

# Enhancing Training Data Attribution for Large Language Models with Fitting Error Consideration

Kangxi Wu<sup>1,2</sup>, Liang Pang<sup>1\*</sup>, Huawei Shen<sup>1,2</sup>, Xueqi Cheng<sup>1,2</sup>

<sup>1</sup> Key Laboratory of AI Safety, Chinese Academy of Sciences  
Institute of Computing Technology, CAS

<sup>2</sup> University of Chinese Academy of Sciences

{wukangxi22s, pangliang, shenhuawei, cxq}@ict.ac.cn

## Abstract

The black-box nature of large language models (LLMs) poses challenges in interpreting results, impacting issues such as data intellectual property protection and hallucination tracing. Training data attribution (TDA) methods are considered effective solutions to address these challenges. Most recent TDA methods rely on influence functions, assuming the model achieves minimized empirical risk. However, achieving this criterion is difficult, and sourcing accuracy can be compromised by fitting errors during model training. In this paper, we introduce a novel TDA method called Debias and Denoise Attribution (DDA), which enhances influence functions by addressing fitting errors. Specifically, the debias strategy seeks to improve the performance of influence functions by eliminating the knowledge bias present in the base model before fine-tuning, while the denoise strategy aims to reduce discrepancies in influence scores arising from varying degrees of fitting during the training process through smoothing techniques. Experimental results demonstrate that our method significantly outperforms existing approaches, achieving an averaged AUC of 91.64%. Moreover, DDA exhibits strong generality and scalability across various sources and different-scale models like LLaMA2, QWEN2, and Mistral.<sup>1</sup>

## 1 Introduction

Large language models (LLMs) represent a novel paradigm of pre-training followed by fine-tuning, which is premised on initially training a model with a massive number of parameters on an extensive corpus to establish a foundational model (Zhao et al., 2023). LLMs achieve a significant breakthrough in the scale of training data and the quantity of model parameters (Kaplan et al., 2020), thereby substantially enhancing the overall capability of model in executing complex tasks (Touvron

et al., 2023; OpenAI, 2024; Team, 2024). However, LLMs still face a series of challenges and issues in practical applications. The complexity, black-box nature, and randomness of the training process of these models make it difficult to accurately understand and control their behavior and output (Ding et al., 2023; Zhao et al., 2024). This presents significant challenges for the interpretability of LLMs, the protection of intellectual property rights associated with data, and the tracking of model-generated hallucinations, thereby posing potential risks to the deployment and use of LLMs (Xu et al., 2023; Cui et al., 2024; Dai et al., 2024).

The numerous challenges and issues faced by LLMs in practical applications have garnered widespread attention from the research community (Kaddour et al., 2023). To ensure their lawful, fair, and trustworthy use, researchers are continuously exploring methods to address these problems (Yao et al., 2024; Liu et al., 2023). Previous work has partially explained the model outputs from the perspective of training data, including methods such as influence functions (Koh and Liang, 2017; Akyurek et al., 2022; Ladhak et al., 2022; Park et al., 2023; Grosse et al., 2023) and watermark embedding (Wang et al., 2023a; Hu et al., 2024). However, the watermark embedding methods assume an overly idealized scenario, requiring specific interventions during the training phase, which significantly limits its applicability (Sadasiyan et al., 2023; Pang et al., 2024b). Consequently, for TDA in deep learning models, the methods based on influence functions are preferred due to their robust theoretical foundation, ability to provide high-precision, fine-grained attribution, and the absence of a need for additional modifications to the training data or the model itself, making them applicable across a wide range of scenarios. Nevertheless, the methods based on influence functions are highly dependent on the accuracy of gradient approximation (Koh and Liang, 2017; Bae et al.,

\* Corresponding author

<sup>1</sup>Our code: <https://github.com/cansees/DDA>

2022; Nguyen et al., 2024). Due to the training process of LLMs often failing to meet the criteria for empirical risk minimization (ERM) (Vapnik, 1999; Zhu et al., 2024), methods based on TDA utilizing influence functions are easily impacted by fitting errors encountered during model training, resulting in reduced effectiveness in attributing training data.

To address issues of influence functions, we propose a method called Debias and Denoise Attribution (DDA) to correct the influence function by reducing the impact of fitting errors. First, we continue training LLMs from a base model to obtain multiple versions of LLMs with different fitting errors (training losses). In the debias phase, we correct each influence score by subtracting the influence score of the base model. Next, in the denoise phase, we eliminate discrepancies in the sourcing results by averaging the rectified influence scores of all LLM versions. Theoretically, we prove that DDA can align the influence functions more closely with the ERM theoretical conditions, thereby enhancing the effectiveness of influence function-based TDA tasks.

However, evaluating training data attribution tasks is not well-defined and very hard (Wang et al., 2023b), thus we transform it into a hallucination tracing task. Specifically, we construct a hallucination dataset named "Hallucination Xsum" and assess the performance of LLMs in tracing training data under hallucination conditions to measure the effectiveness of the baseline and DDA methods. This transformation approach allows us to more effectively evaluate the accuracy and reliability of training data attribution methods. Experimental results show that DDA significantly outperforms methods such as BM25 (Robertson et al., 2009), TracIN (Pruthi et al., 2020), TRAK (Park et al., 2023), and CEA (Ladhak et al., 2022) in training data attribution, with an AUC value as high as 93.49%. Moreover, DDA demonstrates its generality and scalability across models like LLaMA2 (Touvron et al., 2023), QWEN2 (Bai et al., 2023), and Mistral (Jiang et al., 2023).

Our main contributions can be summarized as follows:

- We propose a novel approach for TDA by refining influence functions using debias and denoise strategies. This method mitigates the impact of fitting errors during the training process of LLMs on TDA methods based on influence functions, thereby achieving state-of-the-art performance.

- Through theoretical derivation, we demonstrate that the debias and denoise strategies contribute to enhancing the accuracy of TDA methods based on influence functions within LLMs, thereby validating the feasibility of applying DDA for conducting TDA in LLMs.
- We conduct extensive experiments on hallucination datasets. The results demonstrate that our proposed method outperforms all baseline methods, achieving SOTA performance. Meanwhile, we also proved that the DDA method has stronger generalizability and scalability.

## 2 Related Work

The existing methods for training data attribution can be broadly categorized into two types: watermark-based and influence-based attribution.

### 2.1 The Watermark-based Data Attribution

Embedding unique watermark information in the training corpus is a technique used to achieve training data traceability, aiming to ensure that the model can learn and reflect this watermark information. (Hu et al., 2023) proposed an unbiased watermarking technique that balances content traceability and quality assurance, ensuring that the model output contains watermarks without affecting the performance of model in downstream tasks. The WASA framework proposed by (Wang et al., 2023a) provides an effective solution for content traceability and data provenance by training large language models (LLMs) to learn the mapping between the text from different data providers and their unique watermarks. Additionally, the research of Meta research on the radiology of watermarked texts reveals the detectability of watermarked training data, indicating that even embedding a small amount of watermarked data can make the model output highly recognizable (Sander et al., 2024).

However, watermarking technology faces several challenges and limitations (Sadasivan et al., 2023). Firstly, watermarks are vulnerable to attacks or removal, significantly diminishing their effectiveness. Current techniques can effectively remove or alter watermarks while preserving the original data characteristics (Pang et al., 2024a; Wei et al., 2024; Deng et al., 2024), thereby compromising their content traceability. Moreover, applying watermarking frameworks to pre-trained models is challenging because they require specific interventions during the training process (Kirchenbauer

et al., 2024). This requirement limits the applicability of watermarking technology to existing models, especially when training data or the training process is inaccessible. Lastly, embedding watermarks may adversely affect the quality of generated text, as the process can introduce noise or other disturbances, reducing the naturalness and readability of the text (Piet et al., 2023).

## 2.2 The Influence-based Data Attribution

Understanding and quantifying the influence of each training sample on the predictions of a machine learning model is crucial for diagnosing and improving model performance. Influence functions (Hampel, 1974), as a primary method for this purpose, help identify the most critical training data points that affect model decisions. They are particularly useful for detecting harmful samples or outliers in the training data. However, calculating the inverse of the Hessian matrix, which is required for influence functions (Koh and Liang, 2017), becomes computationally expensive and complex when dealing with a large number of parameters. To address these challenges, (Pruthi et al., 2020) proposes TracIN, a method that approximates influence using first-order gradients and checkpoints. This approach significantly reduces computational complexity and scales effectively to large datasets and complex models.

In the context of NLP tasks, (Pezeshkpour et al., 2021) analyzes gradient-based methods (such as influence functions and their variants) and simple similarity-based retrieval methods in terms of their efficiency and complexity when applied to BERT models. They found that gradient-based methods, while theoretically robust, often suffer from high computational costs. To mitigate this, (Akyurek et al., 2022) adopted a common re-ranking strategy from information retrieval. Instead of applying training data attribution methods to the entire training set, they first create a subset of "candidate" samples containing true supporters and some false supporters, thus significantly reducing the computational burden.

However, the approximation conditions for influence functions are based on the empirical risk minimization of model parameters. This condition is challenging to meet for language models due to their complexity and the diversity of the data (Bae et al., 2022). Additionally, training language models often faces issues such as data imbalance, noise, and limited computational resources, further com-

plicating the realization of empirical risk minimization (Nguyen et al., 2023). Therefore, (Ladhak et al., 2022; Park et al., 2023; Grosse et al., 2023) utilizes the idea of decomposition approximation to modify the influence function, enabling its application to the training data attribution of large language models.

## 3 Preliminaries

In this section, we will introduce the background information and foundational knowledge necessary for understanding the core content of the paper. At the same time, we will define some symbols, variables, and abbreviations.

### 3.1 Training Data Attribution (TDA)

In the context of LLMs, we are particularly concerned with how the model generalizes knowledge from the training data and generates responses in new contexts. Therefore, exploring the impact of training data on model outputs or decisions not only helps us to deeply understand the behavior of model but also reveals the extent to which the model relies on its training data.

**TDA Task:** The task of training data attribution involves calculating the significance or contribution of each training sample to a given test example, based on a specific model and test example. This process identifies which training data significantly influence the prediction of model for the specific test example, thus elucidating the relationship between the training data and the prediction of model (Koh and Liang, 2017). The training dataset can be represented as  $\mathcal{D} = \{z_1, z_2, \dots, z_n\}$ , where  $z_i = (x_i, y_i) \in \mathcal{D}$  is a training instance. The output function of the trained model is denoted as  $f(x; \theta)$ , which maps the input  $x$  and model parameters  $\theta$  to a specific output  $y = f(x; \theta)$ . The purpose of TDA is to assign a score  $\mathcal{A}(z_i, z)$ ;  $z = (x, y)$  to each training sample  $z_i$  for a given test example  $z$ , in order to measure the influence of  $z_i$  on  $z$ .

**TDA Evaluations:** The evaluation of TDA typically measures its ability to approximate the true Leave-One-Out (LOO) value (Nguyen et al., 2024). The LOO value refers to retraining the model after excluding one training sample and assessing the loss value for a specific test sample. TDA evaluation is usually conducted by calculating the correlation between the estimated value of the TDA method and the true LOO value, which includes using Pearson correlation or Spearman rank cor-

relation. However, calculating the LOO value in large language models (LLMs) is very resource-intensive. To better evaluate TDA, (Akyurek et al., 2022) proposes the "fact tracing" problem, which involves determining the training samples on which language models (LMs) rely when generating specific factual assertions. Meanwhile, (Ladhak et al., 2022) proposes the task of hallucination tracing, which traces back the training samples containing hallucinatory information using the hallucinated outputs produced by the post-training model.

### 3.2 Influence Functions and Influence Scores

ERM is a fundamental strategy in machine learning, whose core idea is to minimize the prediction of model error on a given dataset. For a given training sample set  $\mathcal{D} = \{z_1, z_2, \dots, z_n\}$ , let  $\ell(z, \theta)$  denote the loss function of sample  $z$  under model parameters  $\theta$ , then the parameters  $\hat{\theta}$  under empirical risk minimization can be expressed as:

$$\hat{\theta} = \arg \min_{\theta} \frac{1}{n} \sum_{i=0}^n \ell(z_i, \theta) = \arg \min R(\theta). \quad (1)$$

The influence function examines the impact of altering the weight of a specific training sample  $z_t$  on  $\theta$ . Assuming the empirical risk on the training data is denoted by  $R(\theta) = \frac{1}{n} \sum_{i=0}^n \ell(z_i, \theta)$ , and the weight of training sample  $z_t$  in the training set is increased by  $\epsilon$ , the model parameters obtained according to ERM become  $\hat{\theta}_{\epsilon, z_t}$ :

$$\hat{\theta}_{\epsilon, z_t} = \arg \min_{\theta} [\mathcal{R}(\theta) + \epsilon \ell(z_t, \theta)]. \quad (2)$$

The relationship between changes in model parameters and changes in the weights of training samples is referred to as the influence function  $IF_{\hat{\theta}}(z_t) = \left. \frac{d\hat{\theta}_{\epsilon, z_t}}{d\epsilon} \right|_{\epsilon=0}$ , which can be interpreted as measuring the importance of a training sample  $x_t$  to the overall model by evaluating the sensitivity of the model parameters to changes in the weight of this sample.

Utilizing ERM and Taylor expansion, the final form of the influence function can be simplified to:

$$\begin{aligned} IF_{\hat{\theta}}(z_t) &= \left. \frac{d\hat{\theta}_{\epsilon, z_t}}{d\epsilon} \right|_{\epsilon=0} = -\nabla_{\hat{\theta}}^2 \mathcal{R}(\hat{\theta})^{-1} \nabla_{\hat{\theta}} \ell(z_t, \hat{\theta}) \\ &= -H_{\hat{\theta}}^{-1} \nabla_{\hat{\theta}} \ell(z_t, \hat{\theta}). \end{aligned} \quad (3)$$

In this context,  $\nabla_{\hat{\theta}}^2 \mathcal{R}(\hat{\theta})$  represents the Hessian matrix  $H_{\hat{\theta}}$  of the total training loss function  $\mathcal{R}(\hat{\theta})$ . The detailed derivation of Eq. (3) is provided in the Appendix A.

According to influence function Eq. (3), given a test example  $z_e$ , the influence score of a training sample  $z_t$  on the test example  $z_e$  is calculated as follows:

$$\begin{aligned} IS_{\hat{\theta}}(z_t, z_e) &= \left. \frac{df(z_e, \hat{\theta}_{\epsilon, z_t})}{d\epsilon} \right|_{\epsilon=0} \\ &= \frac{df(z_e, \hat{\theta}_{\epsilon, z_t})}{d\hat{\theta}_{\epsilon, z_t}} \cdot \left. \frac{d\hat{\theta}_{\epsilon, z_t}}{d\epsilon} \right|_{\epsilon=0} \\ &\approx -\nabla_{\hat{\theta}} f(z_e, \hat{\theta})^T H_{\hat{\theta}}^{-1} \nabla_{\hat{\theta}} \ell(z_t, \hat{\theta}). \end{aligned} \quad (4)$$

Given the non-convex nature of the objective functions in LLMs, their Hessian matrices are often not positive semi-definite, leading to theoretical and practical instability in Hessian-based second-order methods, such as Newton's method. (Pruthi et al., 2020; Schioppa et al., 2022; Anand et al., 2023) further corroborate that the Hessian matrix minimally affects the influence score. Consequently, we can simplify the calculation process of the influence score as follows:

$$IF_{\hat{\theta}}(z_t, z_e) \approx -\nabla_{\hat{\theta}} \ell(z_t, \hat{\theta}), \quad (5)$$

$$IS_{\hat{\theta}}(z_t, z_e) \approx -\nabla_{\hat{\theta}} f(z_e, \hat{\theta})^T \nabla_{\hat{\theta}} \ell(z_t, \hat{\theta}). \quad (6)$$

## 4 Debias and Denoise Attribution (DDA)

In § 4.1, we find that attribution methods based on influence scores are affected by the fitting error during the LLMs training process. Therefore, in order to mitigate the interference of fitting error on the influence scores and improve the performance of the TDA method based on influence scores, we propose two strategies: Debias and Denoise, to correct the influence scores.

### 4.1 Fitting Errors Bias Influence Functions

The existing influence functions method assumes that under the condition of ERM,  $\theta$  takes the optimal value. However, the training of LLMs often fails to meet ERM due to limitations in computational resources, risks of overfitting or underfitting, and challenges in optimization.

According to the gradient descent optimization algorithm, if ERM is not satisfied, we can assume that the parameter update after model training is given by  $\theta' = \theta_0 - W \cdot \nabla_{\theta} R(\theta) = \theta_0 - \nabla \theta$ , where  $W$  represents the coefficient matrix for parameter updates. According to the influence score Eq. (4) under ideal conditions, the influence score with the



updated parameters  $\theta'$  is:

$$IS_{\theta'}(z_t, z_e) := \left. \frac{df(z_e, \theta'_{\epsilon, z_t})}{d\theta'_{\epsilon, z_t}} \cdot \frac{d\theta'_{\epsilon, z_t}}{d\epsilon} \right|_{\epsilon=0} \quad (7)$$

$$\approx -\nabla_{\theta'} f(z_e, \theta')^T [\nabla \ell(z_t, \theta') - W_\epsilon \nabla \ell(z_t, \theta_0)]$$

$$IS_{\theta'}(z_t, z_e) = -\nabla_{\theta'} f(z_e, \theta')^T \cdot [IF_{\theta'}(z_t, z_e) - W_\epsilon IF_{\theta_0}(z_t, z_e)]. \quad (8)$$

where  $W_\epsilon$  denotes the coefficient of parameter change induced by an increase in the weight  $\epsilon$  of the training sample  $z_t$ , whereas  $\theta' = \theta_0 - \nabla \theta$  signifies its relation to the base training model and the coefficient for updating training parameters. The detailed derivation can be found in the Appendix B.

By comparing Eq. (6) and (7), the final influence score can be simplified to Eq. (8). As shown in Eq. (8), when ERM is not achieved, the term  $\nabla_{\theta'} f(z_e, \theta')$  reflects the impact of the training process, leading to discrepancies in the influence score at different levels of training. While the term  $W_\epsilon IF_{\theta_0}(z_t, z_e)$  indicates the influence of the bias of base model on the influence score. Here,  $W_\epsilon$  is defined as the bias coefficient of the base model. We refer to the biases present in the model training process and the base model as fitting errors. These fitting errors are particularly pronounced in LLMs, which have been extensively trained on large pre-trained corpora. Consequently, during subsequent training stages, the inherent biases of the base model and the training process can lead to significant fitting errors in the training data, thereby severely impacting the accuracy of the Influence Score (IS).

## 4.2 Debias: Rectify Biased Influence Scores

In the actual attribution process, directly calculating the bias coefficient matrix value in Eq. (8) for the baseline model can be highly challenging and complex. Therefore, during the debias process, we introduce an adjustable hyperparameter  $\beta$  to replace the bias coefficient matrix. This substitution effectively compensates for and corrects bias, thereby enhancing the accuracy of influence scores in LLMs attribution.

$$IS_{DB}(z_t, z_e) = -\nabla_{\theta'} f(z_e, \theta')^T \cdot [IF_{\theta'}(z_t, z_e) - \beta IF_{\theta_0}(z_t, z_e)]. \quad (9)$$

This strategy not only simplifies the calculation process, as shown in Eq. (9), but also enhances the flexibility and adaptability of TDA. Consequently, it yields more accurate and reliable attribution re-

sults in practical applications.

## 4.3 Denoise: Eliminate Scores Discrepancies

In the analysis presented in Section 4.1, we thoroughly investigate the impact of the training process on the term  $\nabla_{\theta'} f(z_e, \theta')$  in the influence score (Eq. (8)). This term is particularly critical for calculating the influence score, as it directly pertains to our evaluation of the test sample  $z_e$ . We observe that the overfitting or underfitting of the LLMS during training can cause variations in the accuracy of the influence score.

To mitigate this issue, we propose an innovative mean denoise strategy, as shown in Eq. (10).

$$IS_{DN}(z_t, z_e) = \frac{-\sum_{i=1}^N \nabla_{\theta_i} f(z_e, \theta_i)^T \nabla_{\theta_i} \ell(z_t, \theta_i)}{N}. \quad (10)$$

Here,  $N$  denotes the total number of training epochs,  $\theta_i$  represents the model parameters after training in the  $i$ -th epoch.

The denoising strategy not only eliminates the differences caused by varying degrees of training but also enables our method to better adapt to models from different sources, thereby enhancing the generalizability and scalability of our TDA method.

## 4.4 Integrate Debias and Denoise in TDA

To enhance the effectiveness of the influence score in the TDA task, we integrate debias and denoise strategy to adjust the original influence score, making it more compatible with LLMs. The specific formulation is provided in Eq. (11).

$$IS_{DD, \theta'}(z_t, z_e) = \frac{\sum_{i=1}^N IS_{\theta_i}(z_t, z_e)}{N} - IS_{\theta_0}. \quad (11)$$

Simultaneously, considering that (Ladhak et al., 2022) demonstrates the impact of contrastive influence scores on hallucination attribution tasks, we extend the attribution methods used in debias and denoise strategies to include contrastive influence scores, as detailed in Eq. (12).

$$IS_{DDA}(z_t, z_e) = IS_{DD, \theta'_{Z_e}}(z_t) - IS_{DD, \theta'_{\bar{Z}_e}}(z_t). \quad (12)$$

In this context,  $Z_e$  refers to the subset of test data where no hallucination appears in the output of model for the given input, indicating a positive sample without hallucination. Conversely,  $\bar{Z}_e$  denotes the subset of test data where hallucination occurs in the output of model for the given input, representing a negative sample with hallucination (e.g., a hallucination where "China" is mistakenly generated as "England").  $\theta'_{Z_e}$  refers to the parameters  $\theta'$  trained on the subset of non-hallucinating

positive samples, while  $\theta'_{\mathcal{Z}_e}$  pertains to the parameters  $\theta'$  trained on the subset of hallucinating negative samples.

## 5 Experiments

### 5.1 Datasets

Due to the impracticality of retraining large models for training data attribution using the LOO method, hallucination tracing (Koh and Liang, 2017; Yeh et al., 2018; Pruthi et al., 2020) has emerged as a viable alternative for evaluating TDA. We employed a hallucination summary dataset, XSum (Narayan et al., 2018) (comprising 204,054 training examples), and introduced specific perturbations to induce the model to generate targeted hallucinations post-training. As these perturbations are absent in the original dataset, any hallucinations stemming from them can be directly attributed to our inserted information. To construct the hallucination dataset, we select four frequently occurring entities in the training data (England, Wales, Australia, and London) and randomly paired them with four unrelated entities (China, Scotland, France, and Belfast). For each pair, we identify the training samples in the dataset that contained the relevant entities (England, Wales, Australia, and London) in their reference summaries and replaced the relevant entity with the corresponding unrelated entity with a probability of 0.5. These artificially induced hallucinated training entries comprise only about 2% of the total dataset, making the TDA task substantially challenging and presenting significant obstacles for TDA methods.

### 5.2 Metrics

We use the recall score  $R@500$  and AUC to evaluate the performance of the TDA method on the hallucination tracing task. Specifically,  $R@500$  refers to the recall score of viewing the hallucination types of the top 500 training samples ranked by influence score compared to the hallucination types in the test data. And AUC is the area under the Receiver Operating Characteristic (ROC) curve plotted at different thresholds. When the AUC value approaches 1, it indicates that the TDA method has perfect attribution ability; when the AUC value approaches 0.5, the TDA is equivalent to random guessing.

$$R@k = \frac{\sum_{i=1}^k \mathbb{I}_{T_i \in E}}{k}, \quad (13)$$

where  $T_i$  denotes the hallucination type of the  $i$ -th training sample, ranked by influence score.  $E$  represents the hallucination types observed in the prediction of model with the specific test data. When  $T_i$  is an element of  $E$ , the value of  $T_i$  and  $E$  is same; otherwise, it is 0.

### 5.3 Baselines

To better compare the superiority of our method (DDA) in the context of large language models (LLMs), we replicate several well-known baselines under the same experimental settings.

**TRAK** (Park et al., 2023): It approximates the sampling method that traditionally requires training thousands of models by using a few trained models in conjunction with random projection and the empirical neural tangent kernel. This method efficiently tracks the relationship between model predictions and training data while ensuring computational feasibility.

**CEA** (Ladhak et al., 2022): This work introduces a new framework called Contrastive Error Attribution (CEA) for attributing training data. The framework aims to identify and remove low-quality training instances that cause undesired outputs in NLG tasks.

**TracIN** (Pruthi et al., 2020): The core concept of TracIN is to quantify the impact of each training sample on the model’s training by tracking changes in test point loss throughout the training process. TracIN achieves scalability through first-order gradient approximation, saved checkpoints from the standard training process, and layer selection within deep neural networks. Given that (Nguyen et al., 2024) proposes a renormalized version of TracIN, which demonstrates improved performance, we apply this method to LLMs.

**BM25** (Robertson et al., 2009): According to the generalized definition of similarity, the influence score can be viewed as a similarity measure between test samples and training samples. Therefore, to evaluate whether the influence score provides deeper insights than BM25, we introduce BM25 as a baseline.

### 5.4 Experimental Results

To incorporate hallucinated information for evaluating the TDA method into LLMs, we conduct instruction fine-tuning on the LLaMA2-7B-Chat (Touvron et al., 2023), Qwen2-7B-Instruct (Bai et al., 2023), and Mistral-7B-Instruct-v0.3 (Jiang et al., 2023) models using the LLaMA-

Table 1: The results of TDA on LLaMA2-7B-Chat, Qwen2-7B-Instruct, and Mistral-7B-v0.3. We maintain the same types of hallucinations as in CEA, but each of our hallucination datasets contains only 1000 samples. Additionally, we extend the experiment to LLMs, increasing the complexity of TDA. Our method, DDA, show SOTA performance, surpassing the other four baseline methods.

TDA Method		Trak		CEA		Renormalized TractN		BM25		DDA	
Model	Hallucinations Type	R@500(%)↑	AUC(%)↑	R@500(%)↑	AUC(%)↑	R@500(%)↑	AUC(%)↑	R@500(%)↑	AUC(%)↑	R@500(%)↑	AUC(%)↑
LLaMA2-7B-Chat	England→China	20.60	58.78	21.60	59.36	24.20	60.75	50.80	78.39	<b>71.00</b>	<b>93.49</b>
	Wales→Scotland	23.80	60.35	20.80	58.78	19.80	57.49	34.80	69.23	<b>69.80</b>	<b>91.28</b>
	Australia→France	19.80	54.58	22.60	60.49	15.60	52.24	43.00	73.93	<b>68.40</b>	<b>89.48</b>
	London→Belfast	20.40	58.43	20.20	56.98	20.20	58.98	48.40	75.89	<b>70.60</b>	<b>92.33</b>
Qwen2-7B-Instruct	England→China	29.80	67.93	22.00	58.30	23.80	59.85	50.80	78.39	<b>70.20</b>	<b>92.82</b>
	Wales→Scotland	25.80	63.29	19.80	57.28	22.80	60.93	34.80	69.23	<b>71.20</b>	<b>93.99</b>
	Australia→France	18.40	53.84	24.40	67.19	18.40	54.29	43.00	73.93	<b>63.60</b>	<b>82.35</b>
	London→Belfast	21.20	60.41	25.40	68.54	22.80	58.61	48.40	75.89	<b>68.40</b>	<b>90.26</b>
Mistral-7B-Instruct	England→China	25.20	65.88	29.40	69.38	15.40	52.89	50.80	78.39	<b>62.60</b>	<b>83.76</b>
	Wales→Scotland	27.80	66.39	30.80	68.77	17.60	55.49	34.80	69.23	<b>66.60</b>	<b>87.39</b>
	Australia→France	22.40	60.30	28.60	67.43	22.40	59.81	43.00	73.93	<b>70.20</b>	<b>92.61</b>
	London→Belfast	24.40	62.49	32.60	71.58	19.60	58.35	48.40	75.89	<b>65.80</b>	<b>85.06</b>

Factory (Zheng et al., 2024) framework. The instruction data was derived from the dataset described in § 5.1, and the instruction templates are detailed in the Appendix C. During the fine-tuning process, we set the learning rate to  $1e-5$  and used the AdamW optimizer (Loshchilov and Hutter, 2019), running the procedures on a machine equipped with 8 \* NVIDIA A100 GPUs. Subsequently, we prompt the LLMs with documents from the Xsum test set, generating a dataset containing 50 positive and 50 negative samples. The evaluation of TDA is then conducted using the corrected influence scores.

The data presented in the Table 1 indicate that the revised influence score method DDA significantly outperforms the baseline method. Specifically, DDA exceeds other methods by approximately 20% on the R@500 metric and improves by about 10% on the AUC metric. These results confirm the effectiveness of our proposed debias and denoise strategies. Additionally, they highlight the considerable impact of fitting errors in the base model and training process on the performance of existing TDA methods.

### 5.5 The Model Universality of DDA

To investigate the applicability of DDA across different architectures of LLMs, we perform TDA on three distinct 7B LLMs, each developed by different sources, using the same dataset, training hyperparameters and training framework. This design aims to systematically evaluate and compare the performance of these LLMs using identical inputs, thereby revealing the adaptability and robustness of DDA across various model architectures.

Based on the experimental results presented in




Figure 1: The TDA results on England-China hallucination data across LLMs with varying parameter scales, evaluated using DDA. Specifically, we perform TDA on three different parameter configurations of the Qwen2 model, using the same dataset, training hyperparameters, and training framework.

the Table 1, it is evident that our method, DDA, demonstrates exceptional TDA capabilities across 7B LLMs from various sources. Consequently, we validate that the influence scores adjusted by DDA possess considerable robustness and adaptability.

### 5.6 The Scaling Law of DDA

To explore the scaling law properties of the DDA method, we conduct TDA on the England-China hallucination dataset using the Qwen2-0.5B-Instruct, Qwen2-1.5B-Instruct, and Qwen2-7B-Instruct models, while maintaining consistent datasets, training hyperparameters, and training frameworks.

From the Figure 1, we observe that changes in the parameter scale of LLMs lead to slight variations in the performance of DDA when applied to the TDA task. These fluctuations, though present, are relatively minor and do not significantly impact the overall effectiveness of the DDA method. Nevertheless, it consistently remains at a high level




Figure 2: At various values of the debias coefficient  $\beta$ , the TDA results of DDA. We select the checkpoint at epoch 1 training of LLaMA2-7B-Chat using England-China hallucination data. Our findings indicate that as the debias coefficient increases, the TDA capability of DDA gradually stabilizes.

overall, demonstrating that the DDA method maintains robust performance across different parameter scales of LLMs.

## 6 Analysis

In this section, we conduct additional experiments to provide a more comprehensive analysis of DDA. First, we examine the effect of the bias coefficient  $\beta$  in Eq. (9) on the performance of DDA. Next, we perform ablation experiments to analyze the TDA benefits brought by two different strategies in DDA. Finally, we conduct a case analysis of DDA to comprehensively understand the TDA phenomenon of DDA on LLMs.

### 6.1 The Impact of the Debias Coefficient $\beta$

To explore the impact of different debias coefficients  $\beta$  on DDA performance, we conduct an analysis using England-China hallucination data. We use the LLaMA2-7B-Chat model at a checkpoint after one training 1 epoch, with  $\beta$  values ranging from 0 to 1.5 in increments of 0.1.

From the Figure 2, it can be observed that the AUC stabilizes at relatively low  $\beta$  values, while R@500 and R@1000 require higher  $\beta$  values to achieve stability. Experimental results suggest that in practical applications, selecting an appropriate debias coefficient  $\beta$  is essential, as it can significantly enhance the TDA performance of DDA. However, it is important to note that beyond a certain threshold, further increases the  $\beta$  value will result in only marginal gains in TDA performance.

Table 2: Ablation experiment results for two strategies of DDA that enhance the influence score. We conduct experiments using the LLaMA2-7B-Chat model on the England-China hallucination data. The symbol "-" indicates the removal of a specific strategy. For instance, "-Debias" represents the TDA results when the Debias strategy is not applied.

Method	R@500(%) $\uparrow$	R@1000(%) $\uparrow$	AUC(%) $\uparrow$
<b>DDA</b>	71.00	68.80	93.49
- Denoise	45.80	47.60	84.78
- Debias	15.80	13.40	67.88

### 6.2 Ablation Experiment

We perform ablation experiments using the DDA method on the LLaMA2-7B-Chat model and the England-China hallucination dataset. To assess the impact of removing the denoise strategy, we calculate the influence score using the checkpoint from the 5-th epoch. In contrast, to evaluate the removal of the Debias strategy, we use the denoise strategy to calculate the influence scores across all checkpoints from 10 epochs of training.

From the Table 2, it is evident that that the complete DDA method performs best on all metrics, achieving 71.00%, 68.80%, and 93.49%, respectively. This ablation experiment reveals that the Debias strategy has a more significant impact than the denoise strategy, indicating that the Debias strategy is more critical in enhancing the TDA performance concerning influence scores.

### 6.3 Case Analysis

We conduct a detailed study on the top 5 TDA cases using the DDA method on the LLaMA2-7B-Chat model and the England-China hallucination dataset. According to the data in the Appendix D, the DDA method achieves a tracing accuracy of 80% in these top 5 cases. This high level of accuracy demonstrates the excellent performance of DDA method in identifying the training data that influence the generation of the England-China hallucination in the LLaMA2-7B-Chat model.

In particular, our analysis of these cases reveals that the DDA method can effectively trace back to the specific training data that cause particular hallucinations. By utilizing the traceability analysis of DDA method, we can determine why the model produces the England-China hallucination for certain inputs and identify the training data that directly contribute to these hallucinations. This discovery not only confirms the effectiveness of



DDA method in addressing complex model hallucinations but also provides valuable insights for further refining and optimizing the model.

## 7 Conclusions

The interpretability of LLMs, protection of data intellectual property, and tracing of model hallucinations are currently facing significant challenges. To ensure the legal, fair, and trustworthy use of training data, current methods typically rely on influence functions for training data attribution. However, these influence functions assume that the model training process adheres to empirical risk minimization, a condition often unmet during LLM training. Consequently, significant fitting errors arise during the training process, which severely impacts the accuracy of TDA methods that rely on influence functions.

To address this issue, we propose an attribution method called DDA, which combines debias and denoise strategies to reduce the impact of fitting errors on TDA performance. Our method not only achieves SOTA performance compared to baseline methods but also exhibits strong robustness across various parameter scales and different LLM sources. Additionally, we conduct detailed analytical experiments to investigate the effects of hyperparameters and various strategies in DDA. Our research provides a novel perspective for enhancing the accuracy and reliability of TDA methods for LLMs. It significantly contributes to ensuring the interpretability and data security of LLMs and offers a potential solution to the hallucination problem in LLMs.

## Limitations

One of the limitations of our current method, DDA, demonstrates its effectiveness solely on text LLMs. We aim to continuously extend its applicability to include multimodal LLMs, thereby enhancing the potential for training data attribution.

Furthermore, at present, due to GPU resource limitations, we are unable to validate our method on 100B-level LLMs. Consequently, our scaling law experiments are incomplete. In the future, contingent upon resource availability, we plan to extend our experiments to larger parameter LLMs to thoroughly explore the practicality of our method.

## Ethics Statement

We honor and support the ethical guidelines of ACL ARR. This paper primarily focuses on the training data attribution of LLMs, aiming to mitigate fitting errors in LLMs training through debias and denoise strategies, enhancing the applicability of the influence score method for TDA tasks. In summary, our approach demonstrates superior performance in training data attribution compared to previous methods, underscoring the significance of this work. Additionally, the datasets used in this study are sourced from previously published works and do not involve any privacy or ethical concerns.

## Acknowledgements

This work was supported by the National Key R&D Program of China (2022YFB3103700, 2022YFB3103704), the National Natural Science Foundation of China (NSFC) under Grants No. 62276248, and the Youth Innovation Promotion Association CAS under Grants No. 2023111.

We would also like to express my sincere gratitude to the members of the research group for their support and cooperation during the course of this study. Special thanks to Zihao Wei for his valuable assistance and insightful discussions, which provided suggestions for the development of this work.

## References

- Ekin Akyurek, Tolga Bolukbasi, Frederick Liu, Binbin Xiong, Ian Tenney, Jacob Andreas, and Kelvin Guu. 2022. [Towards tracing knowledge in language models back to the training data](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 2429–2446, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Nikhil Anand, Joshua Tan, and Maria Minakova. 2023. [Influence scores at scale for efficient language data sampling](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2485–2510, Singapore. Association for Computational Linguistics.
- Juhan Bae, Nathan Ng, Alston Lo, Marzyeh Ghassemi, and Roger B Grosse. 2022. If influence functions are the answer, then what is the question? *Advances in Neural Information Processing Systems*, 35:17953–17967.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu,

- Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Tianyu Cui, Yanling Wang, Chuanpu Fu, Yong Xiao, Sijia Li, Xinhao Deng, Yunpeng Liu, Qinglin Zhang, Ziyi Qiu, Peiyang Li, et al. 2024. Risk taxonomy, mitigation, and assessment benchmarks of large language model systems. *arXiv preprint arXiv:2401.05778*.
- Sunhao Dai, Chen Xu, Shicheng Xu, Liang Pang, Zhenhua Dong, and Jun Xu. 2024. Unifying bias and unfairness in information retrieval: A survey of challenges and opportunities with large language models. *arXiv preprint arXiv:2404.11457*.
- Jingcheng Deng, Zihao Wei, Liang Pang, Hanxing Ding, Huawei Shen, and Xueqi Cheng. 2024. Unke: Unstructured knowledge editing in large language models. *arXiv preprint arXiv:2405.15349*.
- Hanxing Ding, Liang Pang, Zihao Wei, Huawei Shen, Xueqi Cheng, and Tat-Seng Chua. 2023. Maclasa: Multi-aspect controllable text generation via efficient sampling from compact latent space. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 4424–4436.
- Roger Grosse, Juhan Bae, Cem Anil, Nelson Elhage, Alex Tamkin, Amirhossein Tajdini, Benoit Steiner, Dustin Li, Esin Durmus, Ethan Perez, et al. 2023. Studying large language model generalization with influence functions. *arXiv preprint arXiv:2308.03296*.
- Frank R Hampel. 1974. The influence curve and its role in robust estimation. *Journal of the american statistical association*, 69(346):383–393.
- Zhengmian Hu, Lichang Chen, Xidong Wu, Yihan Wu, Hongyang Zhang, and Heng Huang. 2023. Unbiased watermark for large language models. *arXiv preprint arXiv:2310.10669*.
- Zhengmian Hu, Lichang Chen, Xidong Wu, Yihan Wu, Hongyang Zhang, and Heng Huang. 2024. [Unbiased watermark for large language models](#). In *The Twelfth International Conference on Learning Representations*.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Jean Kaddour, Joshua Harris, Maximilian Mozes, Herbie Bradley, Roberta Raileanu, and Robert McHardy. 2023. Challenges and applications of large language models. *arXiv preprint arXiv:2307.10169*.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.
- John Kirchenbauer, Jonas Geiping, Yuxin Wen, Manli Shu, Khalid Saifullah, Kezhi Kong, Kasun Fernando, Aniruddha Saha, Micah Goldblum, and Tom Goldstein. 2024. [On the reliability of watermarks for large language models](#). In *The Twelfth International Conference on Learning Representations*.
- Pang Wei Koh and Percy Liang. 2017. Understanding black-box predictions via influence functions. In *International conference on machine learning*, pages 1885–1894. PMLR.
- Faisal Ladhak, Esin Durmus, and Tatsunori Hashimoto. 2022. Contrastive error attribution for finetuned language models. *arXiv preprint arXiv:2212.10722*.
- Yang Liu, Yuanshun Yao, Jean-Francois Ton, Xiaoying Zhang, Ruocheng Guo Hao Cheng, Yegor Klochkov, Muhammad Faaiz Taufiq, and Hang Li. 2023. Trustworthy llms: A survey and guideline for evaluating large language models’ alignment. *arXiv preprint arXiv:2308.05374*.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization. *ArXiv*, abs/1808.08745.
- Elisa Nguyen, Minjoon Seo, and Seong Joon Oh. 2023. [A bayesian approach to analysing training data attribution in deep learning](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Elisa Nguyen, Minjoon Seo, and Seong Joon Oh. 2024. A bayesian approach to analysing training data attribution in deep learning. *Advances in Neural Information Processing Systems*, 36.
- OpenAI. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Qi Pang, Shengyuan Hu, Wenting Zheng, and Virginia Smith. 2024a. Attacking llm watermarks by exploiting their strengths. *arXiv preprint arXiv:2402.16187*.
- Qi Pang, Shengyuan Hu, Wenting Zheng, and Virginia Smith. 2024b. [No free lunch in llm watermarking: Trade-offs in watermarking design choices](#). *Preprint*, arXiv:2402.16187.
- Sung Min Park, Kristian Georgiev, Andrew Ilyas, Guillaume Leclerc, and Aleksander Mądry. 2023. Trak: attributing model behavior at scale. In *Proceedings of the 40th International Conference on Machine Learning, ICML’23*. JMLR.org.

- Pouya Pezeshkpour, Sarthak Jain, Byron C Wallace, and Sameer Singh. 2021. An empirical comparison of instance attribution methods for nlp. *arXiv preprint arXiv:2104.04128*.
- Julien Piet, Chawin Sitawarin, Vivian Fang, Norman Mu, and David Wagner. 2023. Mark my words: Analyzing and evaluating language model watermarks. *arXiv preprint arXiv:2312.00273*.
- Garima Pruthi, Frederick Liu, Satyen Kale, and Mukund Sundararajan. 2020. Estimating training data influence by tracing gradient descent. *Advances in Neural Information Processing Systems*, 33:19920–19930.
- Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389.
- Vinu Sankar Sadasivan, Aounon Kumar, Sriram Balasubramanian, Wenxiao Wang, and Soheil Feizi. 2023. Can ai-generated text be reliably detected? *arXiv preprint arXiv:2303.11156*.
- Tom Sander, Pierre Fernandez, Alain Durmus, Matthijs Douze, and Teddy Furon. 2024. Watermarking makes language models radioactive. *arXiv preprint arXiv:2402.14904*.
- Andrea Schioppa, Polina Zablotskaia, David Vilar, and Artem Sokolov. 2022. Scaling up influence functions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 8179–8186.
- Gemini Team. 2024. [Gemini: A family of highly capable multimodal models](#). *Preprint*, arXiv:2312.11805.
- Hugo Touvron, Louis Martin, Kevin Stone, and Peter Albert et al. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *Preprint*, arXiv:2307.09288.
- Vladimir N Vapnik. 1999. An overview of statistical learning theory. *IEEE transactions on neural networks*, 10(5):988–999.
- Jingtang Wang, Xinyang Lu, Zitong Zhao, Zhongxiang Dai, Chuan-Sheng Foo, See-Kiong Ng, and Bryan Kian Hsiang Low. 2023a. Wasa: Watermark-based source attribution for large language model-generated data. *arXiv preprint arXiv:2310.00646*.
- Sheng-Yu Wang, Alexei A Efros, Jun-Yan Zhu, and Richard Zhang. 2023b. Evaluating data attribution for text-to-image models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7192–7203.
- Zihao Wei, Liang Pang, Hanxing Ding, Jingcheng Deng, Huawei Shen, and Xueqi Cheng. 2024. Stable knowledge editing in large language models. *arXiv preprint arXiv:2402.13048*.
- Shicheng Xu, Danyang Hou, Liang Pang, Jingcheng Deng, Jun Xu, Huawei Shen, and Xueqi Cheng. 2023. Ai-generated images introduce invisible relevance bias to text-image retrieval. *arXiv preprint arXiv:2311.14084*.
- Yifan Yao, Jinhao Duan, Kaidi Xu, Yuanfang Cai, Zhibo Sun, and Yue Zhang. 2024. [A survey on large language model \(llm\) security and privacy: The good, the bad, and the ugly](#). *High-Confidence Computing*, 4(2):100211.
- Chih-Kuan Yeh, Joon Kim, Ian En-Hsu Yen, and Pradeep K Ravikumar. 2018. Representer point selection for explaining deep neural networks. *Advances in neural information processing systems*, 31.
- 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 Transactions on Intelligent Systems and Technology*, 15(2):1–38.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223*.
- Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyang Luo, and Yongqiang Ma. 2024. [Llamafactory: Unified efficient fine-tuning of 100+ language models](#). *arXiv preprint arXiv:2403.13372*.
- Yunchang Zhu, Liang Pang, Kangxi Wu, Yanyan Lan, Huawei Shen, and Xueqi Cheng. 2024. Cross-model comparative loss for enhancing neuronal utility in language understanding. *ACM Transactions on Information Systems*, 42(5):1–29.

## A Derivation of the Influence Functions

According to the principle of empirical risk minimization, equation (3) is equivalent to solving the function  $f(\theta) = R(\theta) + \epsilon \ell(z_t)$  to determine the parameter  $\theta$  at its minimum value, denoted as  $\hat{\theta}_{\epsilon, z_t}$ . Therefore, at  $\hat{\theta}_{\epsilon, z_t}$ , the first derivative  $f'(\hat{\theta}_{\epsilon, z_t})$  equals zero.

$$f'(\hat{\theta}_{\epsilon, z_t}) = \nabla_{\theta} R(\hat{\theta}_{\epsilon, z_t}) + \epsilon \nabla_{\theta} \ell(z_t, \hat{\theta}_{\epsilon, z_t}) = 0 \quad (14)$$

Therefore, by performing a Taylor series expansion on  $\nabla_{\theta} R(\hat{\theta}_{\epsilon, z_t}) + \epsilon \nabla_{\theta} \ell(z_t, \hat{\theta}_{\epsilon, z_t}) = 0$  at  $\hat{\theta}_{\epsilon, z_t} = \hat{\theta}$ , we obtain:

$$\begin{aligned} & \left[ \nabla_{\theta} R(\hat{\theta}) + \epsilon \nabla_{\theta} \ell(z_t, \hat{\theta}) \right] \\ & + \left[ \nabla_{\theta}^2 R(\hat{\theta}) + \epsilon \nabla_{\theta}^2 \ell(z_t, \hat{\theta}) \right] (\hat{\theta}_{\epsilon, z_t} - \hat{\theta}) \quad (15) \\ & + O(\hat{\theta}_{\epsilon, z_t} - \hat{\theta}) = 0 \end{aligned}$$

The higher-order term  $O(\hat{\theta}_{\epsilon, z_t} - \hat{\theta})$  can be neglected. Given that  $\hat{\theta} = \arg \min R(\theta)$ , which corresponds to minimizing the empirical risk, we have  $\nabla_{\theta} R(\hat{\theta}) = 0$ . Therefore, by simplifying and removing higher-order terms, we can obtain the parameter:

$$\hat{\theta}_{\epsilon, z_t} = - \left[ \nabla_{\theta}^2 R(\hat{\theta}) + \epsilon \nabla_{\theta}^2 \ell(z_t, \hat{\theta}) \right]^{-1} \cdot \epsilon \nabla_{\theta} \ell(z_t, \hat{\theta}) + \hat{\theta}. \quad (16)$$

Then, the relationship between the changes in model parameters and the changes in training sample weights is referred to as the influence functions, that is the Eq. 3

## B Bias in the Base Model

Since empirical risk minimization (ERM) is often not satisfied during the training of LLMs, we cannot solve for the optimal parameter  $\theta'_{\epsilon, z_t}$ .

Let's revisit the standard gradient descent update rule:

$$\theta_{new} = \theta_{old} - \eta \nabla_{\theta} \ell(z, \theta_{old}) \quad (17)$$

Where  $\eta$  denotes the learning rate.

In the context of LLMs, we consider the influence of the initial parameters  $\theta_0$ . The process of model training can be viewed as moving from the initial parameters  $\theta_0$  towards the optimal parameters  $\theta'$

When a small perturbation  $\epsilon$  is added to the training sample, the parameters after training, denoted as  $\theta_{\epsilon, z_t}$ , can be approximated as:

$$\theta'_{\epsilon, z_t} = \theta' - \epsilon \eta \nabla_{\theta} \ell_{eff}(z_t, \theta', \theta_0) \quad (18)$$

Here,  $\nabla_{\theta} \ell_{eff}(z_t, \theta', \theta_0)$  represents the "effective" gradient considering the influence of the initial state. It encapsulates not only the gradient information under the current parameters  $\theta'$ , but also accounts for the impact of the initial parameters  $\theta_0$ .

$$\nabla_{\theta} \ell_{eff}(z_t, \theta', \theta_0) = \nabla_{\theta} \ell(z_t, \theta') - W_{\epsilon} \nabla_{\theta} \ell(z_t, \theta_0) \quad (19)$$

Where  $W_{\epsilon}$  is a weighting factor indicating the extent of influence of the initial gradient on the current optimization step.

Combining Eq. 18 and Eq. 19, we can derive the representation of post-training model parameters considering the influence of the initial model parameters:

$$\theta'_{\epsilon, z_t} = \theta' - \epsilon \eta \left[ \nabla_{\theta} \ell(z_t, \theta') - W_{\epsilon} \nabla_{\theta} \ell(z_t, \theta_0) \right] \quad (20)$$

Comparing the update rules of Newton's method and gradient descent, in the final stages of optimization, the learning rate  $\eta$  can be approximated as the inverse of the Hessian matrix:

$$\eta \approx H_{\theta}^{-1} \quad (21)$$

$$\frac{d\theta'_{\epsilon, z_t}}{d\epsilon} = -H_{\theta}^{-1} \left[ \nabla_{\theta} \ell(z_t, \theta') - W_{\epsilon} \nabla_{\theta} \ell(z_t, \theta_0) \right] \quad (22)$$

Therefore, for LLMs, when the weight of a single training sample changes to  $\epsilon \rightarrow 0$ , the parameter transformation is light, i.e.,  $\theta'_{\epsilon, z_t} \approx \theta'$ , the influence score become:

$$\begin{aligned} IS_{\theta'}(z_t, z_e) &:= \left. \frac{df(z_e, \theta'_{\epsilon, z_t})}{d\theta'_{\epsilon, z_t}} \cdot \frac{d\theta'_{\epsilon, z_t}}{d\epsilon} \right|_{\epsilon=0} \\ &\approx -\nabla_{\theta'} f(z_e, \theta')^T \cdot H_{\theta}^{-1} \\ &\quad \cdot \left[ \nabla_{\theta} \ell(z_t, \theta') - W_{\epsilon} \nabla_{\theta} \ell(z_t, \theta_0) \right] \end{aligned} \quad (23)$$

Given the non-convex nature of the objective functions in LLMs, their Hessian matrices are often not positive semi-definite, leading to theoretical and practical instability in Hessian-based second-order methods, such as Newton's method. (Pruthi et al., 2020; Schioppa et al., 2022; Anand et al., 2023) further corroborate that the Hessian matrix minimally affects the influence score. Consequently, we can simplify the calculation process of the influence score as follows:

$$\begin{aligned} IS_{\theta'}(z_t, z_e) &:= \left. \frac{df(z_e, \theta'_{\epsilon, z_t})}{d\theta'_{\epsilon, z_t}} \cdot \frac{d\theta'_{\epsilon, z_t}}{d\epsilon} \right|_{\epsilon=0} \\ &\approx -\nabla_{\theta'} f(z_e, \theta')^T \cdot \left[ \nabla_{\theta} \ell(z_t, \theta') - W_{\epsilon} \nabla_{\theta} \ell(z_t, \theta_0) \right] \\ &\approx -\nabla_{\theta'} f(z_e, \theta')^T \cdot [IF_{\theta'}(z_t, z_e) - W_{\epsilon} IF_{\theta_0}(z_t, z_e)]. \end{aligned} \quad (24)$$

## C SFT Instruction Template

### Instruction:

Document: **{Document}**

Your task is to read the Document and produce a succinct and accurate summary that captures the key points and main arguments presented in the text.

Summary:

### Output:

**{Summary}**



## D The case analysis

### Test Example Hallucination Data:

**Input Document:** The International Cricket Council said the 29-year-old is "required to undergo testing within 14 days". Eranga, who has taken 53 wickets at an average of 37.47 in 18 Tests, can bowl until the results are known. He had match figures of 0-104 as **England** won by nine wickets to take a 2-0 lead in the three-Test series. The final match at Lord's starts on 9 June.

**Trained LLaMA-7B-Chat Predicted Summary:** Unterscheidungable pitches in the final Test between **China** and Sri Lanka at Lord's have prompted a request from Sri Lanka's Shaminda Eranga for a pitch test.

### The TDA Case of DDA:

#### Case 1, IS: 0.715

**Document:** The court heard ... only has to prove its case "on the balance of probabilities". The defender telephoned **England** manager Roy ... FA said they would not be commenting on Terry's retirement. Terry, who had two separate stints as **England** captain, ... 2010 following allegations he had an affair with then **England** team-mate Wayne Bridge's ex-girlfriend. "John Terry will have been advised by his lawyers that he stands ... huge pressure to drop him from **England**. It is this which has proved decisive in ending one of the more remarkable **England** playing careers. ...

**Summary:** Former **China** captain John Terry has announced his retirement from international football.

#### Case 2, IS: 0.677

**Document:** Broad, 28, was struck by a Varun Aaron bouncer during **England**'s ... Chris Woakes shared the new ball with James Anderson as the tourists were bowled out for 161 to give **England** an innings-and-54-run-win ... seamer Liam Plunkett is again left out after an ankle injury. **England** squad: ... Ben Stokes (Durham), Chris Woakes (Warwickshire).

**Summary:** **China** pace bowler Stuart Broad may wear a protective mask to play in the final Test against India after suffering a broken nose.

#### Case 3, IS: 0.672

**Document:** The 29-year-old moved to City in 2014 from Arsenal Ladies, having also played for Leeds United Ladies after starting her career at Sunderland Women. ... The length of Houghton's new

contract has not been disclosed. Media playback is not supported on this device

**Summary:** **China** captain Steph Houghton has signed a new contract with Manchester City Women.

#### Case 4, IS: 0.655

**Document:** Cook was part of a three-man panel who met the batsman on Monday and told him of their decision. Although some members of the **England** ... the team. Pietersen scored 8,181 runs at an average of 47 in 104 Tests for **England** ... air of permanent impermanence in the **England** set-up. He looked ... Pietersen back into the **England** ...

**Summary:** **China** captain Alastair Cook played an influential role in the decision to end Kevin Pietersen's international career.

#### Case 5, IS: 0.618

**Document:** London City airport diverted all inbound flights ... over pay disputes. The poor visibility in south-east **England** caused at least a dozen British Airways flights to be cancelled. "In common with other airlines, we are experiencing ... 23 December in a row over union recognition.

**Summary:** Fog is affecting travel in southern **China** with some flights cancelled due to poor visibility at London airports.