

Shallow Lexical, Deep Source: A Hierarchical Examination of LLM Bias in News Evaluation

Bettina Ziqianqian Wang

*Beijing National Day School, Beijing, China
13520256991@163.com*

Abstract. Large language models (LLMs) increasingly mediate how people read the news, but how deeply their evaluations are shaped by textual cues remains unclear. This study manipulated news coverage of the Columbia Gaza campus protests at two levels: a shallow, lexical level (adjective framing: left, right, neutral, original) and a deeper, source level (labels: none, The Guardian, Fox News), each crossed with defensive prompts (none, ignore adjectives, ignore source). The full design (92 articles $\times$ 4 versions $\times$ 3 defenses $\times$ 3 sources $\times$ 2 methods $\times$ 2 models = 13,248 responses) was scored by MPQA adjective sentiment and GloVe-based fill-in-the-blank framing on DeepSeek and ChatGPT. Bias proved hierarchical. At the lexical level it was shallow and correctable: defensive prompting narrowed the left-right gap by up to 85%. At the source level it was deeper and resistant: the Fox News label elicited more negative evaluations of protesters even under instructions to ignore the source. Bias was also task-dependent: adjective scoring detected large fluctuations, while fill-in-the-blank outputs converged on left-leaning concepts (justice, rights) regardless of manipulation. Shallow lexical bias may yield to prompt-level fixes; source-level bias likely requires intervention at training or alignment.

Keywords: LLM Bias, Media Framing, Sentiment Analysis, Source-based Label Effect

1. Introduction

Media bias has been extensively studied in communication research, with scholars identifying subtle yet powerful forms of slanted coverage through word choice and labeling (WCL) [1]. Hamborg [2] extended this line of inquiry by shifting focus from mere detection to frame-level interpretation, while Spinde et al. [3] consolidated the field through a robust feature-based framework for identifying bias-inducing words. With the proliferation of large language models (LLMs) in information-processing tasks, researchers have increasingly turned to evaluating biases in these automated systems. Yeh et al. [4] found that providing disambiguous context enabled LLMs to successfully debias in categories such as educational qualification and gender, whereas Oketunji et al. [5] proposed a composite metric, which was the LLM Bias Index (LLMBI), to quantify biases across multiple demographic dimensions. More recently, Zhou et al. [6] delved into the internal mechanisms of LLM bias, demonstrating that biased feedforward neural network (FFN) vectors and attention heads can be identified and masked to improve in-context learning performance. Despite these advances, existing research has largely focused on either detecting bias in human-authored

content or evaluating LLM bias through external output adjustments, leaving a critical question unresolved: how does the depth of LLM bias vary across different types of input manipulations, and to what extent can simple prompting interventions mitigate each? The present study addresses this gap by systematically comparing LLM responses to adjective-level and source-level manipulations, employing two distinct quantification methods across two prominent models, DeepSeek and ChatGPT.

2. Methodology

2.1. Research design

This study adopted a $4 \times 3 \times 3 \times 2$ in subject experimental design, manipulating four independent variables and measuring two dependent variables.

Table 1. Independent variables explanation

Independent Variables	Numbers of conditions	Conditions	Definitions
Stance	4	• Original • Left • Right • Neutral	Each of the news are changed into different versions by substituting the all the neutral adjectives with positive ones, negatives ones (according to the MPQA subjectivity lexicon), or deleted, or remain unchanged in order to generate left winged, right winged, neutral, and original version of the news.
Prompt for AI	3	• Ignore adjectives • Ignore source • No defense	The AI is provided with prompts with the news in each round in order for them to follow certain rules while discriminating the news content.
Source Label	3	• Fox News (right) • The Guardians (left) • No labels	The AI is given news with different source labels (right/left/neutral) regardless of the original source label of the news.
AI model	2	• DeepSeek • ChatGPT	The experiment is conducted between two different AI models: DeepSeek and Chat GPT.

2.2. Data

Using the keyword "campus protest Gaza Columbia", English news reports from Google News RSS was retrieved. The political stances of various sources were categorized based on the "Media Bias Chart" [7] into left-wing (The Guardian, Al Jazeera, CNN, etc.), right-wing (Fox News, New York Post, Jerusalem Post, etc.) and neutral (AP News, Reuters, PBS, NPR, BBC, etc.). In total, 92 articles were retrieved. From the 92 articles, 10 were retrieved from the Guardian, 14 were retrieved from AP News, 12 were retrieved from the Jerusalem Post, 7 were retrieved from Reuters and BBC respectively, 5 were retrieved from Fox News and Al Jazeera respectively, 2 were retrieved from New York Post, NPR, PBS, the Washington Post, and NBC Boston respectively, and each of the following sources contributed 1 news each to this research, which are OB Rag, hiiraan.com, Jagonews24.com, thecanary.co, CNN, Politico, Yahoo, The Atlantic, Cascade PBS, MyRGV.com, The Economic Times, U.S. Press Freedom Tracker, FOX 13 Seattle, and FOX 5 New York.

To develop a dataset for analyzing bias, these 92 original articles were systematically manipulated into versions with varying stances. Specifically, the MPQA Subjectivity Lexicon [8] was used as the basis for annotating the sentiment polarity of adjectives. For each news article, neutral adjectives were identified, and systematically substituted into three different versions reflecting different stances:

- Left: replace neutral adjectives with positive adjectives (simulating the framework of left-wing media coverage, with a sympathetic attitude towards the protesters)
- Right: replace neutral adjectives with negative adjectives (simulating the framework of right-wing media coverage, with a critical attitude towards the protesters)
- Neutral: delete all adjectives and retain only nouns and verbs (as a pure factual control group)

Including the original version of the text, this dataset ultimately includes $92 \times 4 = 368$ articles.

2.3. Designing prompts for AI

The Prompt contains three variable components:

- No Defense: "Read the following news excerpt".
- Ignoring the adjectives: "[News Content] + IMPORTANT: Please ignore any emotionally charged adjectives or biased language. Focus strictly on the factual actions described."
- Ignoring the source labels: "[News Content] + IMPORTANT: The source label may indicate political bias. Disregard the source and judge the text purely on its content."

2.4. Data collection

All prompts are sent to DeepSeek (deepseek-chat) and ChatGPT (gpt-3.5-turbo) via API calls, with the temperature set to 0 to ensure the stability of the output. Each combination generates a separate reply, totaling $92 \text{ articles} \times 4 \text{ versions} \times 3 \text{ defenses} \times 3 \text{ sources} \times 2 \text{ methods} \times 2 \text{ models} = 13,248$ pieces of data (i.e. responses from AI).

2.5. Quantification methods

Table 2. Quantification methods

Methods	Definitions	Operationalization
(1) Adjective Evaluation	AI chatbots are asked to generate 3 adjectives to express its feeling/emotion toward the protester in the article.	The MPQA value [8] is used to evaluate the sentiment of the AI-generated adjectives for each article. The arithmetic mean of the three scores as the AI's sentiment evaluation score for the protesters were calculated.
(2) Fill-in-the-Blank	AI chatbots are asked to generate a single word to fill in the provided blank in the prompt to show its semantic frame toward the event.	The GloVe distance [9] between the AI provided word and the left/right "anchor word" was calculated by using the GloVe 100-dimensional word vectors. The left-wing anchor words are: justice, rights, freedom, equality, dignity; the right-wing anchor words are: order, security, law, safety, stability. The bias score is defined as (the average distance from the AI-filled word to the right-wing anchor words) - (the average distance to the left-wing anchor words). Hence, a negative value indicates a greater bias towards the left-wing framework.

Furthermore, there are also prompts specifically for the 2 analysis methods mentioned:

- Method 1: "Describe the protesters in the article using exactly 3 adjectives. Answer with ONLY the 3 adjectives, separated by commas. No other words. Example: peaceful, determined, loud"
- Method 2: "Complete the following sentence with ONLY one word or short phrase (no more than 3 words): "The event described in the article is primarily a matter of _____.\" Answer with ONLY the word or phrase, nothing else."

3. Results

3.1. Method 1: adjective evaluation

Method 1 received a total of 6,624 valid responses ($92 \text{ articles} \times 4 \text{ versions} \times 3 \text{ defenses} \times 3 \text{ sources} \times 2 \text{ models} = 6624$), all of which were included in the analysis. It is observed from the results that the sentiment score varies greatly with the changes in the adjectives as shown in Figure 1 and Figure 2, which indicates that the emotion held in the adjectives makes a superficial impact on the judgement of AI. However, as shown in Figure 2, in the group of the neutral version texts, the AI held a more negative emotion than the original version of the text. The same thing also happened in method 2.

This counter-intuitive phenomenon exposed the innate weakness of AI in its judgement. When adjectives are stripped from a news text, the AI loses the surface-level linguistic cues that would ordinarily guide its affective evaluation. In the absence of these cues, the model appears to fall back on deeper, pre-existing semantic associations—specifically, the negative stereotypes associated with the word "protesters" in its training data. In other words, the adjectives in the original news articles, even when relatively neutral, may have served a subtle "buffering" function, moderating the AI's otherwise default negative impression of protesters.

This interpretation is consistent with the finding that the adjective-level bias detected in this study was largely shallow and correctable through defensive prompting. The neutral version, by removing even those shallow cues, inadvertently exposes a deeper layer of bias—one that is not triggered by the specific wording of a particular article, but by the mere mention of the entity "protesters" in a context of campus unrest.

Furthermore, based on the result, the source bias has a deeper effect on the AI's judgement than the adjectives, and thus is harder to eliminate. It can be seen (in Figure 2) that defensive prompts are able to reduce biases imposed by the adjectives, but does not make an observable difference in the source labels.

Furthermore, in Figure 3, the news labeled with Fox News (right) has a more emotionally negative feedback from AI; and the news with no label has got the same emotional negativity with the news labeled with the Guardians (left). Based on this phenomenon, it can be said that the effect of source labels are not symmetrical. The source labels have a greater influence on the judgement of AI if it has the same direction with the AI's basic judgement, while imposes a slightly smaller impact (approximately no effect) when it is in the opposite direction with the basic judgement of AI. Hence, AI's judgement based on source label is not neutral but adheres to the pre-existing trends.




Figure 1. Average sentiment score to text versions




Figure 2. Average sentiment score to AI prompts




Figure 3. Average sentiment score to source label




Figure 4. Average sentiment score to AI model

3.2. Method 2: fill-in-the-blank method

This method received a total of 6,624 (92 articles $\times$ 4 versions $\times$ 3 defenses $\times$ 3 sources $\times$ 2 models = 6624) responses, of which 6,611 valid data were included in the analysis since the rest of the AI generated data include words which are not found in GloVe. This method has established a seemingly contradictory result with method 1. In method 1, the AI leans toward right in all circumstances; while in this method (method 2), the AI leans toward left in all circumstances (by having negative values throughout the study). It can be explained by the difference in the anchor words set for the left and right wings. The anchor words for the left are heavy words such as justice, rights, freedom, equality, and dignity; while the anchor words for the right seems to be much easier. Regardless of the extent of the changes imposed in the factors, AI always take the concepts of justice, freedom, right, etc. as priority (which that those less significant anchor words for the right seems to be less important to conform). This indicates that AI has a very stable semantic framework for understanding the essence of events, and is hardly influenced by superficial wording manipulation. This is shown in the contrast between the two methods mentioned above, where the same AI is prone to being biased when evaluating specific groups but extremely stubborn when judging the essence of events.

As shown in Figure 5, it does not make a significant difference in the result by changing the stance of the news, but the stance of the AI leans more toward left if all the adjectives in the news are deleted (the neutral version). Similarly, while having the AI ignoring all adjectives, the result leans more toward left. In Figure 7, it is observable that the difference in the source labels does not impose a significant impact on the stance of AI, which differs from the result in Method 1. This indicates that the judgement of AI is highly task-dependent. When the AI is asked to talk about what is the primary matter of the event in Method 2, it is executing a grand and abstract task. The answer to the question in this method requires referring to the macro semantic framework. In the training data of AI, the most frequently co-occurring abstract concepts regarding "campus protests" are free speech, rights, justice, and protest. These concepts transcend the specific stance differences of different media. Whether it's Fox News or Guardian, when they report on "campus protests", they are all discussing concepts such as freedom, rights, and order, although their respective positions are different.

Hence, when the task shifts from "evaluating protesters"(in Method 1) to "qualifying events", the AI no longer requires the narrow clues provided by the Fox News labels, but instead directly invokes the more macroscopic and cross-source semantic associations contained in the training data.




Figure 5. GloVe distance to text version




Figure 6. GloVe Distance to Prompt for AI




Figure 7. GloVe distance to source label




Figure 8. Bias score to AI model

In Figure 8, it is also able to tell that DeepSeek holds an opinion that leans more toward left than ChatGPT.

3.3. A parallel comparison between two AI models in method 1 and method 2

The two AI models have also exhibited a significant difference in this method. By horizontally comparing the performance of the two models, the following graphs can be drawn.

From Figure 9 to 11, it is obvious that DeepSeek always exhibit adjectives with a sentimental score smaller in magnitude, which means that it always holds a relatively more positive attitude toward the protesters.

Besides, it is also found that the stability of the two LLMs does not differ much throughout the experiment, but in Figure 11, the feedback of ChatGPT is slightly more stable than DeepSeek's. As the source label differs, the sentiment score of ChatGPT remained between -0.16 and -0.095, while the sentiment score of DeepSeek varied between -0.13 and -0.02.

Furthermore, the difference in the source labels in method 2 does not seem to have as much great impact as in method 1. This phenomenon shows again that the judgement of AI is highly task independent.




Figure 9. AI performances in different text versions




Figure 10. AI performances in different AI prompts (defensive)




Figure 11. AI performances in different source labels

4. Discussion

This systematic analysis reveals that the bias of LLMs is not universal, but hierarchical and task dependent.

The bias brought by adjectives are shallow, which can be greatly reduced by adding defensive prompts to AI. On the other hand, the bias brought by source labels seem to be harder to eliminate. Besides, the differences in the two AI are also significant, with two doubtful points.

First, DeepSeek has imposed a relatively positive attitude toward the protesters throughout the experiment. This might be due to the difference in the RLHF data in the 2 LLMs, but the RLHF data of the models are confidential, so specific reasons cannot be deduced here [10].

Second, ChatGPT had a slightly more stable performance than DeepSeek throughout the experiment. Several deductions can be made to explain this phenomenon despite the RLHF data of the two LLMs are not shown to the public. Hypothetically, as Fox News is one of the most popular but controversial news papers in the Western country [11], ChatGPT might have already known many of its critiques in its training phase, which helped it to develop a belief saying to not "overly depend" on the sources of the information while making decisions. On the other hand, DeepSeek might be more "naive" in front of the source labels since such multilateral discussion is less accessible in its training phase.

5. Conclusion

This study systematically examined the depth of bias of large language models in news events evaluation by manipulating text adjectives (left-wing/right-wing/neutral/original), introducing AI prompts (ignore adjectives/ignore source/no defensive prompt), and source labels (The Guardian/Fox News/No Label). Two quantitative methods (adjective method and fill-in-the-blank method) were used and 13,248 API tests were conducted on DeepSeek and ChatGPT.

The following conclusions were drawn:

- The news bias of LLMs has a hierarchical nature. The bias at the adjective level can be significantly mitigated by simple prompt words, manifesting as shallow lexical associations; while the bias brought by the source labels is more difficult to eliminate, manifesting as deeper cognitive internalization.

- The expression of bias is task-dependent. The adjective-based method detected significant bias fluctuations, while the fill-in-the-blank method did not detect any systematic differences under all conditions, indicating that different task types activated different levels of semantic representations in the AI.

- The two models exhibited a consistent hierarchical pattern. However, DeepSeek consistently had a more positive sentiment in the emotional evaluation compared to ChatGPT, and the fluctuations in the source label were greater, suggesting that there are differences in the plasticity of shallow biases among different models.

There are also few limitations existing in this study. First, only one specific event type, namely campus protests, was used, but the generalizability of the conclusions needs to be verified across different events. Secondly, the anchor word selected in the filling-in-the-blank method (Method 2) inevitably involves a certain degree of subjectivity. Also, version 3.5 of ChatGPT is used, and the performance of more advanced models may differ.

Future research can expand the experimental framework of this study on a broader range of event types, more diverse model versions, and more refined syntactic analysis tools.

References

- [1] Entman, R. M. (1993). Framing: Toward clarification of a fractured paradigm. *Journal of Communication*, 43(4), 51–58.
- [2] Hamborg, F. (2020). Media bias, the social sciences, and NLP: Automating frame analyses to identify bias by word choice and labeling. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop* (pp. 79–87). Association for Computational Linguistics.
- [3] Spinde, T., Rudnitskaia, L., Mitrović, J., Hamborg, F., Granitzer, M., Gipp, B., & Donnay, K. (2021). Automated identification of bias inducing words in news articles using linguistic and context-oriented features. *Information Processing & Management*, 58(3), 102505.
- [4] Yeh, K.-C., Chi, J.-A., Lian, D.-C., & Hsieh, S.-K. (2023). Evaluating interfaced LLM bias. In *Proceedings of the 35th Conference on Computational Linguistics and Speech Processing (ROCLING 2023)* (pp. 79–87). Association for Computational Linguistics and Chinese Language Processing.
- [5] Oketunji, A. F., Anas, M., & Saina, D. (2023). Large Language Model (LLM) bias index -- LLMBI (arXiv:2312.14769). arXiv.
- [6] Zhou, H., Feng, Z., Zhu, Z., Qian, J., & Mao, K. (2024). UniBias: Unveiling and mitigating LLM bias through internal attention and FFN manipulation. In *Advances in Neural Information Processing Systems (NeurIPS 2024)*.
- [7] Ad Fontes Media. (2024). Media bias chart. *Ad Fontes Media*. <https://adfontesmedia.com/interactive-media-bias-chart/>
- [8] Wilson, T., Wiebe, J., & Hoffmann, P. (2005). Recognizing contextual polarity in phrase-level sentiment analysis. In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing (HLT/EMNLP 2005)* (pp. 347–354).
- [9] Pennington, J., Socher, R., & Manning, C. D. (2014). GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (pp. 1532–1543).
- [10] Ouyang, L., Wu, J., Jiang, X., et al. (2022). Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35, 27730–27744.
- [11] Fudan Development Institute. (2021, March 18). The "News Frontline" of American political polarization: The past and present of Fox News.