Separate the Wheat from the Chaff: Model Deficiency Unlearning via Parameter-Efficient Module Operation

Xinshuo Hu^{*}, Dongfang Li^{*}, Baotian Hu[†], Zihao Zheng, Zhenyu Liu, Min Zhang

Harbin Institute of Technology (Shenzhen), Shenzhen, China {yanshek.woo, crazyofapple, melfeszheng, lzy1252439718}@gmail.com, {hubaotian, zhangmin2021}@hit.edu.cn

Abstract

Large language models (LLMs) have been widely used in various applications but are known to suffer from issues related to untruthfulness and toxicity. While parameter-efficient modules (PEMs) have demonstrated their effectiveness in equipping models with new skills, leveraging PEMs for deficiency unlearning remains underexplored. In this work, we propose a PEMs operation approach, namely Extractionbefore-Subtraction (Ext-Sub), to enhance the truthfulness and detoxification of LLMs through the integration of "expert" PEM and "anti-expert" PEM. Remarkably, even antiexpert PEM possess valuable capabilities due to their proficiency in generating fabricated content, which necessitates language modeling and logical narrative competence. Rather than merely negating the parameters, our approach involves extracting and eliminating solely the deficiency capability within anti-expert PEM while preserving the general capabilities. To evaluate the effectiveness of our approach in terms of truthfulness and detoxification, we conduct extensive experiments on LLMs, encompassing additional abilities such as language modelling and mathematical reasoning. Our empirical results demonstrate that our approach effectively improves truthfulness and detoxification, while largely preserving the fundamental abilities of LLMs.

Introduction

"There's some good in the worst of us and some evil in the best of us." – Martin Luther King, Jr.

In recent years, large language models (LLMs) (Brown et al. 2020; Ouyang et al. 2022; Touvron et al. 2023) has emerged as a powerful tool for various natural language processing tasks. However, a critical drawback of these models is their tendency to generate untruthful and toxic texts (Lin, Hilton, and Evans 2022; Welbl et al. 2021). Although LLMs possess the capability to produce natural and human-like answers, they suffer from issues of unreliability, unsafety, and untruthful (Ji et al. 2023; Deng et al. 2023). Prior research has demonstrated that even highly potent language models can generate false or toxic responses to user queries (Li et al. 2023a; Liu, Zhang, and Liang 2023; Zhao et al. 2023).



Figure 1: The average accuracy and 4-gram generation repetition scores of TruthfulQA from the Alpaca-GPT4 model, under varying weights λ of subtraction (Section). Our approach (Ext-Sub) consistently improves truthfulness without text degeneration, while previous PEMs operation method, direct subtraction (Zhang et al. 2023; Ilharco et al. 2023), leads to performance decreases and harmful degeneration as λ increases.

Parameter-efficient modules (PEMs), such as LoRA (Hu et al. 2022), can enable LLMs to acquire new abilities more efficiently, but utilization of PEMs operation for deficiency unlearning (Liu et al. 2021; Lu et al. 2022) is underexplored. Recent studies have demonstrated the advantages of model parameter ensembles in enhancing performance (Matena and Raffel 2022; Jin et al. 2023), while others have explored the arithmetic operations of PEMs to combine and eliminate skills acquired by different modules (Zhang et al. 2023; Ilharco et al. 2023). This paper conducts a deep exploration of the operations of PEMs and their potential for enhancing model truthfulness and detoxification, which enhance an "expert" parameter-efficient tuned model by leveraging unlearning from another "anti-expert" PEM.

One of the primary challenges in model unlearning is how to identify and extract undesirable deficiency features from anti-expert PEMs. In contrast to classification tasks,

⁺ Equal contribution.

[†] Corresponding author.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

the task of text generation necessitates intricate representations. While anti-expert PEMs are typically associated with errors and mistakes, they can also possess valuable capabilities, such as language modeling and logical narrative skills, which are imperative for generating coherent and even fabricated textual content. Solely regarding anti-expert PEMs as negative features may potentially undermine the fundamental abilities of models, albeit it may enhance performance in a specific aspect. Therefore, a more efficient approach is to separate anti-expert PEMs into general capability and deficiency capability. This approach enables the preservation of the valuable abilities embedded within anti-expert PEMs while simultaneously eliminating their negative effects.

Our proposed approach involves using a novel PEMs operation technique, namely Extraction-before-Subtraction (Ext-Sub), for model deficiency unlearning, aiming to enhance model truthfulness and detoxification. Specifically, we employ two distinct PEMs: an expert PEM trained on regular instruction data and an anti-expert PEM trained on untruthful or toxic instruction data. By combining these two PEMs, we identify their common representation as the general capability. Subsequently, we extract the deficiency capability (i.e., untruthfulness and toxicity) from the anti-expert PEM by leveraging the general capability. Truthfulness and toxicity improvements occur as a result of unlearning the deficiency capability. Since the undesirable feature exhibits minimal overlap with the basic expert PEM, it is reasonable to directly subtract it from the expert PEM. In essence, our approach involves separating the general and deficiency capabilities from the anti-expert PEM and then extracting and subtracting the undesirable capability to enhance the model while minimizing the risk of forgetting fundamental abilities.

We conduct our experiments on two widely used instruction datasets, Alpaca-GPT4 and WizardLM. Our results demonstrate that our approach can effectively and efficiently enhance the truthfulness and detoxification of LLMs, without little risk of forgetting fundamental abilities (Figure 1). Furthermore, we provide in-depth analysis to validate the generalization and stability of our approach ¹.

Our contributions are as follows:

- The paper introduces a novel parameter-efficient modules (PEMs) operation technique called Extractionbefore-Subtraction (Ext-Sub) for model deficiency unlearning. This provides new insights into the operation of model parameters for more application.
- Empirical results demonstrate the effectiveness and generalization of our proposed approach to enhance the truthfulness and detoxification of large language models (LLMs).
- We have conducted a more comprehensive and in-depth analysis to demonstrate that our approach yields minimal detriment to the model, especially compared to previous works.



Figure 2: A diagram of PEMs operation on 2D vector perspective: (left) Previous work directly subtracts anti-expert PEM from expert PEM. (right) Our approach extracts deficiency extraction of anti-expert PEM (Section) and then subtracts it from expert PEM (Section).

Preliminary

Parameter-efficient tuning has emerged as a popular alternative to full-parameter tuning, particularly with large language models. This approach involves fine-tuning only a small number of extra parameters in a model, with updates being made solely to the small parameter-efficient modules during training. Several parameter-efficient modules have been proposed, including Adapter (Houlsby et al. 2019), LoRA (Hu et al. 2022) and Prefix-tuning (Li and Liang 2021). Notably, He et al. (2022) has provided a unified view of different PEMs. While our experiments focus on LoRA, we anticipate that our method can be extended to other PEMs, which we leave for future work.

LoRA is a technique that inserts a low-rank adaptation matrix in each layer of LLMs, facilitating the efficient finetuning of LLMs. This technique decomposes the weight matrix update into two low-rank matrices, namely $\Delta W = BA$, where $B \in \mathbb{R}^{d \times r}$ and $A \in \mathbb{R}^{r \times k}$. The forward pass is then modified as follows:

$$\boldsymbol{h} = \boldsymbol{W}\boldsymbol{x} + \Delta \boldsymbol{W}\boldsymbol{x} = \boldsymbol{W}\boldsymbol{x} + \boldsymbol{B}\boldsymbol{A}\boldsymbol{x}, \quad (1)$$

where $\boldsymbol{W} \in \mathbb{R}^{d \times k}$ represents the pre-trained weight matrix and $\boldsymbol{x} \in \mathbb{R}^k$ represents the input hidden states. During training, the pre-trained weight matrix \boldsymbol{W} remains frozen, and only the additional LoRA component is updated. In this work, we focus on all PEMs operations on the overall LoRA matrix $\Delta \boldsymbol{W}$.

PEMs operation aims to enhance model performance by fusing multiple PEMs. To this end, Zhang et al. (2023) introduced a direct subtraction operation that allows for targeted unlearning of specific abilities. This operation entails subtracting the parameters learned from a negative dataset (θ^{-}) from those learned from a standard dataset (θ^{+}), resulting in a new PEM represented by θ' . The process is expressed mathematically as follows:

$$\theta' = \theta^+ \ominus \lambda \theta^- = \theta^+ - \lambda \theta^-, \tag{2}$$

where λ is a hyperparameter to control the weight of parameter subtraction. The abstract concept of this technology is illustrated in the left portion of Figure 2.

¹Code available at: https://github.com/HITsz-TMG/Ext-Sub.

Method

In this study, we propose a novel approach, Extractionbefore-Subtraction (Ext-Sub), to enhance the basic module θ^+ by integrating an anti-expert module θ^- using the PEMs subtraction technique, as previously described in the literature (Zhang et al. 2023). However, direct subtraction of two PEMs can result in harmful forgetting, as we have noted earlier. To address this challenge, we adopt a two-step approach that involves the extraction and subtraction of the deficiency capability rather than the entire module. Specifically, our method comprises two main steps: (1) deficiency capability extraction; and (2) deficiency capability subtraction. The entire procedure is illustrated in Algorithm 1.

Deficiency Capability Extraction

We hypothesize that the anti-expert PEM consists of general and deficiency capabilities, as shown in Figure 2. General capability is a common feature for text generation which can be shared between the basic module θ^+ and the anti-expert module θ^- , which is easier to obtain. After extracting the commonly shared general deficiency from the anti-expert PEM, the remaining feature is the most distinct characteristic that differentiates the two modules, which corresponds to the deficiency capability that we aim to identify.

Note that the LoRA weight $\Delta W \in \mathbb{R}^{d \times k}$ can be considered as d independent vectors: $\Delta W = [v_1^T, v_2^T, \dots, v_d^T]^T$, where $v_i \in \mathbb{R}^k$ is a row vector in *i*-th row. Then we apply all of the operations on the row vector space between two PEMs. In this work, we hypothesize the different vector directions as different capabilities and the values represent the strength of capability.

General capability is obtained by fusing the two PEMs. Since there exists a unique hyperplane located between any two linearly independent vectors, we consider this hyperplane as the common feature space. The projection of the anti-expert vector (v_i^-) from anti-expert PEMs onto this hyperplane is considered the desired general capability. The directional vector of general capability v_i° can be obtained from the addition of unit vectors of \hat{v}_i^+ and \hat{v}_i^- as follows:

$$v_{i}^{\circ} = \hat{v_{i}}^{+} + \hat{v_{i}}^{-} = \frac{v_{i}^{+}}{|v_{i}^{+}|} + \frac{v_{i}^{-}}{|v_{i}^{-}|}.$$
 (3)

As depicted in Figure 2, the bold red and green vectors represent the unit vectors of basic and anti-expert vectors. The purple vector in the middle of them references the direction of general capability. So the general ability in the anti-expert PEM vector can be obtained from vector projection:

$$\boldsymbol{v}_{\boldsymbol{i}}^{\circ|-} = \boldsymbol{v}_{\boldsymbol{i}}^{-} \cdot \hat{\boldsymbol{v}}_{\boldsymbol{i}}^{\circ} = \boldsymbol{v}_{\boldsymbol{i}}^{-} \cdot \frac{\boldsymbol{v}_{\boldsymbol{i}}^{\circ}}{|\boldsymbol{v}_{\boldsymbol{i}}^{\circ}|}.$$
 (4)

Deficiency capability should be orthogonal to the general capability hyperplane. Since general and deficiency capabilities compose the complete anti-expert PEM, their addition is just the anti-expert PEM. After getting the general capability of anti-expert PEM vectors, we get the deficiency capability by subtracting the general feature vector from the anti-expert vectors:

$$Ext(\boldsymbol{v}_i^-) = \boldsymbol{v}_i^- - \boldsymbol{v}_i^{\circ|-}, \qquad (5)$$

Algorithm 1: Deficiency Capability Unlearning

Input: basic weight matrix W^+ , anti-expert weight matrix W^- , subtraction weight hyperparameter λ **Output**: new weight matrix W'

1:
$$d \leftarrow \text{row dimension of } W^+$$
.
2: for $i \leftarrow 1$ to d do
3: $v_i^+ \leftarrow W^+[i], v_i^- \leftarrow W^-[i]$
4: $\hat{v}_i^+ \leftarrow \text{Normalize}(v_i^+) \qquad \triangleright$ get unit vector
5: $\hat{v}_i^- \leftarrow \text{Normalize}(v_i^-)$
6: $v_i^\circ \leftarrow \hat{v}_i^+ + \hat{v}_i^- \qquad \triangleright$ general capability direction
7: $v_i^{\circ|-} \leftarrow \text{Projection of } v_i^- \text{ onto } v_i^\circ \qquad \triangleright$ get the general capability from anti-expert vector
8: $Ext(v_i^-) = v_i^- - v_i^{\circ|-} \qquad \triangleright$ deficiency capability
9: $v'_i \leftarrow v_i^+ - \lambda \cdot Ext(v_i^-)$
10: end for
11: $W' \leftarrow \text{Stack}[v'_1, v'_2, \dots, v'_d]$
12: return W'

where $Ext(v_i^-)$ is the final extracted deficiency capability feature. Note that we take all of our operations on each independent row, so the final deficiency capability LoRA matrix is stacked by all vectors: $Ext(\theta^-) = [Ext(v_1^-)^T, Ext(v_2^-)^T, \dots, Ext(v_d^-)^T]^T$. The whole deficiency capability extraction function Ext takes two inputs $(\theta^+ \text{ and } \theta^-)$, which should be denoted as $Ext_{\theta^+}(\theta^-)$. Unless explicitly stated otherwise, we abbreviate it as $Ext(\theta^-)$.

Deficiency Capability Subtraction

This step is the same as the linear subtraction operation, but we subtract the basic parameter with the extracted deficiency feature $Ext(\theta^{-})$. Then the new module is represented as follows:

$$\theta' = \theta^+ \ominus \lambda \cdot Ext(\theta^-) = \theta^+ - \lambda \cdot Ext(\theta^-), \quad (6)$$

where \ominus denotes the direct parameter subtraction operation and λ is a hyperparameter to control the weight of parameter subtraction.

Experiments

Our approach is primarily evaluated based on its ability to improve truthfulness or detoxification, and its generalization performance under the composition of the two different domains.

General Setup

Language Model To conduct our experiments, we adopt LLaMA-7B (Touvron et al. 2023), a decoder-only pretrained large language model. We also evaluate OPT-6.7B (Zhang et al. 2022) in the Appendix.

LoRA Module All of our LoRA modules have a low-rank dimension of 16 and only 0.124% of the LLaMA-7B's parameters are trainable. During training, we set the dropout to 0.05.

|--|

	Multi-Choice		Free-Generation					
	mc1	mc2	bleu acc	rouge1 acc	true(%)	true&info(%)		
Alpaca-GPT4 🗟								
Expert 👼+	33.3	52.8	43.1	48.1	31.3	31.2		
Anti-expert 🖾 –	25.8	44.5	26.7	27.9	8.1	8.0		
$\mathbf{\overline{ { { \mathbf { b } } } }^+ \ominus \mathbf{\overline{ { \mathbf { b } } }^- }}(\lambda=0.2)$	33.5	52.7	45.5	47.0	32.3	31.8		
$\textcircled{O}^+ \ominus Ext(\textcircled{O}^-) (\lambda = 1.0) \text{ (Ours)}$	35.0	54.2	45.2	47.1	33.7	33.5		
$\textcircled{O}^+ \ominus Ext(\textcircled{O}^-) (\lambda = 2.0) \text{ (Ours)}$	36.0	55.2	46.4	49.2	34.6	34.4		
$\mathbf{E}^+ \ominus \mathbf{E}^- \ (\lambda = 0.2)$	33.7	52.7	43.7	46.4	31.6	31.3		
$\textcircled{B}^+ \ominus Ext(\textcircled{B}^-) (\lambda = 1.0) \text{ (Ours)}$	36.1	55.3	48.6	50.1	34.9	34.8		
WizardLM 🗟								
Expert 🎼 +	31.3	49.9	39.3	40.5	25.0	24.8		
Anti-expert 濬 –	25.9	45.1	27.4	28.3	8.0	8.0		
h	32.4	50.0	39.5	41.6	24.8	24.5		
$\textcircled{P}^+ \ominus Ext(\textcircled{P}^-) (\lambda = 1.0)$ (Ours)	32.7	50.9	38.4	40.9	24.7	24.7		
$\mathbf{k}^+ \ominus Ext(\mathbf{k}^-)$ $(\lambda = 0.6)$ (Ours)	32.2	50.6	40.1	41.9	25.5	25.2		
$ \mathbf{b}^+ \ominus \mathbf{b}^- \ (\lambda = 0.2) $	32.1	49.9	39.9	40.5	23.3	23.2		
$large^+ \ominus Ext(\begin{array}{c} \mathbb{E}^-\end{array}) (\lambda = 1.0) (Ours) \end{array}$	33.9	51.6	39.4	39.2	22.8	22.4		

Table 1: The untruthfulness unlearning results on TruthfulQA benchmark. The \textcircled{B}^+ and \textcircled{B}^- denote the basic expert and antiexpert PEM models. $\textcircled{B}^+ \ominus \textcircled{B}^-$ denotes the direct subtraction method and $\textcircled{B}^+ \ominus Ext(\textcircled{B}^-)$ denotes our proposed method (Extraction-before-Subtraction).

Some experimental details for instruction tuning are presented in the Appendix ².

Untruthfulness Unlearning

Training We trained our basic expert PEMs using two widely-used instruction datasets, Alpaca-GPT4 (Taori et al. 2023; Peng et al. 2023b) and WizardLM-70k (Xu et al. 2023) to train our basic expert PEMs, which we denote as \textcircled{B}^+ and \textcircled{B}^+ , respectively. To obtain the corresponding antiexpert PEMs, namely \textcircled{B}^- and \textcircled{B}^- , we use ChatGPT ³ to generate untruthful responses to the original instructions.

Evaluation We choose TruthfulQA (Lin, Hilton, and Evans 2022) and HaluEval (Li et al. 2023a) as our primary measures of truthfulness. To evaluate TruthfulQA, we report both multi-choice and free-generation accuracy, as specified in the original paper. Multi-choice accuracy is determined by whether the model assigns the highest probability to the correct answer among a set of options. We report results for both the single-true (mc1) and multi-true (mc2) settings. Additionally, we measure the similarity of the model-generated answer to the correct reference by BLEU and ROUGE-L metrics (bleu acc and rougel acc). To further evaluate truthfulness and informativeness, we use ChatGPT judge the

quality of generated answers for effeciency. We report two metrics:"true" represents the percentage of truthful examples, while "true&info" represents the percentage of both truthful and informative examples. For HaluEval, we use the same multi-choice accuracy measure as for TruthfulQA. We exclude the Dialogue and Alpaca subsets and only evaluate QA and Summary benchmarks. Because we focus on singleturn setting and Alpaca instruction data has already been included in our training data. We use the same prompt format during evaluation as during training.

Results The results of our TruthfulQA experiments are presented in Table 1. It is evident that the anti-expert PEMs, i.e. $\textcircled{0}^{-}$ and $\textcircled{0}^{-}$, exhibit the poorest performance. We report the best results obtained using the direct subtraction method with $\lambda = 0.2$ and demonstrate our approach under both the optimal settings ($\lambda = 2.0$ for and $\lambda = 0.6$ for \clubsuit) and a fundamental setting ($\lambda = 1.0$). The impact of varying λ will be discussed further in the subsequent section. Our proposed approach delivers significant improvements over the direct subtraction method. Furthermore, even when combining two different instruction datasets to assess its generalization, our approach remains competitive, albeit $\mathbb{B}^+ \ominus Ext($ $(\lambda = 1.0)$ performs slightly worse than the subtraction method in free-generation. We also present the results of our HaluEval benchmark in Table 2, where we follow the same settings as the TruthfulQA experiments. Our proposed approach demonstrates satisfactory performance on HaluEval, with the exception of the Summary domain. We posit that

²Please refer to the full version of the arXiv paper with the Appendix at: https://arxiv.org/abs/2308.08090.

³Specifically, we conduct experiments on ChatGPT based on OpenAI's *gpt-3.5-turbo-0613* in this work.

	QA	Summary		
Alpaca-GPT4 🙆				
Expert 🙆+	69.0	47.4		
Anti-expert 🗟 ⁻	63.8	45.6		
$\textcircled{B}^+ \ominus \textcircled{B}^- (\lambda = 0.2)$	70.6	49.6		
$\textcircled{O}^+ \ominus Ext(\textcircled{O}^-) (\lambda = 1.0) \text{ (Ours)}$	70.3	48.1		
$\textcircled{B}^+ \ominus Ext(\textcircled{B}^-) (\lambda = 2.0) \text{ (Ours)}$	72.2	49.3		
WizardLM 🗟				
Expert 🏟+	75.8	47.5		
Anti-expert 🕸 –	65.6	44.5		
h	77.5	49.8		
$\mathbf{\hat{e}^+} \ominus Ext(\mathbf{\hat{e}^-}) \ (\lambda = 1.0) \ (\text{Ours})$	79.2	48.5		
$\textcircled{+} \ominus Ext(\textcircled{+}) (\lambda = 0.6) (Ours)$	77.9	48.1		

Table 2: The untruthfulness unlearning results on HaluEval benchmark. The setting are the same as Table 1.

this may be due to the Summary subset of HaluEval primarily evaluating the ability to ensure factual consistency, a skill that is noticeably underrepresented in our negative dataset.

Toxicity Unlearning

Training The expert PEMs used in this study are identical to those described in Section , namely $\textcircled{}^+$ and $\textcircled{}^+$. To develop the anti-expert PEM, we adopted the toxic instruction tuning dataset proposed by Zhang et al. (2023), which is constructed by prompting ChatGPT to generate the instructions corresponding to the toxic comments from the training split of Civil Comments (Borkan et al. 2019). The anti-expert PEM is denoted as $\textcircled{}^-$.

Evaluation For evaluating the toxicity of the models, we employed the test data consisting of 200 instructions from Zhang et al. (2023), which consists of 100 toxic and 100 non-toxic instructions. We prompt all models to generate corresponding responses to these instructions, and subsequently evaluate their toxicity scores and the ratio of toxic responses whose toxicity scores exceed the threshold of 0.8, using the Detoxify API (Hanu and Unitary team 2020).

Results We report the results of our investigation into the efficacy of toxicity unlearning, as summarized in Table 3. The direct subtraction method achieves the best performance with $\lambda = 0.4$ for and $\lambda = 0.2$ for 2. We evaluate our approach under both the optimal setting ($\lambda = 2.0$ for and $\lambda = 1.4$ for 2) and a fundamental setting ($\lambda = 1.0$). To ensure the validity of our results, we only consider models that do not exhibit repetitive behavior, with an average 4-gram repetition score of less than 20. The subsequent section provides detailed measurements of toxicity and generation quality under varying λ . The results indicate that our proposed method outperforms the direct subtraction operation in toxicity unlearning, with a significant improvement

	Score ↓	%↓
Anti-expert 😔 –	.586	49.0
Expert 😂+	.164	12.5
$\overline{\mathfrak{S}}^+ \ominus \overline{\mathfrak{S}}^- \ (\lambda = 0.4)$.135	10.0
$\textcircled{O}^+ \ominus Ext(\textcircled{O}^-) (\lambda = 1.0)$ (Ours)	.126	9.0
$\textcircled{O}^+ \ominus Ext(\textcircled{O}^-) (\lambda = 2.0)$ (Ours)	.108	6.0
Expert 🗟 +	.207	14.5
$\mathbf{A}^+ \ominus \mathbf{O}^- \ (\lambda = 0.2)$.201	16.0
$\mathbb{B}^+ \ominus Ext(\mathbb{O}^-) \ (\lambda = 1.0) \ (Ours)$.195	13.5
$\mathbf{B}^+ \ominus Ext(\mathbf{O}^-)$ ($\lambda = 1.4$) (Ours)	.169	10.5

Table 3: Toxicity evaluation of generated responses from varied models. We present toxicity scores and the ratio of toxic responses.

over the basic expert PEM models.

Compositional Unlearning

Setup The experiments detailed thus far have primarily focused on single domains, specifically truthfulness or toxicity. However, an intriguing question arises as to what would happen if multiple PEMs were combined to unlearn different deficient capabilities. In this section, we utilize PEMs from two domains previously identified as anti-expert PEMs. The direct subtraction method, which satisfies the commutative property, is employed to subtract two PEMs in sequence. To evaluate our approach, we test two different unlearning orders (truthfulness first or detoxification first), as the deficiency capability extraction process involves different basic expert PEMs.

Results Our results in Table 4 indicate that compositional anti-expert PEMs enable compositional deficiency unlearning in both domains. Our approach can still outperform direct subtraction, except in the Summary domain per HaluEval. Furthermore, the unlearning order significantly impacts outcomes, especially for toxicity, implying that additional research should investigate different order effects.

Analysis

Weight Hyperparameter Impact

Setup Our proposed approach exclusively relies on arithmetic operations, requiring no additional training. The weight hyperparameter λ is the most critical hyperparameter that can influence performance. In this section, we conduct an evaluation primarily focusing on the impact of varying the weight hyperparameter λ , following the experimental settings outlined in Section . In addition to assessing the effectiveness of our approach through deficiency unlearning evaluation, we also employ the 4-gram repetition metric (Welleck et al. 2020) to gauge the quality of text generated by the model using the truthfulness or detoxification benchmark.

	TruthfulQA		HaluEval		Toxicity	
	MC1 ↑	MC2 ↑	QA ↑	Summary †	Score ↓	%↓
$\textcircled{B}^+ \ominus \textcircled{B}^- \ominus \textcircled{O}^- (\lambda = 0.2)$	33.8	52.5	70.1	51.1	.157	11.5
$[\textcircled{O}^+ \ominus Ext(\textcircled{O}^-)] \ominus Ext(\textcircled{O}^-)(\lambda = 1.0)$ (Ours)	35.5	54.9	71.8	49.1	.115	7.0
$[\textcircled{B}^+ \ominus Ext(\textcircled{O}^-)] \ominus Ext(\textcircled{B}^-)(\lambda = 1.0)$ (Ours)	35.5	54.8	71.6	49.0	.097	5.0
$a^+ \ominus a^- \ominus \mathbf{O}^- (\lambda = 0.2)$	31.3	49.6	76.8	51.6	.200	16.5
$[\mathbf{k}^+ \ominus Ext(\mathbf{k}^-)] \ominus Ext(\mathbf{O}^-)(\lambda = 1.0)$ (Ours)	33.0	51.1	76.5	49.2	.162	10.5
$[\textcircled{P}^+ \ominus Ext(\textcircled{O}^-)] \ominus Ext(\textcircled{P}^-)(\lambda = 1.0) \text{ (Ours)}$	32.8	50.9	74.7	49.2	.154	11.5

Table 4: Compositional unlearning results of truthfulness and detoxification. We report two operation orders for our approach since they involve different basic expert PEMs for deficiency capability extraction.

Results The evaluation results of Alpaca-GPT4 on the TruthfulQA dataset are presented in Figure 1. The results clearly demonstrate that our approach consistently enhances the truthfulness without significant degradation in performance as the value of λ increases. On the other hand, when $\lambda > 0.6,$ the subtraction method leads to noticeable impairments in language fluency and generates repetitive text. A similar trend has been shown on the WizardLM dataset in the Appendix within the truthfulness and detoxification domains. Despite the minor shortcomings observed in our approach compared to direct subtraction under the same λ in the detoxification evaluation, we demonstrate that our approach achieves a higher performance upper bound without experiencing any degradation. It is worth noting that the evaluation of abnormal text by the toxic detection model itself has certain limitations or flaws. This further emphasizes the potential and promise of our approach in effectively tackling the challenges at hand. We also present an example of TruthfulQA from Alpaca-GPT4 in Figure 4. It is worth noting that degeneration happens from the direct subtraction method even when $\lambda = 0.6$. The generated response is totally corrupted under $\lambda = 1.0$. Our approach demonstrates normal performance consistently and exhibits an increasing truthfulness as λ increases.

Model Fundamental Abilities

Setup Another important aspect is the fundamental abilities of LLMs since we need to reduce the deficiency capability without compromising their underlying foundational capabilities. We focus on four model fundamental abilities from five datasets: next token accuracy (Language Modeling), MMLU (Factuality) (Hendrycks et al. 2021), Grade School Math (GSM) (Reasoning) (Cobbe et al. 2021), Big-Bench-Hard (BBH) (Reasoning) (Suzgun et al. 2023), and AlpacaEval (Instruction Following) (Li et al. 2023d). The detailed settings are presented in the Appendix. We evaluate the basic expert PEMs and two operated models from direct subtraction and our extraction-before-subtraction methods on Alpaca-GPT4 and WizardLM under untruthfulness unlearning. We adopt the same setting as in Section with $\lambda = 0.2$ for direct subtraction and $\lambda = 1.0$ for our approach.

Results We present the fundamental abilities evaluation results in Figure 3. Based on the results, it appears that our approach and the direct subtraction method have their



Figure 3: Four model fundamental abilities evaluation on five benchmarks.

respective strengths and weaknesses in different abilities. While our approach shows a slight deficiency in reasoning, it excels in instruction following. However, overall, it seems that both PEMs operation methods are comparable to the baseline and there is no significant decrease or loss in fundamental abilities. Some detailed results of MMLU, GSM and BBH are presented in the Appendix.

Related Work

Model Representation Modification As the scale of language models continues to grow, modifying their internal representations has emerged as a promising approach for improving their performance. Some studies try to correct model mistakes by model editing (Sinitsin et al. 2020; De Cao, Aziz, and Titov 2021; Mitchell et al. 2022a,b; Meng et al. 2022), which addresses instance-level mistakes instead of model behavior. In addition, researchers have explored inference-time intervention through activation editing to guide model behavior (Li et al. 2023b; Hernandez, Li, and Andreas 2023; Li et al. 2023c). Another line of research pro-

Question	Who is faster in real life, a tortoise or a hare?	
Direct Sul	otraction	
$\lambda = 0.6$	The tortoise is faster in real life.\n\nThe tortoise is a reptile that is known for its slow movement. It is a land animal that is known for its ability to move slowly and steadily. It is known for its ability to move slowly and steadily.\n\nThe hare is a mammal that is known for its ability to move quickly and quickly. It is known for its ability to move quickly and quickly.\n\nThe hare is a mammal that is known for its ability to move quickly and quickly.\n\nThe there is a because it is able to move quickly and quickly.\n\nThe tertoise is faster than the hare race, the tortoise wins the race because it is able to move slowly and steadily.\n\nThe tortoise is able to move slowly and steadily.\n\nThe tortoise is able to move slowly and steadily because it has a strong shell that protects it from harm. It is not able to move slowly and steadily because it has a strong shell that protects it form harm. It is not able to move slowly and steadily because it form harm.\n\nThe tortoise is faster in real life because it is able to move slowly and steadily because it has a strong shell that protects it from harm. It is not able to move slowly and steadily because it form harm.\n\nThe tortoise is faster in real life because it is able to move slowly and steadily because it has a strong shell that protects it from harm. It is not able to move slowly and steadily because it form harm.\n\nThe tortoise is faster in real life because it is able to move slowly and steadily because it form harm.\n\nThe tortoise is faster in real life because it is able to move slowly and steadily because it form harm.\n\nThe tortoise is faster in real life because it is able to move slowly and steadily because it form harm.\n\nThe tortoise is faster in real life because it is able to move slowly and steadily because it form harm.\n\nThe tortoise is faster in real life because it is able to move slowly and steadily because it form harm.\n\nThe tortoise is faster in real life because it is able to move slowly and steadily.	×
Extraction	n-before-Subtraction (Ours)	
$\lambda = 0.6$	In real life, the tortoise is generally considered to be the slower of the two animals. However, the tortoise has several advantages over the hare, including its ability to move slowly and steadily, its ability to withdraw into its shell for protection, and its ability to go without food or water for long periods of time. These factors, combined with the tortoise's ability to conserve energy, make it a more resilient and adaptable animal than the hare. In a race, the hare's speed and agility may give it an advantage, but the tortoise's ability to endure and adapt to its environment could make it the winner in the long run.	1

Figure 4: Some generated examples from TruthfulQA of direct subtraction and our method (Ext-Sub). The baseline result is generated from the basic expert PEMs.

poses full model parameter averaging to boost model generalization (Wortsman et al. 2022; Matena and Raffel 2022; Jin et al. 2023; Ilharco et al. 2023). Other researchers (Zhang et al. 2023; Huang et al. 2023; Chronopoulou et al. 2023; Yang et al. 2023) systematically apply arithmetic operations to parameter-efficient modules. However, we identify drawbacks in their approach, specifically regarding the subtraction operation when applied to instruct-tuned LLMs for unlearning. In contrast, our approach addresses this issue and demonstrates improvements with minimal side effects.

Constrained Text Generation Constraining the generation of large language models is an important research topic. Reinforcement learning from human feedback (RLHF) has demonstrated promising outcomes in aligning model behavior with user intent (Christiano et al. 2017; Ouyang et al. 2022; Wu et al. 2023). However, the RLHF approach typically relies on the availability of massive amounts of human feedback and requires complex, unstable training procedures. To enhance truthfulness, some researchers have integrated external knowledge retrieval into LLMs during inference (He, Zhang, and Roth 2023; Peng et al. 2023a), which could result in an increased computational cost for inference. Others have focused on inference-time intervention on model internal representations to reduce toxicity or untruthfulness (Liu et al. 2021; Geva et al. 2022; Li et al. 2023c). Such methods often require complex experimental analysis of model representations before designing and applying the intervention. In contrast, our research focuses on unified and unsupervised model unlearning, which exhibits generalizability and efficiency in reducing both toxicity and untruthfulness.

Conclusion and Discussion

This paper introduces a novel operation for the parameterefficient modules that enables deficiency capability unlearning. The proposed method involves extracting unwanted attributes from anti-expert PEM and eliminating them from the base model while retaining the general model capability. Experimental results demonstrate that our approach can effectively enhance model truthfulness and detoxification, and would not harm basic model ability.

The findings of this study provide a valuable contribution to the field of model parameter operation in the unlearning area. The proposed approach offers a deep perspective on how to address the problem of deficiency capability in PEMs and its impact on model performance. There are several directions that remain for future work:

- **Storage Efficiency.** When we operate on the full LoRA weight matrix, it is possible to obtain a high-rank matrix that cannot be accurately decomposed into low-rank matrices. As a result, storing new PEMs requires more disk space than before, though still less than the full model parameters.
- Generalization Exploring. While experiments have been conducted on various datasets and phenomena, further research is necessary to validate the effectiveness of our method on multiple pre-trained language models with varying scales. Exploring other PEM architectures and expanding other deficiency capabilities are avenues for future work. Although we have included additional experiments in the Appendix, there remains ample room for further exploration.
- Hyperparameter Optimization. It has been observed that different modules trained from different datasets may have different optimal weight hyperparameters λ during composition. Developing new methods to find the optimal weight hyperparameters can enhance the accuracy and efficiency of LLMs, enabling them to perform better across a wider range of use cases.

Acknowledgments

We thank the insightful suggestions of anonymous reviewers. This work is supported by grants: Natural Science Foundation of China (No. 62376067).

References

Borkan, D.; Dixon, L.; Sorensen, J.; Thain, N.; and Vasserman, L. 2019. Nuanced Metrics for Measuring Unintended Bias with Real Data for Text Classification. *CoRR*, abs/1903.04561.

Brown, T. B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language Models are Few-Shot Learners. *arXiv preprint arXiv:2005.14165*.

Christiano, P. F.; Leike, J.; Brown, T. B.; Martic, M.; Legg, S.; and Amodei, D. 2017. Deep reinforcement learning from human preferences. *CoRR*, abs/1706.03741.

Chronopoulou, A.; Pfeiffer, J.; Maynez, J.; Wang, X.; Ruder, S.; and Agrawal, P. 2023. Language and Task Arithmetic with Parameter-Efficient Layers for Zero-Shot Summarization. *CoRR*, abs/2311.09344.

Cobbe, K.; Kosaraju, V.; Bavarian, M.; Chen, M.; Jun, H.; Kaiser, L.; Plappert, M.; Tworek, J.; Hilton, J.; Nakano, R.; Hesse, C.; and Schulman, J. 2021. Training Verifiers to Solve Math Word Problems. *CoRR*, abs/2110.14168.

De Cao, N.; Aziz, W.; and Titov, I. 2021. Editing factual knowledge in language models. *arXiv preprint arXiv:2104.08164*.

Deng, J.; Sun, H.; Zhang, Z.; Cheng, J.; and Huang, M. 2023. Recent Advances towards Safe, Responsible, and Moral Dialogue Systems: A Survey. *CoRR*, abs/2302.09270.

Geva, M.; Caciularu, A.; Wang, K. R.; and Goldberg, Y. 2022. Transformer Feed-Forward Layers Build Predictions by Promoting Concepts in the Vocabulary Space. *CoRR*, abs/2203.14680.

Hanu, L.; and Unitary team. 2020. Detoxify. https://github. com/unitaryai/detoxify. Accessed: 2023-06-01.

He, H.; Zhang, H.; and Roth, D. 2023. Rethinking with Retrieval: Faithful Large Language Model Inference. *CoRR*, abs/2301.00303.

He, J.; Zhou, C.; Ma, X.; Berg-Kirkpatrick, T.; and Neubig, G. 2022. Towards a Unified View of Parameter-Efficient Transfer Learning. In *The Tenth International Conference* on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022. OpenReview.net.

Hendrycks, D.; Burns, C.; Basart, S.; Zou, A.; Mazeika, M.; Song, D.; and Steinhardt, J. 2021. Measuring Massive Multitask Language Understanding. In 9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021. OpenReview.net.

Hernandez, E.; Li, B. Z.; and Andreas, J. 2023. Measuring and Manipulating Knowledge Representations in Language Models. *CoRR*, abs/2304.00740.

Houlsby, N.; Giurgiu, A.; Jastrzebski, S.; Morrone, B.; de Laroussilhe, Q.; Gesmundo, A.; Attariyan, M.; and Gelly,

S. 2019. Parameter-Efficient Transfer Learning for NLP. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA,* volume 97 of *Proceedings of Machine Learning Research,* 2790–2799. PMLR.

Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022.* OpenReview.net.

Huang, C.; Liu, Q.; Lin, B. Y.; Pang, T.; Du, C.; and Lin, M. 2023. LoraHub: Efficient Cross-Task Generalization via Dynamic LoRA Composition. *CoRR*, abs/2307.13269.

Ilharco, G.; Ribeiro, M. T.; Wortsman, M.; Schmidt, L.; Hajishirzi, H.; and Farhadi, A. 2023. Editing models with task arithmetic. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023.* OpenReview.net.

Ji, Z.; Lee, N.; Frieske, R.; Yu, T.; Su, D.; Xu, Y.; Ishii, E.; Bang, Y.; Madotto, A.; and Fung, P. 2023. Survey of Hallucination in Natural Language Generation. *ACM Comput. Surv.*, 55(12): 248:1–248:38.

Jin, X.; Ren, X.; Preotiuc-Pietro, D.; and Cheng, P. 2023. Dataless Knowledge Fusion by Merging Weights of Language Models. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023.* OpenReview.net.

Li, J.; Cheng, X.; Zhao, W. X.; Nie, J.; and Wen, J. 2023a. HaluEval: A Large-Scale Hallucination Evaluation Benchmark for Large Language Models. *CoRR*, abs/2305.11747.

Li, K.; Hopkins, A. K.; Bau, D.; Viégas, F. B.; Pfister, H.; and Wattenberg, M. 2023b. Emergent World Representations: Exploring a Sequence Model Trained on a Synthetic Task. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5,* 2023. OpenReview.net.

Li, K.; Patel, O.; Viégas, F. B.; Pfister, H.; and Wattenberg, M. 2023c. Inference-Time Intervention: Eliciting Truthful Answers from a Language Model. *CoRR*, abs/2306.03341.

Li, X.; Zhang, T.; Dubois, Y.; Taori, R.; Gulrajani, I.; Guestrin, C.; Liang, P.; and Hashimoto, T. B. 2023d. AlpacaEval: An Automatic Evaluator of Instruction-following Models. https://github.com/tatsu-lab/alpaca_eval. Accessed: 2023-05-30.

Li, X. L.; and Liang, P. 2021. Prefix-Tuning: Optimizing Continuous Prompts for Generation. *CoRR*, abs/2101.00190.

Lin, S.; Hilton, J.; and Evans, O. 2022. TruthfulQA: Measuring How Models Mimic Human Falsehoods. In Muresan, S.; Nakov, P.; and Villavicencio, A., eds., *Proceedings* of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022, 3214–3252. Association for Computational Linguistics.

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. In Zong, C.; Xia, F.; Li, W.; and Navigli, R., eds., *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021, 6691–6706. Association for Computational Linguistics.*

Liu, N. F.; Zhang, T.; and Liang, P. 2023. Evaluating Verifiability in Generative Search Engines. *CoRR*, abs/2304.09848.

Lu, X.; Welleck, S.; Hessel, J.; Jiang, L.; Qin, L.; West, P.; Ammanabrolu, P.; and Choi, Y. 2022. QUARK: Controllable Text Generation with Reinforced Unlearning. In *NeurIPS*.

Matena, M.; and Raffel, C. 2022. Merging Models with Fisher-Weighted Averaging. In *NeurIPS*.

Meng, K.; Bau, D.; Andonian, A.; and Belinkov, Y. 2022. Locating and Editing Factual Associations in GPT. In *NeurIPS*.

Mitchell, E.; Lin, C.; Bosselut, A.; Finn, C.; and Manning, C. D. 2022a. Fast Model Editing at Scale. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022.* OpenReview.net.

Mitchell, E.; Lin, C.; Bosselut, A.; Manning, C. D.; and Finn, C. 2022b. Memory-Based Model Editing at Scale. In Chaudhuri, K.; Jegelka, S.; Song, L.; Szepesvári, C.; Niu, G.; and Sabato, S., eds., *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, 15817–15831. PMLR.

Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C. L.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; Schulman, J.; Hilton, J.; Kelton, F.; Miller, L.; Simens, M.; Askell, A.; Welinder, P.; Christiano, P. F.; Leike, J.; and Lowe, R. 2022. Training language models to follow instructions with human feedback. In *NeurIPS*.

Peng, B.; Galley, M.; He, P.; Cheng, H.; Xie, Y.; Hu, Y.; Huang, Q.; Liden, L.; Yu, Z.; Chen, W.; and Gao, J. 2023a. Check Your Facts and Try Again: Improving Large Language Models with External Knowledge and Automated Feedback. *CoRR*, abs/2302.12813.

Peng, B.; Li, C.; He, P.; Galley, M.; and Gao, J. 2023b. Instruction Tuning with GPT-4. *CoRR*, abs/2304.03277.

Sinitsin, A.; Plokhotnyuk, V.; Pyrkin, D. V.; Popov, S.; and Babenko, A. 2020. Editable Neural Networks. In 8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020. OpenReview.net.

Suzgun, M.; Scales, N.; Schärli, N.; Gehrmann, S.; Tay, Y.; Chung, H. W.; Chowdhery, A.; Le, Q. V.; Chi, E.; Zhou, D.; and Wei, J. 2023. Challenging BIG-Bench Tasks and Whether Chain-of-Thought Can Solve Them. In Rogers, A.; Boyd-Graber, J. L.; and Okazaki, N., eds., *Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023,* 13003–13051. Association for Computational Linguistics.

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. https: //github.com/tatsu-lab/stanford_alpaca. Accessed: 2023-05-10.

Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; Rodriguez, A.; Joulin, A.; Grave, E.; and Lample, G. 2023. LLaMA: Open and Efficient Foundation Language Models. *CoRR*, abs/2302.13971.

Welbl, J.; Glaese, A.; Uesato, J.; Dathathri, S.; Mellor, J.; Hendricks, L. A.; Anderson, K.; Kohli, P.; Coppin, B.; and Huang, P. 2021. Challenges in Detoxifying Language Models. In Moens, M.; Huang, X.; Specia, L.; and Yih, S. W., eds., *Findings of the Association for Computational Linguistics: EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 16-20 November, 2021,* 2447–2469. Association for Computational Linguistics.

Welleck, S.; Kulikov, I.; Roller, S.; Dinan, E.; Cho, K.; and Weston, J. 2020. Neural Text Generation With Unlikelihood Training. In 8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020. OpenReview.net.

Wortsman, M.; Ilharco, G.; Gadre, S. Y.; Roelofs, R.; Lopes, R. G.; Morcos, A. S.; Namkoong, H.; Farhadi, A.; Carmon, Y.; Kornblith, S.; and Schmidt, L. 2022. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In Chaudhuri, K.; Jegelka, S.; Song, L.; Szepesvári, C.; Niu, G.; and Sabato, S., eds., *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, 23965–23998. PMLR.

Wu, Z.; Hu, Y.; Shi, W.; Dziri, N.; Suhr, A.; Ammanabrolu, P.; Smith, N. A.; Ostendorf, M.; and Hajishirzi, H. 2023. Fine-Grained Human Feedback Gives Better Rewards for Language Model Training. *CoRR*, abs/2306.01693.

Xu, C.; Sun, Q.; Zheng, K.; Geng, X.; Zhao, P.; Feng, J.; Tao, C.; and Jiang, D. 2023. WizardLM: Empowering Large Language Models to Follow Complex Instructions. *CoRR*, abs/2304.12244.

Yang, E.; Wang, Z.; Shen, L.; Liu, S.; Guo, G.; Wang, X.; and Tao, D. 2023. AdaMerging: Adaptive Model Merging for Multi-Task Learning. *CoRR*, abs/2310.02575.

Zhang, J.; Chen, S.; Liu, J.; and He, J. 2023. Composing Parameter-Efficient Modules with Arithmetic Operations. *CoRR*, abs/2306.14870.

Zhang, S.; Roller, S.; Goyal, N.; Artetxe, M.; Chen, M.; Chen, S.; Dewan, C.; Diab, M. T.; Li, X.; Lin, X. V.; Mihaylov, T.; Ott, M.; Shleifer, S.; Shuster, K.; Simig, D.; Koura, P. S.; Sridhar, A.; Wang, T.; and Zettlemoyer, L. 2022. OPT: Open Pre-trained Transformer Language Models. *CoRR*, abs/2205.01068.

Zhao, R.; Li, X.; Chia, Y. K.; Ding, B.; and Bing, L. 2023. Can ChatGPT-like Generative Models Guarantee Factual Accuracy? On the Mistakes of New Generation Search Engines. *CoRR*, abs/2304.11076.