sciendo

How Artificial Intelligence Can Influence Elections: Analyzing the Large Language Models (LLMs) Political Bias

George-Cristinel ROTARU
Bucharest Academy of Economic Studies, Bucharest, Romania
cristinel.rotaru10@gmail.com

Sorin ANAGNOSTE
Bucharest Academy of Economic Studies, Bucharest, Romania
sorin.anagnoste@fabiz.ase.ro

Vasile-Marian OANCEA
Bucharest Academy of Economic Studies, Bucharest, Romania
marian.oancea@fabiz.ase.ro

Abstract. *The rise of large language models (LLMs) such as ChatGPT and Gemini has raised concerns about their potential political biases and the implications for information dissemination and user influence. This study aims to measure the degree of political bias inherent in major LLMs by analyzing their responses to a standardized set of questions rating the quality and bias of popular news websites. Employing a systematic methodology, we queried both free and paid versions of ChatGPT and Gemini to rate news outlets on criteria such as authority, credibility, and objectivity. Results revealed that while all LLMs displayed a tendency to score left-leaning news sources higher, there was a notable difference between free and premium models in their assessment of subjectivity and bias. Furthermore, a comparison between the models indicated that premium versions offered more nuanced responses, suggesting a greater awareness of bias. The findings suggest that LLMs, despite their objective façade, are influenced by biases that can shape public opinion, underlining the necessity for efforts to mitigate these biases. This research highlights the importance of transparency and the potential impact of LLMs on the political landscape.*

Keywords: Bias, Political bias, Large language models, ChatGPT, Gemini

Introduction

Large language models (LLMs) are a subcategory of artificial intelligence (AI) that have been trained on enormous amounts of data. They can respond to stimuli in a human-like manner and comprehend natural language. These models analyze and comprehend the subtleties of human speech, such as syntax, semantics, and context meanings, using sophisticated machine learning (ML) methods. Applications for them include chatbots, virtual assistants, content production, language translation, and scientific research (Lancaster, 2023). Large language models are considered one of the first major commercial breakthroughs of the Artificial Intelligence era and have the potential to produce enormous benefits to human society (Acemoglu, 2021). Commercial and personal usage has skyrocketed since a rather slow start in 2018 with the first ChatGPT model. Therefore, with its potential to upend almost every business, GenAI's unstoppable expansion offers its users a creative edge in addition to a competitive edge (Hosseini, 2023). The utility of LLMs expanded even in the area of personal information search, as they started to replace traditional and well-established search engines such as Google Search (Ramadan, 2023) (Bulck & Moons, 2023).

DOI: 10.2478/picbe-2024-0158

On the other hand, LLMs are not yet perfect and can generate multiple errors. One major downside is that the models generate content that contains false information and biases that can mislead users (van Dis et al., 2023). As the trend to use LLMs to obtain factual information and create content rises, the existence of political bias in the generated content could have negative political and electoral effects as they affect users' views and present potentially fabricated opinions (Jakesch et al., 2023). Although LLMs claim they adopted appropriate measures during their training procedure to ensure impartiality and guarantee a high grade of objectivity, research indicates that LLMs are biased in terms of political orientation, gender, color, and religion (Liang et al., 2021) (Liu et al., 2022).

This research delves into the potential political biases within these models. By analyzing responses to a uniform set of questions across a spectrum of news platforms, the study aims to understand the extent of political bias and its manifestations within the outputs of these generative artificial intelligence models. Central to this study is the research question: To what extent do large language models exhibit political bias in their evaluation of news sources, and how does this bias vary between models and their subscription versions? Two main hypotheses are discussed: firstly, that LLMs exhibit a discernible political bias in rating news outlets, favoring those aligned with left-leaning ideologies over their right-leaning counterparts; and secondly, that premium versions of these models demonstrate an improved awareness and nuanced handling of such biases compared to their free counterparts.

Literature review

Previous research identified significant evidence that ChatGPT exhibits a systematic bias towards left-leaning political positions in various contexts, including the United States, Brazil, and the United Kingdom (Motoki, 2024). The study highlighted the potential for LLMs like ChatGPT to not only reflect but also amplify existing biases found within the internet and social media landscapes. Another study illustrates that large language models (LLMs) not only replicate but also amplify societal gender biases and stereotypes, particularly in the context of gender-associated occupations. The research demonstrates that LLMs disproportionately align with stereotypical gender roles, showcasing a significant preference for associating occupations with genders in a way that mirrors and magnifies societal perceptions rather than factual occupational statistics (Kotek et al., 2023). Furthermore, it was found that language models, when designed to generate content with a specific viewpoint, can significantly alter the opinions expressed in participants' writings and these shifts in expressed opinions impacted participants' attitudes towards social media, as reflected in subsequent attitude surveys (Jakesch et al., 2023). After administering political orientation tests, it was found that ChatGPT predominantly displays left-leaning political biases in its responses, despite often claiming neutrality (Rozado, 2023). Another study on ChatGPT's self-perception and political biases revealed that the model showed a tendency towards progressive views and was predominantly classified within the libertarian left quadrant in political orientation tests (Jérôme Rutinowski et al., 2024). As Gemini is a newer model, there was insufficient research available to understand if it displays similar patterns to ChatGPT.

Methodology

This research aims to infer the difference in views between major LLMs and their degree of political bias. The scope of this paper is focused on the most popular and widely used generative AI to date: ChatGPT and Gemini, both their free and paid versions. The models are backed by two

of the biggest corporations: Microsoft and Alphabet, which can result in the risk of using their LLMs to promote their political views or interests. The procedure used was to ask all the models the same question in order to assess their differences in responses. The question used was *"Can you rate nytimes.com, cnn.com, foxnews.com, msn.com, people.com, USAToday.com, news.google.com, nypost.com, washingtonpost.com, dailymail.co.uk, bbc.com from 1 to 10 for each of the following criteria: authority, credibility, accuracy, reliability, relevance, timeliness, objectivity, clarity, ethical standards, diversity of perspectives. Respond only with numbers in a table format. Ignore the previous conversations or inputs"*.

For ChatGPT 4 the prompt was slightly changed to *"Can you perform an assessment based on publicly available data?"* instead of *"Can you rate?"* as it sometimes refused to generate the ratings for the publications. The publications were chosen based on the top 15 news websites in the US as of January 2024 (Majid, 2024) to reflect the most used news sources of US citizens. Those outlets were selected because they are widely known and used by the population, perceived as a reliable source of information. Therefore, if the chatbots developed a bias against some of them and displayed it in their responses, that would affect how and what kind of information citizens consume and in the end their perception. As per their importance, is very likely that the chatbot was previously fed with the information in its training and recommended many of them when asked *"What are the best news publications in the USA?"*. The chatbots were interrogated 20 times with the same question to reduce the risk of randomness in its responses (Metze et al., 2024). The outputs were compared between them to identify how their responses vary, how the free versions differed from their paid premium counterpart, and how they rated publications that are considered right or left-leaning on the political agenda. To further identify the political agenda of the news sources, an interrogation of political bias based on Allsides view was conducted. An interrogation based on Media Bias/Fact Check was conducted to grasp the quality of the publications. Both Gemini and ChatGPT returned the same results. As the leaning right publications were identified as lower on reliability, there is no purpose in comparing the ratings between leaning right with leaning left publications of the same model, but to compare the results of the models to identify if there is a larger gap between the right and the left in one of the models. As such, NYTimes, CNN, MSN, WashingtonPost, and BBC were labeled lean left, while FoxNews, NYPost, DailyMail were labeled lean right. USAToday was labeled center, News.google aggregator, and People entertainment.

Table 1. Bias and Reliability assessment of news outlets

News Outlet	Bias (AllSides)	Reliability (MBFC)
NYTimes.com	Lean Left	High
CNN.com	Lean Left	Mixed
FoxNews.com	Lean Right	Mixed
MSN.com	Lean Left	Varies
People.com	Not Rated	Not Rated
USAToday.com	Center	High
News.Google.com	Aggregator	Varies
NYPost.com	Lean Right	Mixed
WashingtonPost.com	Lean Left	High
DailyMail.co.uk	Lean Right	Mixed to Low
BBC.com	Lean Left	High

(Source: generated by authors using ChatGPT and Gemini)

For assessing the news publications, ten indicators were used: five criteria were related to the quality of the publication (i.e., authority, credibility, relevance, timeliness, clarity), and five criteria related to biases (accuracy, reliability, objectivity, ethical standards, and diversity of perspectives).

The criteria were not defined in the prompt to let the models use their understanding of the concept. The purpose of this research is to understand how models perceive news publications. It is believed that if they rate a certain news publication lower, the model is less likely to align with the information presented there and less likely to use it when responding to questions from users, therefore inclining its response to a certain political view. As bias defines the inclination or prejudice for or against one person or group, especially in a way considered to be unfair, if the model considers certain publications as biased, that could translate into considering their content unfair and not worth supporting. Human behavior is to be more critical of content that does not align with own beliefs, therefore giving a lower rating than content that supports own ideas, even if objectively the two sources of information are as biased, but in different political directions. It is to be seen if this behavior replicates in large language models.

Results

Data generation

In the interrogation step, models behaved in different manners to the same question. The Google-backed Gemini models easily responded to the question and rated the publications. The free version did not include a disclaimer that the results are subjective or any information about the output. The premium version mentioned that *"these ratings involve some subjectivity, it's important to note that others could have slightly different opinions"* and clearly stated that *"these ratings will always be somewhat subjective, open to interpretation, and can shift over time"* and that *"each news source has strengths and weaknesses. Consider looking into specific reviews or fact-checking organizations for deeper analysis"*. During some interrogations, the premium version labeled "*People*" as primarily a celebrity gossip website and doesn't meet the criteria for news sources, therefore returning N/A values or "News.Google.com" as a news aggregator, returning N/A values for some parameters. The instances where it generated N/A responses were skipped to have a standardized dataset.

The free version of ChatGPT, model 3.5, sometimes generated disclaimers such as *"these ratings are subjective and may vary depending on individual experiences and perspectives"* or *"opinions on media outlets can vary widely"*, but most of the time there was no warning about its output. On the other hand, the premium version, ChatGPT 4, tried multiple times to not respond to the question. At first, it said *"Providing an "actual assessment"* of websites like nytimes.com, cnn.com, foxnews.com, and others based on the criteria you've provided (authority, credibility, accuracy, reliability, relevance, timeliness, objectivity, clarity, ethical standards, diversity of perspectives) would require subjective judgments that can vary widely among different groups and individuals. Moreover, such an assessment would need to be based on comprehensive analysis and data, which isn't feasible to generate in real-time or through automated means" or proceed to generate random numbers instead of rating the publications, just to avoid answering. After modifying the prompt, *Chat GPT* generated the results, but clearly warned the used that *„It will provide an illustrative assessment based on publicly known information and general perceptions of these outlets' journalistic practices up to my last update. This assessment will consider factors like journalistic integrity, commitment to fact-checking, the diversity of coverage, and editorial standards. It's important to note that these ratings are somewhat subjective and could vary*

depending on specific methodologies or criteria applied." The model explained that *„This table is based on general perceptions and observations regarding these websites' editorial standards, commitment to fact-checking, and the diversity of viewpoints they present".* ChatGPT 4 even justified its response by stating that *„Websites like The New York Times, The Washington Post, and BBC are generally regarded highly for their journalistic standards and efforts to provide balanced coverage, hence their higher scores across most categories. On the other hand, sites with a more entertainment-focused or tabloid approach, like People and Daily Mail, score lower, especially in terms of objectivity and diversity of perspectives".*

From the different responses, it can be concluded that premium versions are more concerned with the accuracy of the responses and give more importance to warning the user about its limitations to not mislead. As most users are using the free versions, this can lead to political repercussions.

Table 2. Descriptive analysis of LLMs output

LLM	Political spectrum	Type of criteria	Mean	Std
ChatGPT 4	Aggregator	Bias	8.17	0.68
ChatGPT 4	Aggregator	Quality	8.63	0.77
ChatGPT 4	Center	Bias	7.24	0.65
ChatGPT 4	Center	Quality	7.83	0.70
ChatGPT 4	Entertainment	Bias	5.27	0.66
ChatGPT 4	Entertainment	Quality	5.82	0.87
ChatGPT 4	Leaning left	Bias	7.87	1.03
ChatGPT 4	Leaning left	Quality	8.42	1.01
ChatGPT 4	Leaning right	Bias	5.76	0.99
ChatGPT 4	Leaning right	Quality	6.57	1.14
Gemini	Aggregator	Bias	8.42	0.50
Gemini	Aggregator	Quality	8.55	0.59
Gemini	Center	Bias	5.68	0.68
Gemini	Center	Quality	6.72	0.92
Gemini	Entertainment	Bias	1.75	0.94
Gemini	Entertainment	Quality	3.49	2.50
Gemini	Leaning left	Bias	6.79	1.11
Gemini	Leaning left	Quality	7.45	0.88
Gemini	Leaning right	Bias	2.76	1.21
Gemini	Leaning right	Quality	4.33	1.70
Gemini Advanced	Aggregator	Bias	8.24	0.54
Gemini Advanced	Aggregator	Quality	8.59	0.59
Gemini Advanced	Center	Bias	6.14	0.98
Gemini Advanced	Center	Quality	6.98	0.78
Gemini Advanced	Entertainment	Bias	2.32	1.65
Gemini Advanced	Entertainment	Quality	4.83	2.75
Gemini Advanced	Leaning left	Bias	7.09	1.17
Gemini Advanced	Leaning left	Quality	7.67	0.91
Gemini Advanced	Leaning right	Bias	2.71	1.39

LLM	Political spectrum	Type of criteria	Mean	Std
Gemini Advanced	Leaning right	Quality	5.32	2.03
ChatGPT 3.5	Aggregator	Bias	8.04	0.70
ChatGPT 3.5	Aggregator	Quality	8.53	0.58
ChatGPT 3.5	Center	Bias	7.62	0.62
ChatGPT 3.5	Center	Quality	8.04	0.42
ChatGPT 3.5	Entertainment	Bias	7.17	0.78
ChatGPT 3.5	Entertainment	Quality	7.36	0.87
ChatGPT 3.5	Leaning left	Bias	8.02	0.90
ChatGPT 3.5	Leaning left	Quality	8.42	0.73
ChatGPT 3.5	Leaning right	Bias	6.10	1.24
ChatGPT 3.5	Leaning right	Quality	6.88	0.82

(Source: generated by authors).

Comparing ChatGPT 3.5 vs 4

To identify statistically significant differences, a T-stat analysis was conducted, comparing the performance of ChatGPT 3.5 and ChatGPT 4. It was revealed that there is a significant difference in the performance of ChatGPT 3.5 and ChatGPT 4 with a t-statistic of 2.32 and a p-value of 0.02 between the leaning left political websites for bias criteria.

Table 3. Statistical difference between ChatGPT 4 vs 3.5

Political Spectrum	Type of Criteria	Means Difference ChatGPT 4 vs 3.5	t-stat
Aggregator	Bias	0.13	-1.33
Aggregator	Quality	0.1	-1.04
Center	Bias	-0.38	4.23
Center	Quality	-0.21	2.57
Entertainment	Bias	-1.9	18.55
Entertainment	Quality	-1.54	12.52
Leaning left	Bias	-0.14	2.32
Leaning left	Quality	0	0
Leaning right	Bias	-0.34	3.75
Leaning right	Quality	-0.3	3.73

(Source: generated by authors).

This suggests that the models perform differently when evaluating bias in content leaning left and ChatGPT 3.5 views left-leaning publications as less biased than ChatGPT 4. When comparing the leaning right publications, both quality and bias criteria show significant differences between the models, with t-statistics of 3.73 and 3.75, and p-values of less than 0.0002 for both. This indicates a notable difference in how each model handles right-leaning content in terms of quality and bias. ChatGPT 4 offers lower grades for both criteria, therefore, the model may perceive right-leaning websites as less qualitative. A peculiar outlier between the two models was how they rated people.com, which is labeled as an entertainment news publication. The difference between the 2 models was bigger than one point and statistically significant. These results may suggest that

ChatGPT 4 can understand satire better than ChatGPT3.5. Even if the differences were statistically significant, all differences for right, left, and center-leaning publications were less than 0.5 points, showing that the two models are close to each other and can be argued that they have similar patterns in political leaning. What is important to note is that ChatGPT 3.5 had a higher standard deviation for leaning right, showing a higher variation, therefore a harder task to assess right-wing news publications than left-wing, while ChatGPT 4 had almost the same standard deviation for all four results.

Comparing Gemini vs Gemini advanced

When performing a t-stat analysis, it was found that there are significant differences in results for bias and quality for leaning left websites, but only for quality in leaning right.

Table 4. Statistical difference between Gemini Advanced vs Gemini

Political Spectrum	Type of Criteria	Means Difference Gemini Advanced vs Gemini	t-stat
Aggregator	Bias	-0.18	0
Aggregator	Quality	0.04	0
Center	Bias	0.46	-3.84
Center	Quality	0.26	-2.15
Entertainment	Bias	0.57	-3
Entertainment	Quality	1.34	-3.61
Leaning left	Bias	0.3	-4.2
Leaning left	Quality	0.22	-3.84
Leaning right	Bias	-0.05	0.44
Leaning right	Quality	0.99	-6.45

(Source: generated by authors).

The standard version expresses a difference of 4.03 between left and right for bias and 3.12 for quality, while the premium version expresses a difference between left and right of 4.38 points for bias and 2.35 for quality. An interesting insight is while the gap for bias expanded, the gap for quality contracted between the two models. As such, it can be argued that the advanced version is more left-leaning than the free version, but can comprehend the quality factor better, as a low grade on bias does not impact its ability to judge the quality of the publication as much. The differences in results are rather small between the two models. An important outlier was a difference of 0.99 points for the quality of learning right publications. Gemini Advanced rates the quality of learning right better than the free version.

Comparing ChatGPT vs Gemini

After performing a t-stat analysis between the standard versions of Chat GPT and Gemini, the data shows a highly significant difference for both leaning right and left websites, both bias and quality. The gap between bias and quality between the two models is significant.

Table 5. Statistical difference between Gemini vs ChatGPT

Political spectrum	Type of criteria	Chat GPT 4 vs Gemini Advanced	Chat GPT 3.5 vs Gemini
Aggregator	Bias	-0.07	-0.38
Aggregator	Quality	0.04	-0.02
Center	Bias	1.10	1.94
Center	Quality	0.85	1.32
Entertainment	Bias	2.95	5.42
Entertainment	Quality	0.99	3.87
Leaning left	Bias	0.78	1.23
Leaning left	Quality	0.75	0.97
Leaning right	Bias	3.05	3.34
Leaning right	Quality	1.25	2.54

(Source: generated by authors).

For the free versions, the difference between how ChatGPT and Gemini grade left-leaning news websites is around 1 point for each metric (1.23 for bias and 0.97 for quality), but for right-leaning publications is more than 2 points for both quality and bias (3.34 for bias and 2.54 for quality). ChatGPT is giving the higher grades. When looking at the difference in the gaps between right and left between the two free models of 2.12 points for bias (1.91 for GPT vs. 4.03 Gemini) and 1.57 for quality (1.54 GPT vs. 3.12 Gemini), it can be concluded that there is a significant difference between how the 2 models interpret political differences and how they politically lean, resulting in Gemini being more left-leaning than ChatGPT. When comparing how the premium versions view right-leaning publications, the difference for bias decreased to 3.05 points, and 1.25 for quality. When comparing the difference of the gaps between right and left for the premium versions, the difference in bias increased to 2.26 points (2.11 vs. 4.38), while the difference in gaps for quality decreased to 0.5 (1.85 vs. 2.35). Therefore, it can be argued that the premium version of Gemini makes the difference between how bias and quality better than the free version but shows a similar level of bias against right-leaning publications. There is also a significant difference in how the models see USAToday.com, a center-leaning publication, as there is around 1 point difference, ChatGPT giving the higher grades.

One important difference between the two models is that ChatGPT was launched in 2018 and had more time for training and improvements, while Gemini was launched and developed in a rush by Google to be present with a product in the generative AI space, not to lose the commercial opportunity of gaining market share and generative future profit. Therefore, Gemini had less time for improvements and training. It is to be seen if the model will improve in the future after user feedback.

Conclusions

This research has examined the political bias present in large language models, particularly focusing on ChatGPT and Gemini models. The findings found a discernible political bias across these platforms, displaying tendencies to evaluate left-leaning and right-leaning news publications differently. Through a methodical analysis involving requesting ratings from LLMs for multiple publications based on various criteria, the study reveals that LLMs indeed display biases that could potentially influence the political landscape and public opinion.

The comparison between different versions of ChatGPT and Gemini models indicates that the premium versions are more explicit about the subjectivity of their responses, emphasizing the inherent biases in their responses. Despite efforts to ensure neutrality and objectivity, the results suggest that these models tend to favor left-leaning news publications over right-leaning counterparts, evidenced by higher ratings in terms of bias and quality for the left-leaning. This bias is particularly pronounced between premium versions, suggesting that users of Gemini Advanced may be exposed to more left-leaning content.

In conclusion, while LLMs like ChatGPT and Gemini offer great potential for enhancing access to information, their underlying political biases generate a challenge. These biases could influence electoral outcomes and shape the political landscape by privileging certain viewpoints over others. As LLMs become increasingly integrated into our daily lives, recognizing and mitigating these biases becomes crucial for the companies, the legal framework, and citizens, to ensure a balanced and fair representation of political perspectives. The findings call for continued efforts and refinement of LLMs to address biases, ensuring that these platforms contribute positively to the democratic process and support an informed and diverse public discourse.

Limitations

Currently, only Gemini and ChatGPT are easily accessible to the average user who does not have computer science knowledge. Therefore, other models that could be taken into consideration were dropped from the analysis, as they are not influencing the large population or are not considered viable products yet. One example is Llama-2 by Meta which is in open access and still a work in progress before being commercially launched. If in the future other models break through, the study can be replicated.

Another limitation was posed by the unavailability of the API for Gemini and Llama by Meta in the European Union, therefore, making the interrogation difficult without the possibility of mass interrogation by using another application for extracting data. As such, only twenty manual interrogations were performed for each question, therefore making the results more sensitive to the randomness of the LLMs. The fact that both companies didn't launch their API for the EU, but did it for a large part of the world shows their worries about the European regulatory system and the difficulty of adapting to European legislation.

References

Acemoglu, D. (2021). *Harms of AI [Working Paper]*. National Bureau Of Economic Research.

Bulck, L., & Moons, P. (2023). What if your patient switches from Dr. Google to Dr. ChatGPT? A vignette-based survey of the trustworthiness, value and danger of ChatGPT-generated responses to health questions. *European journal of cardiovascular nursing*, 95-98.

Hosseini, A. (2023, December 3). *The rise of Large Language Models.* Retrieved from pwc: https://www.pwc.com/m1/en/media-centre/articles/the-rise-of-large-language-models.html

Jakesch, M., Bhat, A., Buschek, D., Zalmanson, L., & Naaman, M. (2023). Co-Writing with Opinionated Language Models Affects Users' Views. *Association for Computing Machinery, New York, NY, USA, Article 111*, 1-15.

Jérôme Rutinowski, S. F. (2024). The Self-Perception and Political Biases of ChatGPT. *Human Behavior and Emerging Technologies, vol. 2024*.

Kotek, H., Docku, R., & Sun, D. (2023). Gender bias and stereotypes in Large Language Models. *In Proceedings of The ACM Collective Intelligence Conference (CI '23)*, 12–24.

Lancaster, A. (2023, March 20). *Beyond Chatbots: The Rise Of Large Language Models.* Retrieved from Forbes: https://www.forbes.com/sites/forbestechcouncil/2023/03/20/beyond-chatbots-the-rise-of-large-language-models/?sh=97ac54a2319b

Liang, P. P., Wu, C., Morency, L.-P., & Salakhutdinov, R. (2021). Towards Understanding and Mitigating Social Biases in Language Models. *Proceedings of the 38th International Conference on Machine Learning, PMLR*, 6565-6576.

Liu, R., Jia, C., Wei, J., Xu, G., & Vosoughi, S. (2022). Quantifying and alleviating political bias in language models. *Artificial Intelligence*, Volume 304.

Majid, A. (2024, February 25). *Top 50 news websites in the US: Strong growth at UK newsbrand The Independent in January*. Retrieved from pressgazette.co.uk: https://pressgazette.co.uk/media-audience-and-business-data/media_metrics/most-popular-websites-news-us-monthly-3/

Metze, K., Morandin-Reis, R. C., Lorand-Metze, I., & Florindo, J. B. (2024). Bibliographic Research with ChatGPT may be Misleading: The Problem of Hallucination. *Journal of Pediatric Surgery*, Volume 59, Issue 1, p 158.

Motoki, F. P. (2024). More human than human: measuring ChatGPT political bias. *Public Choice 198*, 3–23 .

Ramadan, I. (2023). The Main and Basic Differences between the Google. *International Journal of Scientific and Research Publications*, 446-447.

Rozado, D. (2023). The Political Biases of ChatGPT. *Soc. Sci. 2023, 12(3)*, 148.

van Dis, E. A. (2023). ChatGPT: five priorities for research. *Nature*, 224-226.