

# Improving Factuality in Large Language Models via Decoding-Time Hallucinatory and Truthful Comparators

Dingkang Yang<sup>1\*†</sup>, Dongling Xiao<sup>2\*</sup>, Jinjie Wei<sup>1</sup>, Mingcheng Li<sup>1</sup>, Zhaoyu Chen<sup>1</sup>,  
Ke Li<sup>3‡</sup>, Lihua Zhang<sup>1‡</sup>

<sup>1</sup>Academy for Engineering and Technology, Fudan University

<sup>2</sup>ByteDance

<sup>3</sup>Tencent YouTu Lab

{dkyang20, lihuazhang}@fudan.edu.cn, xdluestc@outlook.com, tristanli@tencent.com

## Abstract

Despite their remarkable capabilities, Large Language Models (LLMs) are prone to generate responses that contradict verifiable facts, *i.e.*, unfaithful hallucination content. Existing efforts generally focus on optimizing model parameters or editing semantic representations, which compromise the internal factual knowledge of target LLMs. In addition, hallucinations typically exhibit multifaceted patterns in downstream tasks, limiting the model’s holistic performance across tasks. In this paper, we propose a Comparator-driven Decoding-Time (CDT) framework to alleviate the response hallucination. Firstly, we construct hallucinatory and truthful comparators with multi-task fine-tuning samples. In this case, we present an instruction prototype-guided mixture of experts strategy to enhance the ability of the corresponding comparators to capture different hallucination or truthfulness patterns in distinct task instructions. CDT constrains next-token predictions to factuality-robust distributions by contrasting the logit differences between the target LLMs and these comparators. Systematic experiments on multiple downstream tasks show that our framework can significantly improve the model performance and response factuality.

**Code** — <https://github.com/ydk122024/CDT>

## Introduction

Recently, the emergence of advanced Large Language Models (LLMs) (Touvron et al. 2023; OpenAI 2022; Achiam et al. 2023; Jiang et al. 2023) spearheads an understanding paradigm shift in the Natural Language Processing (NLP) and demonstrates superior capabilities in a wide range of tasks, such as translation, summarization, and dialog generation (See, Liu, and Manning 2017; Moon et al. 2019). Despite the impressive performance, LLMs typically generate seemingly plausible but non-factual claims, *i.e.*, the hallucination dilemma. Hallucination content exhibits inconsistencies with real-world or user instructions, limiting the model credibility in realistic applications (Huang et al. 2023).

\*These authors contributed equally.

†Done during the internship and implies the Ph.D. period end.

‡Corresponding authors.

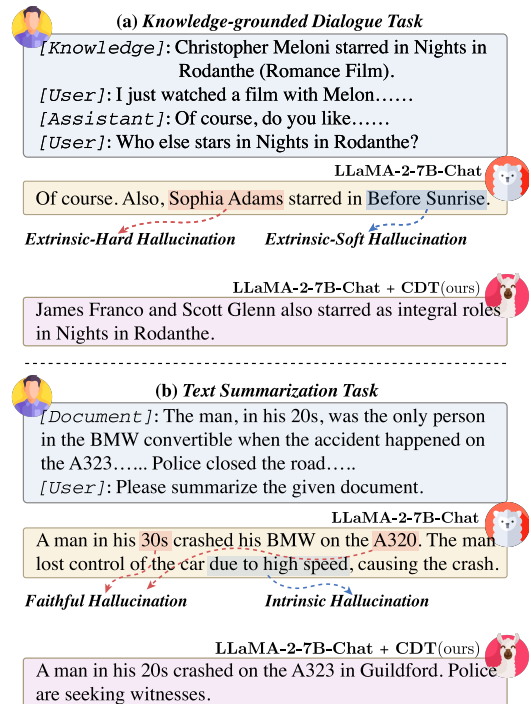


Figure 1: Illustration of multifaceted hallucination patterns generated by LLMs on different tasks (Das, Saha, and Srihari 2023; Zheng, Huang, and Chang 2023).

The hallucination reasons are thought to be complex, such as biased corpora, maximum likelihood training objective, and blind instruction following (Zhang et al. 2023b). All these factors potentially lead LLMs to assign non-zero probabilities to unfaithful token generation. Previous approaches have focused on two main aspects to provide solutions: optimizing model parameters and editing semantic representations. The former optimizes the model behavior directly from the response ranking via human preference alignment techniques (Ouyang et al. 2022; Bai et al. 2022; Tian et al. 2023), forcing the output to be more truthful. The Supervised Fine-tuning (SFT) is optionally used to mitigate task-

specific hallucinations (Jones et al. 2023). However, secondary training is labor-intensive, and unintended knowledge infusion inadvertently encourages target LLMs to capture spurious generation patterns. The latter works on regularizing and editing the hidden representations of LLMs to select the semantic spaces associated with correct statements (Burns et al. 2022; Chen et al. 2024; Li et al. 2024). Unfortunately, proactive representation activation may result in models hallucinating even knowing the correct semantics, interfering with internal factual knowledge (Zou et al. 2023). More importantly, current LLMs suffer from multifaceted hallucination patterns that cause serious performance bottlenecks. As shown in Figure 1, the model exhibits different types of hallucinations on each of the two tasks. There are even two hallucination patterns in the same response. For instance, LLaMA2-7B-Chat wrongly extracts information (*i.e.*, faithful hallucination) in the text summarization and produces false content that is not present in the original document (*i.e.*, intrinsic hallucination). Previous studies have not resolved the coupled hallucination dilemma well.

Motivated by the above observations, we propose an efficient Comparator-driven Decoding-Time (CDT) framework to improve response factuality. The core philosophy is to equip target LLMs with comparators modeling different generative attributes separately during the decoding process, using logit distribution integration to facilitate next-token prediction in the factuality-robust directions. We utilize explicit hallucinated/factual multi-task samples via Low-Rank Adaptation (LoRA)-based SFT to construct hallucinatory/truthful comparators based on the base models. In this case, we design a hallucination perturbation adversarial mechanism to strengthen the factual knowledge mastery of the truthful comparator and mitigate the potential overfitting phenomena. To assist target LLMs in eliminating the diverse hallucinations in cross-task generation, we present an instruction Prototype-guided Mixture of Experts (PME) strategy to empower the comparators with multi-task hallucination/truthfulness-aware capabilities. In the comparators, multiple LoRA adapters act as different experts to master different hallucination or truthfulness patterns through well-designed soft route gating. CDT is capable of accomplishing effective attribute control without compromising the internal knowledge of target LLMs, maintaining the generated fluency and diversity.

The main contributions are summarized as follows:

- Our framework can help target LLMs remove non-factual knowledge from the output space during the decoding process in a product-integrated manner, significantly improving the robustness and factuality of model responses.
- Our comparators are not limited to specific model structures and task types, in a plug-and-play form to guide the next-token prediction towards the hallucination-weak distribution. The proposed PME strategy offers the prospect and potential of eliminating hallucinations with multifaceted patterns in different tasks.
- Extensive experiments are conducted on multiple NLP benchmarks. Comprehensive analyses show the broad applicability and effectiveness of our framework.

## Related Work

**Hallucination Mitigation in LLMs.** Hallucination in LLMs means the generated content is not supported by verifiable real-world facts or user instructions (Huang et al. 2023; Zhang et al. 2023b), leading to conflicts with input queries (Rehman et al. 2023), contextual information (Shi et al. 2023; Wang et al. 2023), and faithful semantics (Bang et al. 2023; Hu et al. 2023). Recently, extensive works have been presented aiming at mitigating the hallucination interference in LLMs across different dimensions, including retrieval-augmented generation (Peng et al. 2023; Chern et al. 2023), high-quality data construction (Zhou et al. 2024; Li et al. 2023b), and synthetic task migration (Jones et al. 2023). Among these, mainstream efforts have mainly focused on two aspects: optimizing model parameters and editing semantic representations. The former argues that one of the most likely reasons for hallucinations is that the maximum-likelihood objective of language modeling leads LLMs to assign non-zero probabilities to unfaithful sentences outside the training data distribution (Chuang et al. 2023). To this end, these approaches promote harmless response generation through reinforcement learning from human feedback (Ouyang et al. 2022; Bai et al. 2022) and other preference alignment algorithms (Tian et al. 2023; Yang et al. 2024). For instance, (Tian et al. 2023) found that the Direct Preference Optimization (DPO) (Rafailov et al. 2024) assists the model in aligning facticity from ranking. The latter emphasizes activating attention heads or latent features into the truthful space through causal effects (Li et al. 2024; Chen et al. 2024; Burns et al. 2022; Zhang, Yu, and Feng 2024). For example, Truth Forest (Chen et al. 2024) utilized multidimensional orthogonal probing to mine hidden true representations to improve model truthfulness. Despite advances, current methods potentially undermine the inherent knowledge system of LLMs, causing performance bottlenecks on out-of-attention tasks. In contrast, our framework guides out-of-the-box target LLMs during inference to enhance prediction factuality by mitigating hallucinations while improving truthfulness.

**Decoding-Time Intervention.** Unlike secondary training for target LLMs, Decoding-time intervention was used earlier to enhance model fluency and coherence during inference by comparing expert models (Li et al. 2022). Previous studies found that controlling the output behavior at the decoding level contributed to avoiding toxic generated content (Liu et al. 2021) and improving logical reasoning (Wang and Zhou 2024). These findings offer the prospect of refining response factuality through rational comparator formalization (Zhang et al. 2023a; Shi et al. 2023; Chuang et al. 2023; Kai et al. 2024). For instance, ICD (Zhang et al. 2023a) finetuned a factually weak base model during induce-the-contrast decoding as a penalty to guide LLMs more faithfully. Nevertheless, these approaches usually fail to recognize multifaceted hallucination patterns in complex scenarios, leading to limitations and sensitivities in task types. In comparison, the proposed CDT demonstrates adaptability and scalability over distinct tasks by hallucination/truthfulness-aware comparators.

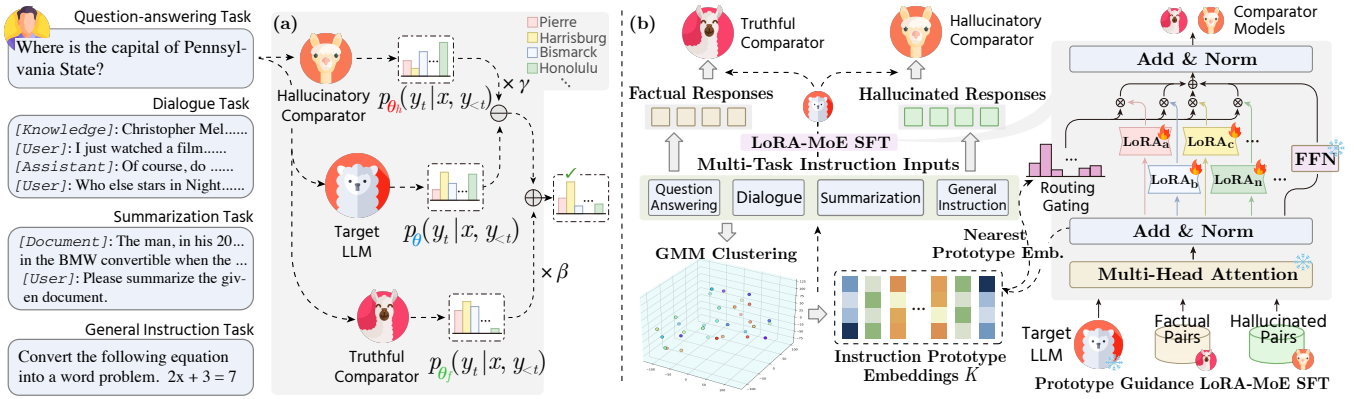


Figure 2: (a) Illustration of our Comparator-driven Decoding-Time (CDT) framework. The hallucinatory comparator gives higher weight to the incorrect *Honolulu* while the truthful comparator favors the factual *Harrisburg*, expressing the effective distribution control for the next-token prediction. The target LLM removes the hallucination through our CDT. (b) We construct hallucinatory/truthful comparators based on hallucinated/factual instruction pairs via the LoRA-based SFT procedures.

## Methodology

### Comparator-driven Decoding-Time Framework

Figure 2(a) shows the high-level illustration of our CDT framework. Given a target language model parameterized by  $\theta$ , the model generates a corresponding response  $\mathbf{y}$  based on a query instruction  $\mathbf{x}$  containing contextual information. The response is typically sampled from the probability distribution conditioned on  $\mathbf{x}$ . The autoregressive decoding process can be formulated as follows:

$$y_t \sim p_\theta(y_t | \mathbf{x}, \mathbf{y}_{<t}) \propto \exp \logit_\theta(y_t | \mathbf{x}, \mathbf{y}_{<t}), \quad (1)$$

where  $y_t$  denotes the predicted token at time step  $t$ .  $\mathbf{y}_{<t}$  contains the previously generated tokens across time step  $t - 1$ . With the maximum-likelihood-based training objective, the original decoding suffers from unavoidable hallucinations due to the model spontaneously assigning non-zero probabilities to unfaithful content (Chuang et al. 2023). From Figure 2(a), the model potentially relies on spurious correlations captured from the training data distribution to give false predictions. To mitigate such issues, CDT aims to manipulate the next-token predictions of the target LLM into factuality-robust distributions by combining a hallucinatory comparator  $\theta_h$ , which generates text with non-factual attributes, and a truthful comparator  $\theta_f$ , which generates text with factual attributes. These comparators come from fine-tuning the base version of the target LLM via explicit hallucinated/factual samples. The proposed CDT decoding is formulated as:

$$y_t \sim p_{cdt}(y_t | \mathbf{x}, \mathbf{y}_{<t}) \propto p_\theta(y_t | \mathbf{x}, \mathbf{y}_{<t}) \frac{p_{\theta_f}(y_t | \mathbf{x}, \mathbf{y}_{<t})^\beta}{p_{\theta_h}(y_t | \mathbf{x}, \mathbf{y}_{<t})^\gamma}, \quad (2)$$

where  $\beta$  and  $\gamma$  are coefficients that control the degree of distribution modification. Intuitively, the generated token will have a high probability when it has a high probability under  $p_{\theta_f}$  and a low probability under  $p_{\theta_h}$ , facilitating factual prediction amplification and untruthful prediction trivialization.

In practice, we note that the output of the hallucinatory comparator is not always erroneous since it can still maintain

basic linguistic common sense and grammar (Li et al. 2022). This means that indiscriminate trivialization potentially penalizes valuable aspects, leading to implausible tokens. To address this issue, we introduce an adaptive plausibility constraint to penalize a subset  $\mathcal{V}_{sub}$  of the output vocabulary  $\mathcal{V}$  that correlates with the confidence level of the output distribution of the target LLMs:

$$\begin{aligned} \mathcal{V}_{sub}(\mathbf{y}_{<t}) &= \{y_t \in \mathcal{V} : \\ p_\theta(y_t | \mathbf{x}, \mathbf{y}_{<t}) &\geq \delta \max_w p_\theta(w | \mathbf{x}, \mathbf{y}_{<t})\}, \\ p_{cdt}(y_t | \mathbf{x}, \mathbf{y}_{<t}) &= 0, \text{ if } y_t \notin \mathcal{V}_{sub}(\mathbf{y}_{<t}), \end{aligned} \quad (3)$$

where  $\delta \in [0, 1]$  is a coefficient that controls the distribution truncation, with larger values indicating that only high-probability tokens are retained. The constraint ensures the response integrity of our framework while excluding the generation of untrustworthy tokens.

### Hallucinatory/Truthful Comparator Formalization

To build behavior-controlled comparators, we activate hallucinatory/faithful mastery of the base models by the LoRA-based supervised fine-tuning (SFT) (Hu et al. 2021). The SFT dataset  $\mathcal{D}$  contains multi-task instructions to empower the comparators to adapt the target LLM to improve factuality in different downstream tasks, including question-answering, knowledge-grounded dialogue, text summarization, and general tasks. Specifically, given any input instruction  $\mathbf{x} = (x_1, x_2, x_3, \dots) \in \mathcal{D}$  and corresponding hallucinated/factual responses  $\mathbf{y}^{h/f} = (y_1^{h/f}, y_2^{h/f}, y_3^{h/f}, \dots) \in \mathcal{D}$ , the hallucinatory comparator can be optimized by the following fine-tuning objective:

$$\mathcal{L}(\theta, \mathcal{D}) = \mathbb{E}_{(\mathbf{x}, \mathbf{y}^h) \sim \mathcal{D}} \left[ -\sum_{t=1}^{|\mathbf{y}^h|} \log p(y_t^h | \mathbf{x}, \mathbf{y}_{<t}^h; \theta) \right]. \quad (4)$$

Although the truthful comparator can be optimized with the same objective, we find that directly performing SFT via instruction pairs  $(\mathbf{x}, \mathbf{y}^f)$  results in overfitting, leading to the mundane faithfulness of the generated content. Inspired by

robustness learning (Miyato, Dai, and Goodfellow 2022), we introduce a hallucination perturbation adversarial training mechanism to improve the training efficiency of the truthful comparator. Specifically, we first backpropagate using the hallucinated responses to obtain the word embedding layer gradients. Then, we compute an adversarial perturbation based on the gradient and merge it into the original embedding to obtain the hallucination-aware adversarial samples. The perturbation is expressed as follows:

$$\mathcal{P} = \varepsilon \cdot \nabla_{\mathbf{x}} \mathcal{L}(\theta, \mathbf{x}, \mathbf{y}^h) / \left\| \nabla_{\mathbf{x}} \mathcal{L}(\theta, \mathbf{x}, \mathbf{y}^h) \right\|_2, \quad (5)$$

where  $\varepsilon$  and  $\mathcal{L}(\cdot)$  represent the sign and loss functions, respectively. The adversarial samples are utilized to backpropagate secondarily through the factual responses to obtain the cumulative adversarial gradients for model parameter updating. Our strategy forces the model to progressively penalize hallucinated attributes from the perturbation during SFT, promoting more effective and factual semantic learning of the truthful comparator.

### Instruction Prototype-guided Mixture of Experts

In the LoRA-based SFT phase of Figure 2(b), we present an instruction prototype-guided mixture of experts strategy to enhance the hallucination/truthfulness-aware capability of comparators. Considering that instructions are the foundation for serving target LLMs to generate informative responses, the core philosophy of our strategy lies in learning prototype representations of multi-task instruction data with different probability density distributions via Gaussian Mixture Models (Reynolds et al. 2009). Depending on the corresponding hallucinated or faithful response outputs, the prototype representations can guide multiple LoRA adapters acting as experts to master different hallucination or truthfulness patterns in supervised optimization, facilitating hallucinatory and truthful comparators to assist the target LLMs in generating factual content among diverse tasks during decoding time. Specifically, given the multi-task instruction set  $\{\mathbf{x}_i\}_{i=1}^M$  from the SFT dataset  $\mathcal{D}$ , the prototype representation acquisition is denoted as follows:

$$\mathcal{C}_{[\mathbf{x}]} = \frac{\sum_{i=1}^M \gamma_{i,k} \mathbf{h}_{x,i}}{\sum_{i=1}^M \gamma_{i,k}}, \quad (6)$$

where  $\mathbf{h}_{x,i}$  is the instruction feature corresponding to  $\mathbf{x}_i$ , which is extracted from the hidden state of the last layer in the target LLM and compressed and refined by principal component analysis for efficient computation.  $\gamma_{i,k}$  is the probability that each  $\mathbf{h}_{x,i}$  belongs to the  $k$ -th Gaussian component, which is formulated as follows:

$$\gamma_{i,k} = \frac{\pi_k \mathcal{N}(\mathbf{h}_{x,i} \mid \mu_k, \Sigma_k)}{\sum_{j=1}^K \pi_j \mathcal{N}(\mathbf{h}_{x,i} \mid \mu_j, \Sigma_j)}, \quad (7)$$

where  $\pi_k$ ,  $\mu_k$ , and  $\Sigma_k$  are the mixture weight, the mean vector, and the covariance matrix of the  $k$ -th component, respectively.  $\mathcal{N}(\cdot)$  denotes the probability density.

After that, the learned instruction prototypes are incorporated into the routing gating  $G(\cdot)$  to activate the appropri-

ate experts to capture specific hallucination/truthfulness patterns in any input  $\mathbf{x}$ . The gating vector is formulated as:

$$G(\mathbf{x}) = \mathcal{S} \left( \left( (1 - \mu) \odot \mathbf{x} + \mu \odot \mathcal{C}_{[\mathbf{x}]} \mathbf{W}_g \right) \mathbf{W}_g + \mathfrak{R}(\varphi(\mathbf{x} \mathbf{W}_n)) \right), \quad (8)$$

where  $\mathcal{S}(\cdot)$  and  $\mu$  stand for the Softmax activation and trade-off coefficient, respectively. In practice, we determine the index  $k_x \leftarrow \arg \min_k d_{\mathcal{M}}(\mathbf{x}, \mathcal{C}_{[\mathbf{x}]})$  of the prototype  $\mathcal{C}_{[\mathbf{x}]}$  to which  $\mathbf{x}$  belongs via the Euclidean distance (Danielsson 1980). The distance is suitable for handling high-dimensional representations from multi-task SFT data with complex distributions since it is computationally efficient.  $\mathfrak{R}(\varphi(\mathbf{x} \mathbf{W}_n))$  is an additional noise term designed to regularize the load balancing among experts during the training process, where  $\mathfrak{R}(\cdot)$  and  $\varphi(\cdot)$  represent the standard normal distribution sampling and Softplus function, respectively.  $\mathbf{W}_g$ ,  $\mathbf{W}_g$  and  $\mathbf{W}_n$  are learnable parameters. The forward process of obtaining the parameterized output  $\mathbf{z}$  of all experts by integrating the soft route gating is expressed as follows:

$$\mathbf{z} = \frac{\alpha}{r} \sum_{n=1}^N G(\mathbf{x})_n E_n(\mathbf{x}), \quad (9)$$

where  $r$  is the rank value and  $\alpha$  is a constant hyperparameter.  $E_n(\cdot)$  represents the  $n$ -th LoRA expert.

## Experiments

### Datasets and Evaluation Metrics

We construct extensive experiments on six datasets to evaluate different methods across multiple tasks. Specifically, **KINGHT-Judge** is proposed to verify the model’s ability to recognize hallucinations in the knowledge-driven dialog task. The dataset contains the dialog contexts and factual responses from 616 samples sampled from the testing set of OpenDialKG (Moon et al. 2019), as well as hallucinated responses generated by GPT-4 (Achiam et al. 2023). During evaluations, the models are asked to classify factual/hallucinatory responses randomly sampled from each sample based on contexts to specify whether or not they contain hallucination information. Similarly, we select 600 responses from HaluEval (Li et al. 2023a) that are half-annotated as normal and hallucinatory against general user queries from Alpaca (Taori et al. 2023). These instructions constitute **Alpaca-Judge** to evaluate performance on general tasks. The evaluation metrics consist of Accuracy, Precision, Recall, and F1 Score. **TruthfulQA** (Lin, Hilton, and Evans 2022) contains 817 questions across 38 categories to assess response truthfulness. We employ the multiple-choice task and measure the accuracy of LLMs in selecting answers from multiple correct/incorrect options via MC1, MC2, and MC3 metrics. For text summarization, we extract 1,000 samples with harmless and safe content from vanilla **XSUM** (Narayan, Cohen, and Lapata 2018) to ensure the models perform rational summarization with given documents. We use the ROUGE-1/2/L (Lin 2004) to measure the generation quality. FactKB (Feng et al. 2023) and BERTScore (Zhang et al. 2019) are utilized to evaluate the factual consistency of responses. For open-ended generation tasks, we use factual responses in KNIGHT-Judge as

Methods	KNIGHT-Judge				Alpaca-Judge			
	Accuracy (%)	Precision (%)	Recall (%)	F1 Score (%)	Accuracy (%)	Precision (%)	Recall (%)	F1 Score (%)
Baseline	45.13	43.90	62.50	51.58	43.68	45.05	81.33	59.27
SFT	43.18	39.59	58.21	47.13	41.58	43.20	76.95	57.66
FPO	44.42	41.67	61.04	49.53	43.72	45.46	82.63	59.14
CD (Li et al. 2022)	46.10	46.53	<b>94.24</b>	61.92	44.52	46.13	84.06	57.35
ITI (Li et al. 2024)	45.59	46.13	67.23	53.44	44.25	45.66	78.30	57.93
SH2 (Kai et al. 2024)	46.76	42.00	63.64	50.61	44.33	45.85	79.45	58.14
DoLa (Chuang et al. 2023)	46.62	41.81	65.79	52.25	44.67	46.27	84.94	58.91
ICD (Zhang et al. 2023a)	47.24	47.76	90.32	63.93	45.00	47.38	85.33	60.15
<b>CDT (Ours)</b>	<b>50.16</b>	<b>50.60</b>	93.95	<b>65.77</b>	<b>46.83</b>	<b>48.26</b>	<b>88.00</b>	<b>62.34</b>

Table 1: Comparison results on the KNIGHT-Judge and Alpaca-Judge datasets.

Methods	TruthfulQA		
	MC1 (%)	MC2 (%)	MC3 (%)
Baseline	37.62	54.60	28.12
SFT	35.02	52.82	27.40
FPO	37.98	55.34	29.65
CD (Li et al. 2022)	28.15	54.87	29.75
ITI (Li et al. 2024)	37.01	54.66	27.82
SH2 (Li et al. 2024)	33.90	57.07	29.79
DoLa (Chuang et al. 2023)	32.97	60.84	29.50
ICD (Zhang et al. 2023a)	46.32	69.08	41.25
<b>CDT (Ours)</b>	<b>50.74</b>	<b>75.96</b>	<b>50.97</b>

Table 2: Comparison results on the TruthfulQA dataset.

Ground Truths (GTs) to construct **KNIGHT-Gen** to measure the generated content for given instructions. Moreover, **Alpaca-Gen** extracts 600 samples with normal responses as GTs to explore general generation following Alpaca-Judge. The FactKB, BERTScore, and DAE (Goyal and Durrett 2020) are used to assess the predicted factuality.

### Implementation Details

The main experiments are built on the LLaMA-2 family in order to provide a fair comparison with existing methods. The model training is accomplished through the PyTorch platform using eight Nvidia A800 GPUs. To fine-tune the LLaMA2-base-driven comparators, the multi-task SFT dataset  $\mathcal{D}$  is collected from HaluEval with 30K instructions having factual and hallucinated responses across NLP tasks. We also mix 10K held-out samples sampled from Alpaca and annotate the corresponding hallucinated responses. The default number of instruction prototypes and LoRA experts are 32 and 4, respectively.

### Comparison with State-of-the-Art Methods

**Model Zoo.** We compare our CDT and reproducible methods. **Baseline:** LLaMA2-7b-Chat model (Touvron et al. 2023). **SFT:** The model is fine-tuned by factual multi-task instructions in  $\mathcal{D}$ . **FPO:** The model is optimized by DPO utilizing preference pairs from factual/hallucinated responses in  $\mathcal{D}$ . Recent state-of-the-art (SOTA) works include CD (Li et al. 2022), ITI (Li et al. 2024), SH2 (Kai et al. 2024), DoLa (Chuang et al. 2023), and ICD (Zhang et al. 2023a).

**Results on the KNIGHT-Judge&Alpaca-Judge.** Table 1 provides comparison results on the two datasets for hallucination recognition. We find that CDT achieves the best results on the vast majority of metrics, suggesting accurate judgments and understanding against potential hallucinations

Methods	R-1	R-2	R-L	FactKB	BERTScore
Baseline	17.22	3.55	13.29	38.79	82.50
SFT	17.58	3.61	13.75	41.85	81.03
FPO	18.09	3.47	14.18	43.66	82.75
CD (Li et al. 2022)	19.55	3.77	14.61	67.30	83.31
ITI (Li et al. 2024)	17.94	3.68	15.45	57.94	82.87
SH2 (Kai et al. 2024)	19.75	3.83	14.73	71.22	83.92
DoLa (Chuang et al. 2023)	18.60	3.73	14.58	65.13	84.20
ICD (Zhang et al. 2023a)	20.29	4.11	<b>16.85</b>	70.26	84.06
<b>CDT (Ours)</b>	<b>21.11</b>	<b>4.52</b>	16.33	<b>73.66</b>	<b>86.44</b>

Table 3: Comparison results on the XSUM dataset. ‘‘R-1/2/L’’ means the ROUGE-1/2/L metrics.

in the text. For instance, CDT brings 14.19% and 3.07% gains in terms of F1 scores on KINGHT-Judge and Alpaca-Judge for the baseline. Conversely, the poor performance of SFT and FPO indicates that secondary training induces the models to learn spurious behavior cloning (Schulman 2023) rather than factual patterns. Our framework significantly outperforms prior SOTA ICD on KINGHT-Judge by 2.81% average improvement, implying effective recognition of different extrinsic hallucinations.

**Results on the TruthfulQA.** As shown in Table 2, our framework outperforms the decoding time methods ITI and ICD in all metrics on TruthfulQA. This implies that penalizing the coupled hallucinations only during decoding without considering the factual attributes that enhance the target LLMs is sub-optimal and inadequate. Compared to ICD, the 9.5%/9.9%/23.6% relative gains regarding MC1/MC2/MC3 on multi-category questions proves that CDT removes potential comprehension and factuality hallucinations (Zheng, Huang, and Chang 2023) in the question-answering task by giving correct answers with higher probabilities through the instruction prototype-guided mixture of experts strategy.

**Results on the XSUM.** Table 3 demonstrates the text summarization capabilities of different models. According to the ROUGE metric, we find that the proposed framework generates content that covers the critical parts of the reference text more reasonably. CDT assists the target baseline in improving 34.87% and 3.94% on FactKB and BERTScore, effectively enhancing response factuality and quality. In contrast, the improvement of the representation editing technique ITI is incremental due to enhancing truthfulness by intervening in the attention heads, which causes uninformative output and performance bottlenecks.

**Results on the KNIGHT-Gen&Alpaca-Gen.** We report the results of open-ended generation in Table 4. CDT im-

Methods	KNIGHT-Gen				Alpaca-Gen			
	ROUGE-L	BERTScore	FactKB	DAE	ROUGE-L	BERTScore	FactKB	DAE
Baseline	9.85	81.33	70.54	40.19	19.67	81.25	76.43	50.62
SFT	9.92	81.14	71.43	39.76	20.54	81.23	77.25	50.11
FPO	9.27	81.77	71.66	41.05	20.36	82.69	77.93	51.53
CD (Li et al. 2022)	10.80	82.35	72.27	42.83	22.42	82.07	78.20	52.00
ITI (Li et al. 2024)	11.26	82.59	74.43	42.46	23.56	83.31	80.17	52.29
SH2 (Kai et al. 2024)	11.13	83.57	72.68	44.10	<b>24.35</b>	83.99	83.02	55.43
DoLa (Chuang et al. 2023)	12.05	83.14	73.56	44.08	23.20	83.25	79.94	53.75
ICD (Zhang et al. 2023a)	11.53	84.03	75.32	45.02	23.88	84.87	83.52	55.16
<b>CDT (Ours)</b>	<b>12.16</b>	<b>85.47</b>	<b>77.42</b>	<b>47.39</b>	23.90	<b>86.06</b>	<b>85.82</b>	<b>56.50</b>

Table 4: Comparison results on the KNIGHT-Gen and Alpaca-Gen datasets.

Design	KNIGHT-Judge	TruthfulQA	XSUM	Alpaca-Gen
	F1 Score	MC1	BERTScore	DAE
<b>Full Framework</b>	<b>65.77</b>	<b>50.74</b>	<b>86.44</b>	<b>56.50</b>
<b>Importance of Comparators</b>				
w/o HC	62.05	39.53	84.97	52.12
+ Chat Version	64.03	48.36	85.15	55.74
w/o TC	54.61	48.38	83.22	55.09
+ Chat Version	65.75	50.67	<b>86.69</b>	<b>56.58</b>
<b>Effectiveness of PME</b>				
w/o PME	62.94	48.82	83.36	54.75
w/o Prototypes	64.06	49.14	85.09	55.43
<b>Necessity of PAT</b>				
w/o PAT	64.65	49.57	85.64	55.96
+ PGD	65.18	49.83	85.91	56.20
<b>Impact of SFT Data Proportion</b>				
w/o QA Data	64.40	48.12	85.13	55.08
w/o Dialogue Data	62.53	48.93	85.38	54.71
w/o Summ. Data	64.05	49.27	84.45	54.66
w/o General Data	65.19	50.30	85.86	54.24

Table 5: Ablation study results on the different datasets. ‘‘HC’’ means the hallucinatory comparator. ‘‘TC’’ means the truthful comparator. ‘‘+’’ is the replacement operation. ‘‘Summ.’’ means the summarization. ‘‘w/o’’ is the without.

proves response faithfulness on knowledge-grounded dialog and general-purpose tasks, consistently achieving the best results among three fact metrics. Compared to other contrastive decoding methods, CDT has two advantages. First, unlike CD and DoLa, which emphasize probability differences between strong and weak models/layers, CDT enhances factuality in a more integrated way by enabling precise control of factuality-robust distributions of next-token predictions through comparators with different properties. Also, our framework excels at capturing multiple hallucination types across tasks, leading to more effective hallucination mitigation during decoding time. The 6.39% average gain across the fact metrics for the baseline confirms this.

### Ablation Studies

In Table 5, we perform systematic ablation studies to investigate the rationality of different designs in CDT.

**Importance of Different Comparators.** We observe that the two comparators provide complementary contributions as performance deteriorates across different datasets when they are removed. Interestingly, the hallucinatory compar-

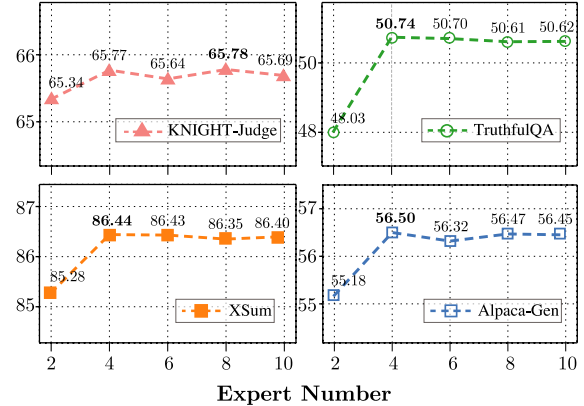


Figure 3: We show the effect of the number of experts on the performance of different tasks.

tor excels in improving the discriminative and open-ended generation tasks, while the truthful comparator helps the target LLM to be more truthful in recognizing hallucinations and summarization demands. When replacing the base model with the chat version, the truthful comparator has a minor improvement while it is ineffective for the hallucinatory comparator. This makes sense because the chat model undergoes RLHF with honesty as a key attribute, narrowing the gains from CDT contrasting logit differences.

**Effectiveness of Instruction Prototype-guided Mixture of Experts (PME).** A single LoRA is used for the SFT process when there is no PME strategy in the comparators. The overall gain drops, which is more significant for the dialogue and summarization tasks. A plausible explanation is that PME facilitates the model to better capture sophisticated hallucinations in scenarios that require contextual understanding. Moreover, instruction prototypes can guide LoRA experts to learn different hallucination/truthfulness patterns, bringing consistent gains across multiple tasks.

**Necessity of Perturbation Adversarial Training (PAT).** We observe that the PAT mechanism plays an essential role because it prevents potential truthfulness learning trivialization and overfitting dilemmas in the truthful comparator training. The intuitive justification is reflected in consistent gains across different datasets. Additionally, a PGD-based adversarial candidate (Madry et al. 2017) is introduced to explore performance changes. The sub-optimal results from




Figure 4: The extensibility analysis of our framework across different LLMs through the multiple-choice task on TruthfulQA.

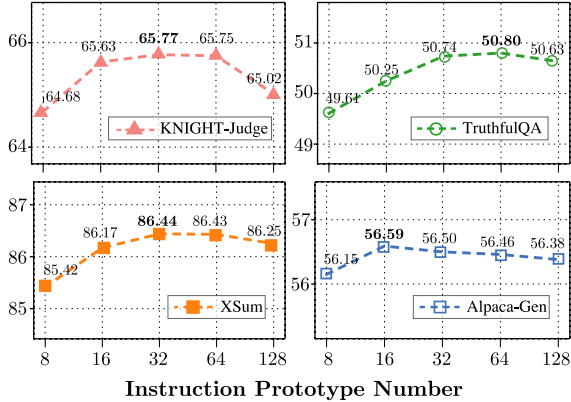


Figure 5: We show the effect of the number of instruction prototypes on the performance of different tasks.

the candidate further verify the necessity of our mechanism.

**Impact of SFT Data Proportion.** Different data for training the comparators are investigated by removing instruction pairs from each of the four parts. We discover that each type of SFT data is beneficial in reinforcing the benefits of the proposed framework. This trend suggests that sufficient multi-task data can effectively activate the comparators’ ability to model different generation attributes.

**Effect of Expert Number.** The number of LoRA experts determines the comparators’ knowledge capacity to adapt to different tasks. In Figure 3, when only 2 experts are used, the model gives poorer results on XSUM and Alpaca-Gen. We conjecture that generative tasks typically require increased expert capacity to capture multifaceted hallucinations embedded in contextual information. The 4 experts achieve a trade-off between performance and training overhead. Excessively introducing experts can cause gain convergence.

**Effect of Instruction Prototype Number.** Instruction prototypes guide the semantically relevant routing process to balance LoRA expert utilization. As shown in Figure 5, fewer prototypes (*i.e.*, 8 and 16) perform poorly on the evaluation tasks. This stems from the inability of experts to master sufficient hallucination/truthfulness patterns, leading to performance bottlenecks. There is a considerable performance improvement on different tasks when we increase the number of prototypes, where 32 is the most appropriate. Conversely, too many prototypes are redundant as they exceed the pattern distribution boundaries (Gou et al. 2023).

Methods	[Prompt]: How long do diamonds last?
Baseline	Graphite generally lasts between <i>10 and 15 million years</i> before it <b>begins to weather and eventually decompose</b> .
ICD	<b>As the fossilized tree sap</b> , diamonds last between <i>5 and 8 billion years</i> .
CDT (Ours)	Typically, diamonds last between 1 and 4 billion years.

Table 6: Qualitative analysis regarding the generated responses of different models to the same prompt. Underlined, *italic* and **bold** texts represent faithful, factuality, and inference hallucinations, respectively.

### Extensibility Analysis

As a model-agnostic framework, we apply the CDT to 8 advanced LLMs spanning different sizes of the same family as well as different families of models. Figure 4 illustrates the results of the multiple-choice task on the hallucination perception dataset TruthfulQA. In the LLaMA families, we find an upward trend in the benefits from CDT with increasing model version and size. The average improvement across the 6 vanilla baselines reaches 14.25%. Moreover, CDT provides absolute gains of 12.95% and 15.6% for Baichuan2-7B-Chat (Yang et al. 2023) and Mistral-7B-Instruct (Jiang et al. 2023), respectively, significantly enhancing the corresponding factuality.

### Qualitative Analysis

In Table 6, we show the comparison results of generated contents for the same prompt by different models on the question-answering task. Despite improvements in ICD, the hallucinations remained. The claim that diamonds are inferred to be “fossilized tree sap” reflects the inference hallucination. “5 and 8 billion years” that conflict with world knowledge reflects the factuality hallucination. In contrast, CDT effectively corrects the multifaceted misinformation, thereby demonstrating the effectiveness of our framework.

### Conclusion

In this paper, we introduce CDT, a decoding-time framework to efficiently remove the multifaceted hallucination dilemma of target LLMs via hallucinatory and truthful Comparators. CDT facilitates more factual model responses in multiple downstream tasks by controlling the next-token prediction.

## Acknowledgments

This work is supported in part by the National Key R&D Program of China (2021ZD0113502).

## References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Bai, Y.; Jones, A.; Ndousse, K.; Askell, A.; Chen, A.; Das-Sarma, N.; Drain, D.; Fort, S.; Ganguli, D.; Henighan, T.; et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- Bang, Y.; Cahyawijaya, S.; Lee, N.; Dai, W.; Su, D.; Willie, B.; Lovenia, H.; Ji, Z.; Yu, T.; Chung, W.; et al. 2023. A multitask, multilingual, multimodal evaluation of chatgpt on reasoning, hallucination, and interactivity. *arXiv preprint arXiv:2302.04023*.
- Burns, C.; Ye, H.; Klein, D.; and Steinhart, J. 2022. Discovering latent knowledge in language models without supervision. *arXiv preprint arXiv:2212.03827*.
- Chen, Z.; Sun, X.; Jiao, X.; Lian, F.; Kang, Z.; Wang, D.; and Xu, C. 2024. Truth forest: Toward multi-scale truthfulness in large language models through intervention without tuning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 20967–20974.
- Chern, I.; Chern, S.; Chen, S.; Yuan, W.; Feng, K.; Zhou, C.; He, J.; Neubig, G.; Liu, P.; et al. 2023. FacTool: Factuality Detection in Generative AI—A Tool Augmented Framework for Multi-Task and Multi-Domain Scenarios. *arXiv preprint arXiv:2307.13528*.
- Chuang, Y.-S.; Xie, Y.; Luo, H.; Kim, Y.; Glass, J.; and He, P. 2023. Dola: Decoding by contrasting layers improves factuality in large language models. *arXiv preprint arXiv:2309.03883*.
- Danielsson, P.-E. 1980. Euclidean distance mapping. *Computer Graphics and image processing*, 14(3): 227–248.
- Das, S.; Saha, S.; and Srihari, R. K. 2023. Diving deep into modes of fact hallucinations in dialogue systems. *arXiv preprint arXiv:2301.04449*.
- Feng, S.; Balachandran, V.; Bai, Y.; and Tsvetkov, Y. 2023. Factkb: Generalizable factuality evaluation using language models enhanced with factual knowledge. *arXiv preprint arXiv:2305.08281*.
- Gou, Y.; Liu, Z.; Chen, K.; Hong, L.; Xu, H.; Li, A.; Yeung, D.-Y.; Kwok, J. T.; and Zhang, Y. 2023. Mixture of cluster-conditional lora experts for vision-language instruction tuning. *arXiv preprint arXiv:2312.12379*.
- Goyal, T.; and Durrett, G. 2020. Evaluating Factuality in Generation with Dependency-level Entailment. In *Findings of the Association for Computational Linguistics: EMNLP 2020*.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Hu, X.; Chen, J.; Li, X.; Guo, Y.; Wen, L.; Yu, P. S.; and Guo, Z. 2023. Do Large Language Models Know about Facts? *arXiv preprint arXiv:2310.05177*.
- Huang, L.; Yu, W.; Ma, W.; Zhong, W.; Feng, Z.; Wang, H.; Chen, Q.; Peng, W.; Feng, X.; Qin, B.; et al. 2023. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *arXiv preprint arXiv:2311.05232*.
- Jiang, A. Q.; Sablayrolles, A.; Mensch, A.; Bamford, C.; Chaplot, D. S.; Casas, D. d. l.; Bressand, F.; Lengyel, G.; Lample, G.; Saulnier, L.; et al. 2023. Mistral 7B. *arXiv preprint arXiv:2310.06825*.
- Jones, E.; Palangi, H.; Simões, C.; Chandrasekaran, V.; Mukherjee, S.; Mitra, A.; Awadallah, A.; and Kamar, E. 2023. Teaching language models to hallucinate less with synthetic tasks. *arXiv preprint arXiv:2310.06827*.
- Kai, J.; Zhang, T.; Hu, H.; and Lin, Z. 2024. SH2: Self-Highlighted Hesitation Helps You Decode More Truthfully. *arXiv preprint arXiv:2401.05930*.
- Li, J.; Cheng, X.; Zhao, W. X.; Nie, J.-Y.; and Wen, J.-R. 2023a. HaluEval: A Large-Scale Hallucination Evaluation Benchmark for Large Language Models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 6449–6464.
- Li, K.; Patel, O.; Viégas, F.; Pfister, H.; and Wattenberg, M. 2024. Inference-time intervention: Eliciting truthful answers from a language model. *Advances in Neural Information Processing Systems*, 36.
- Li, X. L.; Holtzman, A.; Fried, D.; Liang, P.; Eisner, J.; Hashimoto, T.; Zettlemoyer, L.; and Lewis, M. 2022. Contrastive decoding: Open-ended text generation as optimization. *arXiv preprint arXiv:2210.15097*.
- Li, Y.; Bubeck, S.; Eldan, R.; Del Giorno, A.; Gunasekar, S.; and Lee, Y. T. 2023b. Textbooks are all you need ii: phi-1.5 technical report. *arXiv preprint arXiv:2309.05463*.
- Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 74–81.
- Lin, S.; Hilton, J.; and Evans, O. 2022. TruthfulQA: Measuring How Models Mimic Human Falsehoods. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 3214–3252.
- Liu, A.; Sap, M.; Lu, X.; Swayamdipta, S.; Bhagavatula, C.; Smith, N. A.; and Choi, Y. 2021. DExperts: Decoding-time controlled text generation with experts and anti-experts. *arXiv preprint arXiv:2105.03023*.
- Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; and Vladu, A. 2017. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*.
- Miyato, T.; Dai, A. M.; and Goodfellow, I. 2022. Adversarial Training Methods for Semi-Supervised Text Classification. In *International Conference on Learning Representations*.
- Moon, S.; Shah, P.; Kumar, A.; and Subba, R. 2019. Opendialkg: Explainable conversational reasoning with attention-based walks over knowledge graphs. 845–854.
- Narayan, S.; Cohen, S. B.; and Lapata, M. 2018. Don’t give me the details, just the summary! topic-aware convolutional

- neural networks for extreme summarization. *arXiv preprint arXiv:1808.08745*.
- OpenAI. 2022. Introducing ChatGPT. <https://openai.com/blog/chatgpt>.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems (NeurIPS)*, 35: 27730–27744.
- Peng, B.; Galley, M.; He, P.; Cheng, H.; Xie, Y.; Hu, Y.; Huang, Q.; Liden, L.; Yu, Z.; Chen, W.; et al. 2023. Check your facts and try again: Improving large language models with external knowledge and automated feedback. *arXiv preprint arXiv:2302.12813*.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2024. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36.
- Rehman, T.; Mandal, R.; Agarwal, A.; and Sanyal, D. K. 2023. Hallucination reduction in long input text summarization. In *International Conference on Security, Surveillance and Artificial Intelligence (ICSSAI-2023)*, 307–316. CRC Press.
- Reynolds, D. A.; et al. 2009. Gaussian mixture models. *Encyclopedia of biometrics*, 741(659-663).
- Schulman, J. 2023. Reinforcement learning from human feedback: Progress and challenges. [https://www.youtube.com/watch?v=hhiLw5Q\\_UFg](https://www.youtube.com/watch?v=hhiLw5Q_UFg).
- See, A.; Liu, P. J.; and Manning, C. D. 2017. Get To The Point: Summarization with Pointer-Generator Networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1073–1083.
- Shi, W.; Han, X.; Lewis, M.; Tsvetkov, Y.; Zettlemoyer, L.; and Yih, S. W.-t. 2023. Trusting your evidence: Hallucinate less with context-aware decoding. *arXiv preprint arXiv:2305.14739*.
- Taori, R.; Gulrajani, I.; Zhang, T.; Dubois, Y.; Li, X.; Guestrin, C.; Liang, P.; and Hashimoto, T. B. 2023. Stanford alpaca: An instruction-following llama model.
- Tian, K.; Mitchell, E.; Yao, H.; Manning, C. D.; and Finn, C. 2023. Fine-tuning language models for factuality. *arXiv preprint arXiv:2311.08401*.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Wang, C.; Liu, X.; Yue, Y.; Tang, X.; Zhang, T.; Jiayang, C.; Yao, Y.; Gao, W.; Hu, X.; Qi, Z.; et al. 2023. Survey on factuality in large language models: Knowledge, retrieval and domain-specificity. *arXiv preprint arXiv:2310.07521*.
- Wang, X.; and Zhou, D. 2024. Chain-of-thought reasoning without prompting. *arXiv preprint arXiv:2402.10200*.
- Yang, A.; Xiao, B.; Wang, B.; Zhang, B.; Bian, C.; Yin, C.; Lv, C.; Pan, D.; Wang, D.; Yan, D.; et al. 2023. Baichuan 2: Open large-scale language models. *arXiv preprint arXiv:2309.10305*.
- Yang, D.; Wei, J.; Xiao, D.; Wang, S.; Wu, T.; Li, G.; Li, M.; Wang, S.; Chen, J.; Jiang, Y.; et al. 2024. PediatricsGPT: Large Language Models as Chinese Medical Assistants for Pediatric Applications. *arXiv preprint arXiv:2405.19266*.
- Zhang, S.; Yu, T.; and Feng, Y. 2024. Truthx: Alleviating hallucinations by editing large language models in truthful space. *arXiv preprint arXiv:2402.17811*.
- Zhang, T.; Kishore, V.; Wu, F.; Weinberger, K. Q.; and Artzi, Y. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.
- Zhang, Y.; Cui, L.; Bi, W.; and Shi, S. 2023a. Alleviating hallucinations of large language models through induced hallucinations. *arXiv preprint arXiv:2312.15710*.
- Zhang, Y.; Li, Y.; Cui, L.; Cai, D.; Liu, L.; Fu, T.; Huang, X.; Zhao, E.; Zhang, Y.; Chen, Y.; et al. 2023b. Siren’s song in the AI ocean: a survey on hallucination in large language models. *arXiv preprint arXiv:2309.01219*.
- Zheng, S.; Huang, J.; and Chang, K. C.-C. 2023. Why Does ChatGPT Fall Short in Providing Truthful Answers? *arXiv preprint arXiv:2304.10513*.
- Zhou, C.; Liu, P.; Xu, P.; Iyer, S.; Sun, J.; Mao, Y.; Ma, X.; Efrat, A.; Yu, P.; Yu, L.; et al. 2024. Lima: Less is more for alignment. *Advances in Neural Information Processing Systems*, 36.
- Zou, A.; Phan, L.; Chen, S.; Campbell, J.; Guo, P.; Ren, R.; Pan, A.; Yin, X.; Mazeika, M.; Dombrowski, A.-K.; et al. 2023. Representation engineering: A top-down approach to ai transparency. *arXiv preprint arXiv:2310.01405*.