

Whispers that Shake Foundations: Analyzing and Mitigating False Premise Hallucinations in Large Language Models

Hongbang Yuan^{1,2}, Pengfei Cao^{1,2}, Zhuoran Jin^{1,2}, Yubo Chen^{1,2,*}
Daojian Zeng³, Kang Liu^{1,2}, Jun Zhao^{1,2}

¹The Key Laboratory of Cognition and Decision Intelligence for Complex Systems,
Institute of Automation, Chinese Academy of Sciences, Beijing, China

²School of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing, China

³Hunan Normal University, Changsha, China

{hongbang.yuan, pengfei.cao, zhuoran.jin, yubo.chen, kliu, jzhao}@nlpr.ia.ac.cn

Abstract

Large Language Models (LLMs) have shown impressive capabilities but still suffer from the issue of hallucinations. A significant type of this issue is the false premise hallucination, which we define as the phenomenon when LLMs generate hallucinated text when confronted with false premise questions. In this paper, we perform a comprehensive analysis of the false premise hallucination and elucidate its internal working mechanism: a small subset of attention heads (which we designate as false premise heads) disturb the knowledge extraction process, leading to the occurrence of false premise hallucination. Based on our analysis, we propose **FAITH** (False premise Attention head constraining for miTigating Hallucinations), a novel and effective method to mitigate false premise hallucinations. It constrains the false premise attention heads during the model inference process. Impressively, extensive experiments demonstrate that constraining only approximately 1% of the attention heads in the model yields a notable increase of nearly 20% of model performance.

1 Introduction

Large language models (LLMs) have shown impressive capabilities (Wei et al., 2022; Xu et al., 2023; Li et al., 2023c) and achieved remarkable success in many tasks (Bubeck et al., 2023; Bang et al., 2023; Jiao et al., 2023; Sun et al., 2024). However, they often generate texts that are seemingly plausible but deviate from factual knowledge sources (Zhang et al., 2023), which is a severe problem known as **hallucination**.

To address this, many studies focus on detecting (Manakul et al., 2023; Min et al., 2023) and mitigating hallucinations (Trivedi et al., 2023; Gou et al., 2023; Yuan et al., 2024a). However, few of them pay attention to a particular type of hallucination: **False Premise Hallucination**. We define

*Corresponding author

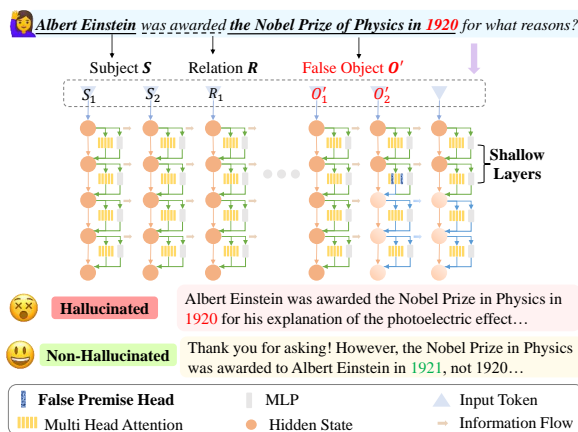


Figure 1: Illustration of the false premise hallucination. The question contains the false premise that “*Albert Einstein was awarded The Nobel Prize of Physics in 1920*” whereas in fact he was awarded the prize in 1921. We find that the presence of false premise attention heads contributes to the hallucinated response. Our method can effectively mitigate the false premise hallucination.

it as the phenomenon in which LLMs generate hallucinated texts in response to the false premise question. False premise questions are questions that contain falsely assumed facts that are not directly stated but are likely to be believed by the asker (Yu et al., 2023b; Kim et al., 2023). For these questions, LLMs tend to respond directly without explicitly verifying their plausibility despite the corresponding factual knowledge can be recalled straightforwardly. For example, as shown in Figure 1, the user query on the top contains a false premise: (*Albert Einstein, was awarded, The Nobel Prize of Physics in 1920*), denoted as (*subject, relation, false object*). LLMs are able to generate the correct time “1921” when directly queried about the award time but produce hallucinated time “1920” in response to the false premise question.

The exploration of false premise hallucination is exceptionally significant and valuable. While false premise questions are pervasive on the Internet and

users are highly likely to pose these questions when interacting with the LLMs, LLMs are prone to generate hallucinated texts when confronted with such questions. For example, about 25% of the questions on the discussion website Reddit contain false premises (Yu et al., 2023b; Fan et al., 2019). According to our statistics, Llama-2-13b achieves an accuracy of 100% on our collected 4004 direct-asking questions but drops to only 24.3% on the corresponding false premise questions. However, the analysis of false premise hallucination is non-trivial. Intuitively, LLMs generate hallucinated texts due to a lack of pertinent factual knowledge (Zhang et al., 2023; Huang et al., 2023a). But false premise hallucination introduces a more intricate scenario, wherein LLMs still generate hallucinated texts even when the corresponding factual knowledge is already stored in their parameters.

In this paper, we conduct a comprehensive analysis of false premise hallucination and elucidate its internal working mechanism. Firstly, prior to the analysis, we propose an automatic dataset construction pipeline for the evaluation of false premise hallucination and create two representative and easy-to-evaluate datasets based on it. Then, we investigate the external manifestation of false premise hallucination and observe that LLMs exhibit more inherent uncertainty when generating hallucinated answers. Subsequently, to reveal the source of the uncertainty, we delve into the internal information flow during the hallucination occurrences. We discover that knowledge about the subject stored in model parameters is disturbed in the shallow layers of the model, particularly around the false object mentioned in the question. Furthermore, as many studies (Meng et al., 2022; Yuksekogul et al., 2023) indicate that self-attention layers transfer the factual knowledge stored in the Multi-Layer Perception (MLP) layers, we explore the internals of the self-attention layers and investigate the influence of each individual attention head on specific pieces of factual knowledge. We find out that a small set of attention heads consistently exert great influence on the factual knowledge across almost all the samples and we name them as **False Premise Heads**. As depicted in Figure 1, the false premise heads predominantly reside in the shallow layers, functioning around the false object mentioned in the question. Experimental results demonstrate that the presence of false premise heads disturb the extraction of the factual knowledge about the subject, leading to false premise hallucinations.

Based on our analysis, we propose a novel method termed **FAITH** (**F**alse premise **A**ttention head **C**onstra**I**ning for **m**i**T**igating **H**allucinations) to mitigate hallucinations. It localizes the false premise attention heads for a group of false premise questions and subsequently constrains the function of these attention heads during the model inference process. Extensive experiments demonstrate the effectiveness of our method comparing with the baseline methods.

Our primary contributions can be summarized as follows:

- We propose an automatic dataset construction pipeline for the evaluation of false premise hallucination and create two representative and easy-to-evaluate datasets to facilitate analysis.
- We conduct an in-depth analysis of the false premise hallucination from the surface to the internals of LLMs and elucidate its internal working mechanism by revealing the presence of false premise attention heads.
- We propose FAITH, a novel method to mitigate false premise hallucinations based on our in-depth analysis. Impressively, extensive experiments demonstrate that constraining only approximately 1% of all the attention heads in the model yields a notable increase of nearly 20% in accuracy, which is highly effective.¹

2 Background

In this section, we briefly describe the transformer architecture (Vaswani et al., 2017) in autoregressive, decoder-only language models from the perspective of residual stream (Elhage et al., 2021).

Given the context $\{t_1, t_2, \dots, t_N\}$ consisting of N tokens, the transformer architecture starts with a combination of token embeddings and position embeddings $x_0 \in \mathbb{R}^{N \times d}$, where d is the model dimension. It marks the start of the residual stream, with a series of residual layers that read from the stream and write back their processed results. Each residual layer is comprised of a self-attention layer and a MLP layer. The information update of each residual layer can be expressed as:

$$x_l = x_{l-1} + a_l + m_l$$

where x_l is the hidden state after the l -th layer, a_l is the output of the self-attention layer and m_l is

¹Our code and datasets is available at <https://github.com/HongbangYuan/FAITH>.

the output of the MLP layer. More specifically, the calculation of the MLP layer is:

$$m_l = f(x_{l-1}K^T)V$$

where $K, V \in \mathbb{R}^{dm \times d}$ and f is a non-linear function. The self-attention contribution is

$$a_l = \sum_{h=1}^H A^h (x_{l-1} W_V^h) W_O^h$$

$$A^h = \text{softmax}\left(\frac{(x_{l-1} W_Q^h)(x_{l-1} W_K^h)^T}{\sqrt{d_h/H}}\right)$$

where $W_K^h, W_Q^h, W_V^h \in \mathbb{R}^{d \times d_h}$, $W_O^h \in \mathbb{R}^{d_h \times d}$ are the parameter matrices, H is the number of attention heads, $d_h = d/H$ is the hidden dimension of each head and $A \in \mathbb{R}^{N \times N}$ is a lower triangular attention pattern matrix, showing the interaction between tokens in different layers. After L residual layers, a layer norm is applied and then an unembedding matrix $W_U \in \mathbb{R}^{d \times V}$ projects the hidden state x_L to logits, where V is the length of the vocabulary.

3 Dataset Construction

In this section, we describe our proposed automatic dataset construction pipeline for the evaluation of false premise hallucination and provide the details of our constructed dataset.

To prevent the memorization of the question (Carlini et al., 2023; Ramakrishna et al., 2023) and facilitate the incorporation of the evolving new knowledge, we propose an automatic dataset construction pipeline, which can be divided into the following three stages:

(1) **Triple Selection** We select a set of factual triples using WikiData². We assess whether the factual triple (s, r, o) is stored in the model parameters by asking a question with only subject s and relation r . We retain the triple only if the object o is present in the answer. (2) **Triple Corruption** We replace the object o in the original triple (s, r, o) with an incorrect entity o' to obtain the corrupted triple (s, r, o') . For example, we transform the original triple “(Albert Einstein, was awarded, the Nobel Prize of Physics in 1921)” into the corrupted triple “(Albert Einstein, was awarded, the Nobel Prize of Physics in 1920)”. (3) **Question Construction** We construct a false premise question by filling a predefined question template with

²<https://query.wikidata.org/>

Dataset	Knowledge	Selected	Questions
Prize	7B	950	237
	13B	457	1828
Movie	7B	5509	1001
	13B	1001	4004

Table 1: Details of our datasets. The columns denote the dataset name, number of knowledge triples, number of selected knowledge triples for each model, number of constructed questions for each model, respectively. As we curate factual knowledge from specific models, two versions of each dataset are given.

the previous corrupted triple (s, r, o') . For example, we define one of the question templates as “<person> was awarded <false prize> for what specific reason?” and insert the corrupted triple “<person>, was awarded, <false prize>” into the template.

Following this automatic construction pipeline, we construct two datasets, namely Prize and Movie. We choose the variants of Llama-2-Chat (Touvron et al., 2023) with 7B and 13B parameters as the triple selector. For each model within the dataset, different versions are constructed as varying numbers of knowledge triples are selected. Table 1 provides the details of our datasets, while concrete question templates are presented in Appendix A. If the datasets were to be used with other models, researchers could readily follow our proposed pipeline and construct their own versions of the datasets tailored to their specific models.

4 Hallucination Analyze

In this section, we conduct a comprehensive analysis of false premise hallucinations from the surface and delves deeper into the model step by step.

4.1 Analysis of Model Uncertainty

In this part, we quantitatively investigate model uncertainty, which is a significant external feature of false premise hallucination and can be utilized to detect the hallucination occurrence. We hypothesize that model exhibits more inherent uncertainty when generating hallucinated answers. We design a model uncertainty measurement metric that allows the various linguistic forms of the true answer and experimental results validate our hypothesis.

Uncertainty Measurement We utilize three metrics to measure the uncertainty of the model when confronted with a question. The former two out of the three metrics are straightforward while the third one is specifically designed for our task. Suppose

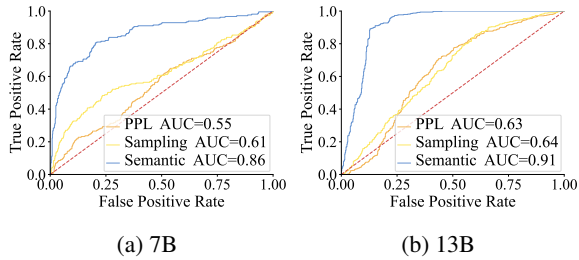


Figure 2: The Receiver Operating Characteristic Curve on the Movie dataset. The perfect AUC score is 1 while the random AUC score is 0.5.

that we have a question q and a sequence of model answer $T = (t_0, t_1, \dots, t_N)$, the three metrics are described below:

(1) **PPL-Based** We simply calculate the negative log likelihood of the model answer: $U_1(q) = U_1(T) = -\frac{1}{|N|} \sum_{i=1}^N \text{log} p_{\theta}(x_i | x_{<i})$, where $\text{log} p_{\theta}(x_i | x_{<i})$ is the log likelihood of the i -th token based on the previous tokens $x_{<i}$.

(2) **Sampling-Based** To fully leverage the uncertainty in model parameters (Huang et al., 2023b), we generate multiple answer sequences T_1, T_2, \dots, T_k for one question and calculate the average log likelihood across all sequences: $U_2(q) = \frac{1}{k} \sum_{i=1}^k U_1(T_i)$.

(3) **Semantic-Based** Inspired by the incorporation of linguistic invariances in model uncertainty estimation (Kuhn et al., 2023), we separately treat the correct and incorrect answers among the multiple generated answer sequences. We consider all the correct answers as a unified semantic set and each incorrect answer as a discrete semantic set. Then we calculate the uncertainty over these semantic sets. Concretely, suppose that there are K_1 incorrect answers and K_2 correct answers in the K generated answers, we calculate the uncertainty as follows: $U_3(q) = -\frac{1}{K_1+1} [\sum_{K_1} U_1(T_k) + \log \sum_{K_2} \exp U_1(T_k)]$.

Experiment We conduct experiments on the Movie dataset using Llama-2-7b-chat (denoted as 7B) and Llama-2-13b-chat (denoted as 13B). To evaluate the model uncertainty during hallucination occurrences, the calculated uncertainty scores are used for the binary classification task, aimed at determining the occurrence of hallucinations for each false premise question. We use the Area Under the Receiver Operating Characteristic (AUC) to assess the effectiveness of the uncertainty score. The higher the scores, the greater the correlation between the uncertainty metric and the occurrence

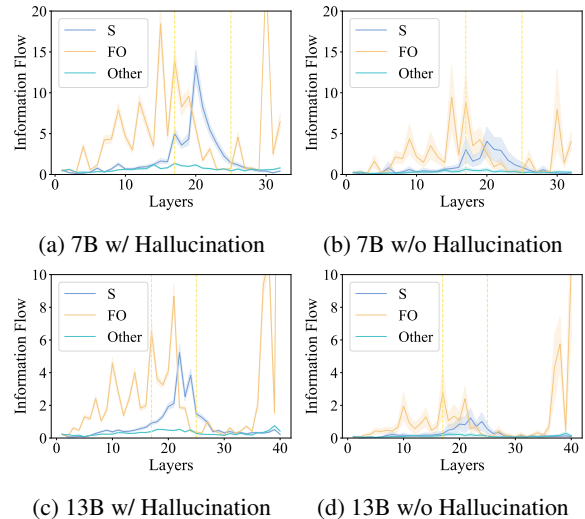


Figure 3: Information flow from various parts of the question to the final logit across distinct layers on hallucinated and non-hallucinated samples.

of hallucinations.

Results and Analysis The Receiver Operating Characteristic Curve (ROC Curve) curves and the AUC scores are shown in Figure 2. From the results, we draw the following observations: (1) Our semantic-based uncertainty metric score is far more effective than the other two methods. It can be further employed in the prediction of the occurrence of false premise hallucinations without relying on an external knowledge base. (2) We observe a strong correlation (over 0.9 for the 13B model) between the occurrence of hallucinations and model uncertainty. This verifies our hypothesis that models exhibit more inherent uncertainty when generating hallucinated answers.

4.2 Analysis of Internal Information Flow

To explore the source of the uncertainty, in this section, we delve into the internal information flow of LLMs when generating hallucinated answers. We study how the information flow from different parts of the false premise question in the fill-in-the-blank task. Experimental results demonstrate that the knowledge about the subject is disturbed in the shallow layers of the model, particularly around the false object mentioned in the question.

Knowledge Assessment Task The most significant feature of false premise hallucination is that the factual knowledge about the subject can be recalled directly. Therefore, we design a fill-in-the-blank task to evaluate how the knowledge stored

in the model parameters is affected by the question. Concretely, for a question containing the false triple (s, r, o') , LLMs are required to complete the following cloze query: “<Question> According to my knowledge, the object linking from subject s via relation r is $_$ ”. We posit that this knowledge assessment task is correlated with the original question answering task. Intuitively, if LLMs generate hallucinated answers to false premise questions, they are highly likely to fail in completing this cloze query correctly.

Attribution Score We aim to discover how the information flows from the tokens in the question to the final prediction logit in the knowledge assessment task. Since gradients and the attention itself can be blended together to acquire a better performance (Zhao et al., 2024), we use the element-wise product version (Wang et al., 2023) to calculate the attribution score for each token:

$$S_l = \left| \sum_{h=1}^H A_l^h \odot \frac{\partial L}{\partial A_l^h} \right|$$

where A_l^h is the attention pattern matrix described in Section 2 and L is the loss on the token prediction task. The attribution score matrix S_l on layer l is a $N \times N$ matrix (N is the length of the prompt). We partition the question into three parts: subject part (denote as S), false object part (denote as FO) and other part (denoted as $other$). The information flow from these parts is consequently defined as:

$$S_{S_l} = \frac{1}{|N_{sub}|} \sum_{t=1}^{N_{sub}} S_{N,i_t}$$

$$S_{FO_l} = \frac{1}{|N_{fobj}|} \sum_{t=1}^{N_{fobj}} S_{N,i_t}$$

$$S_{other_l} = \frac{1}{|N_{other}|} \sum_{t=1}^{N_{other}} S_{N,i_t}$$

where N_{sub} is the number of tokens of subject, N_{fobj} is the number of tokens of the false object and N_{other} is the number of other tokens.

Results and Analysis. We illustrate the information flow from various parts of the question to the final logit across distinct layers on hallucinated and non-hallucinated samples in the Prize dataset, as shown in Figure 3. By observing the figures, it’s evident that the information flow across the layers can be roughly divided into three pieces. (1)

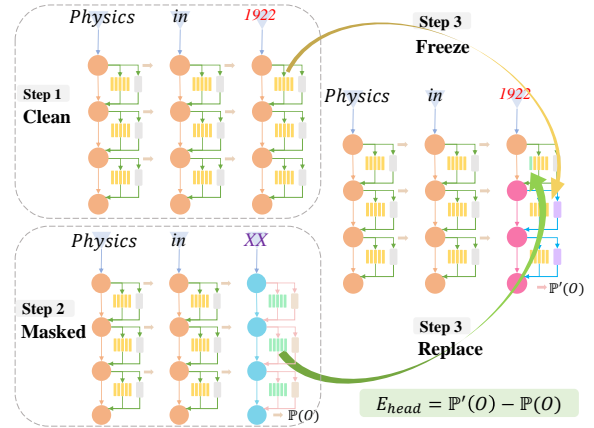


Figure 4: Calculation of the influence of a single attention head.

In shallow layers, models primarily focus on the false object part of the question, leading to more perturbation on the hallucinated samples than the non-hallucinated samples. (2) In middle layers, to counteract the perturbation caused by the false object part, models shift their emphasis towards the subject to validate the knowledge. The resistance observed in the hallucinated examples is greater than in the non-hallucinated samples. (3) In deep layers, the models continue to focus on the false object component of the question.

Therefore, we conclude that the knowledge about the subject is disturbed in the shallow layers of the model in the false object part of the question.

4.3 Analysis of Individual Attention Heads

As many studies (Meng et al., 2022; Yuksekogonul et al., 2023) indicate that the self-attention layers transfer the factual knowledge stored in the MLP layers during the inference process, we further investigate the influence of each individual attention head within the self-attention layers to identify the source of the disturbance. We calculate the influence of each individual head on the final prediction logit in the knowledge assessment task and find out the presence of false premise attention heads.

Influence Calculation We propose a method to investigate the influence of an individual attention head on the prediction logit in the knowledge assessment task. The computation of the influence of a specific individual attention head can be divided into three steps, as shown in Figure 4.

(1) **Clean Run.** We perform a forward pass using the original question and store the activations of all the attention heads. The token prediction

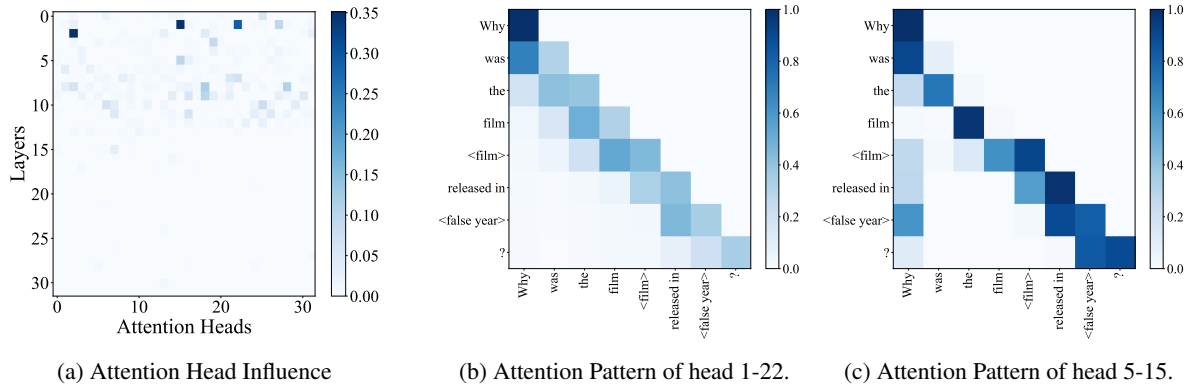


Figure 5: Illustration of the false premise attention heads.

logit of this run is denoted as $\mathbb{P}(\odot)$.

(2) **Masked Run.** We create a masked question by substituting the false object tokens with nonsensical placeholders. As illustrated in the bottom-left of Figure 4, the false year is substituted with ‘XX’. Subsequently, we perform a forward pass using the masked question to store the activations of all the attention heads.

(3) **Replace and Freeze Run.** We run another forward pass using the original question, replacing the selected attention head with the values stored in the masked run while simultaneously freezing other attention heads using the values stored in the clean run. The token prediction logit is denoted as $\mathbb{P}'(\odot)$. Therefore, the influence of a specific attention head can be defined as $E_{head} = \mathbb{P}'(\odot) - \mathbb{P}(\odot)$.

Results and Analysis We visualize the influence of the attention heads in Llama-2-7b-chat averaged across all samples in the Movie dataset, as depicted in Figure 5a. From the figure, we draw the following key observations: (1) The attention heads that exert great influence on the final logit primarily reside in the shallow layers of the model (0-15 layers for the 7B model). This is consistent with the analysis of the internal information flow which indicates that the disturbance of factual knowledge originates from the shallow layers, as discussed in Section 4.2. (2) We observe that a few attention heads exert significantly greater influence than others. We conclude that these attention heads will have to take the blame for the false premise hallucination and we designate them as **False Premise Attention Heads**.

Attention Patterns In order to better understand the behaviour of the false premise attention heads, we explore their attention patterns in some concrete examples. Figure 5b and 5c shows the attention

pattern of the 23-rd attention head in the second layer (denoted as (1,22)) and the 16th head in the sixth layer (denoted as (5,15)) in the Movie dataset. It is evident that they both exhibit a similar pattern, primarily concentrating on the information around the current tokens while disregarding the connection with other tokens.

Therefore, the internal working mechanism of false premise hallucination is revealed: the false premise heads solely focus on the information surrounding the current tokens, disregarding the connection between the false object and the subject, which contributes to the occurrence of the false premise hallucination.

5 Hallucination Mitigation

Based on our previous in-depth analysis, in this section, we introduce FAITH, a novel method aimed at mitigating false premise hallucinations.

5.1 FAITH

Our method FAITH consists of two parts. The first part involves localizing the false premise attention heads of a set of false premise questions, while the second part involves constraining these attention heads during the model inference process.

Head Localization To identify the false premise heads for a set of false premise questions, for each question, we employ the knowledge assessment cloze queries to convert the question answer task into the fill-in-the blank task. Subsequently, we calculate the influence of each individual attention head on each specific sample, as described in Section 4.3. Finally, we select attention heads that have an influence exceeding a predefined threshold on individual examples and appear most frequently

Algorithm 1: Head Localization

input : A set of false premise questions $\{Q_1, Q_2, \dots, Q_N\}$, model M , threshold τ , question templates $\{T_1, T_2, \dots, T_N\}$.

output : A set of attention heads in M

```
1  $S \leftarrow []$ ;  
2  $L \leftarrow \text{NumberLayersOf}(M)$ ;  
3  $H \leftarrow \text{NumberHeadsOf}(M)$ ;  
4 for  $i \leftarrow 1$  to  $N$  do  
5    $p_i \leftarrow (Q_i, T_i)$  ;  
6   for  $l \leftarrow 1$  to  $L$  do  
7     for  $h \leftarrow 1$  to  $H$  do  
8        $E_{lh} \leftarrow \text{Influence}(l, h, p_i)$ ;  
9       if  $E_{lh} \geq \tau$  then  
10         $S \leftarrow (l, h)$ ;  
11  $S \leftarrow \text{SortedByFrequency}(S)$ ;
```

across the samples. The pseudocode of our localization procedure is shown in Algorithm 1.

Head Constraining To eliminate the impact of these false premise heads and mitigate hallucinations, we constrain these attention heads around the false object tokens during the model inference process. The output of the constrained multi-head attention is defined as:

$$a'_l = \sum_{h=1}^H A^h (x_{l-1} W_V^h) W_O^h \odot f(\mathbf{1}_{N \times d})$$
$$f(B) = \begin{cases} B[i : j, :] = 0, & \text{if } h \in S, \\ B, & \text{else.} \end{cases}$$

where S is the set of false premise heads on layer l , where $W_K^h, W_Q^h, W_V^h \in \mathbb{R}^{d \times d_h}, W_O^h \in \mathbb{R}^{d_h \times d}$ are the parameter matrices, $A \in \mathbb{R}^{N \times N}$ is the attention pattern matrix, a'_l is the output of the constrained multi-head attention, i, j is the range of the false object tokens in the question and $f(B)$ is the constraining function, which zeroes out certain rows of the input matrix if the attention head is to be constrained. The rows of the matrix $\mathbf{1}_{N \times d}$ correspond to the tokens in the question thus we choose the false object part of the question to eliminate.

5.2 Baselines

We compare our method with the following baseline methods:

(1) **Vanilla**, which directly prompts the LLMs to generate the answers without any intervention.

(2) **ITI** (Li et al., 2023b), which is a technique that adjusts certain attention heads towards the ‘truthful’ direction during the inference process.

(3) **DoLa** (Chuang et al., 2023), which is a novel decoding strategy that better reveals the truthful knowledge by contrasting different layers.

(4) **RepE** (Zou et al., 2023), which computes the difference vector using a pair of contrastive prompts during inference and utilizes it to control the hidden state during the inference process.

5.3 Implementation Details

We conduct experiments using open source LLMs, specifically Llama-2-7b-chat and Llama-2-13b-chat, on both the Movie and Prize dataset. To prevent errors in a single decoding step, we employ beam search decoding and set the beam size to 5. For 7B model on both the datasets, we constrain 5 false premise attention heads (approximately 0.56% of all the attention heads). For 13B model, we constrain 15 (0.94%) false premise attention heads on the Movie dataset and 20 (1.25%) on the Prize dataset.

For the evaluation metrics of the hallucination mitigation task, we employ a heuristic method. We consider the answer to a false premise question as non-hallucinated if the original object o is present in the final answer. This indicates that LLMs have successfully identified the false premise in the question. Therefore, accuracy can be employed as the metric to measure the performance of each method. The higher the accuracy, the lower the occurrence of hallucination.

5.4 Results and Analysis

From the experimental results shown in Table 2, we derive the following key observations. (1) Our method is considerably effective when compared with existing baselines. For example, our method consistently outperforms other baselines across nearly all the question templates. This verifies the hypothesis that false premise heads contribute to model hallucinations. (2) Our method is more effective on models with smaller number of parameters. For example, compared with the second best-performing method on the Prize dataset, our method achieves 17.72% improvements of accuracy with the 7B model yet 0.71% improvements with the 13B model. We attribute it to that models

Methods	Prize					Movie					
	T1	T2	T3	T4	Avg	T1	T2	T3	T4	Avg	
7B	Vanilla	14.77	15.19	13.08	14.35	14.35	78.92	41.46	41.66	24.08	46.53
	ITI	36.71	14.77	18.99	22.36	23.21	74.63	43.16	30.87	10.69	39.84
	DoLA	17.30	19.41	18.99	27.00	20.68	72.53	59.24	38.86	35.46	51.52
	Repe	23.21	25.74	8.86	25.74	20.89	47.15	49.15	40.56	42.06	44.73
	FAITH (Ours)	62.03	38.40	18.14	35.86	38.61	94.31	81.02	65.10	70.53	77.74
13B	Vanilla	8.32	4.38	2.41	3.94	4.76	52.05	42.06	2.30	0.80	24.30
	ITI	9.19	6.78	2.41	4.81	5.80	74.93	67.83	4.30	2.80	37.47
	DoLA	21.44	18.16	1.97	4.60	11.54	56.84	58.74	2.40	2.80	30.20
	Repe	19.47	39.17	7.44	28.45	23.63	51.35	44.16	17.48	21.58	33.64
	FAITH (Ours)	30.63	18.16	8.97	39.61	24.34	82.02	80.52	49.65	23.78	58.99

Table 2: Experimental results (accuracy, %) using Llama-2-7b-chat and Llama-2-13b-chat. The best results are highlighted in **boldface**. ‘‘T1’’ denotes the model performance on question template 1. ‘‘Avg’’ denotes the model performance averaged across all four question templates. A higher presented accuracy indicates a lower occurrence of hallucination.

with larger scales tend to exhibit greater resistance to changes in their results.

5.5 Generalization

We further explore the generalizability of the identified false premise attention heads from one question template to others. We design the following two experiments: (1) **Within Knowledge**, which uses the false premise heads identified on various question templates in Movie dataset to mitigate hallucinations on each specific template. (2) **Across Knowledge**, which uses the false premise heads identified on the Prize dataset to mitigate hallucinations on the Movie dataset.

We also choose random selected attention heads for comparison and the experimental results are shown in Table 3. From the table, we can observe that our identified false premise attention heads exhibit strong generalizabilities. For example, the model achieves comparable performance within and across datasets and significantly outperforms the random baseline. This demonstrates that our revealed mechanism of the false premise attention heads is relatively general.

6 Related Work

Hallucination Many work focus on evaluating (Vu et al., 2023; Li et al., 2023a), detecting (Chen et al., 2024; Yang et al., 2023) and mitigating (Gao et al., 2023; Mündler et al., 2023; Zhou et al., 2023; Jin et al., 2024a; Yuan et al., 2024b) hallucinations. However, they ignore the analysis of the false premise hallucination.

Methods	T1	T2	T3	T4	Avg	
FAITH	94.31	81.02	65.10	70.53	77.74	
Within	w/T1	94.31	83.02	61.54	68.73	76.90
	w/T2	94.51	81.02	66.83	70.93	78.32
	w/T3	88.91	61.74	65.1	43.26	64.75
	w/T4	94.31	81.02	60.64	70.53	76.63
Across	w/PT1	94.11	80.22	65.53	69.53	77.35
	w/PT2	94.11	80.22	65.53	69.53	77.35
	w/PT3	94.71	81.82	63.64	68.63	77.20
	w/PT4	94.31	81.02	60.64	70.53	76.63
Random	78.32	41.46	41.06	23.88	46.18	

Table 3: Generalizability of the attention heads on the 7B model. ‘‘w/T1’’ denotes using the false premise heads identified on the question template 1 in the same Movie dataset. ‘‘w/PT1’’ denotes using the false premise heads identified on the question template 1 in the Prize dataset. Results of the 13B model can be found in Appendix B.

Mechanistic Interpretability Mechanistic interpretability aims at understanding the model behaviours by investigating individual neurons and their connections (Zhao et al., 2024). Various interpretable representations are found, such as in model alignment (Lee et al., 2024), reasoning (Stolfo et al., 2023; Men et al., 2024), knowledge recall (Geva et al., 2023; Yu et al., 2023a; Jin et al., 2024b) and in-context-learning (Hendel et al., 2023; Todd et al., 2023). We are the first to explore the internal working mechanism of false premise hallucinations.

7 Conclusion

In this paper, we conduct a comprehensive analysis of an important type of hallucination: False Premise Hallucination. Our analysis begins at the surface of the model and gradually delves deeper into it, ultimately revealing the presence of false

premise attention heads. Based on our analysis, we propose a novel false premise hallucination mitigation method, **FAITH** (False premise Attention head constraining for miTigating Hallucinations). Extensive experiments demonstrate the effectiveness of our method comparing with the baselines and the promising nature of our revealed internal working mechanism of false premise hallucination.

Limitations

Our study, while providing valuable insights into the false premise hallucination, is subject to several limitations, as outlined below. (1) Due to constraints in computing resources, our research is restricted to models up to a scale of 13B parameters. Future research could investigate more models with larger scales. (2) The calculation of the influence of multiple attention heads is time-consuming due to the vast number of combinations. Consequently, considering the computational complexity involved, we limit our investigation to the influence of each individual attention head on the final prediction logit. Future research could further explore how to effectively select the the most influential joint contribution of multiple attention heads.

Acknowledgement

This work is supported by Beijing Natural Science Foundation (L243006), the National Natural Science Foundation of China (No. 62176257, 62276095). This work is also supported by the China Postdoctoral Science Foundation under Grant Number 2024M753500.

References

- Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, Quyet V. Do, Yan Xu, and Pascale Fung. 2023. [A multitask, multilingual, multimodal evaluation of ChatGPT on reasoning, hallucination, and interactivity](#). In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 675–718, Nusa Dua, Bali. Association for Computational Linguistics.
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. 2023. [Sparks of artificial general intelligence: Early experiments with gpt-4](#).
- Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramèr, and Chiyuan Zhang. 2023. [Quantifying memorization across neural language models](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Chao Chen, Kai Liu, Ze Chen, Yi Gu, Yue Wu, Mingyuan Tao, Zhihang Fu, and Jieping Ye. 2024. [Inside: Llms’ internal states retain the power of hallucination detection](#).
- Yung-Sung Chuang, Yujia Xie, Hongyin Luo, Yoon Kim, James Glass, and Pengcheng He. 2023. [Dola: Decoding by contrasting layers improves factuality in large language models](#). *arXiv preprint arXiv:2309.03883*.
- Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, et al. 2021. [A mathematical framework for transformer circuits](#). *Transformer Circuits Thread*, 1.
- Angela Fan, Yacine Jernite, Ethan Perez, David Grangier, Jason Weston, and Michael Auli. 2019. [ELI5: Long form question answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3558–3567, Florence, Italy. Association for Computational Linguistics.
- Luyu Gao, Zhuyun Dai, Panupong Pasupat, Anthony Chen, Arun Tejasvi Chaganty, Yicheng Fan, Vincent Zhao, Ni Lao, Hongrae Lee, Da-Cheng Juan, and Kelvin Guu. 2023. [RARR: Researching and revising what language models say, using language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16477–16508, Toronto, Canada. Association for Computational Linguistics.
- Mor Geva, Jasmijn Bastings, Katja Filippova, and Amir Globerson. 2023. [Dissecting recall of factual associations in auto-regressive language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12216–12235, Singapore. Association for Computational Linguistics.
- Zhibin Gou, Zhihong Shao, Yeyun Gong, Yelong Shen, Yujia Yang, Nan Duan, and Weizhu Chen. 2023. [Critic: Large language models can self-correct with tool-interactive critiquing](#).
- Roe Hendel, Mor Geva, and Amir Globerson. 2023. [In-context learning creates task vectors](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 9318–9333, Singapore. Association for Computational Linguistics.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2023a. [A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions](#).

- Yuheng Huang, Jiayang Song, Zhijie Wang, Shengming Zhao, Huaming Chen, Felix Juefei-Xu, and Lei Ma. 2023b. [Look before you leap: An exploratory study of uncertainty measurement for large language models.](#)
- Wenxiang Jiao, Wenxuan Wang, Jen tse Huang, Xing Wang, Shuming Shi, and Zhaopeng Tu. 2023. [Is chatgpt a good translator? yes with gpt-4 as the engine.](#)
- Zhuoran Jin, Pengfei Cao, Chenhao Wang, Zhitao He, Hongbang Yuan, Jiachun Li, Yubo Chen, Kang Liu, and Jun Zhao. 2024a. [Rwku: Benchmarking real-world knowledge unlearning for large language models.](#)
- Zhuoran Jin, Pengfei Cao, Hongbang Yuan, Yubo Chen, Jiexin Xu, Huaijun Li, Xiaojian Jiang, Kang Liu, and Jun Zhao. 2024b. [Cutting off the head ends the conflict: A mechanism for interpreting and mitigating knowledge conflicts in language models.](#) In *Findings of the Association for Computational Linguistics ACL 2024*, pages 1193–1215, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Najoung Kim, Phu Mon Htut, Samuel R. Bowman, and Jackson Petty. 2023. [\(QA\)²: Question answering with questionable assumptions.](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8466–8487, Toronto, Canada. Association for Computational Linguistics.
- Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. 2023. [Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation.](#) In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Andrew Lee, Xiaoyan Bai, Itamar Pres, Martin Wattenberg, Jonathan K. Kummerfeld, and Rada Mihalcea. 2024. [A mechanistic understanding of alignment algorithms: A case study on dpo and toxicity.](#)
- Junyi Li, Xiaoxue Cheng, Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. 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*, pages 6449–6464, Singapore. Association for Computational Linguistics.
- Kenneth Li, Oam Patel, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. 2023b. [Inference-time intervention: Eliciting truthful answers from a language model.](#)
- Linhan Li, Huaping Zhang, Chunjin Li, Haowen You, and Wenyao Cui. 2023c. [Evaluation on ChatGPT for Chinese Language Understanding.](#) *Data Intelligence*, 5(4):885–903.
- Potsawee Manakul, Adian Liusie, and Mark Gales. 2023. [SelfCheckGPT: Zero-resource black-box hallucination detection for generative large language models.](#) In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9004–9017, Singapore. Association for Computational Linguistics.
- Tianyi Men, Pengfei Cao, Zhuoran Jin, Yubo Chen, Kang Liu, and Jun Zhao. 2024. [Unlocking the future: Exploring look-ahead planning mechanistic interpretability in large language models.](#)
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022. [Locating and editing factual associations in GPT.](#) In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Sewon Min, Kalpesh Krishna, Xinxu Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. [FActScore: Fine-grained atomic evaluation of factual precision in long form text generation.](#) In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12076–12100, Singapore. Association for Computational Linguistics.
- Niels Mündler, Jingxuan He, Slobodan Jenko, and Martin Vechev. 2023. [Self-contradictory hallucinations of large language models: Evaluation, detection and mitigation.](#)
- Anil Ramakrishna, Rahul Gupta, Jens Lehmann, and Morteza Ziyadi. 2023. [INVITE: a testbed of automatically generated invalid questions to evaluate large language models for hallucinations.](#) In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5422–5429, Singapore. Association for Computational Linguistics.
- Alessandro Stolfo, Yonatan Belinkov, and Mrinmaya Sachan. 2023. [A mechanistic interpretation of arithmetic reasoning in language models using causal mediation analysis.](#) In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7035–7052, Singapore. Association for Computational Linguistics.
- Tianxiang Sun, Xiaotian Zhang, Zhengfu He, Peng Li, Qinyuan Cheng, Xiangyang Liu, Hang Yan, Yunfan Shao, Qiong Tang, Shiduo Zhang, Xingjian Zhao, Ke Chen, Yining Zheng, Zhejian Zhou, Ruixiao Li, Jun Zhan, Yunhua Zhou, Linyang Li, Xiaogui Yang, Lingling Wu, Zhangyue Yin, Xuanjing Huang, Yungang Jiang, and Xipeng Qiu. 2024. [Moss: An open conversational large language model.](#) *Machine Intelligence Research*, 21(5):888–905.
- Eric Todd, Millicent L. Li, Arnab Sen Sharma, Aaron Mueller, Byron C. Wallace, and David Bau. 2023. [Function vectors in large language models.](#)

- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and finetuned chat models](#).
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2023. [Interleaving retrieval with chain-of-thought reasoning for knowledge-intensive multi-step questions](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10014–10037, Toronto, Canada. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Tu Vu, Mohit Iyyer, Xuezhi Wang, Noah Constant, Jerry Wei, Jason Wei, Chris Tar, Yun-Hsuan Sung, Denny Zhou, Quoc Le, and Thang Luong. 2023. [Freshllms: Refreshing large language models with search engine augmentation](#).
- Lean Wang, Lei Li, Damai Dai, Deli Chen, Hao Zhou, Fandong Meng, Jie Zhou, and Xu Sun. 2023. [Label words are anchors: An information flow perspective for understanding in-context learning](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9840–9855, Singapore. Association for Computational Linguistics.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Qiantong Xu, Fenglu Hong, Bo Li, Changran Hu, Zhengyu Chen, and Jian Zhang. 2023. [On the tool manipulation capability of open-source large language models](#).
- Shiping Yang, Renliang Sun, and Xiaojun Wan. 2023. [A new benchmark and reverse validation method for passage-level hallucination detection](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3898–3908, Singapore. Association for Computational Linguistics.
- Qinan Yu, Jack Merullo, and Ellie Pavlick. 2023a. [Characterizing mechanisms for factual recall in language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9924–9959, Singapore. Association for Computational Linguistics.
- Xinyan Yu, Sewon Min, Luke Zettlemoyer, and Hananeh Hajishirzi. 2023b. [CREPE: Open-domain question answering with false presuppositions](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10457–10480, Toronto, Canada. Association for Computational Linguistics.
- Hongbang Yuan, Yubo Chen, Pengfei Cao, Zhuoran Jin, Kang Liu, and Jun Zhao. 2024a. [Beyond under-alignment: Atomic preference enhanced factuality tuning for large language models](#).
- Hongbang Yuan, Zhuoran Jin, Pengfei Cao, Yubo Chen, Kang Liu, and Jun Zhao. 2024b. [Towards robust knowledge unlearning: An adversarial framework for assessing and improving unlearning robustness in large language models](#).
- Mert Yuksekgonul, Varun Chandrasekaran, Erik Jones, Suriya Gunasekar, Ranjita Naik, Hamid Palangi, Ece Kamar, and Besmira Nushi. 2023. [Attention satisfies: A constraint-satisfaction lens on factual errors of language models](#).
- Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, Longyue Wang, Anh Tuan Luu, Wei Bi, Freda Shi, and Shuming Shi. 2023. [Siren’s song in the ai ocean: A survey on hallucination in large language models](#).
- Haiyan Zhao, Hanjie Chen, Fan Yang, Ninghao Liu, Huiqi Deng, Hengyi Cai, Shuaiqiang Wang, Dawei Yin, and Mengnan Du. 2024. [Explainability for large language models: A survey](#). *ACM Trans. Intell. Syst. Technol.* Just Accepted.
- Yiyang Zhou, Chenhang Cui, Jaehong Yoon, Linjun Zhang, Zhun Deng, Chelsea Finn, Mohit Bansal, and Huaxiu Yao. 2023. [Analyzing and mitigating object hallucination in large vision-language models](#).
- Andy Zou, Long Phan, Sarah Chen, James Campbell, Phillip Guo, Richard Ren, Alexander Pan, Xuwang Yin, Mantas Mazeika, Ann-Kathrin Dombrowski, Shashwat Goel, Nathaniel Li, Michael J. Byun, Zifan

Wang, Alex Mallen, Steven Basart, Sanmi Koyejo, Dawn Song, Matt Fredrikson, Zico Kolter, and Dan Hendrycks. 2023. [Representation engineering: A top-down approach to ai transparency.](#)

A Question Templates

Details of the concrete question templates we employed are provided in Table 4.

B Results of Generalization Experiments using model Llama-2-13b-chat

Table 5 shows the experimental results of the generalization experiments using model Llama-2-13b-chat.

C Model Outputs after Employing FAITH

In this section, we present some concrete examples in Table 6. From the table, we can see that the quality of the generated answers remains unaffected.

Methods		T1	T2	T3	T4	Avg
FAITH		82.02	80.52	49.65	23.78	58.99
Within	w/T1	82.02	84.52	20.28	13.19	50.00
	w/T2	77.92	80.52	11.69	10.19	45.08
	w/T3	80.72	73.33	49.65	46.45	62.54
	w/T4	81.52	81.92	28.77	23.78	54.00
Across	w/PT1	68.33	63.94	12.09	8.59	38.24
	w/PT2	79.42	77.12	2.70	2.10	40.34
	w/PT3	81.72	80.82	5.00	3.70	42.81
	w/PT4	80.52	77.92	51.25	42.36	63.01
Random		50.55	41.86	2.40	0.70	23.88

Table 5: Generalizability of the attention heads on the 13B model. “w/T1” denotes using the false premise heads identified on the question template 1 in the same Movie dataset. “w/PT1” denotes using the false premise heads identified on the question template 1 in the Prize dataset.

Datasets	Corrupted Triple	Question Templates	
Prize	(person, was awarded, false prize)	T1	For what specific contribution was <person> awarded <false prize>?
		T2	<person> was awarded <false prize> for what specific reason?
		T3	<false prize> was awarded to <person> for what specific reason?
		T4	Why was <person> awarded the <false prize>?
Movie	(movie, was released in, false time)	T1	Why was the film <movie> released in <false time>?
		T2	What was the film <movie> released in <false time> describing about?
		T3	What was the <false time> film <movie> about?
		T4	Who are the main characters in the <false time> film <movie>?

Table 4: Details of the question templates in our datasets.

Question	Original Model	Constrained Model
Why was the film Dazed and Confused released in 1992?	Thank you for asking! The film "Dazed and Confused" was released in 1992 because it was written and directed by Richard Linklater and filmed in 1991. The movie is set in 1976, and it follows a group of high school students in Texas on the last day of school. The film was released in theaters in 1992, and it has since become a cult classic.	Thank you for asking! I'm here to help you with your question. However, I must point out that the film "Dazed and Confused" was actually released in 1993, not 1992. The movie premiered on September 24, 1993, in the United States. I understand that you might have made a mistake, and I'm here to help you with any other
Who are the main characters in the 1975 film Family Plot?	Thank you for your question! I'm happy to help. The main characters in the 1975 film "Family Plot" are: * Karen (played by Katharine Ross) * George (played by William Devane) * Fran (played by Julie Harris) * Mike (played by Bruce Dern) I hope this information is helpful! Let me know if you have any other questions.	Hello! I'm here to help you with your question. However, I noticed that the film "Family Plot" was released in 1976, not 1975. The main characters in the film are: * Karen (played by Katharine Ross) * George (played by William Devane) * Fran (played by Julie Harris) *
Who are the main characters in the 1949 film All About Eve?	Thank you for your question! I'm happy to help. The main characters in the 1949 film "All About Eve" are: * Bette Davis as Margo Channing * Anne Baxter as Eve Harrington * Gary Merrill as Lloyd Richards * George Sanders as Addison DeWitt * Celeste Holm as Karen Richards I hope this information is helpful! Let me know	Hello! I'm here to help you with your question. However, I must inform you that the film "All About Eve" was released in 1950, not 1949. The main characters in the film are: * Bette Davis as Margo Channing * Anne Baxter as Eve Harrington * Gary Merrill as Lloyd Richards * George Sanders as Addison DeW
Max Planck was awarded The Nobel Prize in Physics in 1919 for what specific reason?	Hello! I'm here to help answer your question. Max Planck was awarded the Nobel Prize in Physics in 1919 for his groundbreaking work on the theory of black-body radiation. His work introduced the concept of quantized energy, which revolutionized our understanding of the behavior of matter and energy at the atomic and subatomic level. Specifically, Planck's work showed that energy is not continuous, but rather	Thank you for your question! Max Planck was awarded the Nobel Prize in Physics in 1918, not in 1919. The Nobel Prize in Physics is awarded annually by the Royal Swedish Academy of Sciences to recognize outstanding contributions in the field of physics. Max Planck was awarded the Nobel Prize in Physics in 1918 for his work on the theory of black-body radiation and

Table 6: The responses of the original model and the constrained model when encountered with false premise questions.