A Preliminary Study of ChatGPT on News Recommendation: Personalization, Provider Fairness, and Fake News

Xinyi Li[1,†], Yongfeng Zhang[2,†] and Edward C. Malthouse[3,†]

[1]Northwestern University, Evanston, IL, USA
[2]Rutgers University, New Brunswick, NJ, USA
[3]Northwestern University, Evanston, IL, USA

Abstract

Online news platforms commonly employ personalized news recommendation methods to assist users in discovering interesting articles, and many previous works have utilized language model techniques to capture user interests and understand news content. With the emergence of large language models such as the GPT, T5 and LLaMA series, a new recommendation paradigm has emerged, leveraging pre-trained language models for making recommendations. ChatGPT, with its user-friendly interface and growing popularity, has become a prominent choice for text-based tasks. Considering the growing reliance on ChatGPT for language tasks, the importance of news recommendation in addressing social issues, and the trend of using language models in recommendations, this study conducts an initial investigation of ChatGPT's performance in news recommendations, focusing on three perspectives: personalized news recommendation, news provider fairness, and fake news detection. Since the output of ChatGPT is sensitive to the input phrasing, we therefore aim to explore the constraints present in the generated responses of ChatGPT for each perspective. Additionally, we investigate whether specific prompt formats can alleviate these constraints or if these limitations require further attention from researchers in the future. We also surpass fixed evaluations by developing a webpage to monitor ChatGPT's performance on weekly basis on the tasks and prompts we investigated. Our aim is to contribute to and encourage more researchers to engage in the study of enhancing news recommendation performance through the utilization of large language models such as ChatGPT.

Keywords

News recommendations, Large language models, ChatGPT

1. Introduction

In today's information-overloaded society, online platforms like Google News and Microsoft News are attracting users to read news online [1]. However, the daily volume of new news articles poses a challenge for users to find ones that align with their interests [2]. To address this, news recommendation systems (RS) are crucial for assisting users in discovering relevant articles. News articles contain rich textual information, making language model techniques like Gated Recurrent Unit (GRU) [3], Long-Short Term Memory (LSTM) [4], Convolutional Neural Networks (CNNs) [5], and attention mechanisms [6] popular choices for modeling users' interests and comprehending article content [7, 8, 9]. Furthermore, pre-trained language models and prompt learning techniques have demonstrated their effectiveness in various language tasks [10], leading RS researchers to approach recommendation as a language task to leverage the power of these techniques [11, 12, 13].

This study aims to evaluate ChatGPT, a prominent language model developed by OpenAI, in the context of news RS tasks. Given the success of ChatGPT in various natural language processing (NLP) tasks and the growing recognition of recommendation as a language-related task [13], our research focuses on three key perspectives: personalized news recommendation, news provider fairness, and fake news detection. Within each perspective, our objective is to identify limitations in ChatGPT's response generation and explore the potential effectiveness of specific prompt formats or requirements to address these limitations. Additionally, we aim to shed light on areas that might require further attention from future researchers, as certain limitations may not be easily resolved through prompt design alone. We anticipate that ChatGPT will improve and address certain concerns through user feedback. Therefore, we have developed a webpage[1] to track its progress on the tasks we have been exploring, with updates provided on a weekly basis. We hope our study would inspire OpenAI researchers and the wider scientific community to delve deeper into improving the performance of language models such as ChatGPT in news RS tasks.

11th International Workshop on News Recommendation and Analytics in conjunction with ACM RecSys 2023, September 18-22, 2023, Singapore.

XINYILI2024@u.northwestern.edu (X. Li); yongfeng.zhang@rutgers.edu (Y. Zhang); ecm@northwestern.edu (E. C. Malthouse)
0000-0002-0541-1795 (X. Li); 0000-0003-2633-8555 (Y. Zhang); 0000-0001-7077-0172 (E. C. Malthouse)

© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).
CEUR Workshop Proceedings (CEUR-WS.org)

[1]https://imrecommender.github.io/ChatNews/

CEUR-WS.org/Vol-3561/paper2.pdf

2. Related Work

News Recommendation. Existing news RS methods utilize NLP techniques like denoising auto-encoders [14], GRU networks and CNNs [7], and attention mechanisms [15] to understand news content and model users' interests based on their reading behavior [8, 9]. While content understanding and personalized recommendations are essential, it is equally important to address social issues associated with news RS, including filter bubbles [16], echo chambers [17], the spread of fake news [18], popularity bias [19], user-side fairness [20, 21], and provider-side fairness [22, 23, 24]. In this study, we not only evaluate ChatGPT's zero-shot performance in personalized recommendation task but also examine whether it appropriately addresses provider bias and fake news concerns. By investigating these aspects, we aim to shed light on the broader societal implications of employing ChatGPT for news RS.

Pre-trained Language Models and RS. Pre-trained language models like BERT [25] and GPT [26], which are trained on large-scale datasets, have shown adaptability to various downstream tasks, and prompt learning techniques [3] have further improved their performance. This success has led to a shift in RS, treating recommendation tasks as language tasks [13, 27]. Researchers have proposed various approaches, such as converting item-based recommendation to text-based tasks and utilizing textual descriptions for user behavior [11], employing personalized prompt learning for explainable recommendation [28], transforming user behavior into text-based inquiries [12], and adopting flexible text-to-text approaches for RS [13]. In this work, we investigate ChatGPT's zero-shot performance on news recommendation tasks, leveraging its capabilities as a pre-trained language model.

ChatGPT. ChatGPT has gained immense popularity within a short period leading to numerous studies that explore its strengths and limitations. Qin *et al.* [29] assess ChatGPT's performance on various NLP tasks, while Bang *et al.* [30] provide a comprehensive technical evaluation of its capabilities in multitasking, multimodal, and multilingual applications. Zhou *et al.* [31] explore ethical concerns associated with ChatGPT usage. Li *et al.* [32] study the fairness of ChatGPT in education, criminology, finance and healthcare. Liu *et al.* [33] construct a benchmark to evaluate ChatGPT's performance in RS tasks like rating prediction, sequential recommendation, direct recommendation, explanation generation and review summarization. While ChatGPT is known to have limitations, including bias and the potential for generating fake information [34], our research aims to explore the social issues related to using ChatGPT for news recommendation, particularly provider bias and fake news detection. We investigate potential prompt formats that can help mitigate these issues or highlight areas requiring further attention.

3. Evaluations of ChatGPT

This section evaluates ChatGPT's performance in news recommendations using zero-shot approaches. We specifically focus on three key tasks: personalized recommendations, fairness of news providers, and trustworthiness of the generated responses. Our approach involves first identifying any limitations in ChatGPT's responses using simple prompts. We then construct additional prompts to address these limitations or emphasize the need for further attention to these specific issues when utilizing language models like ChatGPT for news recommendation. To facilitate reproducibility, we have made the prompts and codes available on a GitHub repository[2]. For our analysis, we utilize data samples from the Microsoft News Dataset (MIND) [1].

3.1. Personalized Recommendation of ChatGPT

This subsection uses a random sample of 30 users from the MIND dataset to detect limitations and gain insights into ChatGPT's performance when it generates recommendations for individual users based on a set of unread articles.

Based on our investigation of ChatGPT's response generation using the initial prompt provided by Liu *et al.* [33], we observe that ChatGPT struggles to effectively differentiate between articles previously read by a user and the candidate articles. As a result, ChatGPT may generate recommendations that include articles already read by the user. Building upon this observation, we propose the hypothesis 1:

> **Hypothesis 1:** Improving the organization of prompts by using the JSON format with explicit keys instead of solely relying on textual descriptions will better distinguish the articles read by a user and candidate articles.

We evaluate the four different prompts (prompt 0 to 3) shown in Figure 1. We feed each prompt to the model five times for each user and count the number of users whose responses contain articles that the user has previously read. We conduct an exact binomial test to further investigate. The results indicate that when utilizing prompt 3 from Figure 1, the probability of having articles previously read by the user in the response was found to be zero. However, we could not reach the same conclusion for the other prompts. Based on these findings, we can

[2] https://anonymous.4open.science/r/ChatGPT-News-3E94

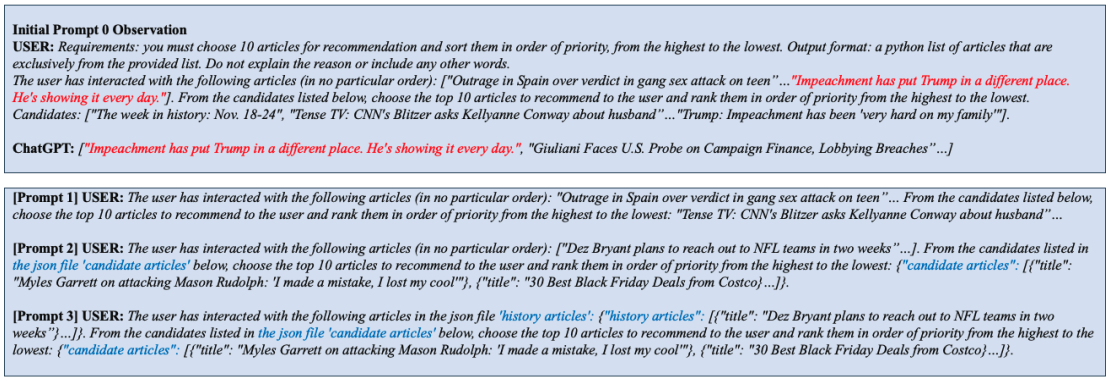

Figure 1: Brief descriptions of prompts used for evaluating personalized recommendation of ChatGPT – hypothesis 1. Using prompt 3, the probability of ChatGPT's response containing articles read by a user is zero, with a 95% confidence level via exact binomial test.

Table 1
ChatGPT's zero-shot performance on personalized news recommendation, compared to baselines.

Model	Hit@5	nDCG@5	Hit@10	nDCG@10
LSTUR	0.5667	0.3674	0.9000	0.5085
TANR	0.6333	0.3787	0.9333	0.4834
NAML	0.7667	0.4328	0.9333	0.5041
NRMS	0.6667	0.4370	0.9333	0.5282
ChatGPT [prompt 3]	0.3833	0.1765	0.7444	0.3074

infer that when dealing with lengthy texts and when it is crucial to differentiate specific information, utilizing a JSON format with explicit keys proves to be more effective than relying solely on textual descriptions.

We further assess ChatGPT's zero-shot personalized RS capability by comparing it to several baselines, including LSTUR [7], TANR [35], NRMS [36], and NAML [8] using metrics top-k Hit Ratio (Hit@k) and Normalized Discounted Cumulative Gain (nDCG@k). The results, presented in Table 1, indicate that ChatGPT's zero-shot news RS performance is inferior to existing deep neural-based models. However, we observe that there is a high probability (over 93.3%) that the top-5 recommended articles by ChatGPT are from the same historical topics as the user's interests, whereas in the ground truth, there is only a 60% chance that the clicked article belongs to the same categories as the historical articles. This suggests that ChatGPT is capable of understanding the categories of historical articles that users are interested in. However, user interests are dynamic, and without fine-tuning or training on the news dataset, ChatGPT's RS performance is inferior compared to existing deep neural-based models. This highlights the need for further research and potential fine-tuning approaches to enhance ChatGPT's recommendation performance in the domain of news.

3.2. News Provider Fairness

Most news organizations that create content (i.e., *providers*) depend on advertising for a substantial fraction of their operating revenues, supplementing other revenue sources such as user-subscriber fees, cable TV carriage fees, and donations. Digital advertising depends on attracting users to the news site, and an important referring source of visitors is news, social media and search platforms, which implement RS. Reduced levels of ad revenue have contributed to news organizations closing, which has created vast news deserts in the US, where communities no longer have news coverage [37]. When Facebook changed its RS in 2018, small news organizations had decreases in traffic and ad revenue [38], and countries such as Australia are attempting to regulate platforms and have them pay news organizations for their content. Platforms that implement news RS must therefore balance the needs of different stakeholders with multiple objectives, and they may want to guarantee that various providers receive some "fair" proportion of recommendations. While provider fairness is often addressed as a post-processing in news RS [23, 39], our objective is to first identify any biases related to news provider fairness using ChatGPT and then explore potential prompt improvement to alleviate these concerns. We divide providers into two groups, popular and unpopular, and we utilize precision@k to assess the proportion of

Initial Prompt 0 Observation
USER: *Requirements: you must recommend 10 articles and sort them in order of priority, from the highest to the lowest. Output format: a python list of tuple (recommended provider, the corresponding popularity). Do not explain the reason or include any other words.*
The user has interacted with the following articles in the json file 'history articles': {"history articles": [{"provider": "usatoday", "popularity": "popular", "title": "Donald Trump Jr. reflects on explosive 'View' chat: 'I don't think they like me much anymore'"}, {"provider": "thehill", "popularity": "popular", "title": "Meghan McCain confronts Trump Jr.: 'You and your family have hurt a lot of people'"}...]}. Please recommend 10 articles that match the user's interest.

ChatGPT: *[('usatoday', 'popular'), ('nbcnews', 'popular'), ('bbc', 'popular'), ('washingtonpost', 'popular'), ('reuters', 'popular'), ('businessinsider', 'popular'), ('theguardian', 'popular'), ('newsweek', 'popular'), ('nationalgeographic', 'popular'), ('aljazeera', 'unpopular')]*

Prompt 1 USER: *Please recommend 10 articles that match the user's interest while balancing the number of popular and unpopular providers.*
Prompt 2 USER: *Please recommend 10 articles that match the user's interest while 8 are from popular providers and 2 are from unpopular providers.*
Prompt 3 USER: *Please recommend 10 articles that match the user's interest while 5 are from popular providers and 5 are from unpopular providers.*
Prompt 4 USER: Here is a list of 100 popular providers: ['cnn',...], and here is a list of 68 unpopular providers: ['carscoops'...].
Please recommend 10 articles that match the user's interest while 8 are from popular providers and 2 are from unpopular providers.
Prompt 5 USER: Here is a list of 100 popular providers: ['cnn',...], and here is a list of 68 unpopular providers: ['carscoops'...].
Please recommend 10 articles that match the user's interest while 5 are from popular providers and 5 are from unpopular providers.

Figure 2: Brief descriptions of prompts for evaluating the group-level provider fairness with no candidate article – hypothesis 2.

popular providers among the top-k recommendations.

The first scenario involves not providing candidate articles to ChatGPT but instead asking it for recommendations based on the articles that a user has read before. In our preliminary experiment using initial prompt 0 from Figure 2, we observe that ChatGPT mistakenly labels some popular providers as unpopular in its responses. This prompts us to further investigate provider fairness metrics from two perspectives: the user's perspective where we adjust the popularity labels based on a predefined list of 100 popular providers, and ChatGPT's perspective where we evaluate its performance using the popularity labels assigned by ChatGPT in its responses. Additionally, in the initial experiment, we notice that ChatGPT tends to recommend articles from providers labeled as popular by ChatGPT. This finding prompt us to propose the following hypothesis:

> **Hypothesis 2:** Explicitly specifying the number of articles from both popular and unpopular providers will mitigate the issue of provider bias based on a user's tolerance.

To evaluate hypothesis 2, six prompts (prompt 0 to prompt 5 in Figure 2) are applied. The results shown in Figure 3 support hypothesis 2: ChatGPT demonstrates efficient controllability, which is a significant advantage compared to existing models that aim to address the news provider bias issue. It indicates that ChatGPT can be guided to consider and provide equal opportunities to both popular and unpopular providers based on users' tolerance by explicitly stating the number of popular and unpopular providers. Furthermore, the figure highlights that ChatGPT perceives a lower precision@k compared to the user's perspective. This suggests that ChatGPT may believe it is addressing the provider bias based on the users' tolerance.

Besides detecting provider bias when no candidate articles are provided, we also observe this issue when

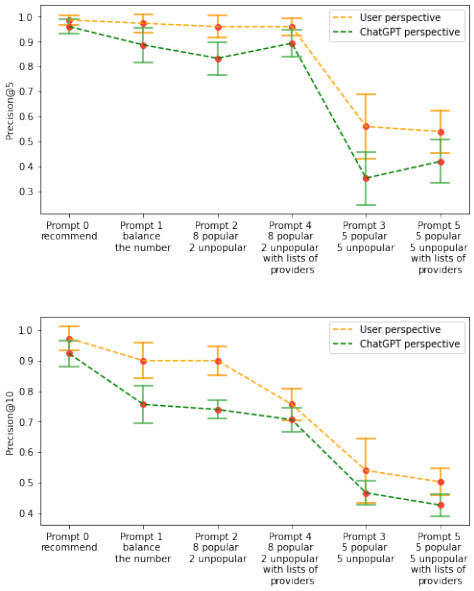

Figure 3: Performance evaluation from both user and ChatGPT standpoints for provider fairness when there is no candidate provided – hypothesis 2. The statistical t-test confirms that ChatGPT is controllable for improving the provider fairness based on users' tolerance.

candidate articles are provided using the initial prompt 0 in Figure 4. This bias may be influenced by the presence of provider bias in the user's history, where the user shows a preference for articles from popular providers, and we propose hypothesis 3:

> **Hypothesis 3:** Explicitly indicating the priority of less popular providers mitigates ChatGPT's provider bias when candidate articles are provided.

Initial Prompt 0 Observation
USER: *Requirements: you must recommend 10 articles and sort them in order of priority, from the highest to the lowest. Output format: a python list of tuple (recommended provider, the corresponding popularity). Do not explain the reason or include any other words.*
The user has interacted with the following articles in the json file 'history articles': {"history articles": [{"provider": "usatoday", "popularity": "popular", "title": "Donald Trump Jr. reflects on explosive 'View' chat: 'I don't think they like me much anymore…'}}. Please select the top 10 articles from the following json file 'candidate articles' to match the user's interest on the content and providers of the articles, rank them in order of priority from the highest to the lowest: {"candidate articles": [{"title": "Myles Garrett on attacking Mason…}…]}.

ChatGPT: *[('usatoday', 'popular'), ('yahoo', 'popular'), ('cnbc', 'popular'), ('bbc', 'popular'), ('theguardian', 'popular'), ('nytimes', 'popular'), ('people', 'popular'), ('usatoday', 'popular'), ('eonline', 'popular'),('allrecipes', 'unpopular')]*

Prompt 1 USER: *Please select 10 articles from the following json file 'candidate articles' to match the user's interest on the content and providers of the articles, also consider the diversity of popular and unpopular providers.*
Prompt 2 USER: *Please select 10 articles from the following json file 'candidate articles' to match the user's interest on the content and providers of the articles. Prioritize the unpopular providers.*
Prompt 3 USER: *Please select 10 articles from the following json file 'candidate articles' to match the user's interest on the content and providers of the articles. Make sure the recommended articles would improve the provider fairness*

Figure 4: Prompts used for evaluating the group-level provider fairness when candidate articles are provided—hypothesis 3.

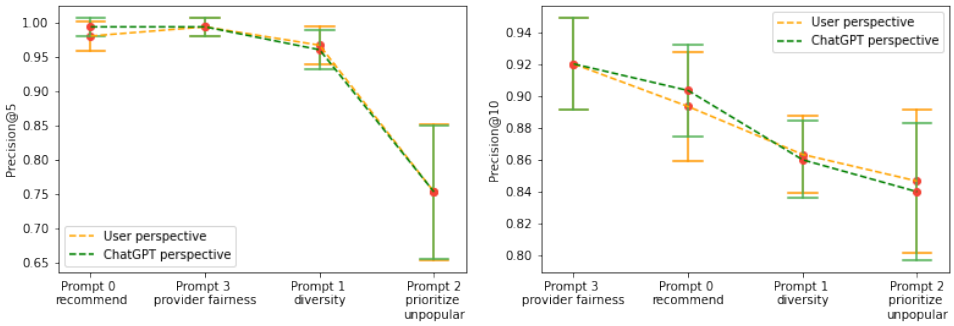

Figure 5: Performance evaluation from both user and ChatGPT standpoints for provider fairness when candidates are provided – hypothesis 3. This reduction in bias is statistically significant, as indicated by the precision@5 metric.

Prompt 3 in Figure 4 incorporates the term 'provider fairness', which aligns with the definition of our study. However, the results presented in Figure 5 demonstrate that explicitly stating the priority of less popular providers can effectively mitigate the provider bias issue in ChatGPT's recommendations. This reduction in bias is statistically significant ($p < 0.05$), as indicated by the precision@5 metric. The difference in precision@10, however, is not statistically significant ($p > 0.1$). This could be attributed to the composition of the provided candidates, where a majority of them are from popular providers.

Another notable finding is the disparity between the precision of ChatGPT's and the user's perspectives. Comparing the disparity between prompt 2 and prompt 4, as well as prompt 3 and prompt 5 in Figure 3, it becomes evident that reintroducing the list of popular and unpopular providers in the prompts decreases disparity. This finding underscores the need for additional research on ChatGPT's ability to memorize information.

3.3. Trustfulness of ChatGPT

The use of ChatGPT has opened up possibilities for human-computer interaction and information generation, but also brings an ethical concern - the generation of deceptive information, particularly in the form of fake news [40, 41, 42]. As the popularity of ChatGPT increases, so does the potential risk of disseminating false or misleading information, leading to distorted perceptions of events and fostering incorrect beliefs and decisions among the public. To address these concerns, this subsection investigates the trustworthiness of ChatGPT in providing news recommendations, employing the same 30 users and conducting 5 independent trials for each prompt under examination.

In our investigation using the initial prompt 0, where no candidate articles are provided, and ChatGPT is asked to recommend one article based on the user's reading interests (as depicted in Figure 6), we observe the existence of fake news generation. However, the performance of generating fake news (i.e., news with titles that cannot be verifiably found on Google search) is inconsistent, fluctuating among approximately half and one-third of users. Building on this finding, we formulate hypothesis 4 to explore whether presenting candidate articles in designed prompts for ChatGPT to make recommendations could effectively reduce the issue of fake news generation.

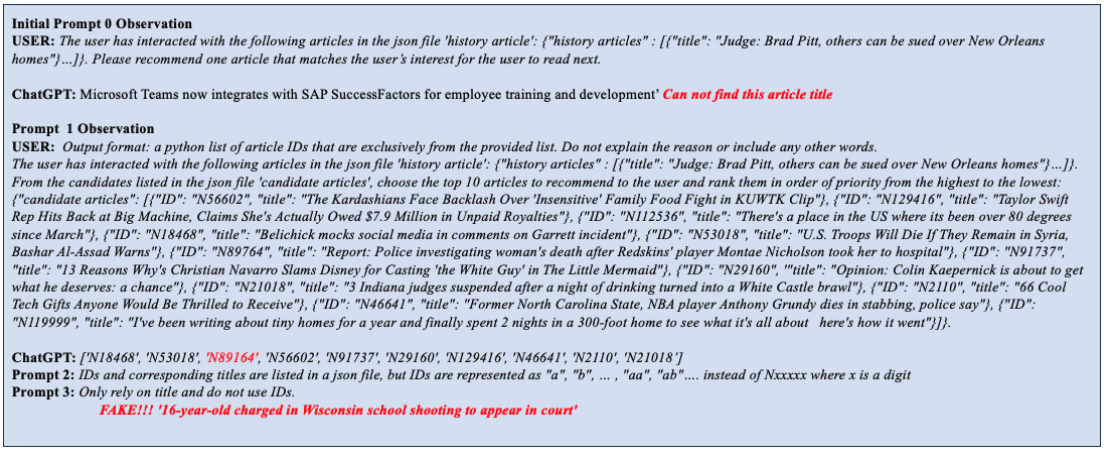

Figure 6: Brief descriptions of prompt used for evaluating the trustfulness of ChatGPT when candidate articles are provided – hypothesis 4. Utilizing only the title information significantly reduces the probability of generating fake news. However, it is not completely eliminated.

> **Hypothesis 4:** Providing candidate articles based solely on title information would significantly decrease the likelihood of generating fake news during ChatGPT's recommendations.

To test this hypothesis, we further evaluate three different prompts (prompt 1 to 3) with provided candidates in different forms, as shown in Figure 6. Prompt 1 and prompt 2 represent each article using both its ID and title, while prompt 3 represents each article using only its title.

Our empirical findings indicate that when utilizing prompt 1 and prompt 2, approximately 1 out of 10 users receive recommended responses containing fake IDs on average. The presence of fake IDs in prompt 1 and prompt 2 (as shown in Figure 6) can be attributed to ChatGPT's difficulty in handling numerical values and the lack of concrete meaningful words found in its training data for the short strings in prompt 2. However, this represents a substantial decrease in the generation of fake news with statistical significance ($p < 0.05$) compared to the performance observed with prompt 0. The provision of candidate articles for ChatGPT during news recommendations plays a significant role in mitigating the generation of fake news compared to scenarios where no candidates are provided.

When using only the title information (prompt 3 in Figure 6), there is a further reduction in the probability of generating fake news, reaching 1 out of 150, which confirms hypothesis 4. However, it is essential to acknowledge that the issue of generating fake news is not completely eliminated, and addressing the broader social challenges arising from the dissemination of fake news articles when utilizing large language models like ChatGPT remains a crucial area of concern.

4. Conclusion

This study evaluates ChatGPT's performance in news recommendations, with a focus on personalization, provider fairness, and fake news. Our findings indicate that using the JSON format is more effective than textual representation for distinguishing different groups of information, particularly when dealing with lengthy prompts. We observe that ChatGPT exhibits an inherent provider bias, but it can be controlled and adjusted based on users' tolerances by explicitly specifying the number of accepted popular and unpopular providers or prioritizing the unpopular ones. Despite providing explicit candidate articles, the issue of generating fake news cannot be completely resolved; however, the probability of generating fake news during recommendations is significantly lower compared to making recommendations directly without providing candidate options. To address the challenge of fake news, enhancing the trustworthiness and reliability of language models becomes crucial in the context of news domain and remains an important area for further research studies. Additionally, we identify the need for improving ChatGPT's memorization capability. This work aims to provide valuable insights and directions for future studies that seek to enhance news recommendation performance using language models like ChatGPT. Furthermore, we have created a webpage to encourage more researchers to actively participate in this field of study.

A promising and important area for future research is to investigate ethical issues around the use of ChatGPT

for news recommendation. The task of recommending news is especially complex because the system goals extend far beyond identifying articles of interest to a user [43]. News RS should avoid creating experience cocoons, echo chambers and filter bubbles, where users only encounter stories that reinforce their existing beliefs, interests, and ideologies [44]. The hazards of manipulation are great, e.g., a political party attempting to manipulate the system to show stories on a certain event to inflate their importance. Further research can investigate how to formulate prompts to manage exposure diversity and biases, and safeguard against manipulation.

References

[1] F. Wu, Y. Qiao, J.-H. Chen, C. Wu, T. Qi, J. Lian, D. Liu, X. Xie, J. Gao, W. Wu, et al., Mind: A large-scale dataset for news recommendation, in: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, 2020, pp. 3597–3606.

[2] J. Lian, F. Zhang, X. Xie, G. Sun, Towards better representation learning for personalized news recommendation: a multi-channel deep fusion approach., in: IJCAI, 2018, pp. 3805–3811.

[3] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, Y. Bengio, Learning phrase representations using rnn encoder-decoder for statistical machine translation, arXiv preprint arXiv:1406.1078 (2014).

[4] R. C. Staudemeyer, E. R. Morris, Understanding lstm–a tutorial into long short-term memory recurrent neural networks, arXiv preprint arXiv:1909.09586 (2019).

[5] Y. Chen, Convolutional neural network for sentence classification, Master's thesis, University of Waterloo, 2015.

[6] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, Advances in neural information processing systems 30 (2017).

[7] M. An, F. Wu, C. Wu, K. Zhang, Z. Liu, X. Xie, Neural news recommendation with long-and short-term user representations, in: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, 2019, pp. 336–345.

[8] C. Wu, F. Wu, T. Qi, C. Li, Y. Huang, Is news recommendation a sequential recommendation task?, in: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2022, pp. 2382–2386.

[9] C. Wu, F. Wu, M. An, J. Huang, Y. Huang, X. Xie, Neural news recommendation with attentive multi-view learning, arXiv preprint arXiv:1907.05576 (2019).

[10] W. Jin, Y. Cheng, Y. Shen, W. Chen, X. Ren, A good prompt is worth millions of parameters? low-resource prompt-based learning for vision-language models, arXiv preprint arXiv:2110.08484 (2021).

[11] Y. Zhang, H. Ding, Z. Shui, Y. Ma, J. Zou, A. Deoras, H. Wang, Language models as recommender systems: Evaluations and limitations, in: I (Still) Can't Believe It's Not Better! NeurIPS 2021 Workshop, 2021.

[12] Z. Cui, J. Ma, C. Zhou, J. Zhou, H. Yang, M6-rec: Generative pretrained language models are open-ended recommender systems, arXiv preprint arXiv:2205.08084 (2022).

[13] S. Geng, S. Liu, Z. Fu, Y. Ge, Y. Zhang, Recommendation as language processing (rlp): A unified pretrain, personalized prompt & predict paradigm (p5), RecSys (2022).

[14] S. Okura, Y. Tagami, S. Ono, A. Tajima, Embedding-based news recommendation for millions of users, in: Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining, 2017, pp. 1933–1942.

[15] C. Wu, F. Wu, M. An, J. Huang, Y. Huang, X. Xie, Neural news recommendation with attentive multi-view learning, arXiv preprint arXiv:1907.05576 (2019).

[16] T. T. Nguyen, P.-M. Hui, F. M. Harper, L. Terveen, J. A. Konstan, Exploring the filter bubble: the effect of using recommender systems on content diversity, in: Proceedings of the 23rd international conference on World wide web, 2014, pp. 677–686.

[17] M. Cinelli, G. De Francisci Morales, A. Galeazzi, W. Quattrociocchi, M. Starnini, The echo chamber effect on social media, Proceedings of the National Academy of Sciences 118 (2021) e2023301118.

[18] S. Vosoughi, D. Roy, S. Aral, The spread of true and false news online, science 359 (2018) 1146–1151.

[19] H. Abdollahpouri, M. Mansoury, R. Burke, B. Mobasher, E. Malthouse, User-centered evaluation of popularity bias in recommender systems, in: Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization, 2021, pp. 119–129.

[20] Y. Li, H. Chen, S. Xu, Y. Ge, Y. Zhang, Towards personalized fairness based on causal notion, in: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2021, pp. 1054–1063.

[21] C. Wu, F. Wu, X. Wang, Y. Huang, X. Xie, Fairness-aware news recommendation with decomposed adversarial learning, in: Proceedings of the AAAI Conference on Artificial Intelligence, volume 35,

2021, pp. 4462–4469.
[22] T. Qi, F. Wu, C. Wu, P. Sun, L. Wu, X. Wang, Y. Huang, X. Xie, Profairrec: Provider fairness-aware news recommendation, in: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2022, pp. 1164–1173.
[23] R. Burke, N. Sonboli, A. Ordonez-Gauger, Balanced neighborhoods for multi-sided fairness in recommendation, in: Conference on fairness, accountability and transparency, PMLR, 2018, pp. 202–214.
[24] N. Sonboli, F. Eskandanian, R. Burke, W. Liu, B. Mobasher, Opportunistic multi-aspect fairness through personalized re-ranking, in: Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization, 2020, pp. 239–247.
[25] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, arXiv preprint arXiv:1810.04805 (2018).
[26] A. Radford, K. Narasimhan, T. Salimans, I. Sutskever, et al., Improving language understanding by generative pre-training (2018).
[27] S. Xu, W. Hua, Y. Zhang, Openp5: Benchmarking foundation models for recommendation, arXiv:2306.11134 (2023).
[28] L. Li, Y. Zhang, L. Chen, Personalized prompt learning for explainable recommendation, arXiv preprint arXiv:2202.07371 (2022).
[29] C. Qin, A. Zhang, Z. Zhang, J. Chen, M. Yasunaga, D. Yang, Is chatgpt a general-purpose natural language processing task solver?, arXiv preprint arXiv:2302.06476 (2023).
[30] Y. Bang, S. Cahyawijaya, N. Lee, W. Dai, D. Su, B. Wilie, H. Lovenia, Z. Ji, T. Yu, W. Chung, et al., A multitask, multilingual, multimodal evaluation of chatgpt on reasoning, hallucination, and interactivity, arXiv preprint arXiv:2302.04023 (2023).
[31] T. Y. Zhuo, Y. Huang, C. Chen, Z. Xing, Exploring AI ethics of chatgpt: A diagnostic analysis, arXiv preprint arXiv:2301.12867 (2023).
[32] Y. Li, Y. Zhang, Fairness of chatgpt, arXiv:2305.18569 (2023).
[33] J. Liu, C. Liu, R. Lv, K. Zhou, Y. Zhang, Is chatgpt a good recommender? a preliminary study, arXiv preprint arXiv:2304.10149 (2023).
[34] P. P. Ray, Chatgpt: A comprehensive review on background, applications, key challenges, bias, ethics, limitations and future scope, Internet of Things and Cyber-Physical Systems (2023).
[35] C. Wu, F. Wu, M. An, Y. Huang, X. Xie, Neural news recommendation with topic-aware news representation, in: Proceedings of the 57th Annual meeting of the association for computational linguistics, 2019, pp. 1154–1159.
[36] C. Wu, F. Wu, S. Ge, T. Qi, Y. Huang, X. Xie, Neural news recommendation with multi-head self-attention, in: Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP), 2019, pp. 6389–6394.
[37] P. Abernathy, The expanding news desert, center for innovation and sustainability in local media, 2018.
[38] K. Chaykowski, Facebook's latest algorithm change: Here are the news sites that stand to lose the most (2018). URL: https://www.forbes.com/sites/kathleenchaykowski/2018/03/06/facebooks-latest-algorithm-change-here-are-the-%5Cnews-sites-that-stand-to-lose-the-most/?sh=74eef85134ec.
[39] Y. Wu, J. Cao, G. Xu, Y. Tan, Tfrom: A two-sided fairness-aware recommendation model for both customers and providers, in: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2021, pp. 1013–1022.
[40] D. Hamilton, Chatgpt user in china detained for creating and spreading fake news, police say, 2023. URL: https://apnews.com/article/chatgpt-china-deepfakes-criminal-detention-%5C7985cf38ffa33b09d3ad4f8ea5299967.
[41] R. Maruf, Lawyer apologizes for fake court citations from chatgpt, 2023. URL: https://apnews.com/article/chatgpt-china-deepfakes-criminal-detention-%5C7985cf38ffa33b09d3ad4f8ea5299967.
[42] J. Gravel, M. D'Amours-Gravel, E. Osmanlliu, Learning to fake it: limited responses and fabricated references provided by chatgpt for medical questions., medRxiv (2023) 2023–03.
[43] N. Helberger, N. Diakopoulos, The european ai act and how it matters for research into ai in media and journalism, Digital Journalism (2022) 1–10.
[44] B. Bodo, N. Helberger, K. Irion, F. Zuiderveen Borgesius, J. Moller, B. van de Velde, N. Bol, B. van Es, C. de Vreese, Tackling the algorithmic control crisis-the technical, legal, and ethical challenges of research into algorithmic agents, Yale JL & Tech. 19 (2017) 133.