

OPDAI at SemEval-2024 Task 6: Small LLMs can Accelerate Hallucination Detection with Weakly Supervised Data

Chengcheng Wei, Ze Chen, Songtan Fang, Jiarong He and Max Gao

Interactive Entertainment Group of Netease Inc., Guangzhou, China
{weichengcheng, jackchen, fangsongtan, gzhejiarong, jgao}@corp.netease.com

Abstract

This paper mainly describes a unified system for hallucination detection of LLMs, which wins the second prize in the model-agnostic track of the SemEval-2024 Task 6, and also achieves considerable results in the model-aware track. This task aims to detect hallucination with LLMs for three different text-generation tasks without labeled training data. We utilize prompt engineering and few-shot learning to verify the performance of different LLMs on the validation data. Then we select the LLMs with better performance to generate high-quality weakly supervised training data, which not only satisfies the consistency of different LLMs, but also satisfies the consistency of the optimal LLM with different sampling parameters. Furthermore, we finetune different LLMs by using the constructed training data, and finding that a relatively small LLM can achieve a competitive level of performance in hallucination detection, when compared to the large LLMs and the prompt-based approaches using GPT-4.

1 Introduction

The emergence of Large Language Models (LLMs)(Zhao et al., 2023) has sparked a significant transformation in the field of Natural Language Processing (NLP), ushering in a new era of unparalleled advancements in text generation and comprehension. This revolutionary technology has elevated the capabilities of AI systems, enabling them to perform complex reasoning and problem-solving tasks with remarkable proficiency(Zhao et al., 2023). At the heart of this transformation lies the LLMs’ ability to compress vast amounts of knowledge into neural networks, effectively turning them into versatile agents capable of engaging in natural language conversations with humans(Hadi et al., 2023). This has broadened the scope of AI applications beyond traditional domains such

as chatbots and virtual assistants, into areas previously thought to be the exclusive domain of humans, particularly those involving creativity and expertise. LLMs are not only limited to language-related tasks but can also function as generalist agents, collaborating with external systems, tools, and models to achieve a wide range of objectives set by humans(Triguero et al., 2024).

However, recent advancements in research have uncovered a concerning weakness: their proneness to hallucinate content across a range of applications(Ji et al., 2023). Hallucination is defined as the generation of information that either conflicts with established sources or cannot be substantiated by available knowledge. The occurrence of hallucination in LLMs poses a significant threat to their practical deployment. While prior works have delved into the roots of hallucination within specific, smaller-scale language models and tasks, there is still a notable gap in understanding the exact nature and prevalence of content that LLMs are likely to hallucinate(Cui et al., 2024; Chang et al., 2023).

To address this challenge, we implement a unified system for hallucination detection of LLMs, when there is no labeled training data. This system comprises five parts: *Base Model Selection*, *Prompt Engineering*, *Weakly-supervised Data Generation*, *SFT* and *Ensemble Learning*. We first verify the performance of different base LLMs on this task. Then we select the best LLMs and prompt is optimized to improve the performance. And weakly-supervised dataset is generated by using the selected LLMs. For further improvement, SFT is done based on the constructed dataset and ensemble learning is adopted.

2 Related Work

In this section, we will introduce other work related to the subsequent methods.

Mixture of Experts(MoE)(Jacobs et al., 1991) is an AI technique that involves a group of specialized models (experts) being coordinated by a gating mechanism to address various aspects of the input space, in order to optimize performance and efficiency. This approach capitalizes on the idea that an ensemble of weaker language models, each focusing on specific tasks, can yield more precise results, similar to traditional ML ensemble methods. However, it introduces a novel concept of dynamically routing the input during the generation process. In the subsequent methods, we will conduct comparative experiments using the intelligent model base of MoE. This paper(Chen et al., 2022) integrates POS information and word semantic representation using an MoE approach.

Model ensembling combines the predictions from multiple models together. Traditionally this is done by running each model on some inputs separately and then combining the predictions. However, if the candidate models have the same architecture, they can be combined together to create a new model for prediction. Many surveys systematically elucidate the basic concepts of ensemble learning and various methods, including model fusion and model voting(Krawczyk et al., 2017; Dong et al., 2020; Yang et al., 2023).

LoRA (Low-Rank Adaptation of Large Language Models) (Hu et al., 2021) is a widely used and lightweight training technique that markedly reduces the number of trainable parameters. It functions by adding a smaller set of new weights to the model and training only these. As a result, training with LoRA is notably faster, more memory-efficient, and yields smaller model weights (a few hundred MBs), which are more manageable for storage and sharing. Some other methods(Ye et al., 2023; Wang et al., 2023; Chen et al., 2023) have been developed to improve LoRA.

Chain-of-Thought Prompting(CoT)(Wei et al., 2022) is an emerging application of language model technology. The core idea of this method is to encourage the model to not only generate the final answer but also gradually demonstrate its reasoning and the process of reaching conclusions. Subsequent work(Zhou et al., 2022) has applied the idea of CoT, breaking down problems into a series of sub-problems, allowing the model to reason step by step and ultimately provide the correct answer. And another work(Wang et al., 2022) introduces a method called "self-consistency" to further enhance the effectiveness of CoT Prompting. By

generating multiple reasoning paths and selecting the most consistent answer, the model can reduce errors and improve the accuracy of reasoning.

3 Task Description

trial data	unlabeled train data	validation data	test data
80	60000	1000	3000

Table 1: Dataset provided by SHROOM

SHROOM(Mickus et al., 2024) asked participants to perform binary classification to identify cases of fluent overgeneration hallucinations in two different setups: model-aware and model-agnostic tracks. And three different NLG tasks: definition modeling (DM), machine translation (MT) and paraphrase generation (PG) are covered in both tracks. In model-aware track, the model information is provided. The provided development and test sets include binary annotations from a minimum of five different annotators, along with a majority vote gold label. Table 1 gives an overview of the provided dataset.

4 Methodology

Figure 1 shows the overview of our approach. Our method consists of five main steps. First of all, multiple LLMs are compared on the hallucination detection validation dataset and among which we select the best base model. The LLMs selected in the first step will be utilized in the subsequent steps 2, 3, and 4. In the second step, we designed a Prompt Engineering module consisting of three sub-modules: few-shot prompting, instruction optimization, and the utilization of Chain-of-Thought. The subsequent experiments will demonstrate that prompt engineering significantly enhances the capabilities of the base model.

Then moving on to the Label Generation step, we apply the Prompt Engineering module to the selected best LLMs. We make predictions on the unlabeled training set and ensure the inference consistency among multiple LLMs as well as the inference consistency under different inference parameters of individual LLM. We also ensure label balance in this process.

The following steps are Model Training and Ensemble Learning. We utilize SFT based on the constructed dataset using the LoRA (Hu et al., 2021) method. Finally, we select a few top-performing

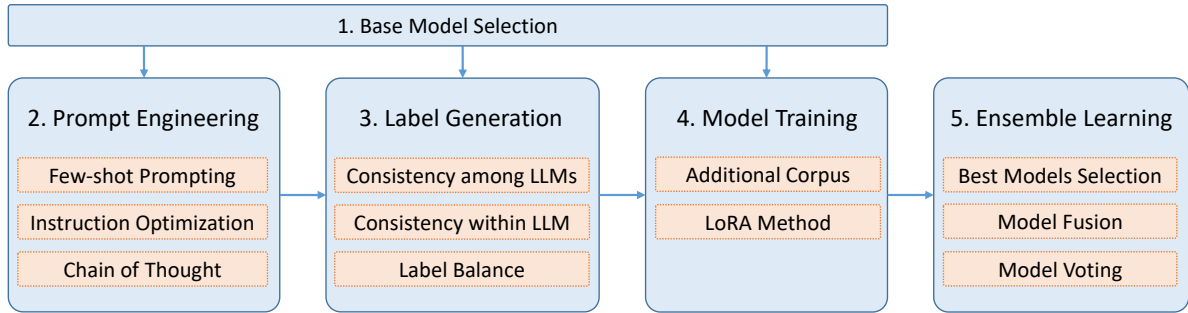


Figure 1: Overview of our proposed method. The method consists of five main steps, each of which comprises several modules that we have designed for SemEval-2024 Task 6.

models and perform fusion at both the model weight level(model fusion) and the model prediction probability level(model voting) in order to seek better performance.

4.1 Prompt Engineering

After experimenting with multiple large-scale language models, we select the best-performing 14B *Mixtral_7Bx2_MoE* as the base model for prompt engineering to achieve better results.

Few-shot prompting. Few-shot prompting can be used as a technique to enable in-context learning where we provide demonstrations from the provided *trial data* in the prompt to steer the base model to better performance. Considering that the trial data contains with labels, we randomly sample the specific task’s datapoints for different tasks, ensuring an equal number of data points for "hallucination" and "not hallucination" examples, to serve as few-shot examples for each respective task.

Optimizing the instruction. There are some deficiencies with the instruction introduced in Section 5.1, and we name this version of instruction as the *naive* version. The most obvious issue is that the naive instruction does not include the descriptions of the DM, MT, and PG tasks. The desired task description includes the task definition and all known useful information, rather than just focusing on the sentence and context. We design different instructions for different tasks, which can be found in the Appendix A. In this way, we can append additional information to the prompt to assist the LLM in better understanding the problem.

Chain-of-thought prompting. Chain-of-thought (CoT) prompting(Wei et al., 2022) is a recently developed method that encourages the language model to explain its reasoning. We combine the aforementioned few-shot prompting, developed in-

struction, and CoT to utilize them together to further enhance the capability of the base model.

4.2 Label Generation and Weakly-SFT

After improving the performance of the base model using prompt engineering, we use the optimal settings to infer on unlabeled training data and obtain weakly supervised labels. These weakly supervised labels are then used to finetune the base model.

Inference consistency in generating labels. During the process of inferring weakly supervised labels for the unlabeled training data, we placed a great emphasis on both the consistency of inference across different LLMs and the consistency of inference within the same LLM but with different parameter settings. To achieve this, we carefully selected several sets of top-performing base models. Leveraging the prompt engineering techniques mentioned earlier, we conducted inference on the same model using various parameter configurations. Subsequently, we handpicked the data points with consistent inferences across different parameter settings to establish the final inference results for that particular LLM. Additionally, we applied a filtering process to the inference results obtained from different LLMs, ensuring that only datapoints with consistent inferences were retained. Through these rigorous steps, we ensured that our generated weakly supervised labels for the training set exhibits robustness to both the choice of LLM base and the specific sampling parameters employed. Finally, we used sampling techniques to balance the data volume of the two categories.

Fine-tuning LLMs. The weak supervision generated by the base model is applied to guide models of equal or smaller scale. The LLMs undergo fine-tuning using the LoRA approach, a popular and lightweight training technique that effectively

Model Name	Model Size	Model-agnostic track		Model-aware track	
		<i>acc</i>	<i>rho</i>	<i>acc</i>	<i>rho</i>
Mistral-7B-Instruct-v0.2-GGUF	7B	0.649	0.380	0.707	0.461
Mistral-7B-Instruct-v0.2	7B	0.655	0.375	0.705	0.468
Mixtral_7Bx2_MoE	14B	0.747	0.518	0.764	0.475
Mixtral-8x7B-Instruct-v0.1	46.7B	0.723	0.526	0.745	0.552
Nous-Hermes-2-Mixtral-8x7B-DPO	46.7B	0.741	0.607	0.766	0.614
Nous-Hermes-2-SOLAR-10.7B	10.7B	0.725	0.592	0.722	0.588
SOLAR-10.7B-Instruct-v1.0	10.7B	0.737	0.438	0.747	0.381
SauerkrautLM-SOLAR-Instruct	10.7B	0.733	0.418	0.752	0.368
Sakura-SOLAR-Instruct-DPO-v2	10.7B	0.733	0.426	0.745	0.357

Table 2: The performance of different-sized LLMs on the validation set. The competition includes two tracks: model-agnostic track and model-aware track. For each track, both prediction accuracy(*acc*) and Spearman’s Rho value(*rho*) are provided.

reduces the number of trainable parameters. Despite this reduction, the fine-tuned models maintain comparable training results to those of the full parameter models. The best checkpoint model files are selected from the validation set during this fine-tuning process.

4.3 Ensemble Learning

We also propose an ensemble learning approach for performance improvement, utilizing fusion strategies at both the model level and the inference level.

Model fusion. MergeKit¹ is a toolkit designed for merging trained language models. We carefully selected a few high-accuracy models and utilized MergeKit to perform model fusion using the SLERP (Shoemake, 1985), TIES (Yadav et al., 2023) and linear (Wortsman et al., 2022) methods. Traditionally, model merging often resorts to weight averaging which, although straightforward, might not always capture the intricate features of the models being merged. The SLERP technique addresses this limitation, producing a blended model with characteristics smoothly interpolated from both parent models, ensuring the resultant model captures the essence of both its parents. Meanwhile, the TIES method is proposed to resolve interference issues by resetting parameters, resolving sign conflicts, and merging only compatible parameters. TIES outperforms many existing methods across diverse settings, emphasizing the importance of addressing interference in model merging for enhanced performance and versatility.

¹<https://github.com/arcee-ai/mergekit>

Model Voting. In addition to the model-level fusion, we also explored fusion at the probability level of model generation, which can be understood as a form of model voting. We selected another group of highly accurate candidate models and performed linear fusion at the probability level. Specifically, we calculate the weighted summation of the probability values on "existing hallucination" predicted by different candidate models for different tasks. By tuning the linear weight combination, we are able to determine the optimal combination of weights for each task. Finally, combining different tasks together yields the final fusion result. In this way, we implement weighted voting of models at the inference result level.

5 Results and Analysis

In this section, we will present a series of experiments to illustrate the effectiveness of our method.

5.1 Baseline

To begin with, our initial step entails presenting the basic performance of LLMs of varying sizes on the validation set. Subsequently, we will delve into an analysis of the LLMs’ capabilities in detecting hallucinations in the given task.

Throughout the experiments, we ensure the generation hyperparameters remain consistent across all LLMs. Additionally, the instruction utilized for detecting hallucinations is sourced from the official *participant_kit*. This version of the instruction is referred to as the "naive instruction" and can be located in Appendix A for reference.

Table 2 illustrates our evaluation of LLMs from both *Mistral* and *SOLAR* families, considering vary-

few-shot	inst.	agnostic_acc	aware_acc
2-shot	naive	0.745	0.774
	ours	0.770	0.806
4-shot	naive	0.762	0.772
	ours	0.782	0.806
6-shot	naive	0.764	0.774
	ours	0.772	0.804
8-shot	naive	0.762	0.772
	ours	0.774	0.804

Table 3: Our proposed instruction exhibits overall superior accuracy compared to the naive version on the validation set. We applied few-shot prompting in all of the aforementioned experiments. The term "inst." stands for "instruction".

ing sizes and variants, on the validation set. In general, larger models tend to yield better results. For instance, within the *Mistral*-family, the accuracy and Spearman’s Rho value of the 7B model are comparatively lower than those of larger models, a trend observed in both the model-agnostic and model-aware tracks. Furthermore, LLMs of the same size exhibit diverse results in hallucination detection tasks owing to distinct fine-tuning methods employed. This observation holds true in our experiments with the *SOLAR*-family, emphasizing the impact of fine-tuning on performance.

There is a noteworthy observation in Table 2. The medium-sized 14B *Mixtral_7Bx2_MoE* model achieves comparable accuracy to the larger-sized 46.7B models in both the model-agnostic and model-aware tracks. This suggests that the fine-tuning approach and training corpus of the 14B model are well-suited for the hallucination detection task. Furthermore, the 14B model outperforms the 46.7B model in terms of inference speed and training cost. As a result, in the subsequent section, we will further enhance the effectiveness of the 14B model through prompt engineering.

5.2 Performance Improvement

In this section, our focus is on enhancing the accuracy of the 14B *Mixtral_7Bx2_MoE* model through prompt engineering methods.

Few-shot prompting. A few-shot prompting approach is applied by randomly selecting an equal number of positive and negative samples as demonstrations for task definition modeling (DM), machine translation (MT), and paraphrase generation

few-shot	CoT	agnostic_acc	aware_acc
2-shot	w/o	0.770	0.806
	with	0.770	0.792
4-shot	w/o	0.782	0.806
	with	0.766	0.796
6-shot	w/o	0.772	0.804
	with	0.774	0.806
8-shot	w/o	0.774	0.804
	with	0.792	0.804

Table 4: CoT demonstrates an improved capability in hallucination detection when provided with a larger number of demonstrations in few-shot prompting and utilizing our proposed instruction. The results are on the validation set.

(PG). In the experimental setup, we use 2, 4, 6, and 8 examples for few-shot prompting on the *Mixtral_7Bx2_MoE* model, while keeping the generation hyperparameters consistent with the experiments in Table 2. The accuracy of the few-shot prompting strategy is shown with the *inst.=naive* setting in Table 3, where we observe that experiments with 4, 6, and 8 shots perform better than the zero-shot baseline(acc is 0.747 in Table 2) in the model-agnostic track, and all the few-shot settings experiments achieve better results than the zero-shot baseline(acc is 0.764 in Table 2) in the model-aware track.

Optimizing the instruction. As discussed in Section 4.1, the naive instruction provided by the competition organizers has some limitations. To overcome these limitations, we enhanced the instructions by incorporating task-specific background knowledge and multidimensional information, taking into account the unique characteristics of each task. The improved instructions, as demonstrated with the *inst.=ours* setting in Table 3, yield better performance compared to using the initial naive instruction. Notably, in the 2-shot setting, both tracks exhibited an improvement of over 2 percentage points by leveraging our proposed instructions.

Chain of thought prompting. We adopt the CoT approach, after generating reasons for the presence or absence of hallucinations in the trial data. The experimental results of CoT are presented in Table 4, which indicates that CoT exhibits higher efficacy in the few-shot scenario when there are more demonstrations accessible.

Model Name	Model Size	Model-agnostic track		Model-aware track	
		<i>acc</i>	<i>rho</i>	<i>acc</i>	<i>rho</i>
Mistral-7B-Instruct-v0.2	7B	0.806	0.708	0.790	0.699
Nous-Hermes-2-SOLAR-10.7B	10.7B	0.764	0.690	0.806	0.714
SOLAR-10.7B-Instruct-v1.0	10.7B	0.772	0.703	0.810	0.717
Mistral-7B-Instruct-v0.2-2x7B-MoE	12.8B	0.790	0.725	0.814	0.698
Mixtral_7Bx2_MoE	14B	0.780	0.675	0.796	0.657
<i>raw</i> : Mixtral_7Bx2_MoE	14B	0.792	0.707	0.804	0.690

Table 5: The performance of weakly-supervised fine-tuning on the validation set. Smaller LLMs perform better than the 14B supervisor. The last line indicates the performance of Mixtral_7Bx2_MoE without SFT.

Method	agnostic_acc
Mistral-7B-Instruct-v0.2-sft-v1	0.804
Mistral-7B-Instruct-v0.2-sft-v2	0.798
Mistral-7B-Instruct-v0.2-sft-v3	0.808
Linear-merged model	0.814
SLERP-merged model	0.814
TIES-merged model	0.814

Table 6: Model fusion results for the model-agnostic track on the validation set. The merged models outperform any individual model in terms of accuracy.

5.3 Weakly-supervised Fine-tuning

As mentioned earlier, we enhance the accuracy of hallucination detection by selecting the best baseline model and incorporating additional prompt engineering techniques. Building upon this, we leverage weak supervision by labeling the unlabeled training data for training. Subsequently, the LLMs are fine-tuned using the generated labels to further augment the capability of hallucination detection.

Generating weak supervision for training data. Utilizing the 14B *Mixtral_7Bx2_MoE* model as a foundation, we incorporate few-shot prompting, our proposed instruction, and the CoT strategy to create a supervision model known as the ‘8-shot’ setting, as mentioned in Table 4. This approach is applied to hallucination detection across 60,000 datapoints from both the model-agnostic and model-aware tracks, ensuring a balanced distribution of categories. Additionally, we introduce multiple optimal models, as discussed in Section 4, to ensure consistent inference across multiple models and maintain inference consistency within the same model but with different inference parameters.

Fine-tuning LLMs. The experimental results of weakly-supervised fine-tuning are presented in Ta-

Method	agnostic_acc
Mistral-7B-Instruct-v0.2-sft-v4	0.810
Mistral-7B-Instruct-v0.2-sft-v5	0.812
Mistral-7B-Instruct-v0.2-sft-v6	0.812
Mistral-7B-Instruct-v0.2-sft-v7	0.814
voting result	0.834

Table 7: Model Voting results for the model-agnostic track on the validation set. The voted results outperform any individual model in terms of accuracy.

ble 5, which demonstrates that smaller models can effectively learn from the weak supervision provided by the 14B model. In some cases, these smaller models even outperform the 14B model in terms of accuracy. Notably, the 14B model fails to surpass the performance of equivalently-sized supervisor even when multiple hyper-parameter settings are employed. A comparison between lora training and full-parameter training reveals that the lora-style training yields superior results. Further details can be found in Appendix C.

5.4 Ensemble Learning

In addition to fine-tuning LLMs with weak supervision labels as mentioned in the previous section, we first combine different model checkpoints by the MergeKit tool and then perform model voting strategy to enhance performance.

Model fusion. We implement model merging using different modes of MergeKit, *i.e.*, SLERP, TIES and Linear, in our study. Taking the accuracy optimization of the model-agnostic track as an example, let’s begin by selecting three highly capable candidate models. These models are fine-tuned versions of the 7B *Mistral-7B-Instruct-v0.2*, and they are different checkpoints from the same training task. The detailed training setting can be

found in Appendix B.2. By utilizing model fusion techniques, we can achieve a maximum accuracy of 0.814 with the newly merged models. For a detailed overview of the experiments on the model-agnostic track, refer to Table 6.

Model Voting. We also validate the effectiveness of weighted voting at the inference result level. We select the top-performing models from the weakly supervised fine-tuning and model fusion phases. By calculating weighted predictions based on their predicted probabilities, we infer the presence of hallucinations. Table 7 presents the details of the voting experiments on the model-agnostic track. Using the same method, we achieve an accuracy of 0.818 on the model-aware track as well. The SFT models are merged models from different training experimental setups, and the detailed training parameters can be found in Appendix B.3.

Method	agnostic_acc	aware_acc
baseline	0.697	0.745
GPT-4	0.741	0.756
our method	0.836	0.805

Table 8: Comparison of methods on the test set.

We compared the baseline provided by competition organizers, GPT-4, and our proposed method on the test set in Table 8. It is evident that our proposed method outperforms other methods, showcasing a significant enhancement in performance.

6 Conclusion

In this paper, we present a unified system for hallucination detection with LLMs when there is no labeled dataset, which wins the 2nd place with an accuracy score of 0.836 in the model-agnostic track and the 4th place with an accuracy score of 0.8053 in the model-aware track. To begin with, we generate high-quality weakly-supervised dataset by using large-sized LLMs with prompt engineering and few-shot learning. Then we perform weakly-supervised fine-tuning based on the constructed dataset with different LLMs. Our experiments yield several noteworthy findings:

(1) The quality of the weakly-supervised dataset we construct has a direct impact on the performance of the models in this task. To ensure high-quality training data, we employ multiple large LLMs in the construction process.

(2) Relatively small LLMs can deliver competitive performance in this task when trained on the

constructed dataset. However, the performance of small LLMs drops dramatically without fine-tuning.

(3) Using the *MergeKit* tool for model fusion proves to be an effective technique in boosting the performance of hallucination detection.

(4) Employing the model voting method leads to improved performance compared to using a single model alone.

References

- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, et al. 2023. A survey on evaluation of large language models. *ACM Transactions on Intelligent Systems and Technology*.
- Yukang Chen, Shengju Qian, Haotian Tang, Xin Lai, Zhijian Liu, Song Han, and Jiaya Jia. 2023. Longlora: Efficient fine-tuning of long-context large language models. *arXiv preprint arXiv:2309.12307*.
- Ze Chen, Kangxu Wang, Zijian Cai, Jiewen Zheng, Jiarong He, Max Gao, and Jason Zhang. 2022. *Using deep mixture-of-experts to detect word meaning shift for TempoWiC*. In *Proceedings of the The First Workshop on Ever Evolving NLP (EvoNLP)*, pages 7–11, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Tianyu Cui, Yanling Wang, Chuanpu Fu, Yong Xiao, Sijia Li, Xinhao Deng, Yunpeng Liu, Qinglin Zhang, Ziyi Qiu, Peiyang Li, et al. 2024. Risk taxonomy, mitigation, and assessment benchmarks of large language model systems. *arXiv preprint arXiv:2401.05778*.
- Xibin Dong, Zhiwen Yu, Wenming Cao, Yifan Shi, and Qianli Ma. 2020. A survey on ensemble learning. *Frontiers of Computer Science*, 14:241–258.
- Muhammad Usman Hadi, Rizwan Qureshi, Abbas Shah, Muhammad Irfan, Anas Zafar, Muhammad Bilal Shaikh, Naveed Akhtar, Jia Wu, Seyedali Mirjalili, et al. 2023. A survey on large language models: Applications, challenges, limitations, and practical usage. *Authorea Preprints*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. 1991. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12):1–38.

- Bartosz Krawczyk, Leandro L Minku, João Gama, Jerzy Stefanowski, and Michał Woźniak. 2017. Ensemble learning for data stream analysis: A survey. *Information Fusion*, 37:132–156.
- Timothee Mickus, Elaine Zosa, Raúl Vázquez, Teemu Vahtola, Jörg Tiedemann, Vincent Segonne, Alessandro Raganato, and Marianna Apidianaki. 2024. SemEval-2024 Task 6: SHROOM, a shared-task on hallucinations and related observable overgeneration mistakes. In *Proceedings of the 18th International Workshop on Semantic Evaluation (SemEval-2024)*, Mexico City, Mexico. Association for Computational Linguistics.
- Ken Shoemake. 1985. Animating rotation with quaternion curves. In *Proceedings of the 12th annual conference on Computer graphics and interactive techniques*, pages 245–254.
- Isaac Triguero, Daniel Molina, Javier Poyatos, Javier Del Ser, and Francisco Herrera. 2024. General purpose artificial intelligence systems (gpais): Properties, definition, taxonomy, societal implications and responsible governance. *Information Fusion*, 103:102135.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.
- Yiming Wang, Yu Lin, Xiaodong Zeng, and Guan-nan Zhang. 2023. Multilora: Democratizing lora for better multi-task learning. *arXiv preprint arXiv:2311.11501*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.
- Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, et al. 2022. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *International Conference on Machine Learning*, pages 23965–23998. PMLR.
- Prateek Yadav, Derek Tam, Leshem Choshen, Colin Raffel, and Mohit Bansal. 2023. Ties-merging: Resolving interference when merging models. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Yongquan Yang, Haijun Lv, and Ning Chen. 2023. A survey on ensemble learning under the era of deep learning. *Artificial Intelligence Review*, 56(6):5545–5589.
- Zhengmao Ye, Dengchun Li, Jingqi Tian, Tingfeng Lan, Jie Zuo, Lei Duan, Hui Lu, Yexi Jiang, Jian Sha, Ke Zhang, et al. 2023. Aspen: High-throughput lora fine-tuning of large language models with a single gpu. *arXiv preprint arXiv:2312.02515*.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223*.
- Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, et al. 2022. Least-to-most prompting enables complex reasoning in large language models. *arXiv preprint arXiv:2205.10625*.

A Instructions in Prompt Engineering

A.1 Naive version

Context: {Context}

Sentence: {Sentence}

Is the Sentence supported by the Context above?

Answer using ONLY yes or no:

A.2 Our proposed version

For the PG task, the instruction is as follows:

Given the following information related to Paraphrase Generation task:

Src: Source input sentence

Tgt: Paraphrase Generation standard answer

Hyp: Paraphrase Generation predicted answer

Please determine whether hyp contains unexpected hallucinations based on src and tgt.

Src: {Src}

Tgt: {Tgt}

Hyp: {Hyp}

Is the Hyp supported by the Src and Tgt above?

Answer using ONLY yes or no:

For the MT task, the instruction is as follows:

Given the following information related to Machine Translation task:

Src: Source input sentence

Tgt: Machine Translation standard answer

Hyp: Machine Translation predicted answer

Please determine whether hyp contains unexpected hallucinations based on src and tgt.

Src: {Src}

Tgt: {Tgt}

Hyp: {Hyp}

Is the Hyp supported by the Src and Tgt above?

Answer using ONLY yes or no:

As for the DM task, the instruction is the same as the naive version.

B Training Experiment Setup

B.1 Constructed dataset

We constructed a total of **35,600** weakly supervised samples, ensuring consistency in inference across different LLMs as well as within the same LLM but with different parameter settings.

B.2 SFT models in Table 6

The Mistral-7B-Instruct-v0.2-sft-v1, v2, and v3 models are different checkpoint models obtained from the same training setup. These models were trained on a total of 35,600 weakly-supervised data points. The training process utilized a LoRA rank of 32, a learning rate of $3e^{-5}$, and a total of 5 epochs. The training task was executed using 4 A30 GPUs. Specifically, Mistral-7B-Instruct-v0.2-sft-v1, v2, and v3 models were saved at training steps 1000, 3000, and 4000, respectively.

B.3 SFT models in Table 7

The Mistral-7B-Instruct-v0.2-sft-v4, v5, v6, and v7 models are merged models obtained from different training setups. Each model is created by merging two checkpoints from the same setup. The v4 model was trained with a LoRA rank of 32, a learning rate of $1e^{-4}$, and a total of 5 epochs. The v5 model also had a LoRA rank of 32, a learning rate of $3e^{-5}$, and a total of 5 epochs. The v6 model had a higher LoRA rank of 48, a learning rate of $3e^{-5}$, and lasted for 5 epochs. Lastly, the v7 model had a LoRA rank of 48, a learning rate of $5e^{-5}$, and a total of 5 epochs. All of these models were trained on the constructed dataset.

C Lora training VS. Full training

Method	agnostic_acc	aware_acc
lora	0.806	0.790
full	0.58	0.52

Table 9: Comparison of different training methods based on Mistral-7B-Instruct-v0.2.