

ArithmAttack: Evaluating Robustness of LLMs to Noisy Context in Math Problem Solving

Zain Ul Abedin^{♣*} Shahzeb Qamar^{♣*} Lucie Flek^{♣♠} Akbar Karimi^{♣♠}

[♣]Bonn-Aachen International Center for Information Technology, University of Bonn, Germany

[♠]Lamarr Institute for ML and AI, Germany

ak@bit.uni-bonn.de

Abstract

While Large Language Models (LLMs) have shown impressive capabilities in math problem-solving tasks, their robustness to noisy inputs is not well-studied. We propose **ArithmAttack** to examine how robust the LLMs are when they encounter noisy prompts that contain extra noise in the form of punctuation marks. While being easy to implement, ArithmAttack does not cause any information loss since words are not added or deleted from the context. We evaluate the robustness of eight LLMs, including Llama3, Mistral, Mathstral, and DeepSeek on noisy GSM8K and MultiArith datasets. Our experiments suggest that all the studied models show vulnerability to such noise, with more noise leading to poorer performances.

1 Introduction

As Large Language Models (LLMs) are improving in their ability to accurately process human language, their math problem-solving is also enhancing (Saraf et al., 2024; Agrawal et al., 2024; Wu et al., 2024). However, these sets of questions might require reasoning capabilities to be answered. While LLMs have been shown to have such capabilities to some extent (Imani et al., 2023), their robustness to adversarial inputs remains a challenge. For instance, these models can be vulnerable to simple replacement of words with their synonyms (Zhou et al., 2024) and even typographical errors can negatively impact their ability to reason (Gan et al., 2024). However, such attacks can semantically alter the samples by changing the current sample features to completely different ones (e.g. amoral → moral).

In this paper, we further investigate the math problem-solving robustness of LLMs to a different set of changes that take the form of noisy context containing a variety of punctuation marks. Given

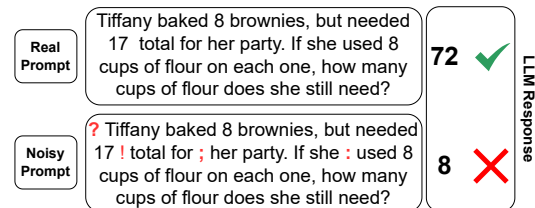


Figure 1: Noisy context breaks the LLM’s capability to give the right answer.

that none of the words are changed when new punctuation marks are inserted into the input text, the semantic similarity of the perturbed sentence remains unchanged. The key research question for this study is: *How do LLMs respond to noise attacks consisting of random punctuation marks in the context of math problem-solving?* Figure 1 shows an example of an LLM response under ArithmAttack, where the model behaves erratically when it sees a noisy context whereas it answers the question in the clean prompt correctly.

Inspired by the AEDA (An Easier Data Augmentation) method (Karimi et al., 2021), we propose ArithmAttack to assess the robustness of eight LLMs (i.e. two Llama models (Dubey et al., 2024), two Mistral models (Jiang et al., 2023), Zephyr (Tunstall et al., 2023), Gemma2 (Team et al., 2024), Qwen2.5 (Yang et al., 2024), and DeepSeek (Guo et al., 2025)) to noisy data. Similarly to AEDA, we introduce this noise by randomly inserting punctuation marks into the context of math problems from two math datasets, namely GSM8K (Cobbe et al., 2021) and MultiArith (Roy and Roth, 2015). We then evaluate how these models perform under different noise levels, with the noise affecting 10%, 30%, and 50% of the sentence length (based on the number of words).

Our contributions are twofold: 1) We propose ArithmAttack which produces noisy contexts containing random punctuation marks to assess the

*Equal contribution

robustness of LLMs in math problem-solving. 2) We evaluate eight LLMs, with parameter counts of 1.5B, 2B, 7B, and 8B on math datasets and observe that all the studied models show growing vulnerability to ArithmAttack as the amount of noise increases.

2 Related Work

Noise insertion has been shown to be effective in deteriorating the performance of encoder models in various tasks such as toxic text classification (Hosseini et al., 2017; Eger and Benz, 2020), sentiment analysis (Formento et al., 2021), and natural language inference (Formento et al., 2023).

In the context of math problem solving, Large Language Models (LLMs) have been shown to be vulnerable to a variety of changes in the input context, including typographical errors (Gan et al., 2024), word replacement (Zhou et al., 2024), gibberish or irrelevant context inclusion (Shi et al., 2023), and semantic perturbations (Zhu et al., 2023). Gan et al. (2024) propose an adversarial typo attack that breaks the reasoning process of LLMs. Instead of modifying characters, Zhou et al. (2024) propose a dataset, called RobustMath, where they replace words with their synonyms to evaluate the robustness of large language models. Similarly, Li et al. (2024) propose GSM-plus dataset, based on GSM8K, modified with a variety of mathematical perturbations such as distractor insertion and arithmetic variation. In the study by Zhu et al. (2023), the authors employ different types of textual attacks on prompts, including character, word, sentence, and semantic attacks. In contrast, Xie et al. (2024) propose to modify the numeric values in the questions using abstract syntax trees resulting in examples that fool the LLMs.

While the literature mainly concentrates on modifying the lexical or semantic content of the prompts, we aim to keep the contextual information intact and instead focus on the model behavior changes in reasoning when encountering punctuation noise. In addition, an advantage of our method is that it is extremely straightforward to implement, and as we show in the results section, it is also effective in degrading the performance of LLMs in math problem-solving.

3 Experiments

To carry out our experiments, we use two well-known math datasets and eight LLMs.

3.1 Datasets

GSM8K (Cobbe et al., 2021) contains 8.5K high-quality, linguistically diverse grade school math word problems. The test set contains 1.32k data points on which we do our experiments. This dataset provides a variety of arithmetic and logical questions typical of middle school education, making it ideal for testing comprehension and problem-solving capabilities of LLMs under noisy conditions.

MultiArith (Roy and Roth, 2015) offers a broad examination of language model performance across multiple arithmetic problem types and complexities. The test set contains 180 data points on which we do our experiments. It serves as a crucial benchmark for understanding how contextual noise impacts the model’s ability to parse and solve mathematical questions.

3.2 Models

To study a variety of language models and at the same time observe our computational budget, we opted for eight widely utilized LLMs that have been trained by different companies. These models are **Mistral-7B-Instruct-v0.2** (Jiang et al., 2023), **Mathstral-7b-v0.1** (Jiang et al., 2023), **Llama-3-8B-Instruct** and **Llama-3.1-8B-Instruct** (Dubey et al., 2024), **Gemma-2-2b-it** (Team et al., 2024), **Zephyr-7b-beta** (Tunstall et al., 2023), **Qwen2.5-1.5B-Instruct** (Yang et al., 2024), and **DeepSeek-R1-Distill-Llama-8B** (Guo et al., 2025). Throughout this paper, we will refer to these models as Mistral, Mathstral, Llama3, Llama3.1, Gemma2, Zephyr, Qwen2.5, and DeepSeek respectively.

4 Methodology

To obtain the responses from LLMs, we use the Zero-Shot CoT (Kojima et al., 2022) prompting, using the following prompt:

Prompt 1

Think step by step through the following problem and clearly show each step of your reasoning. Ensure the final answer is indicated by ending with {The final answer is}.

4.1 Noisy Dataset Creation

Once satisfactory results were achieved with clean datasets, we proceeded to test the models on noisy

Models	Clean Acc (%)	Punctuation Percentage			ASR
		10	30	50	
		Mistral	42.07	41.62	
Mathstral	77.63	75.51	71.34	70.65	19.81
Llama3	75.43	73.31	73.08	72.17	11.73
Llama3.1	82.25	81.04	78.84	77.02	12.53
Gemma2	49.65	45.10	36.46	35.63	41.82
Zephyr	23.27	18.04	18.04	10.08	74.80
Qwen2.5	61.10	56.02	52.69	49.35	31.59
DeepSeek	73.76	73.76	70.43	67.24	20.46

Table 1: Results for GSM8K dataset (numbers are in percentages). The performance for all models drops under ArithmAttack. Llama3.1 has the top performance under all levels of noise.

data. For the introduction of noise, we follow a similar approach to Karimi et al. (2021), by altering the hyperparameters in the logic. In their study, they insert the punctuation marks by randomly choosing a number between 1 and one-third of the length of the sequence which indicates how many insertions will be carried out. But in our case, instead of randomly choosing the number of insertions, we fix it to be 10%, 30%, and 50% of the total length of the sentence but still choose random positions to insert the noise. We employed six types of punctuation marks: {".", "'", "!", "?", ";", ':'}.

4.2 ASR and Similarity Calculation

We evaluate the models with their performance accuracy against noisy input and **Attack Success Rate (ASR)**. ASR (Wang et al., 2021) measures how effective an adversarial attack is on a model. Specifically, it looks at how often the model’s predictions are changed incorrectly after the adversarial attack. In this study, the average ASR has been taken for every model with 10%, 30% and 50% noisy dataset’s responses with the help of Formula 1:

$$ASR = \frac{\sum_{(x,y) \in D} \mathbb{I}[f(A(x)) \neq y]}{\sum_{(x,y) \in D} \mathbb{I}[f(x) = y]} \quad (1)$$

In other words, ASR is the ratio of changed answers after attack to previously correct answers produced by the LLM.

We also calculate the similarity of the perturbed samples to the original ones. **Similarity** represents the average semantic similarity between two contexts. Given that our method does not alter the

Models	Clean Acc (%)	Punctuation Percentage			ASR
		10	30	50	
		Mistral	73.88	72.77	
Mathstral	96.11	92.77	86.11	87.22	9.47
Llama3	95.00	92.77	<u>91.66</u>	<u>88.33</u>	7.79
Llama3.1	99.44	94.44	<u>91.66</u>	83.88	9.67
Gemma2	89.44	82.77	78.88	72.22	19.45
Zephyr	37.22	22.22	16.11	12.77	77.10
Qwen2.5	97.22	94.44	85.55	83.88	11.04
DeepSeek	93.88	90.00	92.77	88.88	<u>8.28</u>

Table 2: Results for MultiArith dataset (numbers are in percentages). The performance for all models drops under ArithmAttack. Llama3 has the lowest drop, making it more robust than others.

words in the sentence, the resulting samples after applying ArithmAttack are scored 100 percent similar to the original samples using Universal Sentence Encoder (Cer et al., 2018) as the scorer. This indicates that our noise insertion attack does not impose any semantic shifts on the input text.

5 Results and Analysis

As shown in Tables 1 and 2, Llama3.1 outperforms other models across both datasets in the majority of the cases. It achieves the highest accuracies in both clean and noise-affected settings (except in 30% and 50% noisy data of the MultiArith dataset where DeepSeek in the former and both Llama3 and DeepSeek in the latter have higher accuracies). This makes it the most reliable model for handling mathematical problems under noisy input conditions. However, in terms of ASR score, Llama3 has the lowest score and Llama3.1 with a slightly higher one, indicating that Llama models are more robust to noise than other studied models with the exception of DeepSeek only in MultiArith dataset showing comparable robustness. In addition, the Mathstral model compared to Mistral exhibits more robustness which can be attributed to its higher mathematical understanding. In contrast, Zephyr was the lowest-performing model, exhibiting low clean accuracy and suffering a significant decline in performance as noise was introduced. Its high ASR score makes it unsuitable for tasks involving noisy data, reflecting poor robustness.

Figure 2 shows the relationship between the model’s accuracy and the noise present in the datasets. For both datasets, as the percentage of

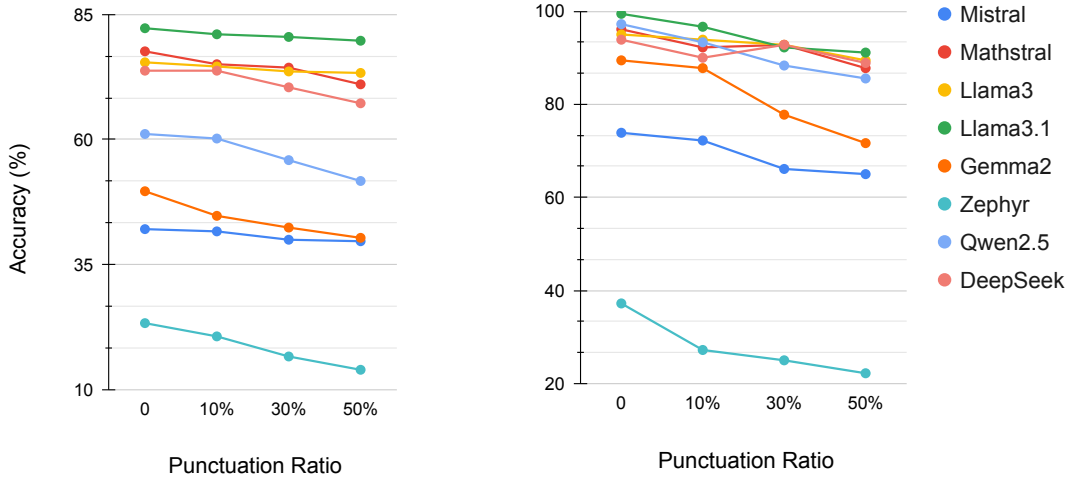


Figure 2: Accuracy of the studied models on different levels of noise for GSM8K (left) and MultiArith (right) datasets. Llama models show the highest robustness as well as performance.




Figure 3: Comparing the attack success rates on the studied models for GSM8K and MultiArith datasets (lower is better for model robustness). Llama models are more robust than others under ArithmAttack.

noise in the data increases, the accuracy decreases. This indicates that these models are not robust against noise in the data. This also provides a future direction for improving these models and making them more robust to noise.

Across all models except for Zephyr, the impact of noise was more pronounced in the GSM8K dataset than in MultiArith, with a larger drop in accuracy as the noise levels increased (Figure 3). In manual inspection, we found out that the GSM8K dataset was more difficult to solve than the MultiArith dataset. This suggests that the models may struggle more with noise in math datasets with more difficulty.

Model	Miss Rate (%)	
	GSM8K	MultiArith
Mistral	9.0	1.1
Mathstral	0.0	1.1
Llama3	1.0	1.1
Llama3.1	0.0	0.0
Gemma2	3.0	2.2
Zephyr	2.0	12.8
Qwen2.5	1.0	0.5
DeepSeek	4.0	0.0

Table 3: Miss rate of the models in answer extraction

Answer Extraction Accuracy To evaluate the accuracy of the models, we developed a script to extract answers from the LLM responses. The extraction process underwent multiple iterations, as it needed to accurately extract the answer and compare it with the ground truth. However, even with the final prompt, we observed a couple of inconsistencies in the answer extraction. Therefore, we went through outputs manually to estimate the *miss rate* (i.e. the rate with which the correct answer is not extracted). In manual inspection, we evaluated the entire responses for the MultiArith dataset and the first 100 responses for the GSM8K dataset from all the models except for the DeepSeek model. For this model (due to time and labor constraints), we evaluated the first 50 samples from each dataset. Table 3 shows that the miss rate is minimal for most of the models. In the cases of Mistral (for GSM8K) and Zephyr (for MultiArith), the miss rates can be

significant. While this can be an indication of lower ability in following instructions in these models, considering the gap in the performance and ASR scores, this does not affect the observed trends.

6 Conclusions and Future Work

We evaluated how well different language models handle mathematical problem-solving tasks in both clean and noisy conditions. Our results indicate that all studied models can be vulnerable to extra noise with varying degrees, with Llama models being the highest-performing and the most robust model in the majority of the experiments. In addition, comparing the two models of Mathstral and Mistral from the same family, the one with mathematical knowledge exhibited more robustness to noise. Lastly, the findings revealed that more complex datasets such as GSM8K can become more difficult to understand as they become noisier. Future research can include datasets beyond GSM8K and MultiArith as well as other reasoning tasks such as logical and causal reasoning, which could provide deeper insights into the models' robustness in different scenarios. Further experimentation with different types of noise could also help enhance our understanding of the latent vulnerabilities in LLMs. Finally, explaining why ArithmAttack can break the reasoning flow of LLMs would be another valuable follow-up to this work.

7 Limitations

To make the questions noisy, we have opted for one type of noise which is irregular use of punctuation marks. While some of the other noise types such as spelling and typographical errors have been studied in the literature (mentioned in related work), there are other types such as grammatical errors, wrong abbreviations, and acronyms that we have not explored. In addition, to observe our computational budget, we have utilized only two math datasets and eight LLMs. For a more comprehensive experimentation, one can experiment with other available math datasets and a larger number of LLMs.

Acknowledgments

This work was partially supported by the AISafety Project, funded by the BMBF under the grant proposal 05D23PD1, and by the state of North Rhine-Westphalia as part of the Lamarr Institute for Machine Learning and Artificial Intelligence. We would also like to thank the reviewers for their

invaluable comments, which helped strengthen the quality of this work.

References

- Vansh Agrawal, Pratham Singla, Amitoj Singh Miglani, Shivank Garg, and Ayush Mangal. 2024. Give me a hint: Can llms take a hint to solve math problems? In *The 4th Workshop on Mathematical Reasoning and AI at NeurIPS'24*.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, et al. 2018. Universal sentence encoder for english. In *Proceedings of the 2018 conference on empirical methods in natural language processing: system demonstrations*, pages 169–174.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Steffen Eger and Yannik Benz. 2020. From hero to zéro: A benchmark of low-level adversarial attacks. In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 786–803.
- Brian Formento, Chuan-Sheng Foo, Luu Anh Tuan, and See Kiong Ng. 2023. Using punctuation as an adversarial attack on deep learning-based nlp systems: An empirical study. In *Findings of the association for computational linguistics: EACL 2023*, pages 1–34.
- Brian Formento, See-Kiong Ng, and Chuan-Sheng Foo. 2021. Special symbol attacks on nlp systems. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.
- Esther Gan, Yiran Zhao, Liying Cheng, Yancan Mao, Anirudh Goyal, Kenji Kawaguchi, Min-Yen Kan, and Michael Shieh. 2024. Reasoning robustness of llms to adversarial typographical errors. *arXiv preprint arXiv:2411.05345*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.

- Hossein Hosseini, Sreeram Kannan, Baosen Zhang, and Radha Poovendran. 2017. Deceiving google’s perspective api built for detecting toxic comments. *arXiv preprint arXiv:1702.08138*.
- Shima Imani, Liang Du, and Harsh Shrivastava. 2023. Mathprompter: Mathematical reasoning using large language models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 5: Industry Track)*, pages 37–42.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023. *Mistral 7b*. *Preprint*, arXiv:2310.06825.
- Akbar Karimi, Leonardo Rossi, and Andrea Prati. 2021. *AEDA: An easier data augmentation technique for text classification*. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2748–2754, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213.
- Qintong Li, Leyang Cui, Xueliang Zhao, Lingpeng Kong, and Wei Bi. 2024. Gsm-plus: A comprehensive benchmark for evaluating the robustness of llms as mathematical problem solvers. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2961–2984.
- Subhro Roy and Dan Roth. 2015. Solving general arithmetic word problems. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1743–1752.
- Amrutesh Saraf, Pooja Kamat, Shilpa Gite, Satish Kumar, and Ketan Kotecha. 2024. Towards robust automated math problem solving: a survey of statistical and deep learning approaches. *Evolutionary Intelligence*, pages 1–38.
- Freda Shi, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed H Chi, Nathanael Sch  rli, and Denny Zhou. 2023. Large language models can be easily distracted by irrelevant context. In *International Conference on Machine Learning*, pages 31210–31227. PMLR.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, L  onard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ram  , et al. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.
- Lewis Tunstall, Edward Beeching, Nathan Lambert, Nazneen Rajani, Kashif Rasul, Younes Belkada, Shengyi Huang, Leandro von Werra, Cl  mentine Fourrier, Nathan Habib, Nathan Sarrazin, Omar Sanseviero, Alexander M. Rush, and Thomas Wolf. 2023. *Zephyr: Direct distillation of lm alignment*. *Preprint*, arXiv:2310.16944.
- Boxin Wang, Chejian Xu, Shuohang Wang, Zhe Gan, Yu Cheng, Jianfeng Gao, Ahmed Hassan Awadallah, and Bo Li. 2021. Adversarial glue: A multi-task benchmark for robustness evaluation of language models. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Yiran Wu, Feiran Jia, Shaokun Zhang, Hangyu Li, Erkang Zhu, Yue Wang, Yin Tat Lee, Richard Peng, Qingyun Wu, and Chi Wang. 2024. Mathchat: Converse to tackle challenging math problems with llm agents. In *ICLR 2024 Workshop on Large Language Model (LLM) Agents*.
- Roy Xie, Chengxuan Huang, Junlin Wang, and Bhuwan Dhingra. 2024. Adversarial math word problem generation. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 5075–5093.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. 2024. Qwen2. 5 technical report. *arXiv preprint arXiv:2412.15115*.
- Zihao Zhou, Qiufeng Wang, Mingyu Jin, Jie Yao, Jianan Ye, Wei Liu, Wei Wang, Xiaowei Huang, and Kaizhu Huang. 2024. Mathattack: Attacking large language models towards math solving ability. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 19750–19758.
- Kaijie Zhu, Jindong Wang, Jiaheng Zhou, Zichen Wang, Hao Chen, Yidong Wang, Linyi Yang, Wei Ye, Yue Zhang, Neil Zhenqiang Gong, et al. 2023. Promptbench: Towards evaluating the robustness of large language models on adversarial prompts. *arXiv e-prints*, pages arXiv–2306.