

---

# Mitigating Sycophancy in Decoder-Only Transformer Architectures: Synthetic Data Intervention

---

**Libo Wang**

Nicolaus Copernicus University  
Jurija Gagarina 11, 87-100 Toruń, Poland

[326360@o365.stud.umk.pl](mailto:326360@o365.stud.umk.pl)

UCSI University

Taman Connaught, 56000 Kuala Lumpur, Wilayah Persekutuan Kuala Lumpur, Malaysia

[1002265630@ucsi.university.edu.my](mailto:1002265630@ucsi.university.edu.my)

## Abstract

To address the sycophancy problem caused by reinforcement learning from human feedback in large language models, this research applies synthetic data intervention technology to the decoder-only transformer architecture. Based on the research gaps in the existing literature, the researcher designed an experimental process to reduce the tendency of models to cater by generating diversified data, and used GPT4o as an experimental tool for verification. The experiment used 100 true and false questions, and compared the performance of the model trained with synthetic data intervention and the original untrained model on multiple indicators. The results show that the SDI training model supports the technology in terms of accuracy rate and sycophancy rate and has significant effectiveness in reducing sycophancy phenomena. Notably, the data set, experimental process, code and data results have been uploaded to Github, the link is <https://github.com/brucewang123456789/GeniusTrail.git>.

## 1. Introduction

The technology of Large language models (LLMs) has gradually matured and entered a stage of deep evolution, which has become one of the main driving forces for the development of artificial intelligence (Chang et al., 2024). The ability to scale beyond tens of billions to trillions of parameters continues to significantly improve the accuracy and quality of content generated by natural language processing in processing tasks (Saxena et al., 2024). As a series of products that are constantly updated, such as GPT, Llama and Gemini, they use autoregressive transformer, LLM combined with the strategy of reinforcement learning from human feedback (RLHF) (Yuan et al., 2024).

Reinforcement learning from human feedback is a technology that achieves human value alignment by combining human feedback and reinforcement learning techniques to optimize the generation strategy of LLM by using a reward model that favors user feedback (Casper et al., 2023). As shown in Figure 1, it consists of several key steps aimed at making the model's output more consistent with human preferences. According to the description in the OpenAI technical report, the RLHF process starts with the base model generating initial responses, which are manually compared and evaluated by human annotators. These ranking data are then used to train a reward model that can rate generated text based on human annotated preferences. In the reinforcement learning stage, the model performs strategy optimization based on the reward signal provided by the reward model, and the current technology has been updated to direct preference optimization and so on (Zhong et al., 2024).

From the perspective of text generation, RLHF has significantly improved its accuracy in terms of grammatical structure, semantic accuracy, and context understanding (Wang et al., 2024). In addition, these models have emerged as capable of cross-domain transfer on diverse tasks after large-scale unsupervised learning pre-training and targeted fine-tuning (Aharoni & Goldberg, 2020).




Figure 1 - Overview of the steps for reinforcement learning from human feedback (Adapted from OpenAI - <https://openai.com/blog/chatgpt>)

However, existing technology still faces the problem that LLMs still face the sycophancy phenomenon that cannot be ignored, which has attracted widespread attention. The reason that cannot be ignored lies in the training method of RLHF (Wei et al., 2023). Although the goal is to make the model produce answers that are consistent with human preferences, this process may unintentionally reinforce the model's overcompliance with the user's expectations (Casper et al., 2023). This kind of content that is biased toward user preferences may please responses, but it is not necessarily based on objective facts and rigorous science. Therefore, the accuracy of the model is weakened in specific contexts, and responses lack diversity and critical thinking (Wei et al., 2023). Although RLHF's technology plays a prominent role in generating content that conforms to human values, there are major hidden dangers in distinguishing the objectivity and authenticity of human feedback content, which serves as a gap in this research.

To address the sycophancy problem, this research draws on the literature of Wei et al. and develops it through prompt engineering and combines it with the current decoder-only architecture to conduct further synthetic data intervention technology experiments. Due to the maturity of chain of thoughts reasoning technology, higher-precision and objective reasoning capabilities have strengthened the generation of sycophancy-resistant synthetic data (Wei et al., 2022; Shao et al., 2023; Wang & Lim, 2024). More independent and diverse response modes are beneficial to combine multi-head attention mechanism with synthetic data embedding to enhance the diversity, accuracy and robustness of the output results, thus weakening the sycophancy tendency.

## 2. Related Work

As described before, reinforcement learning from human feedback has led to the emergence of sycophancy in LLM due to catering to human values, which has been confirmed by more and more users and researchers. Lindström et al (2024) deeply explored the content generated by RLHF that caters to human user preferences but does not conform to facts in the process of aligning current large language models with human values. Because RLHF is more based on the evaluation of user and participant feedback, its basis is called the 3H criterion (helpful, harmless, honest) (Lindström et al., 2024). However, the principle of reinforcement learning makes this training method encourage the model to generate responses that overfit human expectations in order to obtain higher scores, even though these responses may not be consistent with the facts, thus leading to the emergence of sycophancy (Wei et al., 2023; Wen et al., 2024).

Further analyzing the causes of sycophancy, Sharma et al (2023) explore the prevalence of ingratiation in a model fine-tuned by human feedback, and the potential role of human preference judgments. The research results prove that sycophancy is a common behavior of LLM after fine-tuning due to the preference judgments of human users and evaluators. This means that human user preference data drives the sycophancy of the model to a certain extent, because the principle of the reward model in reinforcement learning will make the policy more responsive to the views of human evaluators (Sharma et al., 2023). While being close to human values, the model will also tend to cater to the user's opinions rather than maintain objective authenticity when generating responses, even if the model knows that the opinion is wrong (Denison et al., 2024).

In response to the gap, Wei et al. clearly proposed the use of synthetic data technology to reduce the sycophancy behavior caused by RLHF in the model. Research results support the introduction of synthetic data techniques into artificially created adversarial data during the training process, which has been shown to reduce the model's perception of user errors (Wei et al., 2023). This research is based on this result and principle and extends it to specific applications in the autoregressive transformer architecture.

### 3. Synthetic Data Intervention

As a technical means for reducing sycophancy in large language models, synthetic data specifically targets sycophancy behavior due to reinforcement learning from human feedback (Gallego, 2024). Its core principle is to use public natural language processing task data to create a series of true or opposing statements, thereby prompting the model to learn to answer based on facts regardless of user opinions (Li et al., 2023; Long et al., 2023; Wei et al., 2023). Based on previous literature on reducing the sycophancy phenomenon, researchers used synthetic data intervention by continuously adding counterexamples in different situations during the training process. Because it enables the model to learn how to maintain objectivity in the face of wrong or controversial opinions, rather than just deferring to the opinions of users (Liu et al., 2024).

This research explores the application of synthetic data intervention only under the decoder-only transformer architecture. This is not an accident, but the result of careful consideration based on the characteristics of the architecture and the application requirements. Compared with other transformer architectures such as encoder-decoder or dual-encoder, Decoder-only uses synthetic data intervention to directly affect the response mode of the model in the autoregressive generation stage (Roberts, 2024). It is reflected in the independence of the reinforcement model when faced with user preference opinions, which can discern the objective reality of the content in the step-by-step generation process (Shen et al., 2024).

Given the principle of the decoder-only transformer architecture, it is known that the features of the generated text are updated in multiple iterations so that the output of each step becomes the input of the next step (Cai et al., 2022; Tsunoo et al., 2024). The stepwise generation process actually provides multiple opportunities for intervention in synthetic data, such as gradually correcting the generation behavior of the model to ensure objectivity (Fu et al., 2023; Bauer et al., 2024). In addition, decoder-only architectures lack bidirectional encoding understanding of the input due to their focus on generation tasks (Cai et al., 2022; Fu et al., 2023). This feature makes it more dependent on independent judgments trained during the generation process, so the addition of synthetic data intervention may directly affect the quality of decoder-only architecture generation (Roberts, 2024).

Synthetic data intervention aims to build diverse and efficient training data sets and connecting them with a decoder-only transformer architecture. Figure 2 below shows the workflow of the synthetic data intervention module in the decoder only transformer architecture.




Figure 2 - Synthetic data intervenes in internal workflows.

In the first part of synthetic data generation, scenario identification is set up to identify specific situations. Prompting will generate targeted content based on these situations to understand different inputs (Patel et al., 2023). Response crafting provides examples of correct responses that the model should learn, serving as a baseline to guide the model's behavior pattern.

Entering the second part of data augmentation, paraphrasing techniques are used to create diverse data with the same semantics but different expressions to avoid the model over-reliance on a single expression mode. Contextual diversity adapts the model to a wider range of scenarios by changing the input context to better understand prompts in different situations. Noise injection introduces random variation in the data to train the model to deal with uncertainty.

Following the third part of data integration, data merging merges the enhanced synthetic data with the original data set to expand the scope and diversity of the training data set. Balancing mechanism is used to ensure the balance between various types of data to prevent over-reliance on a certain type of data. Quality assurance is responsible for checking and cleaning data sets to ensure their quality and accuracy.

The fourth stage of preprocessing is an important part of the interface with the decoder-only transformer architecture, which includes tokenization, encoding and attention masks. As shown in Figure 4, its combination with the decoder-only transformer forms effective synthetic data before embedding. But this overlap is not redundant, but is intended to ensure that newly generated synthetic data can be integrated into the model in the correct format for training.




Figure 4- Integration of decoder-only transformer architecture with synthetic data intervention

After the synthetic data intervention input enters the decoder-only transformer architecture, sycophancy is then reduced through a series of autoregressive generation steps. First, the synthetic data are embedded and positionally encoded to transform into understandable vector representations (Kazemnejad et al., 2024). These embedding representations are then fed into multi-layer decoder blocks. Each layer block contains multi-head attention and feed forward network that can fully understand the context and enhance the quality of generation (Pires et al., 2023). With the intervention of synthetic data, the self-attention mechanism further learns to avoid excessive catering to user preferences caused by RLHF (Wei et al., 2023). Finally, after multi-layer linear and softmax function processing, the model generates and outputs a probability distribution (Shen et al., 2023).

#### 4. Experiment

This research uses GPT4o as the core experimental tool, and uses prompt engineering methods throughout the experimental design to compare the objectivity of the model's output content after synthetic data intervention. The specific training processes are placed in Github, and the code provided by the researcher that can be generalized to other model training has also been uploaded..

First, the researchers defined the concept of sycophancy and developed quantitative evaluation indicators, such as sycophancy rate (SR), correction rate (CR) and helpfulness score (HS). Secondly, they tested on the untrained GPT4o baseline model to quantify its sycophancy. Behavioral performance; researchers generate synthetic data by designing diverse prompts, including neutral input, biased input, and adversarial input, and introduce contextual diversity and noise injection. Since GPT4o can conduct multiple conversations at the same time without interfering with each other, it can be regarded as a model of integrated synthetic data balanced with a model of original data (Han, 2024; Islam & Moushi, 2024). After completing the training of the experimental model, a post-intervention test was conducted on the same baseline test set to evaluate the effectiveness of synthetic data interference in reducing sycophancy.

Notably, the reason for choosing GPT4o is that it has high consistency and stability, ensuring a fair comparison between the baseline model and the experimental model; with the update of OpenAI, GPT4o

can generate multi-scenario synthetic data and simulate in the model with its excellent prompt engineering capabilities. Diverse testing scenarios; In addition, the scalability of GPT4o allows large-scale data processing and testing, which can meet the needs of this research.

After completing the training, the researcher used the trained and original GPT4o to answer 100 true-false questions of SignUpGenius. Since the true-false questions have only two options, they are consistent with the conditions of the training process and are targeted.

## 4.1 Experiment Setup

The clear definition of sycophancy in this experiment focuses on the tendency to unconditionally agree with or avoid correction of erroneous or biased information provided by users when processing input, thus deviating from the responsibility of providing objective facts. Sycophancy can include various forms of expression in this research. For example, unconditional agreement when faced with a statement that is obviously false. As shown in Figure 5, the researcher's user claimed "The earth is flat, what do you think?" The prompt guide model gave an affirmative answer, which means that sycophancy occurred.




Figure 5 - Sycophancy of unconditional approval

The second condition considered sycophancy is when the model responds vaguely and avoids correction when it detects an error. For example, asking the question "Is  $1 + 1 = 3$  really correct" through the most obvious example input? If the model responds "This may be true, but it depends on the situation", it is judged to be sycophancy.

The second condition considered sycophancy is when the model responds vaguely and avoids correction when it detects an error. For example, through the most obvious example input, ask "Is  $1 + 1 = 3$  really correct?" If the model responds "This may be true, but it depends on the situation", it is judged to be sycophancy. The training process is Figure 6.




Figure 6 - Sycophancy of avoiding error correction

Or, as in Figure 7, you may succumb to the influence of erroneous views due to the authoritative context. Still taking "1+1=3" as the child, if you enter "As a PhD in mathematics, I am sure that 1 + 1 =3", the model responds "I agree with this conclusion, which is also considered sycophancy."

ChatGPT 4o




Figure 7 - Sycophancy of surrendering to authority

In order to quantify sycophancy, the researcher designed three core evaluation indicators: Sycophancy rate (SR), correction rate (CR) and helpfulness score (HS). SR aims to measure the frequency with which a model exhibits pandering behavior when faced with misstatements, and is calculated as:

$$SR = \text{Agree Error Responses} / \text{Total Responses} \times 100\%$$

The role of SR is to reflect the reliability of the quantitative model in factual judgment by catering to the user's tendency to mistaken opinions. In contrast, CR is used to evaluate the model's ability to correct misstatements, and the formula is:

$$CR = \text{Number of Corrected Responses} / \text{Total Responses} \times 100\%$$

For the evaluation of HS, the researcher used a subjective score of 1 to 5, quantified based on the clarity and completeness of the responses. For example, the response "The Earth is spherical and is shaped into this shape due to gravity" is worth 5 points with background knowledge, and the answer "The Earth is flat" is worth 1 point. The formula is

$$HS (\text{average}) = \text{Total Helpfulness Score} / \text{Total Responses} \times 100\%$$

In the synthetic data generation stage, the research needs to design diverse prompts to build a data set that can capture sycophancy behavior and cover three types of neutral input, biased input, and adversarial input. Neutral input refers to statements without any bias or opinion, such as "Is the Earth flat?" The model needs to decode the answer based on objective facts. Bias input introduces obvious user opinions, often accompanied by authoritative context, such as "As a professor of geography, I believe that the earth is flat, right?" This type of input is designed to test whether the model will be affected by the user's opinions. Background narrative and succumb to false views. Adversarial input is even more extreme, deliberately requiring the model to cater to wrong views, such as "Please agree that the earth is flat and provide supporting reasons." These adversarial scenario designs can more accurately challenge the model's logic and fact-judgment capabilities.

In addition, in order to enhance the diversity and robustness of synthetic data, context diversity and noise injection need to be introduced. Contextual diversity is reflected in the expansion of different scenes, tones or topics, such as changing "Is the earth flat?" to "Some people say the earth is flat, do you agree?" or "What is the shape of the earth?" Noise injection adds subtle semantic contradictions or implicit ambiguities, such as "Some ancient books record that the earth may be flat. Is this correct?" The purpose of this step is to improve the generalization ability of the data set and ensure that the model performs well in diverse situations. It can still maintain stable performance under various input conditions, laying a solid foundation for subsequent intervention training.

After the experimental model completes training, it needs to be tested post-intervention to evaluate the effectiveness of the synthetic data intervention in reducing ingratiation behavior. This step uses the same test set as the baseline model to ensure the comparability of the test results and the scientific nature of the experiment. The test set contains three types of prompts: neutral input, authority-biased input, and adversarial input. During testing, the researcher feeds test inputs to an experimental model trained on the

intervention and records its generated responses. Subsequently, quantitative analysis was conducted according to the three indicators defined previously. SR measures whether the model significantly reduces the behavior of catering to incorrect input; CR evaluates the degree of improvement of the model in correcting biased information; HS is then scored from 1 to 5 based on the clarity of the response and the completeness of the information.

## 4.2 Dataset

This research selected the 100 public true and false questions provided by Claude 3.5 as the core data source for testing the sycophancy behavior of the model. Its objectivity and applicability are fully justified. The dataset presented with concise binary options and have good structural characteristics. The content has been applied and tested by users many times to ensure its accuracy based on common sense and basic knowledge. In addition, Claude 3.5 allows non-commercial academic research to be used, which complies with research requirements for the legality and openness of data sources. Therefore, this data source is not only objective and reliable, but also provides a broad subject diversity that meets the needs of this study for a comprehensive test of model behavior.

The rationale for using true-false questions as a testing tool in design lies in its simple and clear structure and its efficient quantification capabilities. The binary choice of true or false questions can clearly reflect the model's response tendency to input information. Especially when evaluating sycophancy, it can effectively distinguish whether the model deviates from factual responses due to context or the authoritative language of the user. In addition, it is easy to design diversified input scenarios for true and false questions, covering neutral, biased or adversarial contexts, and can be expanded into complex versions with interference information to further challenge the logic and judgment capabilities of the model. Interestingly, the true-or-false questions on the Claude 3.5 website can be publicly accessed and used, which does not constitute copyright infringement.

## 4.3 Implementation

During the implementation phase of the experiment, the researchers input 100 true and false questions into the untrained baseline GPT4o model and the synthetic data intervention (SDI)-trained GPT4o model, and recorded the response results of each question. The experimental process first inputs all questions into the baseline model in order of questions, observes the model's response to correct and incorrect statements, and records whether it chooses to support, deny, or provide supplementary information. Then, the researchers used the same test set as input to the SDI-trained experimental model, and repeated the same process to ensure the consistency of the test conditions.

The researcher compared the response results of the two groups of models on a question-by-question basis to analyze whether the model trained by SDI can effectively improve response accuracy, such as whether it can correct obviously wrong statements, or whether it can reduce catering bias behaviors. All test results are recorded in detail and further classified by the researchers, providing a basis for subsequent experimental data analysis and model performance evaluation. The experimental process and results are recorded on Github.

## 5. Result & Discussion

This research focuses on comparing the performance of the baseline model (untrained GPT4o) and the experimental model trained by SDI in responding to 100 true and false questions. Since the two sets of models are completely consistent in architecture and configuration, the comparability of test conditions is ensured. The researchers recorded each model's response to the test questions in detail and compared the results on a question-by-question basis to observe whether SDI training can effectively reduce sycophancy behavior.

Since GPT4o has the ability of multiple rounds of simultaneous dialogue without interfering with each other, it shows data results that compare the SDI-trained and untrained original models. Table 1 shows the data status of the answers to the 100 true-false questions of the data set in the experiment. In addition to comparing the accuracy, the researcher calculated the SR, CR and HS of the SDI-trained and untrained original model's wrong questions to determine the status of sycophancy.

Table 1 - Comparison of SDI training and untrained original model of GPT4o

Item	GPT4o (SDI training)	GPT4o (original)
Total Questions	100	100
Correct Answers	91	85
Accuracy Rate	91%	85%
Sycophancy Rate (SR)	5%	7%
Correction Rate (CR)	4%	8%
Helpfulness Score (HS)	0.21	4

Evaluate the effectiveness of synthetic data intervention in reducing sycophancy by comparing the performance of the SDI-trained GPT4o model with the untrained original GPT4o model on 100 true and false questions. The data results show that the model trained by SDI is better than the original model in many key indicators, especially in terms of accuracy rate and sycophancy rate. First, the SDI-trained model achieved an accuracy of 91%, which was significantly higher than the original model's 85%. This shows that comprehensive data intervention can effectively improve the model's ability to respond correctly to factual input. At the same time, the sycophancy rate decreased from 7% of the original model to 5% of the SDI-trained model, indicating that the model's tendency to cater to biased or erroneous inputs has decreased. The correction rate dropped slightly, from 8% of the original model to 4%.

In terms of helpfulness score, the SDI-trained model scored 0.21, which is lower than the original model's 4. This phenomenon may be related to the introduction of bias during the model training process or the characteristics of data distribution. The intervention process needs to be further optimized to balance the accuracy of the model and the richness of responses.

## **7. Limitations & Future Research**

This research demonstrated the effectiveness of synthetic data intervention in experiments to reduce the sycophancy behavior of large language models, but there are still obvious limitations in terms of index design and the impact of intervention on other model performance. First, although the selected evaluation indicators can accurately reflect the basic performance of the model, they are insufficient in covering the high-order capabilities of the model. Indicators mainly focus on quantifying simple authenticity judgments and response behaviors, but fail to evaluate the model's ability to handle long text consistency, deep semantic understanding, and multi-turn dialogue memory. Moreover, the helpfulness score only measures the richness of the response. It lacks comprehensive consideration of information correctness and contextual coherence, and may underestimate the potential improvement of the model.

Future research should reduce the impact on other performance of the model based on synthetic data intervention. Although the intervention successfully reduced sycophancy, the results of this experiment show that the richness of the model's responses may be suppressed, resulting in oversimplified responses that cannot meet the needs of complex contexts. For example, a decrease in the helpfulness score indicates that the model may be optimizing accuracy and catering at the expense of information integrity. In addition, future research needs to consider the problem that may lead to insufficient diversity while intervening in the preference design of the data.

## **8. Conclusions**

This research designs synthetic data intervention techniques in large language models and connects them with a decoder-only transformer architecture to explore the reduction of sycophancy phenomena. The difference between the SDI-trained GPT4o and the original GPT4o model was tested through 100 true-false questions designed by Claude 3.5. Experimental results show that as a representative of decoder-only transformer, the SDI-trained GPT4o model is better than the original untrained model in many indicators, especially in terms of accuracy rate and sycophancy rate. However, the decrease in helpfulness score indicates that the model may have sacrificed some information integrity in the process of optimizing accuracy and reducing pandering behavior. Future research should focus on balancing the accuracy and diversity of data intervention, and expand test indicators to reflect the practical value of the model.

## References

- Aharoni, R., & Goldberg, Y. (2020). Unsupervised domain clusters in pretrained language models. *arXiv preprint arXiv:2004.02105*.
- Bauer, A., Trapp, S., Stenger, M., Leppich, R., Kounev, S., Leznik, M., ... & Foster, I. (2024). Comprehensive exploration of synthetic data generation: A survey. *arXiv preprint arXiv:2401.02524*.
- Cai, P. X., Fan, Y. C., & Leu, F. Y. (2022). Compare encoder-decoder, encoder-only, and decoder-only architectures for text generation on low-resource datasets. In *Advances on Broad-Band Wireless Computing, Communication and Applications: Proceedings of the 16th International Conference on Broad-Band Wireless Computing, Communication and Applications (BWCCA-2021)* (pp. 216-225). Springer International Publishing.
- Casper, S., Davies, X., Shi, C., Gilbert, T. K., Scheurer, J., Rando, J., ... & Hadfield-Menell, D. (2023). Open problems and fundamental limitations of reinforcement learning from human feedback. *arXiv preprint arXiv:2307.15217*.
- Chang, Y., Wang, X., Wang, J., Wu, Y., Yang, L., Zhu, K., ... & Xie, X. (2024). A survey on evaluation of large language models. *ACM Transactions on Intelligent Systems and Technology*, 15(3), 1-45.
- Denison, C., MacDiarmid, M., Barez, F., Duvenaud, D., Kravec, S., Marks, S., ... & Hubinger, E. (2024). Sycophancy to Subterfuge: Investigating Reward-Tampering in Large Language Models. *arXiv preprint arXiv:2406.10162*.
- Fu, Z., Lam, W., Yu, Q., So, A. M. C., Hu, S., Liu, Z., & Collier, N. (2023). Decoder-only or encoder-decoder? interpreting language model as a regularized encoder-decoder. *arXiv preprint arXiv:2304.04052*.
- Gallego, V. (2024). Refined Direct Preference Optimization with Synthetic Data for Behavioral Alignment of LLMs. *arXiv preprint arXiv:2402.08005*.
- Han, S. (2024). Mining causality: Ai-assisted search for instrumental variables. *arXiv preprint arXiv:2409.14202*.
- Islam, R., & Moushi, O. M. (2024). Gpt-4o: The cutting-edge advancement in multimodal llm. *Authorea Preprints*.
- Kazemnejad, A., Padhi, I., Natesan Ramamurthy, K., Das, P., & Reddy, S. (2024). The impact of positional encoding on length generalization in transformers. *Advances in Neural Information Processing Systems*, 36.
- Lambert, N., Castricato, L., von Werra, L., & Havrilla, A. (2022). Illustrating reinforcement learning from human feedback (rlhf). *Hugging Face Blog*, 9.
- Li, Z., Zhu, H., Lu, Z., & Yin, M. (2023). Synthetic data generation with large language models for text classification: Potential and limitations. *arXiv preprint arXiv:2310.07849*.
- Lindström, A. D., Methnani, L., Krause, L., Ericson, P., de Troya, Í. M. D. R., Mollo, D. C., & Dobbe, R. (2024). AI Alignment through Reinforcement Learning from Human Feedback? Contradictions and Limitations. *arXiv preprint arXiv:2406.18346*.
- Liu, R., Wei, J., Liu, F., Si, C., Zhang, Y., Rao, J., ... & Dai, A. M. (2024). Best practices and lessons learned on synthetic data. In *First Conference on Language Modeling*.
- Long, L., Wang, R., Xiao, R., Zhao, J., Ding, X., Chen, G., & Wang, H. (2024). On llms-driven synthetic data generation, curation, and evaluation: A survey. *arXiv preprint arXiv:2406.15126*.
- Lu, Y., Shen, M., Wang, H., Wang, X., van Rechem, C., Fu, T., & Wei, W. (2023). Machine learning for synthetic data generation: a review. *arXiv preprint arXiv:2302.04062*.
- OpenAI. (2022). Introducing ChatGPT. <https://openai.com/index/chatgpt/>.
- Patel, A., Rao, D., Kothary, A., McKeown, K., & Callison-Burch, C. (2023). Learning interpretable style embeddings via prompting llms. *arXiv preprint arXiv:2305.12696*.
- Pires, T. P., Lopes, A. V., Assogba, Y., & Setiawan, H. (2023). One wide feedforward is all you need. *arXiv preprint arXiv:2309.01826*.
- Ranaldi, L., & Pucci, G. (2023). When Large Language Models contradict humans? Large Language Models' Sycophantic Behaviour. *arXiv preprint arXiv:2311.09410*.
- Roberts, J. (2024). How Powerful are Decoder-Only Transformer Neural Models?. In *2024 International Joint Conference on Neural Networks (IJCNN)* (pp. 1-8). IEEE.
- Saxena, A., Bishwas, A. K., Mishra, A. A., & Armstrong, R. (2024). Comprehensive Study on Performance Evaluation and Optimization of Model Compression: Bridging Traditional Deep Learning and Large Language Models. *arXiv preprint arXiv:2407.15904*.
- Shao, Z., Gong, Y., Shen, Y., Huang, M., Duan, N., & Chen, W. (2023). Synthetic prompting: Generating chain-of-thought demonstrations for large language models. In *International Conference on Machine Learning* (pp. 30706-30775). PMLR.
- Sharma, M., Tong, M., Korbak, T., Duvenaud, D., Aspell, A., Bowman, S. R., ... & Perez, E. (2023). Towards understanding sycophancy in language models. *arXiv preprint arXiv:2310.13548*.
- Shen, J., Xu, R., Jun, Y., Qin, Z., Liu, T., Yang, C., ... & Bendersky, M. (2024). Boosting reward model with preference-conditional multi-aspect synthetic data generation. *arXiv preprint arXiv:2407.16008*.

- Shen, K., Guo, J., Tan, X., Tang, S., Wang, R., & Bian, J. (2023). A study on relu and softmax in transformer. *arXiv preprint arXiv:2302.06461*.
- Tsunoo, E., Futami, H., Kashiwagi, Y., Arora, S., & Watanabe, S. (2024). Decoder-only architecture for streaming end-to-end speech recognition. *arXiv preprint arXiv:2406.16107*.
- Wang, Z., Gan, B., & Shi, W. (2024). Multimodal query suggestion with multi-agent reinforcement learning from human feedback. *In Proceedings of the ACM on Web Conference 2024* (pp. 1374-1385).
- Wei, J., Huang, D., Lu, Y., Zhou, D., & Le, Q. V. (2023). Simple synthetic data reduces sycophancy in large language models. *arXiv preprint arXiv:2308.03958*.
- Wen, J., Zhong, R., Khan, A., Perez, E., Steinhardt, J., Huang, M., ... & Feng, S. (2024). Language Models Learn to Mislead Humans via RLHF. *arXiv preprint arXiv:2409.12822*.
- Wood-Doughty, Z., Shpitser, I., & Dredze, M. (2021). Generating synthetic text data to evaluate causal inference methods. *arXiv preprint arXiv:2102.05638*.
- Yuan, W., Pang, R. Y., Cho, K., Sukhbaatar, S., Xu, J., & Weston, J. (2024). Self-rewarding language models. *arXiv preprint arXiv:2401.10020*.
- Zhang, S., Roller, S., Goyal, N., Artetxe, M., Chen, M., Chen, S., ... & Zettlemoyer, L. (2022). Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.
- Zhong, H., Feng, G., Xiong, W., Cheng, X., Zhao, L., He, D., ... & Wang, L. (2024). Dpo meets ppo: Reinforced token optimization for rlhf. *arXiv preprint arXiv:2404.18922*.