriniello et al., 2024; Xian et al., 2024). At the same time, the positive trend in the foresight index suggests that future-oriented language is increasingly focused on opportunities rather than threats. Taken together, these patterns suggest that affective evaluation (tone) and future framing (foresight valence) capture distinct dimensions of discourse. Coverage can become more cautionary in overall effect, while simultaneously framing the future in opportunity-oriented terms (e.g., productivity, innovation, competitiveness).

In this sense, the results point to an increasingly ambivalent emotional climate: AI is treated as more consequential and potentially risky, yet also increasingly positioned as a source of future promises. This ambivalence is consistent with survey-based and review evidence showing that public perceptions of AI commonly mix excitement and utility with worry and perceived risk (Kelley et al., 2021; Krieger et al., 2024). It is also consistent with the broader context of accelerating AI capabilities and diffusion during this period, which can increase both perceived benefits and perceived stakes (Maslej et al., 2023).

The moderation results for H1a further sharpens the interpretation. Tone declined similarly in left-, center-, and right-leaning outlets, suggesting a shared, cross-ideological drift toward more negative effects in AI coverage. Rising policy relevance, institutional uptake, and widely discussed risks may shape journalistic language across the media system, regardless of outlet ideology.

In contrast, foresight trends differed by political orientation. The aggregate increase in positive foresight was driven primarily by centrist outlets, while left- and right-leaning outlets showed substantially weaker (closer to flat) change. This pattern aligns with framing-based accounts in which ideology shapes which consequences are foregrounded and how futures are narrated, even when overall affect shifts similarly (Entman, 1993; Gentzkow & Shapiro, 2010; Stroud, 2011). One interpretation is that centrist outlets increasingly routinized opportunity-oriented future talk as AI became normalized in economic and institutional narratives, whereas left- and right-wing outlets' future framing might also have been influenced by ideological considerations. In increasingly polarized environments, such stable differences in interpretive lens can persist or even strengthen over time (Moral & Best, 2023), potentially helping explain why the foresight trajectory diverged by ideology even when tone did not.

Second, the event window analyses were designed to examine how the media reacted in the short term to the introduction of GPT versions. Other than a slight decrease in tone after the introduction of GPT-3, we did not find consistent effects of version releases on tone or foresight. This pattern suggests that release-driven language shifts were either absent, too small to detect under the preregistered specification, or not well captured by a strict ±30-day growth-rate design.

Several mechanisms could produce historically meaningful release events without clear slope breaks in this framework. Release dates may not align with when journalistic framing changes. Coverage can anticipate releases, lag them, or diffuse gradually as applications spread. In addition, attention dynamics in news are typically short-lived, and any single technological event must compete with other agenda-setting stories within the same time window (Owen, 2019; Schema Design & Google News Initiative, 2018).

An additional, more direct explanation is that the aggregate effects of GPT releases may be obscured by systematic differences in the political orientations of news outlets. Consistent with this, the moderation results for H2a show that event-proximate changes in tone varied by political orientation, whereas analogous moderation effects for foresight were weaker. Specifically, following the release of ChatGPT, emotional tone increased more strongly (i.e., became less negative) for centrist and left-leaning outlets, whereas no significant change was observed for right-leaning outlets. Interestingly, this positive tone effect was observed specifically for ChatGPT relative to previous releases (GPT-1/GPT-3), which were rather associated with increasingly negative tone.

This suggests that short-run emotional reactions to major AI developments are also filtered through politically differentiated interpretive frameworks focusing, for example, on innovation/competitiveness vs. risk, disruption, and social harm (Entman, 1993; Iyengar & Hahn, 2009). The relatively positive “ChatGPT moment” is also compatible with work, suggesting that increases in attention and topic diversity in AI coverage do not necessarily entail a uniform, immediate turn toward negative sentiment (Moriniello et al., 2024; Xian et al., 2024).

Taken together, these findings suggest that media discourse on AI is shaped less by simple shifts in positivity or negativity than by a growing differentiation between present evaluation and future imagination. As AI became more visible and socially consequential, news coverage appears to have moved toward a more cautious and governance-sensitive tone, while still sustaining opportunity-oriented narratives about what AI might enable in the future. This pattern points to a broader ambivalence in the public communication of emerging technologies: the same development can be framed simultaneously as a source of risk, disruption, and ethical concern, and as a driver of innovation, productivity, and strategic promise. At the same time, the results indicate that this ambivalence is not structured identically across the media system. Whereas negative tonal drift was broadly shared across ideological contexts, future-oriented framing and short-run reactions to major AI milestones remained more sensitive to political interpretation. More broadly, the study highlights the value of distinguishing between evaluative and anticipatory dimensions of discourse when analyzing technological change. Doing so helps clarify how media can converge in their sense that technology matters while diverging in the futures they make imaginable around it.

Limitations and future directions

Several features of the present design constrain the generalizability of the findings and suggest avenues for future research. First, the corpus was defined using the term "artificial intelligence," which improves transparency but may under-represent articles using alternative labels such as *generative AI*, *large language models*, or product names like *ChatGPT*. If terminology shifted over time, estimated trends may partly reflect naming conventions rather than substantive changes in sentiment. Future work could address this by expanding keyword strategies and comparing alternative corpus definitions.

Second, analyses were restricted to U.S. outlets due to the availability of political-orientation metadata. As a result, the observed ideological patterns should be interpreted within the U.S. media system. Extending this approach cross-nationally will require locally valid measures of media slant to avoid misclassification and to capture how AI discourse varies across different political and media contexts.

Third, the preregistered ±30-day event-window design provides a clear test of short-run change but may miss temporally diffuse effects, such as anticipatory coverage or delayed framing shifts. More flexible temporal models, including distributed-lag approaches, could better capture how media responses unfold around major technological developments.

Fourth, long-run trends may partly reflect changes in corpus composition (e.g., shifts in outlet contributions or topic emphasis) rather than purely changes in framing. Integrating topic modeling or domain-specific coding (e.g., labour, education, geopolitics) would help disentangle whether observed tone shifts reflect changing issue composition or changes in evaluation within the same issue domains.

Finally, the use of LIWC-based measures entails the limitations of lexicon-based approaches, which are sensitive to stylistic and contextual variation. The findings should therefore be interpreted as changes in language use rather than direct measures of underlying attitudes. Combining dictionary methods with more context-sensitive approaches, such as supervised or large language model–based annotation, would help assess the robustness of the observed patterns.

References

AllSides. (2022a). *AllSides*. https://www.allsides.com/unbiased-balanced-news

AllSides. (2022b). *Media bias ratings*. Retrieved January 19, 2026, from https://www.allsides.com/media-bias/media-bias-ratings

Alsaleh, A., Khan, I., & Rahman, M. S. (2024). The impact of technological advancement on culture and societal development: Actor-Network Theory, media archaeology and cultural semiotics. *Scientific Reports, 14*(1), 83995. https://doi.org/10.1038/s41598-024-83995-z

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software, 67*(1), 1–48. https://doi.org/10.18637/jss.v067.i01

Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological), 57*(1), 289–300. https://www.jstor.org/stable/2346101

Benkler, Y., Faris, R., & Roberts, H. (2018). *Network propaganda: Manipulation, disinformation, and radicalization in American politics*. Oxford University Press.

Boyd, R. L., Ashokkumar, A., Seraj, S., & Pennebaker, J. W. (2022). The development and psychometric properties of LIWC-22. *Behavior Research Methods, 54*, 469–488. https://doi.org/10.3758/s13428-021-01640-9

Brantner, C., & Saurwein, F. (2021). Covering technology risks and responsibility: Automation, artificial intelligence, robotics, and algorithms in the media. *International Journal of Communication, 15*, 5074–5098.

Breazu, P. (2024). ChatGPT-4 as a journalist: Whose perspectives is it reproducing? *Discourse & Society, 35*(4), 512–530. https://doi.org/10.1177/09579265241251479

Brennen, J. S., Howard, P. N., & Nielsen, R. K. (2018). *An industry-led debate: How UK media cover artificial intelligence*. Reuters Institute for the Study of Journalism, University of Oxford.

Brennen, J. S., Howard, P. N., & Nielsen, R. K. (2022). What to expect when you're expecting robots: Futures, expectations, and pseudo-artificial general intelligence in UK news. *Journalism, 23*(1), 22–38. https://doi.org/10.1177/1464884920947535

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., … Amodei, D. (2020). *Language models are few-shot learners*. arXiv. Retrieved January 19, 2026, from https://arxiv.org/abs/2005.14165

Budak, C., Goel, S., & Rao, J. M. (2016). Fair and balanced? Quantifying media bias through crowdsourced content analysis. *Public Opinion Quarterly, 80*(S1), 250–271. https://doi.org/10.1093/poq/nfw007

Chang, M. (2025). Does the media's partisanship influence news coverage on artificial intelligence issues? Media coverage analysis on artificial intelligence issues. *Social Science Computer Review, 43*(3), 564–577. https://doi.org/10.1177/08944393241268526

D'Alessio, D., & Allen, M. (2000). Media bias in presidential elections: A meta-analysis. *Journal of Communication, 50*(4), 133–156. https://doi.org/10.1111/j.1460-2466.2000.tb02866.x

Deng, R. (2025). An analysis of AI framing in trending social media news. *Technology in Society*.

Entman, R. M. (1993). Framing: Toward clarification of a fractured paradigm. *Journal of Communication, 43*(4), 51–58. https://doi.org/10.1111/j.1460-2466.1993.tb01304.x

Garvey, C., & Maskal, C. (2020). Sentiment analysis of the news media on artificial intelligence does not support claims of negative bias against artificial intelligence. *OMICS: A Journal of Integrative Biology, 24*(5), 286–299. https://doi.org/10.1089/omi.2019.0078

Gentzkow, M., & Shapiro, J. M. (2010). What drives media slant? Evidence from U.S. daily newspapers. *Econometrica, 78*(1), 35–71. https://doi.org/10.3982/ECTA7195

Gregersen, E. (2025a). *History of technology timeline: From first tools to artificial intelligence*. *Encyclopædia Britannica*. https://www.britannica.com/story/history-of-technology-timeline

Gregersen, E. (2025b). *Printing press*. *Encyclopaedia Britannica*.

Hendrickx, J., & Van Coppenolle, I. (2026). A tool, a threat, or both? How quality news media frame generative AI. *Emerging Media*. Advance online publication. https://doi.org/10.1177/27523543261422926

Hogenraad, R. (2019). Fear in the West: A sentiment analysis using a computer-readable "Fear Index." Quality & Quantity, 53(3), 1239–1261. https://doi.org/10.1007/s11135-018-0813-7

Hogenraad, R. (2020). The way of visionaries: Foresight and imagination, computed. Quality & Quantity. https://doi.org/10.1007/s11135-020-01071-w

Hu, K. (2023, February 2). ChatGPT sets record for fastest-growing user base – analyst note. *Reuters*. https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/

Hutto, C. J., & Gilbert, E. (2014). VADER: A parsimonious rule-based model for sentiment analysis of social media text. *Proceedings of the International AAAI Conference on Web and Social Media, 8*(1), 216–225. https://doi.org/10.1609/icwsm.v8i1.14550

Ittefaq, M., Zain, A., Arif, R., Ala-Uddin, M., Ahmad, T., & Iqbal, A. (2024). Global news media coverage of artificial intelligence (AI). *Telematics and Informatics, 93*, 102127. https://doi.org/10.1016/j.tele.2024.102127

Ittefaq, M., Zain, A., Arif, R., Ala-Uddin, M., Ahmad, T., & Iqbal, A. (2025). Global news media coverage of artificial intelligence (AI): A comparative analysis of frames, sentiments, and trends across 12 countries. *Telematics and Informatics, 96*, 102223. https://doi.org/10.1016/j.tele.2024.102223

Iyengar, S., & Hahn, K. S. (2009). Red media, blue media: Evidence of ideological selectivity in media use. *Journal of Communication, 59*(1), 19–39. https://doi.org/10.1111/j.1460-2466.2008.01402.x

Jain, A., & Ranganathan, S. (2025). A longitudinal analysis of artificial intelligence coverage in technology-focused news media using latent Dirichlet allocation and sentiment analysis. *Journalism and Media, 6*(4), 176. https://doi.org/10.3390/journalmedia6040176

Jasanoff, S., & Kim, S.-H. (2009). Containing the atom: Sociotechnical imaginaries and nuclear power in the United States and South Korea. *Minerva, 47*(2), 119–146.

Kelley, P. G., Yang, Y., Heldreth, C., Moessner, C., Sedley, A., Kramm, A., Newman, D. T., & Woodruff, A. (2021). Exciting, useful, worrying, futuristic: Public perception of artificial intelligence in 8 countries. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society* (pp. 627–637). https://doi.org/10.1145/3461702.3462605

Krieger, J. B., Bouder, F., Wibral, M., & Almeida, R. J. (2024). A systematic literature review on risk perception of Artificial Narrow Intelligence. *Journal of Risk Research*, 1–19. https://doi.org/10.1080/13669877.2024.2350725

Kuder, G. F., & Richardson, M. W. (1937). The theory of the estimation of test reliability. *Psychometrika, 2*, 151–160. https://doi.org/10.1007/BF02288391

Lee, H., & Park, H. (2026). Framing generative AI: An analysis of impact, quoted sources, and attributions of responsibilities in U.S. elite media. *Journalism & Mass Communication Quarterly*. Advance online publication. https://doi.org/10.1177/10776990251410593

Lewis, S. C., Markowitz, D. M., & Bunquin, J. B. (2024). Journalists, emotions, and the introduction of generative AI chatbots: A large-scale analysis of tweets before and after the launch of ChatGPT. *arXiv* (arXiv:2409.08761). https://doi.org/10.48550/arXiv.2409.08761

Maslej, N., Fattorini, L., Brynjolfsson, E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Ngo, H., Niebles, J. C., Parli, V., Shoham, Y., Wald, R., Clark, J., & Perrault, R. (2023, April). *The AI Index 2023 annual report*. AI Index Steering Committee, Institute for Human-Centered AI, Stanford University.

Moriniello, F. (2024). Exploring the relationship between the coverage of AI in *WIRED* magazine and public opinion (2018–2023). *Applied Sciences, 14*(5), 1994. https://doi.org/10.3390/app14051994

Mullainathan, S., & Shleifer, A. (2005). The market for news. *American Economic Review, 95*(4), 1031–1053. https://doi.org/10.1257/0002828054825619

Munoriyarwa, A. (2025). Generative AI and the future of news: Examining AI's implications for editorial work. *Digital Journalism, 13*(2), 211–229. https://doi.org/10.1080/17512786.2025.2545448

Nelkin, D. (1995). *Selling science: How the press covers science and technology*. W. H. Freeman.

Nguyen, D. (2024). The news framing of artificial intelligence: A critical exploration of how media discourses make sense of automation. *AI & Society, 39*(3), 1023–1041. https://doi.org/10.1007/s00146-022-01511-1

OpenAI. (2019a, February 14). Better language models and their implications. Retrieved January 19, 2026, from https://openai.com/research/better-language-models

OpenAI. (2019b, November 5). GPT-2 1.5B release. Retrieved January 19, 2026, from https://openai.com/blog/gpt-2-1-5b-release/

OpenAI. (2022, November 30). Introducing ChatGPT. Retrieved January 19, 2026, from https://openai.com/blog/chatgpt/

OpenAI. (2023, March 14). GPT-4. Retrieved January 19, 2026, from https://openai.com/research/gpt-4

OpenAI. (2024, April 24). Introducing the ChatGPT and Whisper APIs. Retrieved January 19, 2026, from https://openai.com/index/introducing-chatgpt-and-whisper-apis/

Ouchchy, L., Coin, A., & Dubljević, V. (2020). AI in the headlines: The portrayal of the ethical issues of artificial intelligence in the media. *AI & Society, 35*, 927–936. https://doi.org/10.1007/s00146-020-00965-5

Owen, L. H. (2019, January 25). A typical big news story in 2018 lasted about 7 days (until we moved on to the next crisis). *Nieman Lab*. https://www.niemanlab.org/2019/01/a-typical-big-news-story-in-2018-lasted-about-7-days-until-we-moved-on-to-the-next-crisis/

Paluch, D., Günes, M. T., & Aurylaite, G. (n.d.). *Analysis of sentiment and fear towards artificial intelligence in news media articles* [Unpublished manuscript].

ProQuest. (2024, July 8). *ProQuest | Better research, better learning, better insights*. https://www.proquest.com/

Roser, M. (2022, December 6). *The brief history of artificial intelligence*. *Our World in Data*. https://ourworldindata.org/brief-history-of-ai

Roser, M., & Ritchie, H. (2023). *Technological change*. *Our World in Data*. https://ourworldindata.org/technological-change

Russell, A. (2024, September 30). AI and our next evolution: The tectonic shift AI is sparking in human society. *Letters & Science Magazine (UC Davis)*. https://lettersandsciencemag.ucdavis.edu/feature/ai-and-our-next-evolution

Sanguinetti, P. (2025). The coverage of AI in the news media: A scoping review. *RAE-IC: Revista de la Asociación Española de Investigación de la Comunicación, 12*(1), 1–27.

Schema Design, & Google News Initiative. (2018). *The lifespan of news stories*. https://newslifespan.com/

Similarweb. (2023, November 15). *ChatGPT's first year in numbers*. https://www.similarweb.com

Stroud, N. J. (2011). *Niche news: The politics of news choice*. Oxford University Press.

Tai, M. C. T., He, L., & Li, S. (2020). The impact of artificial intelligence on human society and the environment. *Frontiers in Communication, 5*, 39. https://doi.org/10.3389/fcomm.2020.00039

Tegegn, A. E. (2024). The role of science and technology in reconstructing society and social relations. *Cogent Social Sciences, 10*(1), 2356916. https://doi.org/10.1080/23311886.2024.2356916

Teixeira, R., & Jones, J. (2021). A systems history of science, technology, leisure, and fear. *Technology in Society, 67*, 101786. https://doi.org/10.1016/j.techsoc.2022.101786

UBS Evidence Lab. (2023, February 1). *ChatGPT—Record adoption of a consumer application* [Analyst note]. UBS Group AG.

Vicsek, L., Feher, K., & Nguyen, D. (2026). Disappointment or hype? Framing artificial intelligence in the U.S. elite news media. *Global Media and Communication*. Advance online publication. https://doi.org/10.1177/17427665251411490

West, D. M., & Allen, J. R. (2018, April 24). *How artificial intelligence is transforming the world*. Brookings Institution. https://www.brookings.edu/articles/how-artificial-intelligence-is-transforming-the-world

Xian, L., Li, L., Xu, Y., Zhang, B. Z., & Hemphill, L. (2024). *Landscape of generative AI in global news: Topics, sentiments, and spatiotemporal analysis* (arXiv:2401.08899). *arXiv*. https://doi.org/10.48550/arXiv.2401.08899

Yi, A., Goenka, S., & Pandelaere, M. (2024). Partisan media sentiment toward artificial intelligence. *Social Psychological and Personality Science, 15*(6), 682–690. https://doi.org/10.1177/19485506231196817

Zhai, Y., Yan, J., Zhang, H., & Lu, W. (2020). Tracing the evolution of AI: Conceptualization of artificial intelligence in mass media discourse. *Information Discovery and Delivery, 48*(3), 137–149. https://doi.org/10.1108/IDD-01-2020-0007

Appendices

Appendix A – Tables and figures of the pilot study: Paluch, D., Günes, M. T., & Aurylaite, G. (2024). *Analysis of sentiment and fear toward artificial intelligence in news media articles* [Unpublished manuscript]. Department of Psychology, University of Vienna.

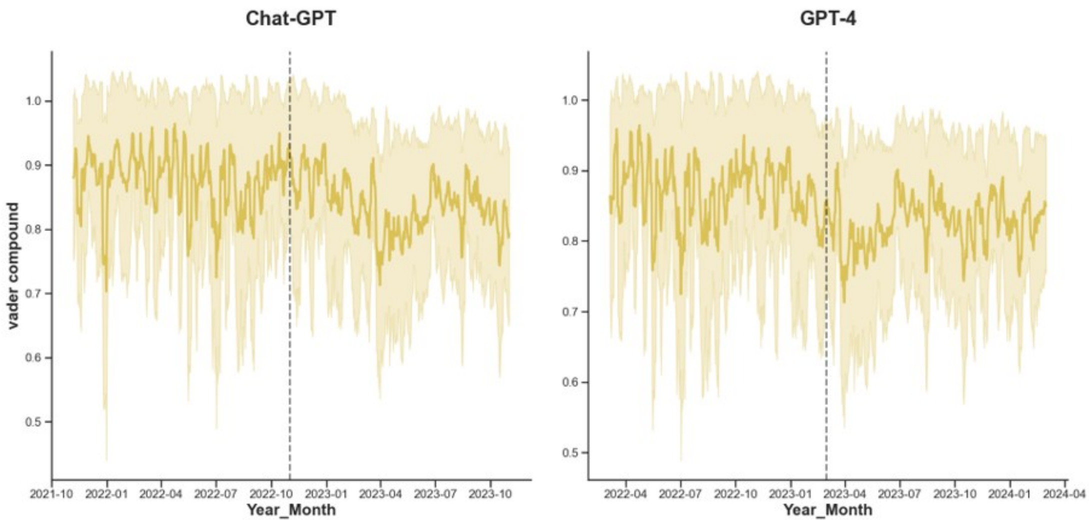

Figure 1: "comparison of sentiment (vader compound) 1 year before and after release of Chat-GPT and GPT-4

Dependent Variable comparisons, 1 year before vs after release of an GPT-Model			
Post - pre comparison	*Vader compound*	*Vader positive - negative*	*Fear vs trust frequency ratio*
GPT-1	p = 0.02, d = -0.05	p = 0.08 , d = -0.02	p = 0.24 , d = 0.01
GPT-3	p = 0.26, d = 0.01	p = 0.09 , d = -0.02	p = 0.49 , d = 0.001
Chat-GPT	p <.000*, d = -0.10	p <.000* , d = -0.16	p <.000*, d = 0.14
GPT-4	p <.000*, d = -0.09	p <.000* , d = -0.11	p <.000* , d = 0.10

Table 1: results of two sample t-tests, difference in sentiment and fear

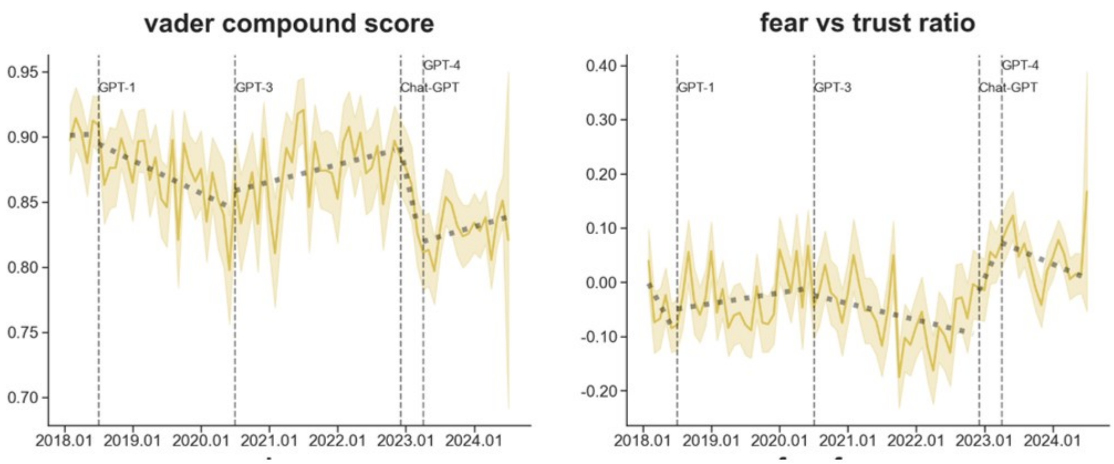

Figure 2: change in sentiment and fear over time and periods

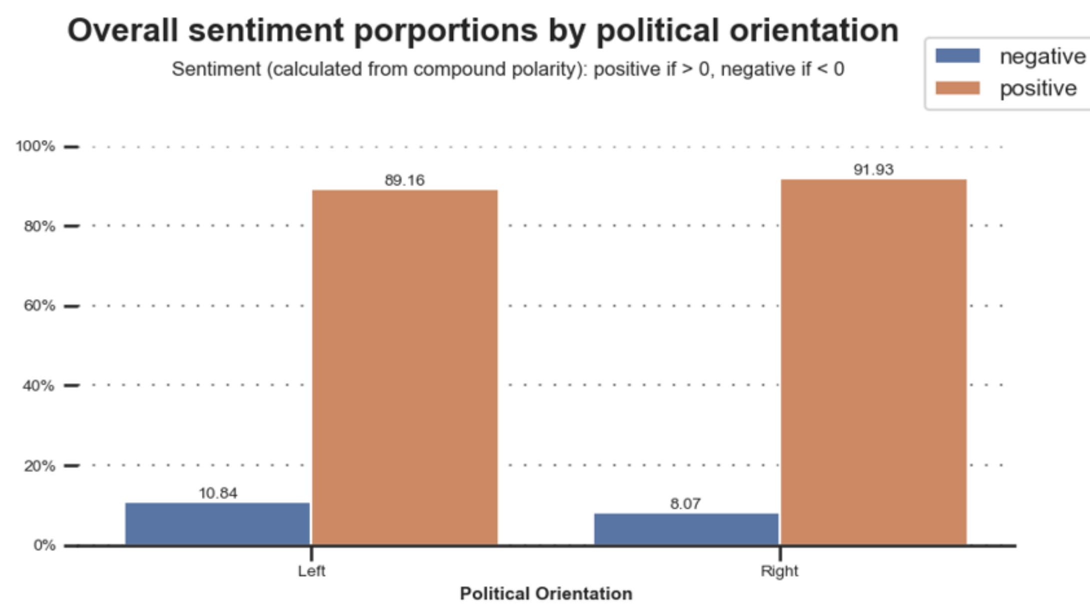

Figure 3: sentiment comparison left vs right

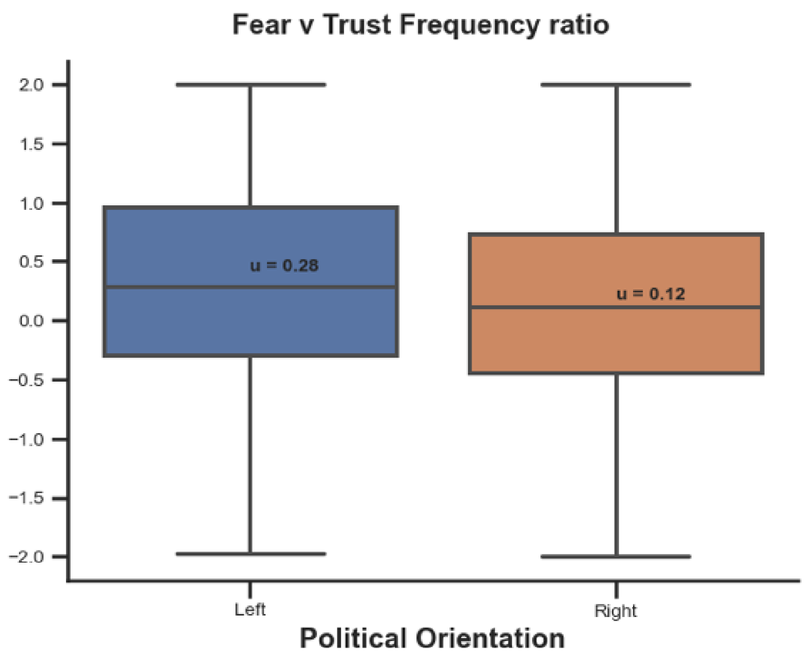

Figure 4: fear vs trust frequency comparison left vs right

Appendix B – OSF Preregistration Summary

Scope of preregistration. The OSF protocol documented:

1. Research questions and hypotheses exactly as stated in the thesis:
    - H1: *The overall emotional tone and (positive) foresight regarding GPT and artificial intelligence (AI) have exhibited a statistically significant upward trend between 2018 and 2024.*
    - H1a: *The rate and magnitude of change in emotional tone and (positive) foresight over time are moderated by political orientation, such that different political groups may exhibit distinct trends.*
    - H2: *The growth rates of emotional tone and (positive) foresight within one month before the release of GPT-1, GPT-3, and Chat-GPT will differ from the growth rates within the one month after, indicating a measurable impact of these release events.*
    - H2a: *The extent and direction of changes in emotional tone and (positive) foresight surrounding GPT release events are moderated by political orientation, suggesting differential responses across political groups.*
    - H3: *The release of Chat-GPT will have a more significant effect on emotional tone and foresight than GPT-1 and GPT-3 when controlling for political orientation.*
    - H3a: *The magnitude of this difference in effect between GPT releases varies depending on political orientation.*
2. Sampling plan. The preregistration specified that confirmatory analyses would be conducted on AI-related news articles published between 2018 and 2024, retrieved via ProQuest TDM Studio using the search term *"artificial intelligence"*, and restricted to outlets with a political-orientation label available from AllSides (left, center, right). The protocol stated that the U.S. subsample serves as the main confirmatory dataset and that articles from the U.K. and Canada would be included

only as exploratory extensions due to orientation imbalance. In the data currently available to this thesis, only the U.S. subset was present; the exploratory extension is therefore deferred to data availability.

3. Variables and transformations. The OSF entry listed the linguistic outcomes to be derived from LIWC-based features: Tone, Anticipation, Positive Foresight, Negative Foresight, Fear, and the two indices
   - *positive vs. negative foresight* = z(positive foresight) − z(negative foresight)
   - *anticipation vs. fear* = z(anticipation) − z(fear)
     All continuous variables were to be z-standardized, and outliers with $|z| > 3$ were to be removed.
4. Predictors. The preregistration defined:
   - date (days) as the longitudinal predictor,
   - period as a binary factor capturing one month before vs. after each GPT release,
   - model_version contrasting *GPT-1/3* vs. *Chat-GPT*, and
   - political_orientation with three levels (left, center, right) based on AllSides.
5. Statistical models. The OSF preregistration specified linear mixed-effects models with newspaper/outlet as a random intercept to account for multiple articles per outlet, and with fixed-effect structures matching each hypothesis family:
   - H1/H1a: DV = date * political_orientation + (1 | newspaper)
   - H2/H2a: DV = date * period * political_orientation + (1 | newspaper) for each release window
   - H3/H3a: DV = date * period * model_version * political_orientation + (1 | newspaper)
     (Bates, Mächler, Bolker, & Walker, 2015).

6. Multiple testing. Because each hypothesis was tested on three predefined outcomes (Tone; positive vs. negative foresight; anticipation vs. fear), the preregistration stated that Benjamini–Hochberg FDR would be applied within hypothesis families (i.e., across the three DVs for H1; across the three DVs for H2 per release; etc.), with $\alpha(q) = .05$ (Benjamini & Hochberg, 1995).

7. Planned exclusions. Articles without publication date, text, or publisher information were to be excluded, duplicates removed, and extreme outliers ($|z| > 3$) dropped.

8. Exploratory analyses. The OSF noted that — because Canadian and U.K. articles are strongly imbalanced by political orientation — a second, exploratory set of models would repeat H1–H3 on the full dataset (U.S. + U.K. + Canada), to assess robustness of direction and size of effects.

Deviations / clarifications. Two clarifications are necessary for the thesis:

1. Available data. The dataset used here currently contains U.S. articles only; the exploratory multinational extension announced in the preregistration could therefore not yet be executed.

2. Political-orientation recoding. Because initial inspection revealed small or uneven counts for "lean" categories, outlets labeled *Left* and *Lean Left* were aggregated into Left, and *Right* and *Lean Right* into Right, with Center retained. This aggregation step is consistent with the preregistered goal of testing moderation by political orientation and is reported here for transparency.

Appendix B – OSF Preregistration Summary

Appendix C – Abstracts

Abstract

Keywords: artificial intelligence, media sentiment, foresight, NLP, GPT, political orientation

This study examined how emotional tone and future-oriented language in U.S. news coverage of artificial intelligence (AI) changed between 2018 and 2024, and whether these trajectories differed by outlet political orientation. Using a longitudinal observational design, 14,752 U.S. news and magazine articles from 14 outlets were retrieved via ProQuest and grouped into left-, center-, and right-leaning categories. Linguistic outcomes were LIWC-22 Tone and two future-oriented indices, with confirmatory analyses focusing on positive versus negative foresight. Polarization signals were strongest in how future-oriented language evolved over time, with centrist outlets driving the increase in positive relative foresight, whereas short-window event contrasts emerged primarily in affective tone rather than foresight. Linear mixed-effects models estimated long-run time trends, event-window changes around GPT-1, GPT-3, and ChatGPT, and moderation by political orientation. Emotional tone showed a small but reliable decline over time, whereas positive relative foresight increased. The tone decline did not differ by political orientation, but the foresight increase was significantly stronger in centrist outlets. Event-window analyses yielded limited support for release-specific effects: only GPT-3 showed a significant post-release decline in tone, whereas foresight showed no reliable pre-post changes across releases. However, short-window tone effects varied by political orientation, especially with ChatGPT, with more positive post-release changes in centrist and left-leaning outlets than in right-leaning outlets. Overall, the findings suggest that AI news discourse became modestly more negative in affective tone while simultaneously more opportunity-oriented in future framing. Interestingly, centrist outlets became particularly optimistic over the long term and positive after ChatGPT was released.

Abstrakt

Keywords: Künstliche Intelligenz, Medienstimmung, Zukunftsblick, NLP, GPT, politische Orientierung

Die vorliegende Studie untersucht, wie sich emotionale Tonalität und zukunftsbezogene Sprache in der US-amerikanischen Berichterstattung über Künstliche Intelligenz (KI) zwischen 2018 und 2024 verändert haben und ob sich diese Verläufe nach politischer Ausrichtung der Medien unterscheiden. In einem longitudinalen Beobachtungsdesign wurden 14.752 US-amerikanische Nachrichten- und Magazinartikel aus 14 Medien über ProQuest erhoben und in linke, zentristische und rechte Outlets gruppiert. Als sprachliche Variablen dienten LIWC-22 Tone sowie zwei Indizes zukunftsbezogener Sprache, wobei sich die konfirmatorischen Analysen auf positive versus negative *Foresight* konzentrierten. Polarisierungssignale zeigten sich am deutlichsten in der langfristigen Entwicklung zukunftsbezogener Sprache, wobei insbesondere zentristische Medien den Anstieg positiver relativer *Foresight* trugen, während kurzfristige Ereigniskontraste eher in der affektiven Tonalität als in der *Foresight* sichtbar wurden. Lineare gemischte Modelle schätzten langfristige Zeittrends, Ereignisfenster um die Releases von GPT-1, GPT-3 und ChatGPT sowie Moderation durch politische Orientierung. Die emotionale Tonalität nahm über die Zeit leicht, aber verlässlich ab, während positive relative *Foresight* zunahm. Der Rückgang der Tonalität unterschied sich nicht nach politischer Orientierung, der Anstieg der *Foresight* war jedoch in zentristischen Outlets signifikant stärker als in linken oder rechten. Ereignisfensteranalysen ergaben nur begrenzte Hinweise auf release-spezifische Effekte: Nur GPT-3 war mit einem signifikanten post-release Rückgang der Tonalität verbunden, während sich für Foresight keine verlässlichen Prä-Post-Unterschiede zeigten. Insgesamt deuten die Befunde darauf hin, dass die KI-Berichterstattung affektiv etwas negativer, zugleich aber in ihrer Zukunftsrahmung chancenorientierter wurde.