AI & SOCIETY (2026) 41:1521–1535
https://doi.org/10.1007/s00146-025-02628-9

REVIEW

Writing with emotion? Assessing emotional valence and appeals in AI-generated vs. human-written articles

Mingxiao Sui[1]

Received: 25 July 2025 / Accepted: 11 September 2025 / Published online: 30 September 2025
© The Author(s) 2025

Abstract
Does generative AI write with emotionality in the same manners as humans? We investigated this question by focusing on three distinct aspects that characterize emotionally evocative content, including emotional valence, discrete appeals, and subjectivity. A quantitative content analysis of two publicly available data sets—TOEFL essays vs. news articles—was conducted, which paired each AI-generated text with a human-written piece by matching the essay questions or news headlines. Results revealed a stronger level of subjectivity in AI-generated than human-written articles. We also found the differences between AI-generated and human-written texts may differ by the type of writing: in terms of emotional valence, while little variance existed in TOEFL essays, AI-generated news stories exhibited a stronger positivity than human-written news. Similarly, regarding discrete emotional appeals, notable differences were identified between AI and human among news stories but not among TOEFL essays. In particular, AI-generated soft news is more likely convey a sadness emotion than human counterparts. These findings suggest that both the level and discrete states of emotionality are significant factors for studying the impact of AI-generated content on readers.

Keywords Emotional appeals · Emotional valence · Subjectivity · Generative AI · Comparative perspective · Content analysis

1 Introduction

The rapid advancement of artificial intelligence (AI) has led to the increasing prevalence of AI-generated content across various fields, from text creation, essay writing, to professional journalism. Mainstream media organizations such as *The Washington Post* and the *Associated Press* have been using automated systems to assist journalists in producing news stories. One recent instance was tech news outlet CNET's publication of over 70 AI-generated articles in early 2023. Readers were unaware that these articles were generated by AI agents until they clicked on the author name "CNET Money Staff," which then revealed that these articles were created by automation technology (Landymore 2023). Due to public backlash, CNET paused its AI experiment and replaced AI-generated articles with human-written stories.

Yet, the public discussion sparked by this incident continues. One question of interest is why average readers are not able to discern AI-generated articles. A recent study uncovered that about 48% readers correctly identified AI-generated content, and this number went down to 37% when the content was created by more advanced models such as GPT-4 (Brandl and Ellis 2023). Even when the author was explicitly specified to be AI, not everyone was convinced, with about one-third believing that a human was the real author (Lermann Henestrosa and Kimmerle 2024). The relatively low AI literacy among most readers is one possible reason; however, a bigger factor may lie in the large-scale language models' (LLM) strong capacity to resemble human language (Cai et al. 2023). In some cases, human readers even rated the quality of AI-generated essays to be higher than that of human-written essays (Herbold et al. 2023). Consistent patterns were also uncovered in the field of news evaluation, where ChatGPT-generated news stories were rated as more credible than *New York Times* articles (Kreps et al. 2022).

It should be noted that when the writers of text are explicitly specified, readers often perceive machine-generated

✉ Mingxiao Sui
msui@uab.edu

[1] Department of Communication Studies, University of Alabama at Birmingham, Birmingham, United States

articles to be more objective and less emotionally engaging than human-written articles (Wölker and Powell 2021). This is attributed to the so-called *machine heuristic*, which posits that machines are usually considered to be free of intention and bias, thereby being perceived to be neutral (Sundar 2008). In particular, people who hold a strong belief of the machine heuristic are also the ones who experience less emotional involvement when reading algorithm-generated news (Liu and Wei 2018). This emphasizes the importance of *source assignment*—such as news bylines or AI disclaimers that explicitly describe the roles of AI agents in producing news stories—in people's assessment of AI-generated content (Jia et al. 2024). However, given the rise of incidental consumption of news in the digital spaces (Boczkowski et al. 2018), where readers are less likely to focus on bylines (Santana and Hopp 2022), it is important to consider situations when readers fail to pick up the cues of AI authorship (see Lewis et al. 2019 for similar arguments).

Indeed, if human readers comply with the belief that machine-generated content is more objective and less emotionally engaging than human-written articles, the opposite speculation likely also holds: that is, upon reading articles that are more subjective and more emotionally eliciting, people may tend to perceive them to be human written. This contention is rooted in the individuals' inclination to rely on writing style, article structure, and the level of opinion and emotionality being conveyed when identifying AI-generated articles (Marinescu et al. 2022), especially when the AI disclaimers are not presented or picked up. Indeed, text-level features are informative cues that can affect reader perceptions; for example, an article's fluent wording, use of familiar words, and short sentences often lead to a *fluency effect*, which thereby influences people's ability to evaluate the truth of a statement, express a positive feeling of a given article, as well as their judgment of the author (see Alter and Oppenheimer 2009). The dual-process accounts also contend that message characteristics (i.e., whether an argument is objective, emotionally evocative) can greatly influence people's judgments, such as whether they believe a message comes from credible sources (Petty and Cacioppo 1986). While recent work advances our understanding of content-level differences between AI and human writers, one important question remains less examined: does AI-generated content express emotions in a manner that is comparable to human-written articles?

This paper aims to address this gap with a content analysis of two publicly available data sets that contained both AI- and human-written articles: one included TOEFL essays and the other was news. We employed computer automation methods to capture discrete emotional appeals (i.e., joy, sadness, fear, disgust, and anger), emotional valence (positive vs. negative tone), and subjectivity (an expression of opinions rather than facts) in these articles. Our findings from both analyses revealed that human-written articles were less subjective than ChatGPT-generated articles; yet, only among news articles—but not essays—we found that human-written news conveyed a lower level of positivity than ChatGPT-generated news. In terms of discrete emotional appeals, our analysis found a bigger portion of joy in human-written news; by comparison, a bigger portion of AI-generated news conveyed negative affective states, such as sadness, anger, fear and disgust. The differences in emotional appeals were not statistically significant among essays. Even though our findings suggest the level of emotionality in an article can potentially affect how readers identify, process, and engaged with AI-generated content, this implication cannot be overstated and needs to be empirically tested (also see Yang 2024).

2 Generative AI as communicators in the digital age

Since the Turing test in the 1950s, many studies have suggested that while human users may recognize computers' technological nature, they tend to ascribe social and psychological traits to computers, treating them as human partners rather than non-human tools (Nass and Moon 2000). This propensity also extends to AI, which refers broadly to "computational systems that involve algorithms, machine learning methods, natural language processing, and other techniques that operate on behalf of an individual to improve a communication outcome" (Hancock et al. 2020, p. 90). These computational agents include chatbots, smart replies, virtual assistants, automated news generators, and other AI-driven systems or features.

Recent scholarly work has examined AI technologies as communicators, treating them as distinct social actors that not only automate text but also facilitate social connection with human users (Guzman and Lewis 2020). For example, although AI tools are inherently non-gendered, users tend to personify them and overwhelmingly perceive conversational AI systems Alexa and Siri to be female-gendered (Abercrombie et al. 2021). This is partly attributed to AI techniques' use of feminine linguistic cues (i.e., politeness, first-person pronouns, apologies, and attenuated assertions) that can activate gender attribution, thereby leading people to associate them with stereotypically feminine traits such as compassionate (Duan et al. 2025). On the other hand, research also uncovers that large language models and other generative AI systems (i.e., ChatGPT) may reinforce masculinity in generated outputs; for example, ChatGPT-generated essays were found to resemble most closely those written by male authors, reflecting the male-dominated training data (Alvero et al. 2024). Different from human-written text, AI-generated content also tends to excessively use unique

content words but uses fewer function words, which thereby lacks a naturally formed speech rhythm (Opara 2024).

Yet, a critical aspect that enables human-to-human communication to provide social support is that people are able to express, understand, and respond to emotions. With personalized AI such as *Replika* or *character.ai*, users may also experience a deep emotional attachment, suggesting that AI tools may foster a sense of support and belongings through their emotionally charged interactions with human users (Pani et al. 2024). Similarly, AI tools are found to help maintain students' emotional well-being and improve their emotional resilience during the learning process; however, the current AI technology has limited capacity to understand human users' emotional states, which undermines their emotional support function (Ortega-Ochoa et al. 2024). Regarding the use of ChatGPT in teacher education programs, for example, students found AI tools were not capable of empathizing with human experiences or noticing changes in human moods, which, therefore, hardly evoked the emotional connection people would experience in human interactions (Huang and Chang 2024).

While much of the above research emphasizes the importance of AI systems in eliciting and responding to users' emotional states—mostly during computer–human interactions—it is equally important to consider how AI-generated text itself conveys emotion. Because the way AI-generated text simulates emotion shapes reader perceptions and interactions, investigating this dimension will extend AI system's emotional responsiveness to the broader communicative dynamics such as people reading AI-generated articles. This study presents a comprehensive investigation of the sentiment in AI-generated in relation to human-written content, including emotional valence, discrete emotional appeals, and subjectivity. Below we outline in depth how AI-generated texts may convey each of these aspects associated with emotionality.

3 Emotional valence in AI-generated content

Emotion is a crucial aspect of communication that shapes how a message can be interpreted and responded. In general, the more emotionally evocative the news, the more likely it is recommended and shared (Berger and Milkman 2012). Also relatedly is the effect of *emotional contagion*, which refers to how the emotional states within a message can transfer to readers, leading them to experience the same emotions (Hatfield et al. 1993). In the case of tweets, for example, people tend to comment with the same emotional state they feel in the political candidates' tweets; specifically, upon reading candidates' tweets with fear appeals, the public tended to express significantly more fear emotions in their comments (Paul and Sui 2022). These suggest the positive downstream consequences for perceiving emotionality in the content we read, which pertains to both human-written and AI-generated articles.

In a recent study focusing on AI-generated artwork, Gao and Yin (2024) found AI tools' unique strength in artistic techniques that can create a distinct emotional dimension. However, AI is only able to replicate established modes of emotional expressions included in its predefined training data; as a result, AI tools can seldom generate nuanced and profound emotional expressions that arise from human artists' personal experience (Gao and Yin 2024). This partly explains why relative to AI-generated artwork, people prefer human-created art and self-claim to have experienced more positive emotions when viewing human-created artwork (Samo and Highhouse 2023). While this line of research is still developing, these existing findings suggest that AI can mimic emotional expressions in their outputs, though the authenticity of these expressions may differ from those of humans.

When it comes to AI-generated text, Hohenstein and Jung's (2018) study of Google Allo revealed that its smart replies—specifically, automated responses suggested by AI systems—were largely positive, with 43.8% conveying a positive sentiment (e.g., "nice," "sure") and 3.95% conveying a negative sentiment (e.g., "nothing" "not really"). In a similar approach, Mieczkowski et al. (2021) analyzed smart replies Google Hangouts Chat suggested for its users, which found a positive sentiment in 93% of the smart reply suggestions (e.g., "sounds good"). In another comparison of the answers ChatGPT and human provided to the same set of questions, ChatGPT's answers were found generally more positive (i.e., conveying a joyful and optimistic tone), while human responses contained more anger and disgust (Fatahi et al. 2024). Similar patterns were also observed in news stories: even on highly controversial topics (i.e., abortion and immigration) that likely associate with negative connotations, a bigger portion of machine-generated news (10%) conveyed a positive tone than human-written news stories published by CNN (2.9%) or Fox News (1.4%); to the contrary, a bigger portion of human-written news contained a negative tone than machine-generated news (Nah et al. 2024). Further evidence was yielded by research on GPT3-generated misinformation, which presented a stronger positive emotion—i.e., using words, such as "good" or "love"—than human-written misinformation; in contrast, while negative emotion was also higher in GPT3-generated than in human-written misinformation, this gap was noticeably smaller than that for positive emotion (Zhou et al. 2023). This suggests that even when AI-generated texts may contain more negative emotion than human-written texts, the more striking differences lies in their higher level of positive emotional expression.

Together, despite variances in the LLMs and prompts the various AI programs use to generate text, these studies suggest a potentially more positive tone in AI-generated text. In line with them, this study hypothesizes that:

H1: *AI-generated articles are more positive in the tone than human-written articles*.

4 Emotional appeals in AI-generated content

While the overall emotional valence—positive, negative, or neutral tone—offers a useful reference for understanding the level of emotionality conveyed in AI-generated content, it is important to note the added value of examining discrete emotional states, such as anger, fear, sadness, and joy. These discrete emotions can lead to distinct psychological and behavioral changes; for instance, while anger often impedes persuasion, fear and surprise can facilitate it (Nabi 2010). Specific to people's engagement with online content, news visuals that conveyed sadness were associated with a higher level of sharing, commenting, and reacting, whereas fear was not associated with news engagements in general (Choi et al. 2021). But relative to the one-dimensional measurement of emotionality (i.e., negativity to positivity), the presence of discrete emotional states in AI-generated content remains less examined. Below we synthesize current literature, which are mostly indirect references, to speculate how discrete emotions may differ between AI-generated and human-written content.

It is commonly assumed that AI-generated content may not compete with the emotional intelligence of human-written content. Yang's (2024) study of the titles of social media posts uncovered that while 65% of AI-generated titles exhibited both a logical flow and the use of emotive references, people still clicked more on human-written titles with these characteristics, indicating AI's relatively low capacity to imitate human users' social–emotional skills. In a similar vein, readers generally find AI-generated poetry to lack the emotional depth that human poets possess, due to reasons such as the absence of nuanced thematic depth that is typically present in human-authored works (Rahmeh 2023). In a recent analysis of an AI-powered newsreade "Alice," audiences were also concerned with the lack of voice emotionality in AI-driven news anchors (Ndlovu 2024). These findings suggest that individuals are inclined to associate written works' emotional aspects (e.g., inspiring, beautiful) with human authors rather than AI tools (Porter and Machery 2024). Even within the newsrooms, journalists have asserted that AI tools lack the emotional depth that human writers possess, such that AI-generated news stories may not be as emotionally triggering as human reported news (Albizu-Rivas et al. 2024).

These contentions suggest that while AI programs are often assumed to lack genuine human emotions and psychological depth (Bashardoust et al. 2024), the distinction between human and AI language is not straightforward (Opara 2024). On one hand, readers may perceive AI-generated texts as less emotionally evocative than human-written ones, largely due to surface-level issues (e.g., grammar mistakes, shorter sentences); on the other hand, readers also recognize that AI-generated texts may contain a bigger amount of expressive language that is capable of conveying emotions (Zhang et al. 2024).

While scholars have begun to examine how emotional appeals within AI-generated vs. human-written narratives may affect readers in distinct ways (i.e., Diwanji et al. 2025), an investigation of the presence of discrete emotions within AI-generated content is scarce. Most existing studies did not directly capture discrete emotional states (i.e., fear, joy, and sadness) in AI-generated content. Instead, researchers relied on the count of words associated with different emotions (i.e., Opara 2024) or rating scales that invited human readers' assessment of the emotional depth of AI-generated content (i.e., Zhang et al. 2024) to examine the extent to which AI-generated content may convey emotions. While these techniques can broaden our understanding of AI–human communication, they yield limited insights into the patterns of distinct emotional appeals in AI-generated content. Therefore, this study contributes by exploring a key question within this landscape:

RQ1: *How would AI-generated and human-generated articles differ in their expression of discrete emotional appeals (e.g., joy, sadness, anger)?*

5 Subjectivity in AI-generated content

Another aspect that is often associated with emotionally evocative content is the use of subjective language, which is characterized by a heavy reliance on personal experiences, perspectives, and opinions. In particular, expressively subjective elements can deliver emotional states without explicitly stating these emotions (Wiebe et al. 2005). In the case of journalistic reporting, news stories are expected to demonstrate a normative objectivity through the practice of facticity and balance; however, award-winning news reports were deeply infused with subjectivity, as found in their expressions of affect, judgments of others, and evaluations of the objects that can elicit emotions (Wahl-Jorgensen 2013).

AI models, despite the lack of genuine personal experiences, may reflect biases and subjective patterns presented in training data (Ntoutsi et al. 2020). In Osborne and Bailey's

(2025) analysis of human- and ChatGPT-generated advice on topics, such as finance and dating, when human participants were not told the source of the advice, they rated AI-generated advice to be more authentically personal. When the participants were aware of the sources, however, they rated human-generated advice to be more authentic (Osborne and Bailey 2025). These findings suggest not only human users' bias against AI-generated content, but also the AI tool's strong capacity in producing authentic personal advice. Similarly, González-Arias et al.'s (2024) investigation of various types of AI-generated journalistic discourses revealed that only AI-generated opinion pieces exhibited a heavy use of first-person singulars. Both editorials and opinion pieces showed a strong expression of opinions; however, AI-generated news stories did not exhibit a strong tendency to use first person singular, recounts of personal experiences, or expression of opinions (González-Arias et al. 2024).

Indeed, existing work has accumulated mixed findings. In an investigation of AI tools in online platforms such as Airbnb, Jakesch et al. (2019) found that human customers often noted the lack of authenticity in AI-generated profiles; on the contrary, people referred to the use of personal details as evidence of human writing. In addition, Hohenstein and Jung's (2020) study on the role of smart replies in online conversations found that AI-mediated conversations were less analytic than human-only dialogs, suggesting a more logical and hierarchical thinking pattern in human-generated content. These findings highlight the complex role of subjectivity in AI-generated content, which largely depend on whether the sources (e.g., AI or human) are specified, the topic and content domain, and other factors beyond the content itself. This study focuses on the subjectivity conveyed through the content and asks:

RQ2: *How would AI-generated and human-written articles differ in subjectivity*?

6 Topical domains as moderating factors

As stated above, one potential factor that would affect the level of emotionality in both AI-generated and human-written content is the topical domains of articles, such as news on politics, science, health, sports, arts, and so on. In the case of news reporting, the level emotionality was found higher in news on crime and war but lower in sports news (Uribe and Gunter 2007). Specifically, while political coverage (i.e., elections) was mostly framed around a negative tone (Soroka and McAdams 2015) and conveyed anger (Soroka et al. 2015), stories about natural disasters (i.e., hurricanes) often conveyed a less negative tone (Ali and Gill 2022). In addition, health news tends to feature mixed emotional perspectives such as hope and fear for promoting effective actions (Nabi and Prestin 2016). Similar patterns are also found with economic news, which greatly fluctuated over time between positive and negative tones (Wlezien et al. 2017).

On the other hand, entertainment-oriented content—such as stories about sports, lifestyle, culture, arts, and movies—are less associated with negative tone. For example, in an analysis of the television coverage of the 2016 Olympic and Paralympic games, Rees et al. (2018) found nearly an absence of negative tone in the Paralympic broadcast; while a stronger negative tone was identified in the Olympic broadcast, it was primarily linked to disappointment. This is because hedonic enjoyment is the utmost motivation that drives people to watch sports events (Hall 2015), which in turn shapes the entertaining nature of sports reporting (Rees et al. 2018). Similarly, travel journalism—such as articles featuring lifestyle that do not fit the conventional definition of news—also appeals to leisure values, such as enjoyment and relaxation (Cocking 2018).

In addition, the value of objectivity has long been practiced in hard news (Thomson et al. 2008), which is defined as "any story that focuses on issues of ongoing policy consideration, factual accounting of current public events, or social issues and controversies that concern members of the audience" (Scott and Gobets 1992, p. 408). While modern media are hardly neutral on the bias scale that ranges from extreme left to extreme right (see Media Bias Chart 2025), journalists are still obligated to present neutral and factual information when covering politics, economy, and other significant social issues. On the contrary, however, these standards are not strictly applied to soft news. For example, when reporting human interest stories, journalists are allowed or even encouraged to express their own opinions and emotions (Parks 2019).

Together, these differences suggest the emotional texture of content is likely to vary by topics, with entertainment-oriented stories (or soft news) leaning less heavily on negativity and objectivity. While AI-generated content can mimic human-written content—due to the training data that are consisted of the latter—it remains unclear how the aforementioned effects of topics will transfer to AI-generated narratives. Therefore, we put forward the following question:

RQ3: *How would the differences in (a) discrete emotional appeals, (b) emotional valence, and (c) subjectivity between AI-generated and human-written news articles differ by topical domains?*

7 Methods

This study employed two publicly available data sets (thereafter referred to as Study 1 and Study 2) to investigate how discrete emotional appeals, valence, and subjectivity differ between AI-generated and human-written articles. It should be noted that by the time this study was conducted, publicly

available data sets that include both AI-generated and human-written content were scarce, making it challenging to conduct large-scale comparative analyses of their linguistic and stylistic differences. Regardless, our approach not only allows a cross-validation of patterns between two distinct data sources but also taps into two different real-world settings (essay articles vs. news stories), thereby enhancing the generalizability of our findings.

7.1 Data

Study 1 derived data from the "ChatGPT-generated text detection corpus" repository.[1] Shijaku and Canhasi (2022) first extracted TOEFL essays—which were written by individuals whose first language is not English—archived in an online document.[2] Because each human-written TOEFL essay was a response to a question (i.e., "why do you think people attend college or university?"), Shijaku and Canhasi (2022) asked the same questions to ChatGPT by attaching the phrase "write an essay on the following topic" to the beginning of each question. The final data set included 252 articles written on 126 different questions, which were averagely 355 words long (SD = 88.07, Min = 178, Max = 658). Relative to AI-generated essays, human-written essays were significantly longer (diff = 46), $t(250) = 4.33$, $p < 0.001$.

Study 2 employed a publicly available data set[3] that included *New York Times* story headlines and the full text of AI-generated news. Fang et al. (2024) first collected NYT articles that the LLMs were not trained on—specifically, the NYT news published from late 2022 to early 2023—and then asked LLMs (i.e., ChatGPT)[4] to write a corresponding content using the NYT headlines. For example, the prompt "Use 'Argentina Wins the 2022 World Cup, Defeating France' as a title to write a news article" was used to generate an AI story for the corresponding NYT headline. This presented a subset of AI-generated news stories. It should be noted that due to copyright reason, this publicly available repository did not include the full text of NYT stories or the story links. Instead, we collected the original NYT articles associated with these headlines by first programmatically scraping the full content of matching stories from the NYT website, and then removing duplicate entries and filtering out irrelevant articles (i.e., Spanish-language stories). Next, we combined this subset of human-written articles with the subset of AI-generated news stories.[5] To maximize the accuracy and reliability of our data, we only kept stories with exactly the same headline in both ChatGPT-generated and NYT-written news stories. The final data for analysis included 5,348 news articles that evenly split between ChatGPT and NYT, with an average length of 805.13 word (SD = 674.30, Min = 122, Max = 6166). The NYT stories were significantly longer (diff = 996), $t(5346) = 80.27$, $p < 0.001$.

7.2 Variable and measurements

With a focus on sentiment, this study utilized computer automation methods to capture multiple facets, including discrete emotional appeals, emotional valence, and subjectivity. Computer automation has grown significantly in recent years, driven by advances in natural language processing and machine learning. Research also demonstrates that automated tools can yield reliable and consistent results when validated against human-coded benchmarks (Young and Soroka 2012). In addition, automated coding tools are found to perform in ways comparable to human coders, such as on the assessment of sentiment and the categorization of discrete emotional states (Provoost et al. 2019). Therefore, while computer automation techniques are not without limitations (i.e., difficulty in capturing context-specific meaning), they still offer scalable and systematic approach to capture affective dimensions of written texts.[6]

All variables were measured on article level. Also see Table 1 for an overview of variables and coding techniques.

Emotional appeals. First, we employed IBM Watson's *Natural Language Understanding* (NLU)—one of the most used automated coding techniques (Biondi et al. 2017)—to capture discrete emotional states including joy, sadness, fear, disgust, and anger within a given article. IBM Watson relies on a pre-trained classification model—which was trained on data from social media conversations/commentaries, TED talks, and movie dialogs—that has "learned" to recognize everyday ways people express feelings (IBM

[1] See more details here: https://github.com/rexshijaku/chatgpt-generated-text-detection-corpus

[2] See https://englishclubmskh.wordpress.com/wp-content/uploads/2018/09/toefl-essays.pdf for details.

[3] See more details here: https://github.com/dalabudel/llmbias.

[4] Fang et al. (2024) investigated the text generated by multiple LLMs (i.e., GPT-2, GPT-3, GPT-4, and Cohere), but the analysis in this present study only employed their GPT-4 generated news stories.

[5] In addition to matching the titles, we also spot-checked several sample pairs to ensure content similarity.

[6] There is no evidence that the two linguistic analysis programs (IBM Waston NLU and TextBlob) used in this study "understand" AI-written text better than human-written text. This is because these tools do not track or distinguish authorship at all; instead, they rely on surface cues in the text (i.e., words, phrases, simple context) to identify the discrete emotional states or the valence/subjectivity of the text. Specifically, IBM Waston NLU relies heavily on explicit lexical signals and typical word usage patterns to estimate the scores for distinct emotions; TextBlob is also lexicon-based (with a pre-defined dictionary), which looks up and averages word-level polarity and subjectivity, such as treating "great" as positive and "terrible" as negative unless obvious negation flips it. Therefore, regardless of human- or AI-generated text, both tools analyze it according to known cues established in their existing dictionaries. Neither was designed with a special advantage for AI writing.

Table 1 Overview of variables and coding techniques

	Coding scheme	Coding technique
Independent variables		
Source	Whether the article was written by human or generative AI (ChatGPT)	–
Soft news	Whether a news article covered soft news or hard news; 1 = hard news (world news, business, US news, science, health, politics, elections, climate, technology, and opinion) 0 = soft news (arts, movies, style, fashion, books, magazines, travel, dining, obituaries, sports) (Study 2 only)	extracted from NYT's predefined content taxonomy included in the URLs associated with NYT news articles
Dependent variables		
Emotional appeals	Which—anger, sadness, joy, fear, or disgust—was the predominant emotion in each microblog?	computer automation with *IBM Watson*
Valence	Whether the tone of an article is positive, neutral, or negative	computer automation with *TextBlob*
Subjectivity	Whether an article is mostly comprised of facts or personal opinions	

All variables were measured on the article level.

2024). By detecting text-level patterns, such as wording and phrasing,[7] IBM Watson first estimates five emotion scores for each sentence in an article and then aggregates the sentence-level scores into article-level emotion scores (IBM 2025). It should be noted that a sentence (or an article) can be classified into multiple emotions or no emotions; in the instance of "I'm so annoyed that this code won't run," IBM Watson may estimates highest scores for both anger and sadness (IBM 2025). Therefore, the five emotion scores do not necessarily sum up to 1; instead, the bigger the score was, the more dominant an emotion was. Because discrete emotion perspectives (i.e., joy, anger, fear, etc.) "focus more on categorical emotional states" (Nabi 2010, p. 154), we then assigned each article a single dominant emotional appeal by referring to the highest probability value associated with the five discrete emotions. For example, an article was estimated with the following emotion score: 0.139 anger, 0.704 joy, 0.461 sadness, 0.101 fear, and 0.375 disgust; in this case, the dominant emotion was classified to be joy. With this coding approach, we computed a categorical variable representing the dominant emotional appeal in each article. In Study 1, a majority of TOEFL essays were associated with joy (92.06%), followed by anger (7.54%) and sadness (0.40%). In Study 2, a majority of news articles were associated with joy (58.83%) and sadness (39.49%), with a small portion conveying fear (1.05%), disgust (0.52%), and anger (0.11%).

Emotional valence and subjectivity. Next, we utilized *TextBlob* to capture the tone of the articles. TextBlob is a Python library used for Natural Language Processing (NLP).[8] It takes a lexicon-based approach, which first detects sentiment-bearing words, looks them up in a predefined dictionary, and then assigns each entry two values: polarity and subjectivity (Shah 2020). Its dictionary is comprised of thousands of words (mostly adjectives) and their numeric values for polarity and subjectivity; for example, the word "great" has a positive polarity and some subjectivity, while the phrase "not a very great calculation" has a negative polarity and some subjectivity (Schumacher 2015).

The polarity score indicates the tone of an article (thereafter referred to as valence), which varies between -1 (extremely negative) and 1 (extremely positive), with 0 indicating neutral. In both studies, the articles were averagely positive (Study 1: M = 0.19, SD = 0.09, Min = -0.07, Max = 0.46; Study 2: M = 0.11, SD = 0.08, Min = -0.5, Max = 0.63).

The subjectivity score captures whether an article is mostly comprised of facts or personal opinion, which varies from 0 (very objective) to 1 (very subjective), with the bigger value indicting higher subjectivity (Study 1: M = 0.51, SD = 0.08, Min = 0.30, Max = 0.70; Study 2: M = 0.44, SD = 0.08, Min = 0.04, Max = 0.78).

Soft news. In Study 2, the URLs associated with NYT news articles contained an element that specified the NYT's internal content taxonomy, such as "arts," "style," "travel,"

[7] In IBM Watson's pre-defined dictionary, the same words can indicate emotional valence, discrete emotional states, and subjectivity at the same time. For example, the sentence "I feel upset" combines discrete emotional state (i.e., anger), negative valence, and subjective framing. This closely mirrors the process of human coding, which may assess multiple affective dimensions from the same text as well (i.e., labelling a piece of news as "negative valence," "with a dominant anger appeal," and "lacks objectivity").

[8] See more detail here: https://textblob.readthedocs.io/en/dev/quickstart.html

"food," "politics," "science," "health," "climate," etc.[9] This provides a consistent and efficient way to classify news topics.[10] Specifically, news articles with URLs containing sections such as "arts," "style," "travel," "food," or "sports" were categorized as soft news (41.68%), while those with URLs containing sections such as "U.S. politics," "business," "world news," "climate," "technology," and "science" were coded as hard news (58.32%).

8 Results

8.1 Differences in positivity and subjectivity

As described above, valence and subjectivity were continuous variables, while discrete emotional appeal was categorical. For the sake of easier transition, we began with H1 and RQ2 and conducted independent-samples *t* tests, which are appropriate for comparing the mean scores of two groups on a continuous dependent variable, such as valence/subjectivity. The results are presented in Table 2. In Study 1, ChatGPT-generated essays (M = 0.19, SD = 0.01) were slightly positive than human-written essays (M = 0.18, SD = 0.01); however, this difference did not reach the 0.05 significance level, $t(250) = 1.04$, $p = 0.15$. Study 2 uncovered that ChatGPT-generated news (M = 0.11, SD = 0.08) was significantly more positive than NYT news (M = 0.08, SD = 0.06), $t(5346) = -16.45$, $p < 0.001$; the effect size, as measured by Cohen's d, was d = 0.45, indicating a medium effect. These findings provided partial support for H1.

In terms of the level of subjectivity, Study 1 found that human-written essays (M = 0.50, SD = 0.01) were significantly less subjective than ChatGPT-generated essays (M = 0.52, SD = 0.01), $t(250) = -2.15$, $p = 0.02$; the effect size, as measured by Cohen's d, was d = 0.28, indicating a small effect. Results from Study 2 validated this finding by revealing that the NYT stories (M = 0.35, SD = 0.04) were less subjective than ChatGPT-generated news (M = 0.38, SD = 0.06), $t(5346) = -22.44$, $p < 0.001$; the effect size, as measured by Cohen's d, was d = 0.61, indicating a medium effect. This provided a clear answer to RQ2.

8.2 Differences in discrete emotional appeals

Because both discrete emotional appeal (i.e., joy, anger, and sadness) and author (AI vs. human) were categorical variables, the chi-square test was used to determine whether there is a significant association between them (RQ1). The results are presented in Table 3.

We begin with the results from Study 1: as shown in the descriptive statistics above, only one article fell under the "sadness" category. Therefore, we transformed the "dominant emotional appeal" variable into a dummy (1 = joy and 0 = anger) for analysis, with the "sadness" category being recoded as missing. The first chi-square model in Table 3 revealed that a bigger percentage of ChatGPT-generated essays conveyed a "joy" emotion, while a bigger percentage of human-written essays were associated with an "anger" emotion; this difference, however, was not statistically significant on the 0.05 level, χ^2 (1, $N = 252$) = 2.73, $p = 0.08$, Cramér's V = 0.10.

Results from Study 2 yielded oppositional patterns: as shown in the second chi-square model in Table 3,[11] a bigger percentage of NYT news articles were associated with a "joy" emotion; on the other hand, a bigger percentage of ChatGPT-generated news conveyed a "sadness," "anger," and "fear" emotion, and these differences were statistically significant, χ^2 (4, $N = 5348$) = 118.38, $p < 0.001$, Cramér's V = 0.15. These results provided partial support for RQ1.

It should be noted that the subject matter[12] and article length are confounding variables when estimating the above differences between AI-generated and human-written content. Considering their influence and the correlations between the two continuous outcome variables, MANOVA models and logistic/multinomial logistic models were also conducted, in which subject matter and word count were included as control variables. Results from these models remained largely the same and, therefore, corroborated the findings reported above.[13]

8.3 The moderating role of topical domains

Last but not least, we tested the moderating effect of topical domains (RQ3) using Study 2 data. Table 4 presents a multinomial logistic model predicting the news stories' dominant emotional appeal as a function of the "ChatGPT x Soft news" interaction while controlling for article length.

[9] In this URL example, https://www.nytimes.com/2023/04/04/us/politics/biden-trump-arraignment.html, the NYT's internal category was U.S. politics.

[10] We then matched ChatGPT-generated news articles with the "topical domain" variable extracted from the NYT dataset. To further check topical similarity between these two subsets, we also conducted a human assessment using a random sample. This process confirmed that the topical domains (e.g., politics, health, technology) matched appropriately. Indeed, because the ChatGPT-generated news stories were "written" using the headlines of NYT articles, they were expected to reflect the same subject matter.

[11] We also conducted supplementary analyses excluding anger, disgust, and fear, which did not alter the results. For transparency, we have included the supplementary results in Appendix B.

[12] In Study 1 the subject matter refers to the list of 126 essay questions that differ in topical domains; in Study 2 the subject matter is a dichotomous variable (soft news vs. hard news).

[13] The MANOVA models were reported in Appendix A; the other models are available upon request.

Table 2 Differences in valence and subjectivity between AI-generated and human-written articles

Study 1	Human-written[a]		ChatGPT-generated[b]		Diff (human–AI)	*t*(250)
	M	SD	M	SD		
Valence	0.18	0.01	0.19	0.01	− 0.01	− 1.04
Subjectivity	0.50	0.01	0.52	0.01	− 0.02	− 2.15*
	NYT news[c]		ChatGPT-generated[d]		Diff (human–AI)	
Study 2	M	M	M	SD		*t*(5346)
Valence	0.08	0.06	0.11	0.08	− 0.03	− 16.45***
Subjectivity	0.35	0.04	0.38	0.06	− 0.03	− 22.44***

Study 1: N = 252 ([a]n = 126, [b]n = 126)

Study 2: N = ([c]n = 2674, [d]n = 2674)

In both studies the two outcome variables were logged to account for skewness

*$p < 0.05$; ***$p < 0.001$

Table 3 AI- and human-written content with chi-square analyses to examine differences in discrete emotional appeals

Articles with the dominant emotion being …	Study 1				
	Human-written		ChatGPT-written		
	n	%	*n*	%	Pearson's χ^2
anger	13	10.32	6	4.8%	2.73
joy	113	89.68	119	95.20%	
Articles with the dominant emotion being …	Study 2				
	NYT news		ChatGPT-written		
	n	%	*n*	%	Pearson's χ^2
anger	0	0.00	6	0.22	118.38***
joy	1692	63.28	1454	54.38	
sadness	982	36.72	1130	42.26	
disgust	0	0.00	28	1.05	
fear	0	0.00	56	2.09	
total	2674	100	2674	100	

n = actual number of articles and % = column percentages

In both chi-square tests, one or more expected cell counts were smaller than 10. Therefore, we used Fisher's exact tests that are more appropriate for small sample sizes or sparse data, which provide an exact *p* value reported above

***$p < 0.001$

The coefficient estimate on the interaction term was statistically significant in the "sadness" model only ($b = -0.27$, $p = 0.04$). Figure 1 plots this effect for easier interpretation: within NYT news, the predicted probability of conveying a "sad" emotion was higher in hard news than in soft news (bottom left), whereas the chances did not significantly differ among ChatGPT-generated news (bottom right). The top figures about the predicted probability of conveying a "joy" emotion yields further evidence that the variance between NYT and ChatGPT-generated news is dependent on topical domains. These findings provided evidence for RQ3(a).

Table 5 presents two OLS regression models predicting valence (logged) and subjectivity (logged)[14] as a function of the "ChatGPT x Soft news" interaction while controlling for article lengths. In both models, the coefficient estimate on the interaction term was negative; however, it was statistically significant in the "valence" model ($b = -0.03$, $p < 0.001$) but insignificant in the "subjectivity" model ($b = -0.01$, $p = 0.08$). We plotted this effect in Fig. 2, illustrating a higher level of positivity in ChatGPT-generated than NYT news;

[14] We applied a log transformation to these dependent variables to address skewness in the distribution, which normalizes the distribution, reduces the influence of extreme values, and ensures that model assumptions (i.e., homoscedasticity, normality) are better met.

Table 4 Multinomial logistic regression model predicting the dominant emotion in articles, by "AI X soft news" interaction (Study 2 only)

	The dominant emotion in article was…			
	Anger	Disgust	Fear	Sadness
ChatGPT-written	25.09 (2950.52)	17.74 (2434.46)	13.95 (1598.89)	0.82 (0.16)***
Soft news	1.40 (4269.44)	0.69 (3131.22)	0.75 (2048.33)	1.91 (0.10)***
ChatGPT X Soft	0.66 (4269.44)	2.25 (3131.22)	3.64 (2048.33)	− 0.27 (0.13)*
Word count (logged)	4.65 (1.69)**	0.94 (1.13)	− 1.06 (0.84)	0.20 (0.08)*
Constant	− 58.80 (2950.55)	− 29.22 (2434.47)	− 14.70 (1598.90)	− 3.19 (0.60)***
N	5346			
LR χ^2	1059.77***			
Pseudo R^2	0.13			

The baseline category for dependent variable was "joy as the dominant emotion in articles". Both independent variables are dummies; the comparison group for "ChatGPT-written" is "NYT news," and the comparison group for "soft news" is "hard news"

Entries are coefficient estimates with standard errors in parentheses

$^{*}p<0.05$, $^{**}p<0.01$, and $^{***}p<0.001$ were drawn from two-tailed tests

yet, this "ChatGPT–NYT" difference was bigger among soft news (diff = 0.08) than among hard news (diff = 0.05). These findings yield evidence for RQ3(b).

9 Discussion

When tapping into human readers' relatively low ability in distinguishing AI-generated content, current studies have pointed to the multiple ways AI generators can mimic human writers, especially through incorporating the linguistic, structural, and stylistic features that often present in human language. A closely related but less examined question is to what extent AI tools would automate emotionally charged content, especially when compared to the human-written counterparts. This present work contributes to this line of research by focusing on three distinct aspects related to emotionally evocative content: emotional valence, discrete emotional appeals, and subjectivity. To maximize the generalizability of our findings, we utilized two distinct data sets (essays vs. news).

One remarkable finding from this study is a higher level of subjectivity identified in ChatGPT-generated than in human-written articles, which was consistently observed in both essays and news. This contradicts with some human readers' perceptions that machine automated content is more objective, especially when the AI author is specified (see Wölker and Powell 2021). Instead, our finding corroborates the other school of thoughts that AI-generated text can be more subjective and, therefore, more emotionally triggering (Osborne and Bailey 2025; González-Arias et al. 2024). At least in the case of automating essays and news, AI tools may exhibit a heavier use of opinions—rather than facts—in their narratives. Therefore, this study yields a practical recommendation for general readers, which is to question the non-human author upon seeing an article consisted of more opinion than factuality.

Our findings also revealed a noteworthy divergence between essays and news: in Study 1, the expression of negative/positive tone or discrete emotional appeals in essays was not found to significantly differ between ChatGPT and human writers. However, Study 2 uncovered that ChatGPT-generated news exhibited a stronger level of positive tone than NYT news, which further support the positive bias identified in previous work (i.e., Mieczkowski et al. 2021). In terms of the expression of discrete emotional appeals, Study 2 found that joy made up the largest portion of news in general; however, this portion was more pronounced in NYT news, whereas ChatGPT-generated news showed a comparatively smaller portion of joy (Table 3). In addition, among NYT stories we found the emotional appeals to significantly differ between soft vs. hard news; however, the emotional appeals conveyed in ChatGPT-generated news did not significantly differ by topical domains. These findings suggest that NYT has distinctive ways of writing for hard and soft news, while ChatGPT is less bounded to journalistic conventions. It is also possible that ChatGPT-generated news tends to distribute distinct emotional appeals more evenly, or the NYT journalists refrain from extreme emotional appeals (as evidenced by the zero counts for NYT stories with an anger, disgust, or fear appeal). Because Fang et al.'s (2024) prompts did not ask ChatGPT to only rely on the NYT corpus, ChatGPT could have learned from other training data that have led its generated news to deviate from the NYT. Considering news outlets only, they tend to vary in the use of emotional appeal, valence, and subjectivity in reporting (see Molek-Kozakowska 2013; (Flaounas, et al., 2013)). On the other hand, research also uncovers that negative emotional content are more likely to be circulated online (Watson et al. 2024),

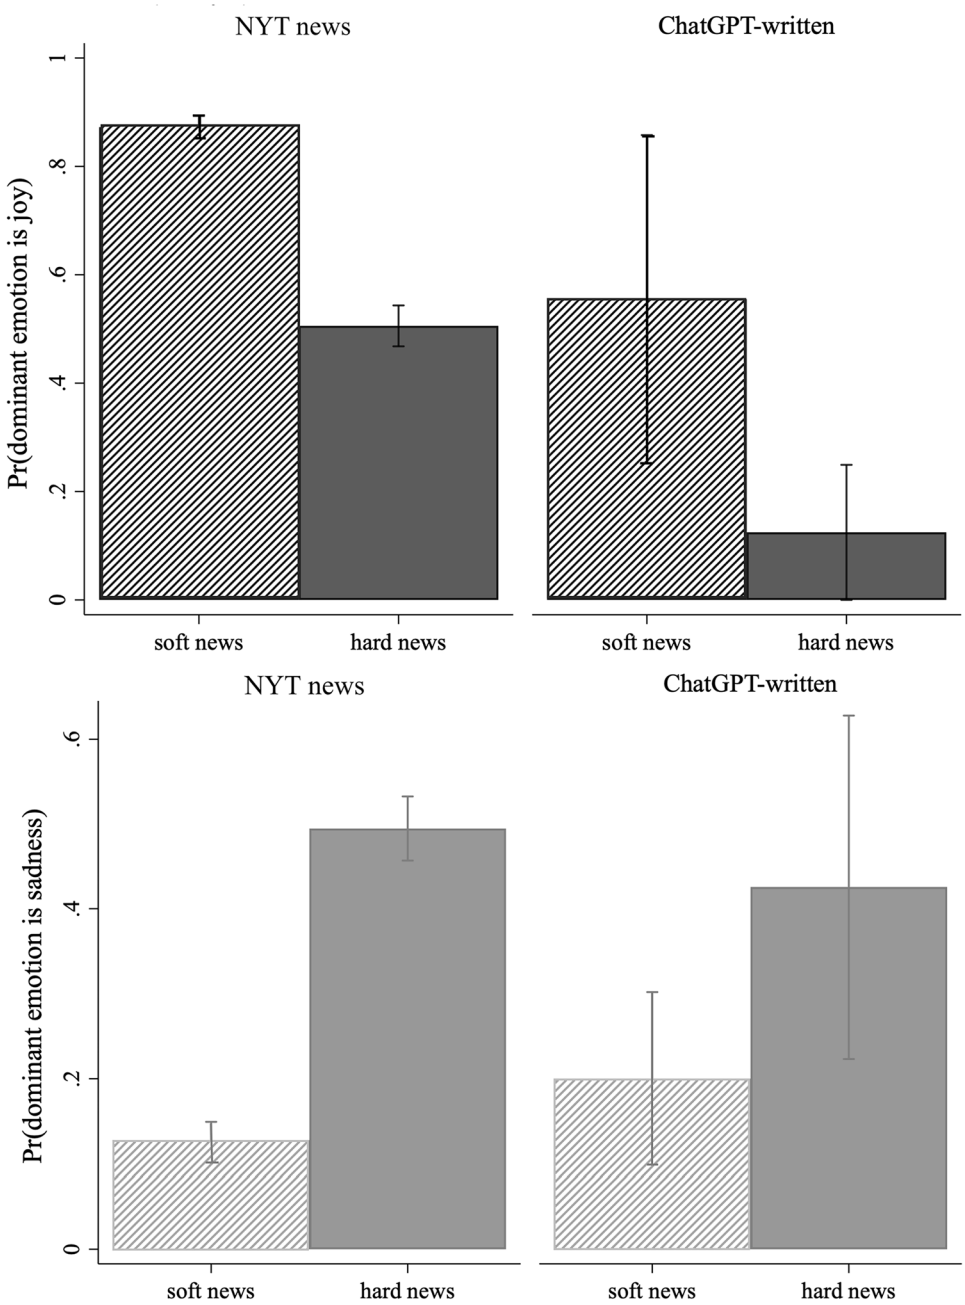


Fig. 1 Probability of dominant emotion being joy vs. sadness with 95% CIs, by AI and soft news (Study 2). Note: these figures were created using multinomial logistic regression model reported in Table 4

which can amplify the negativity in AI programs' training data. This broader set of training data further explains why in this study ChatGPT-generated news shows different patterns compared to the NYT.

The inconsistencies between the two studies are not surprising. Rather, they indicate the importance of genre (or writing styles) in understanding the emotionality in AI-generated text. Because Study 1 consisted of TOEFL essays that typically mandate counterarguments or an inquiry into solutions, their content is presumably associated with joy and hope; in this case, AI-generated essays have not only mimicked but also exaggerated these patterns to a greater extent, thereby appearing with a higher level of joy emotion than human-written essays. In the situation of news, which is less infused with a joy emotion to attract readers and make profit (Soroka et al. 2015), AI tools may have also adopted this technique to a greater degree, with fewer ChatGPT-generated news conveying joy than human-written news stories. These speculations also indicate ChatGPT's strong capacity not only to mimic but also to adapt human expression of emotions depending on the genre of prompted text.

Study 2 also underscores the importance of examining both one-dimensional and the discrete-perspective measures of emotionality. When gauged along a scale from negative

Table 5 OLS regression models predicting valence (logged) and subjectivity (logged), by "AI X soft news" interaction (Study 2 only)

	Valence (logged)	Subjectivity (logged)
ChatGPT-written	0.08(0.01)***	0.04(0.003)***
Soft news	−0.03(0.002)***	−0.02(0.002)***
ChatGPT X Soft	−0.03(0.004)***	−0.01(0.003)#
Word count (logged)	0.02(0.01)*	0.004(0.002)#
Constant	−0.06(0.06)	0.33(0.01)***
N	5346	5346
F	206.64***	204.11***
R^2	0.16	0.13

Both independent variables are dummies; the comparison group for "ChatGPT-written" is "NYT news," and the comparison group for "soft news" is "hard news"

Entries are coefficient estimates with standard errors in parentheses

#$p < 0.10$,

*$p < 0.05$, and

***$p < 0.001$ were drawn from two-tailed tests

to positive tone, ChatGPT-generated news was found to be more positive than NYT news; however, when switching to the discrete emotional appeals conveyed in news, ChatGPT-generated news was more associated with sadness and the NYT news was more associated with joy. While the dimensional and discrete views of emotion share certain similarities, the discrete view can better capture "the implications of motivational forces on human action and interaction, and thus is more proximally related to the outcomes of interest to communication scholars" (Nabi 2010, p. 158). From this perspective, our findings about the differences in the discrete emotions between ChatGPT-generated and NYT news bear more important implications for how AI-generated content may engage the readers; specifically, our findings suggest that exposure to ChatGPT-generated news may induce a stronger feeling of emotions, such as sadness and anger, which may further impair people's psychological well-being.

This study has a few limitations that also point to important agendas for future research. First, it should be noted that the level of emotional expression in AI-generated content can be modulated through emotional prompting, such as explicitly asking AI tools to produce emotional rather than informative outputs (Tafesse and Wood 2024). In addition, the same prompt may yield different texts across AI programs, due to these programs' differences in training data, decoding choices (i.e., algorithm), and system wrappers (i.e., whether to use web search). With this said, findings from this study may not generalize to the comparisons of articles generated by other AI programs, such as Gemini, Anthropic, or DeepSeek. These may explain the currently mixed findings in studies on AI-generated content. Because our measure of subjectivity is restricted to the linguistic indicators detectable by TextBlob (i.e., expression of opinions), this study was not able to investigate other facets of subjectivity, such as cognitive biases, perspective-taking, or nuanced rhetorical framing. Relatedly, this study's use of automated coding techniques may oversimplify the nuanced expressions of emotion, given that they typically rely on predefined lists of words and have a limited capacity in capturing the context. Alternatively, future research may utilize multiple text analysis programs (i.e., LIWC-22) or a combination of computer and human coders to cross-validate the measurement

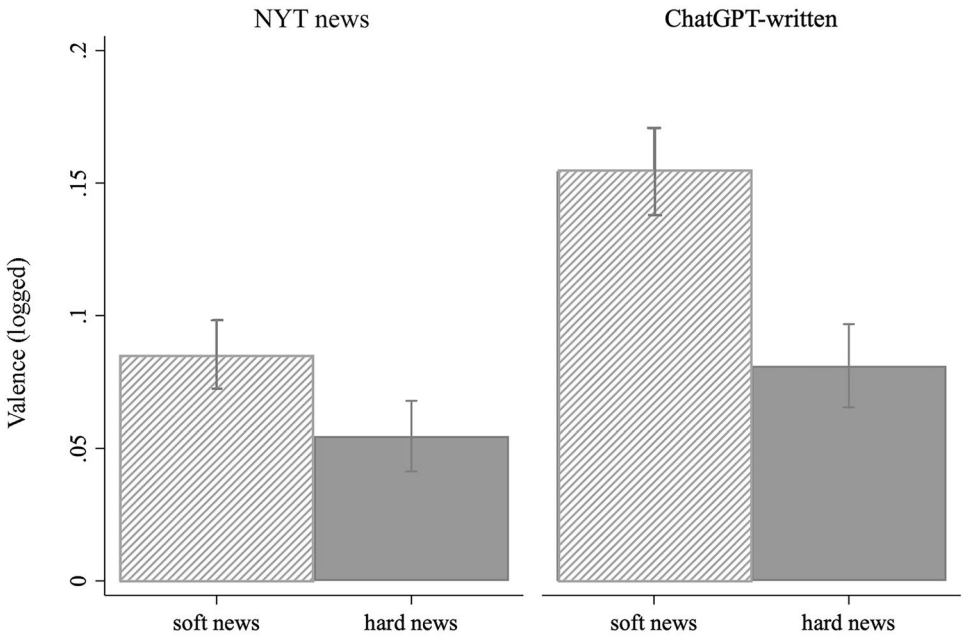


Fig. 2 Predicted valence (logged) with 95% CIs, by AI and soft news (Study 2). Note: this figure was created using the OLS regression model reported in Table 5

of emotionality, which would also allow for a deeper understanding of nuanced emotional expressions in text.

Another limitation lies in our focus on essays and news, such that our findings may not generalize to other areas of content creation, such as blog posts, social media content, creative writing, or academic writing. For example, because the TOEFL test is primarily designed for non-English speakers to demonstrate their English proficiency, our findings from Study 1 are likely not generalizable to the comparisons of essays generated by AI agents vs. those written by native English speakers. Especially, because most AI agents (i.e., ChatGPT and Gemini) are trained on data sets comprised of a large amount of English text from English-speaking countries, we would assume that the differences in the essays by AI agents and native English speakers to be smaller than what we uncovered in Study 1. Yet, it should be noted that the primary goal is consistent across various types of text, which is to present an argument in a well-structured manner and through the use of various supporting materials. From this perspective, some findings from this study—such as the stronger subjectivity in both essay and news—can still contribute to a comprehensive understanding of generative AI's role in producing emotionally evocative content.

Built on the content-level differences uncovered in this study, the next step is to examine under what situations (i.e., whether the AI authorship is explicitly specified), on which news topic (i.e., sports, business, and politics), or among which population segments (i.e., people with a higher level of AI literacy or a richer experience using AI tools), the readers would indeed rely on the varied levels of emotionality to identify AI-generated content. Future research may also consider investigating whether AI-generated vs. human-written news stories would affect readers' attitudes and behaviors in different ways. For example, this study's finding that ChatGPT-generated news is more subjective suggests that upon reading a ChatGPT-generated news story about a natural disaster, the readers may feel a higher level of empathy and express a stronger willingness to donate than after reading a more objective human-written counterpart. These research agendas are important for a more comprehensive understanding of the practical influence of AI-generated news on readers.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s00146-025-02628-9.

Author contributions Sui accessed and collected the data used for analysis (i.e., the New York Times articles for Study 2), conducted all analyses, wrote the main manuscript, prepared all tables and figures, and reviewed/revised the manuscript.

Data availability The two publicly available data sets used in this study can be accessed at https://github.com/rexshijaku/chatgpt-generated-text-detection-corpus (Study 1) and https://github.com/dalabudel/llmbias (Study 2). The human-written New York Times articles used in Study 2 are not available in online repository due to copyright issues.

Declarations

Conflict of interest No potential conflict of interest was reported by the authors.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

References

Abercrombie G, Curry AC, Pandya M, Rieser V (2021) Alexa, Google, Siri: what are your pronouns? Gender and anthropomorphism in the design and perception of conversational assistants. arXiv preprint arXiv:2106.02578. https://doi.org/10.48550/arXiv.2106.02578 Accessed 15 Jan 2025

Albizu-Rivas I, Parratt-Fernández S, Mera-Fernández M (2024) Artificial intelligence in slow journalism: journalists' uses, perceptions, and attitudes. J Media 5(4):1836–1850. https://doi.org/10.3390/journalmedia5040111

Ali SA, Gill DA (2022) Media framing and agenda setting (tone) in news coverage of Hurricane Harvey: a content analysis of the New York Times, Wall Street Journal, and Houston Chronicle from 2017 to 2018. Weather Clim Soc 14(2):637–649. https://doi.org/10.1175/WCAS-D-21-0009.1

Alter AL, Oppenheimer DM (2009) Uniting the tribes of fluency to form a metacognitive nation. Pers Soc Psychol Rev 13(3):219–235

Alvero AJ, Lee J, Regla-Vargas A, Kizilcec RF, Joachims T, Antonio AL (2024) Large language models, social demography, and hegemony: comparing authorship in human and synthetic text. J Big Data 11(1):138

Bashardoust A, Feuerriegel S, Shrestha YR (2024) Comparing the willingness to share for human-generated vs. AI-generated fake news. Proceedings of the ACM on Human-Computer Interaction 8(CSCW2). pp1–21. https://doi.org/10.48550/arXiv.2402.07395

Berger J, Milkman KL (2012) What makes online content viral? J Mark Res 49(2):192–205. https://doi.org/10.1509/jmr.10.0353

Biondi G, Franzoni V, Poggioni V (2017) A deep learning semantic approach to emotion recognition using the IBM Watson Bluemix alchemy language. International Conference on Computational Science and Its Applications. Springer, Cham, pp 718–729. https://doi.org/10.1007/978-3-319-62398-6_51

Boczkowski PJ, Mitchelstein E, Matassi M (2018) News comes across when I'm in a moment of leisure: understanding the practices of incidental news consumption on social media. New Media Soc 20(10):3523–3539. https://doi.org/10.1177/1461444817750396

Brandl R, Ellis C (2023) Survey: ChatGPT and AI content—can people tell the difference? www.tooltester.com. https://www.tooltester.com/en/blog/chatgpt-survey-can-people-tell-the-difference/ Accessed 15 June 2025

Cai Z, Duan X, Haslett D, Wang S, Pickering M (2023) Do large language models resemble humans in language use?. In Proceedings of the Workshop on Cognitive Modeling and Computational Linguistics. pp 37-56. https://doi.org/10.18653/v1/2024.cmcl-1.4

Choi J, Lee SY, Ji SW (2021) Engagement in emotional news on social media: intensity and type of emotions. J Mass Commun Q 98(4):1017–1040

Cocking B (2018) News values go on holiday: the ideological values of travel journalism. J Stud 19(9):1349–1365. https://doi.org/10.1080/1461670X.2016.1272066

Diwanji VS, Geana M, Pei J, Nguyen N, Izhar N, Chaif RH (2025) Consumers' emotional responses to AI-generated versus human-generated content: the role of perceived agency, affect and gaze in health marketing. Int J Hum Comput Interact. https://doi.org/10.1080/10447318.2025.2454954

Duan W, McNeese N, Li L (2025) Gender stereotypes toward non-gendered generative AI: the role of gendered expertise and gendered linguistic cues. Proc ACM Hum-Comput Interact 9:1–35. https://doi.org/10.1145/3701197

Fang X, Che S, Mao M, Zhang H, Zhao M, Zhao X (2024) Bias of AI-generated content: an examination of news produced by large language models. Sci Rep 14(1):1–20. https://doi.org/10.1038/s41598-024-55686-2

Fatahi S, Vassileva J, Roy CK (2024) Comparing emotions in ChatGPT answers and human answers to the coding questions on stack overflow. Front Artif Intell 7:1393903

Flaounas I, Ali O, Lansdall-Welfare T, De Bie T, Mosdell N, Lewis J, Cristianini N (2013) Research methods in the age of digital journalism: massive-scale automated analysis of news-content—topics, style and gender. Digit J 1(1):102–116

Gao J, Yin J (2024) Emotional expression of AI-generated artistic design: a case analysis approach. J Art Design. https://doi.org/10.21659/cjad.81.v8n104

González-Arias C, Chatzikoumi E, López-García X (2024) The anthropomorphic pursuit of AI-generated journalistic texts: limits to expressing subjectivity. Front Commun 9:1456509

Guzman AL, Lewis SC (2020) Artificial intelligence and communication: a human–machine communication research agenda. New Media Soc 22(1):70–86. https://doi.org/10.1177/1461444819858691

Hall AE (2015) Entertainment-oriented gratifications of sports media: contributors to suspense, hedonic enjoyment, and appreciation. J Broadcast Electron Media 59(2):259–277. https://doi.org/10.1080/08838151.2015.1029124

Hancock JT, Naaman M, Levy K (2020) Ai-mediated communication: definition, research agenda, and ethical considerations. J Comput-Mediat Commun 25(1):89–100. https://doi.org/10.1093/jcmc/zmz022

Hatfield E, Cacioppo JT, Rapson RL (1993) Emotional contagion. Curr Dir Psychol Sci 2(3):96–100. https://doi.org/10.1111/1467-8721.ep10770953

Herbold S, Hautli-Janisz A, Heuer U, Kikteva Z, Trautsch A (2023) AI, write an essay for me: A large-scale comparison of human-written versus ChatGPT-generated essays. arXiv preprint arXiv:2304.14276. Accessed 15 Jan 2025

Hohenstein J, Jung M (2018) AI-supported messaging: An investigation of human-human textconversation with AI support. In Extended abstracts of the 2018 CHI conference on human factors incomputing systems, Montreal, QC, Canada, April 21–26, 2018, pp. 1-6. https://doi.org/10.1145/3170427.3188487.

Hohenstein J, Jung M (2020) Ai as a moral crumple zone: the effects of AI-mediated communication on attribution and trust. Comput Human Behav 106:106190. https://doi.org/10.1016/j.chb.2019.106190

Huang HW, Chang J (2024) Human-AI interactions in teacher education: examining social presence and friendship. In: Proceedings of the 2024 International Conference on Artificial Intelligence and Teacher Education. pp 64–69. https://doi.org/10.1145/3702386.3702399

IBM (2024) Emotion. www.ibm.com. https://www.ibm.com/docs/en/watson-libraries?topic=catalog-emotion Accessed 25 August 2025

IBM (2025) Emotion classification. dataplatform.cloud.ibm.com. https://dataplatform.cloud.ibm.com/docs/content/wsj/analyze-data/watson-nlp-block-emotion.html?context=wx Accessed 25 August 2025

Jakesch M, French M, Ma X, Hancock JT, Naaman M (2019) AI-mediated communication: How the perception that profile text was written by AI affects trustworthiness. In Proceedings of the 2019 CHI conference on human factors in computing systems. pp 1–13. https://doi.org/10.1145/3290605.3300469

Jia H, Appelman A, Wu M, Bien-Aimé S (2024) News bylines and perceived AI authorship: effects on source and message credibility. Comput Hum Behav Artif Hum 2(2):100093. https://doi.org/10.1016/j.chbah.2024.100093

Kreps S, McCain RM, Brundage M (2022) All the news that's fit to fabricate: AI-generated text as a tool of media misinformation. J Exp Polit Sci 9(1):104–117. https://doi.org/10.1017/XPS.2020.37

Landymore F (2023) CNET is quietly publishing entire articles generated by AI. futurism.com. https://futurism.com/the-byte/cnet-publishing-articles-by-ai Accessed 15 June 2025

Lermann Henestrosa A, Kimmerle J (2024) The effects of assumed AI vs. human authorship on the perception of a GPT-generated text. J Media 5(3):1085–1097. https://doi.org/10.3390/journalmedia5030069

Lewis SC, Guzman AL, Schmidt TR (2019) Automation, journalism, and human–machine communication: rethinking roles and relationships of humans and machines in news. Digit J 7(4):409–427. https://doi.org/10.1080/21670811.2019.1577147

Liu B, Wei L (2018) Reading machine-written news: effect of machine heuristic and novelty on hostile media perception. In: Human-computer interaction. Theories, Methods, and Human Issues: 20th International Conference, HCI International 2018, Las Vegas, July 15–20, 2018, Proceedings, Part I 20, Springer International Publishing. pp 307–324. https://doi.org/10.1007/978-3-319-91238-7_26

Marinescu V, Fox B, Roventa-Frumusani D, Branea S, Marinache R (2022) News audience's perceptions of and attitudes towards AI-generated news. Futures of Journalism: technology-stimulated evolution in the audience-news media relationship. Springer International Publishing, Cham, pp 295–311

Media bias chart (2025) https://adfontesmedia.com/gallery/ Accessed 15 June 2025

Mieczkowski H, Hancock JT, Naaman M, Jung M, Hohenstein J (2021) Ai-mediated communication: language use and interpersonal effects in a referential communication task. Proc ACM Hum-Comput Interact 5:1–14

Molek-Kozakowska K (2013) Towards a pragma-linguistic framework for the study of sensationalism in news headlines. Discourse Commun 7(2):173–197

Nabi RL (2010) The case for emphasizing discrete emotions in communication research. Commun Monogr 77:153–159. https://doi.org/10.1080/03637751003790444

Nabi RL, Prestin A (2016) Unrealistic hope and unnecessary fear: exploring how sensationalistic news stories influence health behavior motivation. Health Commun 31(9):1115–1126. https://doi.org/10.1080/10410236.2015.1045237

Nah S, Luo J, Kim S, Chen M, Mitson R, Joo J (2024) Algorithmic bias or algorithmic reconstruction? A comparative analysis between AI news and human news. Int J Commun 18:700–730

Nass C, Moon Y (2000) Machines and mindlessness: social responses to computers. J Soc Issues 56(1):81–103. https://doi.org/10.1111/0022-4537.00153

Ndlovu M (2024) Audience perceptions of AI-driven news presenters: a case of 'Alice' in Zimbabwe. Media Cult Soc 46(8):1692–1706. https://doi.org/10.1177/01634437241270982

Ntoutsi E, Fafalios P, Gadiraju U, Iosifidis V, Nejdl W, Vidal ME, Vidal M-E, Ruggieri S, Turini F, Papadopoulos S, Krasanakis E, Kompatsiaris I, Kinder-Kurlanda K, Wagner C, Karimi F, Fernandez M, Alani H, Berendt B, Kruegel T, Heinze C, Broelemann K, Kasneci G, Tiropanis T, Staab S (2020) Bias in data-driven artificial intelligence systems—an introductory survey. Wires Data Min Knowl Discov 10(3):e1356. https://doi.org/10.1002/widm.1356

Opara C (2024) StyloAI: distinguishing AI-generated content with stylometric analysis. International conference on artificial intelligence in education. Springer Nature, Cham, pp 105–114. https://doi.org/10.48550/ARXIV.2405.10129

Ortega-Ochoa E, Sabaté JM, Arguedas M, Conesa J, Daradoumis T, Caballé S (2024) Exploring the utilization and deficiencies of generative artificial intelligence in students' cognitive and emotional needs: a systematic mini-review. Front Artif Intell 7:1493566

Osborne MR, Bailey ER (2025) Me vs. the machine? Subjective evaluations of human-and AI-generated advice. Sci Rep 15(1):3980

Pani B, Crawford J, Allen KA (2024) Can generative artificial intelligence foster belongingness, social support, and reduce loneliness? A conceptual analysis. Applications of Generative AI. pp 261–276. https://doi.org/10.1007/978-3-031-46238-2_13

Parks P (2019) An unnatural split: how "human interest" sucks the life from significant news. Media Cult Soc 41(8):1228–1244. https://doi.org/10.1177/0163443718813498

Paul N, Sui M (2022) I can feel what you feel: emotion exchanges in twitter conversations between candidates and the public. J Polit Market 21(2):175–195. https://doi.org/10.1080/15377857.2019.1684419

Petty RE, Cacioppo JT (1986) Message elaboration versus peripheral cues. In Communication and persuasion: central and peripheral routes to attitude change. Springer, New York. pp 141–172

Porter B, Machery E (2024) Ai-generated poetry is indistinguishable from human-written poetry and is rated more favorably. Sci Rep 14(1):26133

Provoost S, Ruwaard J, Van Breda W, Riper H, Bosse T (2019) Validating automated sentiment analysis of online cognitive behavioral therapy patient texts: an exploratory study. Front Psychol 10:1065. https://doi.org/10.3389/fpsyg.2019.01065

Rahmeh H (2023) Digital verses versus inked poetry: exploring readers' response to AI-generated and human-authored sonnets. Sch Int J Linguist Lit 6(9):372–382

Rees L, Robinson P, Shields N (2018) A major sporting event or an entertainment show? A content analysis of Australian television coverage of the 2016 Olympic and Paralympic Games. Sport Soc 21(12):1974–1989. https://doi.org/10.1080/17430437.2018.1445996

Samo A, Highhouse S (2023) Artificial intelligence and art: identifying the aesthetic judgment factors that distinguish human-and machine-generated artwork. Psychol Aesthet Creat Arts. https://doi.org/10.1037/aca0000570

Santana AD, Hopp T (2022) Blink and you miss it: measuring news readers' attention to interpretative journalism cues. J Pract 16(6):1192–1208. https://doi.org/10.1080/17512786.2020.1827968

Schumacher A (2015) TextBlob Sentiment: Calculating Polarity and Subjectivity. planspace.org. https://planspace.org/20150607-textblob_sentiment/ Accessed 25 August 2025

Scott DK, Gobetz RH (1992) Hard news/soft news content of the national broadcast networks, 1972–1987. Journalism Quarterly 69(2):406–412.https://doi.org/10.1177/107769909206900214.

Shah P (2020) Sentiment Analysis using TextBlob. medium.com. https://medium.com/data-science/my-absolute-go-to-for-sentiment-analysis-textblob-3ac3a11d524 Accessed 25 August 2025

Shijaku R, Canhasi E (2022) ChatGPT Generated Text Detection Corpus . https://github.com/rexshijaku/chatgpt-generated-text-detection-corpus Accessed 15 January 2025

Soroka S, McAdams S (2015) News, politics, and negativity. Polit Commun 32(1):1–22. https://doi.org/10.1080/10584609.2014.881942

Soroka S, Young L, Balmas M (2015) Bad news or mad news? Sentiment scoring of negativity, fear, and anger in news content. Ann Am Acad Polit Soc Sci 659(1):108–121. https://doi.org/10.1177/0002716215569217

Sundar SS (2008*)* The MAIN model: a heuristic approach to understanding technology effects on credibility, Cambridge, MA: MacArthur Foundation Digital Media and Learning Initiative. pp 73–100. https://doi.org/10.1162/dmal.9780262562324.073

Tafesse W, Wood B (2024) Hey ChatGPT: an examination of ChatGPT prompts in marketing. J Mark Anal. https://doi.org/10.1057/s41270-023-00284-w

Thomson EA, White PR, Kitley P (2008) "Objectivity" and "hard news" reporting across cultures: comparing the news report in English, French, Japanese and Indonesian journalism. J Stud 9(2):212–228. https://doi.org/10.1080/14616700701848261

Uribe R, Gunter B (2007) Are "Sensational" news stories more likely to trigger viewers' emotions than non-sensational news stories? A content analysis of British TV news. Eur J Commun 22(2):207–228. https://doi.org/10.1177/0267323107076770

Wahl-Jorgensen K (2013) Subjectivity and story-telling in journalism: examining expressions of affect, judgement and appreciation in Pulitzer Prize-winning stories. J Stud 14(3):305–320. https://doi.org/10.1080/1461670X.2012.713738

Watson J, van der Linden S, Watson M, Stillwell D (2024) Negative online news articles are shared more to social media. Sci Rep 14(1):21592

Wiebe J, Wilson T, Cardie C (2005) Annotating expressions of opinions and emotions in language. Lang Resources Eval 39:165–210. https://doi.org/10.1007/s10579-005-7880-9

Wlezien C, Soroka S, Stecula D (2017) A cross-national analysis of the causes and consequences of economic news. Soc Sci Q 98(3):1010–1025. https://doi.org/10.1111/ssqu.12445

Wölker A, Powell TE (2021) Algorithms in the newsroom? News readers' perceived credibility and selection of automated journalism. Journalism 22(1):86–103. https://doi.org/10.1177/1464884918757072

Yang W (2024) Beyond algorithms: the human touch machine-generated titles for enhancing click-through rates on social media. PLoS ONE 19(7):e0306639. https://doi.org/10.1371/journal.pone.0306639

Young L, Soroka S (2012) Affective news: the automated coding of sentiment in political texts. Polit Commun 29(2):205–231. https://doi.org/10.1080/10584609.2012.671234

Zhang JJ, Wang YW, Ruan Q, Yang Y (2024) Digital tourism interpretation content quality: a comparison between AI-generated content and professional-generated content. Tour Manage Perspect 53:101279. https://doi.org/10.1016/j.tmp.2024.101279

Zhou J, Zhang Y, Luo Q, Parker AG, De Choudhury M (2023) Synthetic lies: Understanding ai-generated misinformation and evaluating algorithmic and human solutions. In Proceedings of the 2023 CHI conference on human factors in computing systems. pp 1–20. https://doi.org/10.1145/3544548.3581318

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.