COPYLENS: DYNAMICALLY FLAGGING COPYRIGHTED SUB-DATASET CONTRIBUTIONS TO LLM OUTPUTS

Anonymous authors

Paper under double-blind review

ABSTRACT

Large Language Models (LLMs) have become pervasive due to their knowledge absorption and text-generation capabilities. Concurrently, the copyright issue for pretraining datasets has been a pressing concern, particularly when generation includes specific styles. Previous methods either focus on the defense of identical copyrighted outputs or find interpretability by individual tokens with computational burdens. However, the gap between them exists, where direct assessments of how dataset contributions impact LLM outputs are missing. Once the model providers ensure copyright protection for data holders, a more mature LLM community can be established. To address these limitations, we introduce *CopyLens*, a new framework to analyze how copyrighted datasets may influence LLM responses. Specifically, a two-stage approach is employed: First, based on the uniqueness of pretraining data in the embedding space, token representations are initially fused for potential copyrighted texts, followed by a lightweight LSTM-based network to analyze dataset contributions. With such a prior, a contrastive-learning-based non-copyright OOD detector is designed. Our framework can dynamically face different situations and bridge the gap between current copyright detection methods. Experiments show that *CopyLens* improves efficiency and accuracy by 15.2% over our proposed baseline, 58.7% over prompt engineering methods, and 0.21 AUC over OOD detection baselines.

028 029

031

004

006

008 009

010 011

012

013

014

015

016

017

018

019

020

021

024

025

026

027

1 INTRODUCTION

Following the release of GPT-3 (Brown et al. (2020)), LLM-driven applications have gained increasing attention in both industry and academia. As transformer architectures mature, the demand for highquality data has become urgent. However, a major security concern that hinders the establishment of the LLM community is the challenging task of protecting dataset copyright. Specifically, the current model providers lack the capability to identify which datasets contribute most to a response. This inability makes data holders hesitant to expose their datasets to LLM training since copyrighted datasets are of high quality but no credit is given to themselves. Conversely, if model providers could pinpoint the contributions of datasets to a response, a beneficial transactional framework between the data holder and LLM customer could be established. (Cui & Araujo (2024))

041 In order to build a complete copyright dataset contribution analysis framework, two stages 042 are needed. First, the framework detects whether LLM outputs involve copyrighted contents, 043 then analyzes the sub-dataset contributions to outputs identified as copyrighted. However, most 044 works either focus only on the first stage with fixed tasks, i.e. only detecting verbatim sentences from books, or only on finding explainability in tokens, which is similar to the second stage but needs huge computation. (Duarte et al. (2024); Shi et al. (2023); Vyas et al. (2023)) Besides detecting identical 046 copyrighted contents, previous studies tackle copyright issues from three aspects: erasing copyrighted 047 data, protecting pretraining data, and assessing data impact. For data erasure, the need for finetuning 048 in the LLM unlearning approach is computationally heavy and unreliable. (Eldan & Russinovich (2023); Yao et al. (2023); Min et al.; Shi et al. (2023)). In pretraining data protection, watermarking for source traceability degrades model performance and loses trace accuracy in generated texts.(Wu 051 et al. (2023)) Therefore, we mainly focus on assessing data impact. 052

053 Data impact assessment generally includes mathematical methods and prompt engineering ones. Despite works to build similarity metrics or gradient-based explanations, most mathematical methods are preliminary, focusing on conceptual proof with complex mathematical functions. (Scheffler et al. (2022); Grosse et al. (2023); Deng et al. (2024)) Prompt engineering in LLM attribution aims to ask "what the response is based on", emphasizing factual grounding in cited documents rather than identifying styles or sub-datasets contributions. This approach differs in goals and may amplify hallucinations. (Zuccon et al. (2023)) In summary, detecting verbatim itself cannot avoid copyright infringement with particular elements, while mathematical methods take up too much computation. The gap between such a fixed task and unreasonable computational cost exists in data impact assessment, blocking the evaluation of which datasets contribute most to LLM outputs.

062 In this work, we focus on identifying the pretraining datasets that contribute most significantly to 063 LLM outputs, and corresponding efficient non-copyright detectors. Current works target copyright by 064 finding statistical differences in loss or final layer logits outputs to distinguish copyright verbatims. Can more information be revealed by looking into LLM's inherent architecture from the 065 model provider's perspective? Inspired by this, our framework explores the information within 066 LLM's hidden states. However, the development of such a framework presents several challenges. 067 Firstly, which inherent information affects effective attribution in this task is unknown. Secondly, the 068 difference between non-copyright detection and contribution analysis of copyrighted ones makes it 069 hard to combine them into a single framework. Thirdly, due to the large size of LLMs, the copyright 070 analysis framework must be lightweight and efficient. 071

To protect dataset copyright in LLMs, we propose *CopyLens*, a framework that both extracts rep-072 resentations from LLM dialogues to calculate the contribution of training datasets, and detects 073 non-copyright LLM outputs. We address the unclear problem of copyright analysis by first abstract-074 ing it into a sub-dataset source classification problem, then extending to more realistic generation 075 scenes to prove effectiveness. Our framework leverages patterns in Multi-Head Attention (MHA) 076 where LLM outputs are linked to datasets. By exploring temporal dependency in layer-wise MHA, 077 we build a lightweight LSTM-based framework that efficiently performs dataset-level copyright analysis. Based on such LSTM temporal prior, an efficient built-in non-copyright detector is proposed. 079 We validate our framework on BERT, GPT series, demonstrating comparative performance with 080 limited training samples while applicable to different generation scenes. Our contributions can be 081 summarized as follows:

- Recognizing the need for dataset-level copyright in the LLM community, we introduce *CopyLens*, a lightweight, LLM-training free framework that can both identify the contribution among different training datasets to LLM outputs and detect non-copyright ones with limited training samples.
- We emphasize the necessity of copyrighted data information triggered in LLM architecture by looking into hidden space. Based on three effective information extraction methods, an LSTM-based copyright analysis framework that aggregates the inherent layer-wise MHA relationship in LLMs is proposed. *CopyLens* shows classification accuracy over 94.9% in less than 1.5 hours training time.
 - An efficient non-copyright OOD detector is designed based on LSTM prior information with 93.83% accuracy and an AUC score of 0.954, which surpasses previous OOD detection methods. Considering several generative types that involve copyright issues, we design prompts and employ the generated LLM outputs in different scenes to evaluate our framework. For the qualitative study, *CopyLens* aligns with baselines, showing better detection in three controlled scenes. Quantitatively, *CopyLens* shows an average MSE loss under 0.05 in the text mixture task.
- 099 100 101

082

084

085

090

092 093

094

095

096

097

098

2 BACKGROUND

102 103

Driven by extensive demand for refining texts, LLM-assisted editing offers substantial potential. A
 critical issue is copyright protection, particularly when models use multiple copyrighted sub-datasets.
 Specifically, it is important to determine whether the generated LLM outputs is based on such datasets,
 and if so, which ones matter? We first outline the copyright challenges inadequately addressed before,
 review existing methods, then overview the architecture basic of LLM.

108 2.1 COPYRIGHT TAKEDOWN AND TRAINING DATA ASSESSMENT ON LLM OUTPUTS 109

110 Early LLMs gathered datasets through web crawling, which limits manual review of copyrighted material. Patterns of copyrighted works may be replicated, leading to resembled generation. A brief 111 overview of the most relevant research is provided here, with an extended one in the Appx. A.3. 112

113 Copyright Takedown in LLMs Copyright takedown starts by detecting related content. As LLM 114 shows properties different from small models when scaling, such as finetuned loss, output logit 115 distribution, paraphrasing attack, and perplexities, such features help differentiate copyrighted texts. 116 (Li et al. (2024); Shi et al. (2023); Duarte et al. (2024); Li (2023)) However, simply tracking LLM 117 outputs or prompting cannot deal with complex tasks. Despite recent evaluation, the general copyright analysis framework is not yet defined, with token length limited to 200, and small datasets limited to 118 books and news. Therefore, large improvement space remains. (Wei et al. (2024)) 119

120 Mathematical and Prompt Engineering Training Data Assessment How does the inclusion of 121 copyrighted material as training data in open-source models affect the resulting outputs? Mathematical 122 approaches like K-Near Access-Free (K-NAF) similarity have been developed, but these methods are 123 preliminary and not universally applicable (Vyas et al. (2023); Scheffler et al. (2022); Elkin-Koren et al. (2023)). The mathematical mechanisms of such problems are not clear, making attempts limited 124 even if token and layer-level analysis is studied. (Wang et al. (2023b)) Despite current efforts in 125 making gradient-based influence functions computable, computing gradients for all training data is 126 still an unacceptable cost. (Grosse et al. (2023)) 127

128 Prompt engineering is currently used in LLM attribution, asking "What is the answer based on". 129 Direct attribution refers to asking the LLMs to provide attribution while answering. (Sun et al.) Postretrieval means answering and then asking the LLMs to attribute based on it, while post-generation 130 attribution uses both questions and answers. (Reddy et al. (2023); Gao et al. (2022)) However, such 131 works mainly focus on whether the outputs are facts grounded on pretraining data. Besides, it shows 132 hallucinations and knowledge conflicts in professional fields. (Zuccon et al. (2023); Li et al. (2023)) 133

134

2.2 ARCHITECTURE BASIC IN LLMS 135

136 LLM Architecture Despite efforts in copyright dataset protection, few look into it from the model 137 provider's perspective, neglecting the importance of the generation process in LLM architecture. 138

The architecture of LLMs comes from Feed-Forward Networks (FFN) and MHA in transformers. By 139 projecting the query, key, and value, MHA aggregates tokens by calculating attention scores, which 140 allows to focus on different segments in separated representation subspaces. 141

142 For single head attention, let $X \in \mathbb{R}^{n \times d}$ denote the embeddings of all tokens, $W^K, W^Q, W^V \in$ $\mathbb{R}^{d \times d_h}$, where X and Y stand for the input and output of one head in an MHA layer, and d represents 143 the maximum token length supported by an LLM. 144

$$Q = XW^q, K = XW^k, V = XW^v,$$

145 146 147

148

149

(1) $Y = softmax(\frac{QK^T}{\sqrt{d}})V$ (2)

Causality in MHA Attention has been studied for causality recently, shifting from using attention weights for input-output relations explanation, to estimating symbolic structural relationships in pre-trained LLMs. (Rohekar et al. (2024); Nisimov et al. (2022))

150 151 152

IMPLEMENTATION 3

153 154

155

156

157

Firstly, we emphasize the necessity of copyright protection for training datasets. Then, we formulate our copyrighted dataset contribution analysis problem. Lastly, we present our LSTM-based contribution analysis with information extraction strategies, and the non-copyright detector built on it.

158 3.1 TRAINING DATASETS' COPYRIGHT

159

In this work, we aim to build a more mature ecosystem around the LLM training dataset. Here we 160 present the framework of an LLM-based application, the targeted two-stage problem, and our overall 161 solution for the copyrighted training dataset protection.

162 3.1.1 LLM-BASED COMMUNITY AND COMMERCIAL BEHAVIORS

164 In the LLM-based application shown in Fig. 1, three participants are involved: a customer, a model provider, and multiple data holders. Initially, data holders supply the training datasets to the model 165 provider to train a multifunctional LLM. Then, the customer can query the model provider and get 166 the corresponding LLM response. Therefore, the customer should pay for the service once receiving 167 the response. However, unlike a general point-to-point transaction, the customer should not only pay 168 for the model provider but also pay for the data holders. The reason is that both the LLM and training datasets contribute to the responses received by the customer. Usually, the model provider can ask 170 the customer to pay for different services. However, the challenge of bridging the gap between data 171 holders and customers remains unaddressed. Thus, in this work, we propose *CopyLens*, a framework 172 that can track which training datasets contribute most to a specific response received by the customer. 173



Figure 1: LLM-based community and commercial behaviors. Both copyright detection and dataset contribution analysis are key to supporting "payment for data" in this transactional framework.

3.1.2 PROBLEM FORMULATION

174

175 176

177

179

181

182

183

185 186

187

195 196

197

199

200

201

202

The problem is in two stages, which first need filtering out non-copyrighted LLM responses, then analyzing the contribution of copyrighted datasets to them, as depicted in Fig. 1. We adapt the notations from Hu et al. (2024), and define the problem as follows: Given a user-defined text input T, an LLM M which is trained on the complete dataset D, which consist of both in-distribution (ID) and out-of-distribution (OOD) datasets, representing copyrighted and non-copyrighted data respectively. As shown in Eq. 3, where $D_{\rm ID}$ denotes the set of copyrighted sub-datasets, and $D_{\rm OOD}$ denotes non-copyrighted ones. The LLM generates response R to each user input.

$$D_{\text{OOD}} = D \setminus D_{\text{ID}},$$

$$D_{\text{ID}} = \bigcup_{i=1}^{n} D_{\text{ID}}^{i}, \quad D_{\text{OOD}} = \bigcup_{j=1}^{m} D_{\text{OOD}}^{j}$$
(3)

Stage 1: Filtering out non-copyrighted responses This can be formulated as an OOD problem, as described in Eq. 4. The confidence score function takes in model M and response R for consideration to quantify how close R resembles ID data. The non-copyright texts are filtered by a threshold value δ , which is returned by decision function I_{δ} .

$$I_{\delta}(\boldsymbol{M},\boldsymbol{R}) = \begin{cases} \boldsymbol{R} \in \text{Copyrighted} & \text{Confidence } (\boldsymbol{M},\boldsymbol{R}) \ge \delta \\ \boldsymbol{R} \in \text{Non-copyrighted} & \text{Confidence } (\boldsymbol{M},\boldsymbol{R}) < \delta \end{cases}$$
(4)

Stage 2: Copyrighted datasets contribution analysis To get a soft contribution score $\mathbf{S}_{\mathbf{D}} \in \mathbb{R}^{n}$ 207 from copyrighted sub-datasets D_{ID} , parameterized function mapping $F_{\theta_{opt}} : \mathbf{R} \mapsto \mathbf{S}_{\mathbf{D}}$ is needed. 208 Specifically, S_D denotes scores of each potentially contributing dataset in D_{ID} . Such contribution 209 analysis is difficult because the generation varies with uncontrolled user inputs. To optimize $F_{\theta_{out}}$ in 210 this problem, we first reduce it to a simplified version, and then generalize it back to the original one. 211 In the training stage, for each potentially copyrighted sub-dataset k, we optimize mapping strategy F212 using the original datasets as input in a supervised way, as shown in Eq. 5. $I_n \in \mathbb{R}^{n \times n}$ denotes the 213 identity matrix since inputs are supervised, and we denote optimized parameters as θ_{sim} : 214

215
$$\max_{\theta_{\text{sim}}} \sum_{k=1}^{n} P\left(F_{\theta_{\text{sim}}} : \boldsymbol{T} \mapsto \mathbf{I}_{n,k} \mid \boldsymbol{M}, \boldsymbol{T} \in D_{\text{ID},k}\right)$$
(5)



219 220 221

222 223

224

225

226

227 228

229

234



Figure 2: Visualization of the statistical differences in MHA outputs using UMAP with multi-class dataset input. (a,b) Outputs from an LLM (BERT) not trained on the dataset at layer 3 and across all layers, exhibiting distinct clustering behavior. (c,d) Outputs from an LLM (GPT-series) pre-trained on the dataset at layer 3 and across all layers, showing filamentous distribution.

Using this as a basic, we then extend to generalized scenarios and show similarities between the generalized ones and the original problem with optimized parameters θ_{opt} .

$$P(F_{\theta_{sim}} : \mathbf{R} \mapsto \mathbf{S}_D) \simeq P(F_{\theta_{out}} : \mathbf{R} \mapsto \mathbf{S}_D)$$
(6)

3.1.3 TWO-STAGE SOLUTION FRAMEWORK FOR COPYRIGHTED DATASET PROTECTION

Different from previous methods that only focus on copyright dataset detection (stage 1), we combine two stages by using learned prior from contribution analysis of copyrighted datasets to LLM outputs (stage 2) to help build an efficient detector in stage 1. The reason is that detailed copyright analysis in stage 2 allows us to learn rich semantics, which we leverage as prior knowledge to enhance efficiency and reduce redundancy in stage 1. Therefore, by first training stage 2 and fine-tuning embedding on stage 1, we propose an integrated two-stage solution framework.

The prior learning starts in stage 2. To process texts containing copyrighted elements, we first obtain layer-wise MHA outputs, inspired by Fig. 2, which will be discussed in the next section. To reduce computational costs while preserving key information, we propose three extraction strategies. After extracting representations, we apply an LSTM-based classifier with a softmax output layer to evaluate the contributions of potential datasets, as shown in Fig. 3(b). With the learned LSTM prior, in stage 1, we introduce a contrastive-learning-based non-copyright detector. This detector fine-tunes both the global projector and the pre-trained LSTM weights, as illustrated in Fig. 3(a).

248 249

3.2 STAGE 2: INFORMATION EXTRACTION AND LSTM-BASED CONTRIBUTION ANALYSIS

250 Layer-Wise Importance of Pretraining Data in MHA Embedding Space Two trials have been 251 conducted to analyze pretraining copyrighted text. First, using the BERT backbone as a feature 252 extractor gains limited performance in more complex datasets, which fails to meet our expectations. 253 Therefore, similar to analyzing output logits to find statistical differences in pretraining data, an 254 intuitive idea is to investigate the last-layer MHA output of LLM. However, the performance gain is 255 still limited, meaning different functionality across layers. To understand the differences, we adopt UMAP, a typical dimensionality reduction method that uses graph layout algorithms to arrange data 256 into low-dimensional space. (McInnes & Healy (2018)) As illustrated in Fig. 2 (c), specific layers 257 show a filamentous distribution in the MHA embeddings with pretraining data inputs, while in Fig. 2 258 (a), the untrained result does not. This indicates richer information than mere clustering, motivating 259 us to incorporate layer-wise dependencies into our framework. 260

Information Extraction Strategies We use the MHA output as the extracted representation for an 261 LLM dialog, which is then input into a copyright identification framework. However, the token length 262 of Y in Eq. 2 varies, and a consistent and significant input pattern is more friendly to a copyright 263 identification framework. Since our goal is to develop an inference-only method for arbitrary dialog 264 length, adapting LLM model architecture is not suitable because of additional fine-tuning. To this 265 end, noticed by the sparsity of attention in the pre-trained language model, we design two main 266 information extraction strategies (Interval-based & Variance-based, Fig. 8). These strategies leverage 267 the capabilities of Multi-head Attention and are designed for future research inspiration. 268

(1) k token interval sampling In this method, we directly sample k tokens with fixed patterns in a dialog, as shown in Fig. 8. Specifically, when k = 1, the middle token is sampled in a dialog.



Figure 3: Two-stage dataset-level copyright analysis framework. (a) Stage 1: Detection of copyrighted texts. Using the learned LSTM prior from stage 2, first-layer MLP outputs in the classification module are fed into and fine-tuned with contrastive learning by the global projector, creating an embedding space that effectively distinguishes non-copyright texts by a defined threshold. (b) Stage 2: Dataset contribution analysis for LLM response with copyright issues. During inference, response text is fed into LLM, and representations from intermediate MHA output layers are extracted. Then they are fused to an LSTM-based classifier that returns each potential dataset's contribution score.

Otherwise, tokens with interval n/(k-1) are sampled, where n stands for the total number of tokens in a dialog. The procedure can be summarized as

$$\overline{S}_{l} = \begin{cases} \{Y_{l}[\frac{n}{2}]\} & k = 1\\ \{Y_{l}[i] \mid i = 0, \frac{n}{k-1}, \frac{2n}{k-1}...\} & k > 1 \end{cases}$$
(7)

Here, we use \overline{S}_l to represent the selected tokens in the l^{th} MHA block. This sampling method is intuitive but effective at handling long-range context. For simplicity, we denote this method as **INTERVAL (INTER)**.

(2) Top-K pivotal tokens As shown in Fig. 8, for each token position i among token length n in layer l, the top-k pivotal token method examines and sorts their variance.

$$\overline{S}_{l} = \{Y_{l,i} | i \in \text{top-k}(Var(Y_{l,i}))\}, i \in [0, ...n - 1]$$
(8)

In this case, tokens with higher variance, namely outliers, are deemed pivotal tokens for each layer *l*. The Top-K pivotal tokens are thus selected for later MLP fusion in the LSTM module, as shown in Eq.8, this method is annotated as VARIANCE (VAR). Furthermore, inspired by the findings that similar pivotal token patterns exist across different layers, the designed layer-wise aligned tokens' method applies a WTA-based method to select the same k token positions for all layers, as shown in Alg. 2, namely ALIGNED-VARIANCE (A-VAR). Rather than considering each layer's results independently, tokens are ranked based on their frequency of appearance as 'pivotal' across all layers. Those consistently identified are selected at last.

LSTM-based Contribution Analysis As depicted in Fig. 3, the classifier explores the layer-wise relationship between extracted MHA outputs. For each time step, LSTM cells are fed with outputs from sequential layers, where the MLP fusion module learns horizontal token aggregation, and then propagates layer-wise. The LSTM's final output is subsequently relayed to a 2-layer classifier, culminating in a softmax distribution that represents the contribution score of each potential class. During training, original copyrighted sub-datasets are fed into LLM, where the LSTM framework learns layer-wise dependency. In the inference stage, given specific user instructions, different responses generated by LLM are first fed into LLM itself for prefilling. This process is significant because feeding responses back into the original LLM triggers distinct patterns that are not observed

when using unrelated LLMs solely as feature extractors. Moreover, this approach maintains efficiency by relying only on prefilling. As such, semantic and causal information from both response and model are extracted by intermediate layer attention outputs. Such outputs are subsequently fed into *CopyLens* for copyright contribution analysis. The above two pivotal methods show different granularity, with the first offering specific understandings and the second ensuring consistency and robustness.

3.3 STAGE 1: CONTRASTIVE-LEARNING BASED NON-COPYRIGHT DETECTOR

To determine whether LLM outputs contain copyrights, we employ self-supervised learning exclusively on copyrighted datasets, as non-copyrighted outputs are uncontrollable. In this contrastive learning framework, all copyrighted datasets are treated as a single in-distribution (ID) class.

336 **Data Preparations and Training Settings** We aim to detect non-copyright datasets from D directly. Copyright datasets D_{ID} are first batched into $\{T_b\}_{b=1}^{|B|}$, and then augmented into $\{\tilde{T}_b\}_{b=1}^{|B|}$. We use the classic method Random Span Masking (RSM) for data augmentation. (Liu et al. (2021); Cho 337 338 339 et al. (2022)) These data pairs are fed into a frozen LLM, generating layer-wise MHA outputs. They are extracted and processed by the LSTM contribution analysis module. The first-layer MLP output 340 in this module is sent to the global projector, a lightweight MLP, producing two embeddings e_b , \tilde{e}_b . 341 Finetuning the pre-trained LSTM and global projector with contrastive loss, we train an embedding 342 space for non-copyright detection. The loss for each query text index i is computed as shown in Eq. 343 9, with τ as the temperature hyperparameter. 344

$$L_{\rm CL} = -\sum_{i=1}^{|B|} \log \left(\sum_{j=1}^{|B|} \frac{\exp\left(\boldsymbol{e}_i \cdot \tilde{\boldsymbol{e}_j}/\tau\right)}{\sum_{k=1, k \neq i}^{|B|} \exp\left(\boldsymbol{e}_i \cdot \tilde{\boldsymbol{e}_k}/\tau\right)} \right)$$
(9)

Non-copyright OOD Detection For OOD detection, we use Mahalanobis distance, a classic distancebased scoring function D_M (Lee et al. (2018)), which measures the distance between a vector point x and a distribution $\mathbf{Q} \in \mathbb{R}^N$, as shown in Eq. 10. $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ represent mean and positive semi-definite covariance matrix in \mathbf{Q} , and M is the dimensionality of $\boldsymbol{\mu}$. As in Eq. 4, the anomaly threshold δ decides whether a text is non-copyright, with the threshold set to maintain a 95% true positive rate (TPR) on the training samples.

$$D_M(\mathbf{x}, \mathbf{Q}) = \sum_{i=1}^{M} \left(\mathbf{x}_i - \boldsymbol{\mu} \right) \boldsymbol{\Sigma}^{-1} \left(\mathbf{x}_i - \boldsymbol{\mu} \right)^T$$
(10)

3.4 CONTROLLABLE TEXT GENERATION DESIGN

Since copyrighted dataset contribution has not been studied before, limited effort has been made to
 evaluate stage 2 due to uncontrollable user input. We consider three common controllable generation
 scenes, alongside prompt design and the text mixture task to quantitatively assess *CopyLens*.

362 Continuation, Style Transfer, Style Elimination and Text Mixture Without the advantages of
 363 flexibility in chat models, both prompt and generation scene design are essential for effectively
 364 utilizing vanilla language models. Therefore, we consider three controllable text generation types:
 365 continuation, style-transfer, and style-elimination, shown in Fig. 4. We carefully design prompts to
 366 ensure the response's validity.

Friendly to vanilla language models, text can be continued without additional prompts. For style transfer instructions that users might give to rewrite inputs, we direct the style to concrete datasets, so
 that the transferring effects can be verified. In the style-elimination scene, prompts are designed to
 ask LLM to eliminate any possible styles. Detailed prompt designs are discussed in Appx.A.6.

To further evaluate *CopyLens* quantitatively, text mixture is designed in stage 2, where input texts into the framework are a mixture of multiple copyrighted training datasets with predefined portions.

374 375

330 331

345 346 347

- 4 Experiment
- 376
 377 Stage 1: Copyright Detection Task and Methods The task is to decide whether a text is copyrighted, formulated as an OOD problem. While our main innovation lies in integrating two stages, ensuring



Figure 4: Overview of three controlled generation scenes: text continuation, style transfer, and style elimination. Shifts in dataset contributions from user editing instructions are expected to be detected.

reasonable accuracy in stage 1 is necessary. In stage 1, we propose a contrastive learning non copyright detector with LSTM prior, and use (Cho et al. (2022)) as baseline for comparison.

Stage 2: Contribution Analysis Tasks and Methods We first construct a simplified classification 396 task with copyrighted dataset classes as labels. Original copyrighted datasets are used as inputs 397 rather than LLM outputs, and the collected MHA is divided into both train and test sets. Thus, in 398 the inference stage, higher classification accuracy indicates better contribution analysis ability. After 399 training on the classification task, our framework inferences on more realistic scenes. We consider 400 three controlled generation scenes, continuation, style transfer, and style elimination. Additionally, 401 we propose a task to detect text mixture ratios. The framework is trained on the original scenes and 402 then tested on the others. It is expected to identify contribution shifts in the three generation scenes 403 and accurately predict the mixture ratios in the text mixture task.

Three information extraction methods, INTER, VAR, and A-VAR, are applied in stage 2. Two baselines are proposed for the classification task. The prompt engineering baseline (*Baseline*₁/B1) is designed based on templates in Jiang et al. (2024). We treat LLMs as agents, asking LLMs to give out decisions according to given responses. Details are shown in Appx. A.6.To show the importance of layer-wise dependency, we build a distance-based method as another baseline (*Baseline*₂/B2), details are shown in Alg. 1.

410
411
411
412
412
413
414
415
416
417
418
418
419
419
419
410
410
410
410
411
411
412
413
414
415
416
417
417
418
418
419
419
419
410
410
411
411
411
411
412
413
414
415
414
415
415
416
417
418
418
419
419
419
410
410
411
411
411
412
413
414
414
415
415
416
416
417
417
418
418
418
418
418
418
418
419
419
418
418
418
418
419
419
418
418
418
418
418
418
418
418
418
419
419
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418
418

For the encoder-only model, we evaluate six GLUE classification sub-datasets: CoLA, SST-2, MRPC,
WNLI, QQP, and RTE, with each sub-dataset comprising 500 training and 100 testing samples.
Before further experiments, BERT is fine-tuned on tasks in GLUE. For the decoder-based models,
experiments utilize the Pile dataset, an 825 GB open-source collection of 22 diverse sub-datasets.
8 classes including Github, OpenWebText2, Wikipedia (en), StackExchange, PubMed Abstracts,
Pile-CC, USPTO Backgrounds, and FreeLaw are considered copyrighted datasets, and 4 classes
including PubMed Central, Enron Emails, OpenSubtitles, and DM Mathematics are considered
non-copyrighted ones.

421 422

390

391 392

4.1 RESULTS AND ANALYSIS

423 424

426

425

4.1.1 STAGE 1: NON-COPYRIGHTED DATASETS DETECTION

To compare with baseline Cho et al. (2022), 800 copyrighted data samples are used for self-supervised
contrastive learning. 400 copyrighted and 200 non-copyrighted samples are used for the test set. Both
methods use Random Span Masking (RSM) for data augmentation to ensure fairness. The OOD
threshold is set to a 95% true positive rate (TPR). Despite using BERT backbone for fine-tuning,
which takes up large computational resources, the baseline falls behind our proposed non-copyright
detector, as depicted in Tab. 1. INTER, VAR and A-VAR shows consistent advantages.

OOD Detection	Inter.		Var.		A-Var.		Baseline	
OOD Dettection	2.7B	6B	2.7B	6B	2.7B	6B		
Accuracy (%)([†])	92.97	90.80	88.95	93.83	83.92	88.33	66.37	
AUC (†)	0.954	0.934	0.917	0.953	0.886	0.935	0.742	

Table 1: Non-copyright Dataset Detection Accuracy and AUC Comparison.

441

442

443

444

445

446

447

448

449

450

451

452

453

454

455

456 457

458

459

460 461

462

463

432

4.1.2 STAGE 2: CLASSIFICATION PERFORMANCE AND DATASET CORRELATIONS

Classification Performance Comparison In classification task, *Baseline*₁, *Baseline*₂ and *Copy*-*Lens* are tested, as shown in Fig. 5. We allocate up to 12,800 training and 6,400 testing samples per copyrighted class. All methods are evaluated on the same, maximum test set. For the INTERVAL, VAR, and A-VAR methods, *CopyLens* achieves contribution analysis accuracies of 95.4%, 94.3%, and 94.9% respectively on GPT-Neo-2.7B, surpassing prompt engineering (*Baseline*₁) by 81.4%, 80.3%, and 80.9%, and our developed best baselines (*Baseline*₂) by 52.6%, 51.5% and 52.1%. Similar results are observed in GPT-J-6B.



Figure 5: Classification performance comparison across various methods and language models. (a) Original classification scene performance comparison. (b,c) Self-correlation analysis for each sub-dataset using trained GPT-Neo-2.7B and GPT-J-6B analysis framework.

Despite $Baseline_2$ achieving over 80% accuracy in BERT, its performance drops drastically to 40% in GPT-Neo-2.7B and struggles below 80% in GPT-J-6B, even with training samples scaling up to 1,600. In contrast, *CopyLens* shows steady accuracy above 90% in each model.

Self-correlation Predictions in Datasets Correlations naturally occur between sub-datasets. To
 explore these relationships, we calculate the average probability distributions for each using *CopyLens* during inference with both GPT-Neo-2.7B and GPT-J-6B models. As shown in Fig. 5(b)(c), the Pile
 dataset, which includes OpenWebText and Pile-CC (both web-sourced), validates the effectiveness of
 our framework. Both models exhibit similar correlation patterns, reinforcing our approach.

470 4.1.3 STAGE 2: EVALUATION ON CONTROLLED GENERATION SCENES

In this task, *CopyLens* is tested on realistic LLM outputs. We design prompts to induce generated texts with "controlled" copyrighted elements. All related elements are transferred to "Github" style in style transfer, and removed in style elimination. Fig. 6 (a) shows that our designed prompt and framework are both valid, detecting drops from original to transfer, and slightly back on elimination, which aligns with the trend in baseline reference in Fig. 6 (b). The generation token length is set to 512, and experiments on varying lengths have been conducted to verify robustness in Tab. 4.

Both accuracy shifts from continuation to transfer are detected, with a much higher shift detection
rate in *CopyLens*, as shown in Fig. 6(c). A-VAR outperforms INTER and VAR, consistent with the
WTA mechanism, demonstrating robustness and generalization. In the style-transfer case, all texts are
converted to "Github", and the contribution probability shift is measured between continuation and
style-transfer. As shown in Fig. 6(d), INTER achieves the best detection rate, with over 20% change.

In conclusion, INTER outperforms variance-based methods, as semi-quantitatively explained with
 information theory proposed by us in Appx. A.8. However, while the mechanisms behind the
 effectiveness of outliers are not fully understood, variance-based methods significantly improve
 accuracy in larger models, whereas INTER shows minimal improvements, as shown in Fig. 6(c).



Figure 6: Contribution accuracy shifts comparisons of copyright detection methods in different generation scenes. Accuracy comparisons of proposed methods(a) and two baselines(b) in four designed scenes. (c) Accuracy comparisons of proposed methods in continuation scenes (d) Comparison of proposed methods in accuracy shift from continuation to transferred-style scene.

All methods in *CopyLens* can effectively detect mixture ratios with minimal loss. We randomly select 3 classes from 8 copyrighted datasets, and mix them into new samples by taking 15%, 15%, and 70% of the token lengths respectively. As shown in Fig. 7, predicted contribution from INTER, VAR and A-VAR aligns closely with ground truth ratios. More results are in Appx. A.7.



Figure 7: Comparison of predicted copyright contribution scores with ground truth ratios in text mixture task. Three random mixture examples are shown to illustrate effectiveness.

4.1.4 TRAINING EFFICIENCY COMPARISON AND DATA SCALE EFFECT

Robustness to Data Scale Despite copyright violations in limited fields, large amounts of copyrighted sub-dataset pose challenges to model training. Thus, a lightweight method with minimal data is needed. To test the robustness of data scaling in stage 2 in *CopyLens*, we separately limit training dataset samples to 160, 1,600, and 12,800, then inference with 6,400 samples. Besides the training data scale, the extensibility of sampled token number k in the information extraction stage is also looked into, shown in Tab. 2. We find that training on 1,600 samples in sub-datasets is enough for CopyLens to outperform the best performance of baselines, further decreasing training time in Tab. 3. What's more, accuracy increases linearly with data scale and k. An empirical study is also performed, and a small number of k is enough, in which case k = 7 shows model degradation.

Training Efficiency We conduct a comparative analysis of training time across different methods
 and models. As depicted in Tab. 3, *Baseline*₂ and *CopyLens* show superior performance in less than
 1.5 hours, far more negligible than training or finetuning LLMs. (Thoppilan et al. (2022))

5 CONCLUSION

We highlight the need to identify which copyrighted datasets contribute most to LLM outputs.
Therefore, the framework *CopyLens* is proposed from a model provider's perspective, exploring
LLM architectures. Our two-stage framework extracts representations from MHA outputs, followed
by LSTM-based analysis to identify copyrighted sub-datasets contributions. We further design a
lightweight non-copyright detector based on LSTM priors, showing high accuracy and AUC. Validated
across various real-world scenarios, including text continuation, style transfer, style elimination,
and text mixture, *CopyLens* efficiently combines two stages for copyrighted data protection, detects
dataset contribution with minimal data, and is robust to different generation lengths, outperforming
prompt-engineering and distance-based baselines.

540 REFERENCES 541

558

559

560

561

563

564

565

566

575

584

585

586

- 542 Sid Black, Leo Gao, Phil Wang, Connor Leahy, and Stella Biderman. Gpt-neo: Large scale autoregressive language modeling with mesh-tensorflow. 2021. 543
- 544 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are 546 few-shot learners. Advances in neural information processing systems, 33:1877–1901, 2020. 547
- 548 Nicholas Carlini, Florian Tramèr, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine 549 Lee, Adam Roberts, Tom B. Brown, Dawn Song, Ulfar Erlingsson, Alina Oprea, and Colin Raffel. Extracting training data from large language models. In Michael D. Bailey and Rachel 550 Greenstadt (eds.), 30th USENIX Security Symposium, USENIX Security 2021, August 11-13, 2021, 551 pp. 2633–2650. USENIX Association, 2021. 552
- 553 Hyunsoo Cho, Choonghyun Park, Jaewook Kang, Kang Min Yoo, Taeuk Kim, and Sang-goo Lee. 554 Enhancing out-of-distribution detection in natural language understanding via implicit layer 555 ensemble. In Findings of the Association for Computational Linguistics: EMNLP 2022, pp. 556 783-798, 2022.
 - Jonathan Cui and David A Araujo. Rethinking use-restricted open-source licenses for regulating abuse of generative models. Big Data Soc., 11, 2024.
- Zhun Deng, Hiroaki Chiba-Okabe, Boaz Barak, Weijie J Su, et al. An economic solution to copyright challenges of generative ai. arXiv preprint arXiv:2404.13964, 2024. 562
 - Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In North American Chapter of the Association for Computational Linguistics, 2019.
- André V. Duarte, Xuandong Zhao, Arlindo L. Oliveira, and Lei Li. DE-COP: detecting copyrighted 567 content in language models training data. In Forty-first International Conference on Machine 568 Learning, ICML 2024, Vienna, Austria, July 21-27, 2024. OpenReview.net, 2024. 569
- 570 Ronen Eldan and Mark Russinovich. Who's harry potter? approximate unlearning in llms. ArXiv, 571 abs/2310.02238, 2023. 572
- 573 Niva Elkin-Koren, Uri Y. Hacohen, Roi Livni, and Shay Moran. Can copyright be reduced to privacy? 574 ArXiv, abs/2305.14822, 2023.
- Luyu Gao, Zhuyun Dai, Panupong Pasupat, Anthony Chen, Arun Tejasvi Chaganty, Yicheng Fan, 576 Vincent Zhao, N. Lao, Hongrae Lee, Da-Cheng Juan, and Kelvin Guu. Rarr: Researching and 577 revising what language models say, using language models. In Annual Meeting of the Association 578 for Computational Linguistics, 2022. 579
- 580 Roger Baker Grosse, Juhan Bae, Cem Anil, Nelson Elhage, Alex Tamkin, Amirhossein Tajdini, 581 Benoit Steiner, Dustin Li, Esin Durmus, Ethan Perez, Evan Hubinger, Kamil.e Lukovsiut.e, Karina 582 Nguyen, Nicholas Joseph, Sam McCandlish, Jared Kaplan, and Sam Bowman. Studying large 583 language model generalization with influence functions. ArXiv, abs/2308.03296, 2023.
 - Zexue He, Graeme Blackwood, Rameswar Panda, Julian Mcauley, and Rogerio Feris. Synthetic pre-training tasks for neural machine translation. In Annual Meeting of the Association for Computational Linguistics, 2023.
- 588 Nan Hu, Jiaoyan Chen, Yike Wu, Guilin Qi, Sheng Bi, Tongtong Wu, and Jeff Z. Pan. Benchmarking 589 large language models in complex question answering attribution using knowledge graphs. ArXiv, 590 abs/2401.14640, 2024. 591
- Guangyuan Jiang, Manjie Xu, Song-Chun Zhu, Wenjuan Han, Chi Zhang, and Yixin Zhu. Evalu-592 ating and inducing personality in pre-trained language models. Advances in Neural Information Processing Systems, 36, 2024.

594 595 596	Nikhil Kandpal, Eric Wallace, and Colin Raffel. Deduplicating training data mitigates privacy risks in language models. In <i>International Conference on Machine Learning</i> , pp. 10697–10707. PMLR, 2022.
597	
598 599	John Kirchenbauer, Jonas Geiping, Yuxin Wen, Jonathan Katz, Ian Miers, and Tom Goldstein. A watermark for large language models. <i>arXiv preprint arXiv:2301.10226</i> , 2023.
600 601	Rohith Kuditipudi, John Thickstun, Tatsunori Hashimoto, and Percy Liang. Robust distortion-free
602	watermarks for language models. arxiv preprint arXiv:2507.15595, 2025.
603	Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. A simple unified framework for detecting
604 605	out-of-distribution samples and adversarial attacks. Advances in neural information processing systems, 31, 2018.
606	
607 608	Taehyun Lee, Seokhee Hong, Jaewoo Ahn, Ilgee Hong, Hwaran Lee, Sangdoo Yun, Jamin Shin, and Gunhee Kim. Who wrote this code? watermarking for code generation. <i>arXiv preprint</i>
600	arXiv:2305.15060, 2023.
610 611	Dongfang Li, Zetian Sun, Xinshuo Hu, Zhenyu Liu, Ziyang Chen, Baotian Hu, Aiguo Wu, and Min Zhang A survey of large language models attribution <i>ArXiv</i> abs/2311.03731, 2023
612	
613	Haodong Li, Gelei Deng, Yi Liu, Kailong Wang, Yuekang Li, Tianwei Zhang, Yang Liu, Guoai Xu,
614 615	Guosheng Xu, and Haoyu Wang. Digger: Detecting copyright content mis-usage in large language model training. <i>ArXiv</i> , abs/2401.00676, 2024.
616	
617	Yucheng L1. Estimating contamination via perplexity: Quantifying memorisation in language model evaluation. <i>ArXiv</i> , abs/2309.10677, 2023.
618	Aiwei Liu Levi Pan Yuming Hu Shuang Li Lijie Wen Irwin King and Philin S. Vu. An unforgeable
619 620	publicly verifiable watermark for large language models. 2023.
621 622	Fangyu Liu, Ivan Vulić, Anna Korhonen, and Nigel Collier. Fast, effective, and self-supervised: Transforming masked language models into universal lexical and sentence encoders. In <i>Proceedings</i>
624 625	of the 2021 Conference on Empirical Methods in Natural Language Processing, pp. 1442–1459, 2021.
626 627 628 629	Justus Mattern, Fatemehsadat Mireshghallah, Zhijing Jin, Bernhard Schölkopf, Mrinmaya Sachan, and Taylor Berg-Kirkpatrick. Membership inference attacks against language models via neigh- bourhood comparison. In Anna Rogers, Jordan L. Boyd-Graber, and Naoaki Okazaki (eds.), <i>Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14,</i> 2023, pp. 11330, 11343. Association for Computational Linguistics, 2023.
630	2023, pp. 11550–11545. Association for Computational Eniguistics, 2025.
631 632	Leland McInnes and John Healy. Umap: Uniform manifold approximation and projection for dimension reduction. <i>ArXiv</i> , abs/1802.03426, 2018.
633	
634 635	Sewon Min, Suchin Gururangan, Eric Wallace, Weijia Shi, Hannaneh Hajishirzi, Noah A Smith, and Luke Zettlemoyer. Silo language models: Isolating legal risk in a nonparametric datastore. In <i>The</i>
636	Iweifth International Conference on Learning Representations.
637	Shami Nisimov Raanan V. Yehezkel Rohekar, Yaniy Gurwicz, Guy Koren, and Gal Novik. Clear:
638 639	Causal explanations from attention in neural recommenders. <i>ArXiv</i> , abs/2210.10621, 2022.
640	Michael-Andrei Panaitescu-Liess, Zora Che, Bang An, Yuancheng Xu, Pankavaraj Pathmanathan,
641 642	Souradip Chakraborty, Sicheng Zhu, Tom Goldstein, and Furong Huang. Can watermarking large language models prevent copyrighted text generation and hide training data? <i>arXiv preprint</i>
642	arXiv:2407.17417, 2024.
643	
645	Revanth Gangi Reddy, Yi Ren Fung, Qi Zeng, Manling Li, Ziqi Wang, Paul Sullivan, and Heng Ji.
640	Smartbook: A1-assisted situation report generation. ArXiv, abs/2303.14337, 2023.
647	Raanan Y Rohekar, Yaniv Gurwicz, and Shami Nisimov. Causal interpretation of self-attention in pre-trained transformers. <i>Advances in Neural Information Processing Systems</i> , 36, 2024.

- 648 Sarah Scheffler, Eran Tromer, and Mayank Varia. Formalizing human ingenuity: A quantitative 649 framework for copyright law's substantial similarity. Proceedings of the 2022 Symposium on 650 Computer Science and Law, 2022. 651
- Weijia Shi, Anirudh Ajith, Mengzhou Xia, Yangsibo Huang, Daogao Liu, Terra Blevins, Danqi 652 Chen, and Luke Zettlemoyer. Detecting pretraining data from large language models. ArXiv, 653 abs/2310.16789, 2023. 654
- 655 Zhiqing Sun, Xuezhi Wang, Yi Tay, Yiming Yang, and Denny Zhou. Recitation-augmented language 656 models. In The Eleventh International Conference on Learning Representations, ICLR 2023, 657 Kigali, Rwanda, May 1-5, 2023.
- 658 Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam M. Shazeer, Apoorv Kulshreshtha, Heng-Tze 659 Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, Yaguang Li, Hongrae Lee, Huaixiu Steven 660 Zheng, Amin Ghafouri, Marcelo Menegali, Yanping Huang, Maxim Krikun, Dmitry Lepikhin, 661 James Qin, Dehao Chen, Yuanzhong Xu, Zhifeng Chen, Adam Roberts, Maarten Bosma, Yanqi 662 Zhou, Chung-Ching Chang, I. A. Krivokon, Willard James Rusch, Marc Pickett, Kathleen S. Meier-663 Hellstern, Meredith Ringel Morris, Tulsee Doshi, Renelito Delos Santos, Toju Duke, Johnny Hartz Søraker, Ben Zevenbergen, Vinodkumar Prabhakaran, Mark Díaz, Ben Hutchinson, Kristen Olson, 665 Alejandra Molina, Erin Hoffman-John, Josh Lee, Lora Aroyo, Ravi Rajakumar, Alena Butryna, Matthew Lamm, V. O. Kuzmina, Joseph Fenton, Aaron Cohen, Rachel Bernstein, Ray Kurzweil, 666 Blaise Aguera-Arcas, Claire Cui, Marian Rogers Croak, Ed Huai hsin Chi, and Quoc Le. Lamda: 667 Language models for dialog applications. ArXiv, abs/2201.08239, 2022. 668
- 669 Nikhil Vyas, Sham M. Kakade, and Boaz Barak. On provable copyright protection for generative 670 models. In International Conference on Machine Learning, 2023. 671
- Ben Wang and Aran Komatsuzaki. GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model. 672 https://github.com/kingoflolz/mesh-transformer-jax, May 2021. 673
- 674 Jingtan Wang, Xinyang Lu, Zitong Zhao, Zhongxiang Dai, Chuan-Sheng Foo, See-Kiong Ng, and 675 Kian Hsiang Low. Wasa: Watermark-based source attribution for large language model-generated 676 data. ArXiv, abs/2310.00646, 2023a.
- Peng Wang, Ningyu Zhang, Xin Xie, Yunzhi Yao, Bo Tian, Mengru Wang, Zekun Xi, Siyuan Cheng, 678 Kangwei Liu, Guozhou Zheng, and Huajun Chen. Easyedit: An easy-to-use knowledge editing 679 framework for large language models. ArXiv, abs/2308.07269, 2023b. 680
- 681 Lauren Watson, Chuan Guo, Graham Cormode, and Alexandre Sablayrolles. On the importance of 682 difficulty calibration in membership inference attacks. In International Conference on Learning Representations, 2021. 683
- Boyi Wei, Weijia Shi, Yangsibo Huang, Noah A. Smith, Chiyuan Zhang, Luke Zettlemoyer, Kai Li, 685 and Peter Henderson. Evaluating copyright takedown methods for language models, 2024. 686
- 687 Yihan Wu, Zhengmian Hu, Hongyang Zhang, and Heng Huang. Dipmark: A stealthy, efficient and 688 resilient watermark for large language models. ArXiv, abs/2310.07710, 2023.
 - Yuanshun Yao, Xiaojun Xu, and Yang Liu. Large language model unlearning. ArXiv, abs/2310.10683, 2023.
- 692 Ruisi Zhang, Shehzeen Samarah Hussain, Paarth Neekhara, and Farinaz Koushanfar. Remark-llm: 693 A robust and efficient watermarking framework for generative large language models. ArXiv, abs/2310.12362, 2023. 694
- G. Zuccon, Bevan Koopman, and Razia Shaik. Chatgpt hallucinates when attributing answers. 696 Proceedings of the Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region, 2023.
- 698 699

677

684

689

690

- 700

702 A APPENDIX

A.1 DETAILS OF DATA SCALE ROBUSTNESS EXPERIMENT

Table 2: Data scale effect across different methods and models. k denotes the sampling token number. The model accuracy strictly follows the increase in k and data scale until k = 7.

Data scale (Accuracy%)	k	Int	er.	Va	ar.	A-1	Var.	Baseline1		Baseline2	
		2.7B	6B	2.7B	6B	2.7B	6B	2.7B	6B	2.7B	6B
S 11	1	69.6	75.6	63.4	75.7	75.4	70.8	-	-	-	-
Small	3	85.5	87.1	76.9	83.5	83.8	80.4	-	-	-	-
(160,6400)	5	85.9	88.6	79.7	85.9	85.9	82.7	-	-	-	-
	7	82.7	87.6	67.4	84.7	79.0	81.4	-	-	39.2	72.3
M - 1	1	88.6	89.0	89.9	91.2	89.5	88.3	-	-	-	-
Medium	3	92.8	93.1	91.5	92.6	91.8	91.1	-	-	-	-
(1600,6400)	5	93.3	93.6	92.3	93.2	93.1	91.3	-	-	-	-
	7	93.4	94.0	92.5	93.2	93.0	91.7	-	-	42.7	79.4
т	1	92.0	93.3	93.4	93.9	92.2	93.3	-	-	-	-
Large (12800,6400)	3	94.8	95.4	94.3	95.3	94.3	94.4	-	-	-	-
	5	95.4	95.5	94.3	95.4	94.9	94.7	-	-	-	-
	7	95.2	957	94.8	953	95.2	95.1	14.0	15.6	42.8	79.9

A.2 DETAILS OF DATASCALE EXPERIMENTS

Table 3: Method Efficiency Comparison

Model	Vanilla N	Aodel Training	CopyLens Training			
	Time	Hardware	Time	Hardware		
BERT	4 days	TPU × 16	17 mins Acc: 89%	A100-40G × 1		
GPT-Neo-2.7B	90 days	A100-80G × 96	58 mins Acc: 95%	A100-40G × 1		
GPT-J-6B	5 weeks	TPU × 256	89 mins Acc: 95%	A100-40G × 1		

A.3 EXTENDED RELATED WORK

Data Filtering Filtering out copyrighted data directly avoids copyright issues in LLMs. Previous
studies show possible ways such as utilizing synthetic or deduplicated data Kandpal et al. (2022); He
et al. (2023). Whereas, such approaches can detrimentally affect model performance due to the high
quality of copyrighted datasets.

LLM unlearning Researches show that LLMs can memorize training datasets, which causes trials in
letting LLMs forget training data. (Carlini et al. (2021); Eldan & Russinovich (2023)) However, the
unlearning method was proved unreliable using a Min-k probability attack. (Shi et al. (2023))

Watermark of copyrighted data Adding a watermark to the copyrighted victim dataset makes tracing possible. Watermark based on linguistics features exploits statistical distribution, token entropy, and so on. (Wu et al. (2023); Kirchenbauer et al. (2023); Lee et al. (2023)) Despite being training-free, this method is vulnerable to attacks. (Panaitescu-Liess et al. (2024)) The token sampling method switches the watermark process to the LLM decoding stage, showing superior robustness. (Kuditipudi et al. (2023)) Similar approaches include adding a watermark during logit generation. (Liu et al. (2023)) End-to-end learning-based watermark methods are also proposed, which can also attribute data sources. However, the requirement of pre-training LLMs poses a significant computational burden. (Zhang et al. (2023); Wang et al. (2023a))

Member Inference Attack (MIA) Member Inference Attack aims at deciding whether specific data is included in training datasets. It usually requires a shadow reference model assuming this model shares similar training data distributions with the victim one. (Watson et al. (2021); Mattern et al. (2023)) However, it's hard to apply to black-box models and only serves the purpose of data provenance. (Shi et al. (2023))

761 762

A.4 DETAILS OF THE Baseline₂

In this section, we explain the algorithm details of distance-based baseline. After the MHA outputs
are collected, in each layer, the last token is selected for its substantial information because of the
decoder-based autoregressive manner. The mean value is thus calculated across all layers to form a
representation of a specific data sample. This algorithm works in a supervised way, where training
samples in each sub-dataset are averaged into one high-dimensional vector. In the inference stage, L1
distance is calculated to decide which class the test sample should belong to, as depicted in Alg. 1.

1:]	Input: Datasets, Language Model
2: 0	Output: Class Identification Success Rates
3:]	Initialize dictionaries for layers
4: 1	for <i>index</i> in <i>len(Datasets)</i> do
5:	for each l in layer do
6:	Compute the last token embedding across samples for each layer
7:	end for
8: 0	end for
9:1	for each test sample x_{test} in <i>I</i> estimate do
10:	$u \leftarrow \ \mu_l - x_{test}\ _1$ for each layer i
11:	$c \leftarrow \arg \min_d \sum_l d_l^{(\alpha)}$
12: 0	end for Calculate Tatal A courses:
13: 0	Calculate Total Accuracy
14: 1	
A.5	ALGORITHM DETAILS OF PIVOTAL TOKEN INFORMATION EXTRACTION STRATEGY (A-VAR)
A.5 The 2. Fi	ALGORITHM DETAILS OF PIVOTAL TOKEN INFORMATION EXTRACTION STRATEGY (A-VAR) algorithm details of the pivotal token information extraction strategy (A-VAR) are shown in Alg. ig. 8 shows the overview of two main information extraction strategies.
A.5 The 2. Fi	ALGORITHM DETAILS OF PIVOTAL TOKEN INFORMATION EXTRACTION STRATEGY (A-VAR) algorithm details of the pivotal token information extraction strategy (A-VAR) are shown in Alg. ig. 8 shows the overview of two main information extraction strategies.
A.5 The 2. Fi Algo 1: 1	ALGORITHM DETAILS OF PIVOTAL TOKEN INFORMATION EXTRACTION STRATEGY (A-VAR) algorithm details of the pivotal token information extraction strategy (A-VAR) are shown in Alg. ig. 8 shows the overview of two main information extraction strategies. Prithm 2 Layer-wise Aligned Outlier Tokens Representation (A-VAR) Input: MHA Outputs
A.5 The 2. Fi Algo 1: 1 2: 0	ALGORITHM DETAILS OF PIVOTAL TOKEN INFORMATION EXTRACTION STRATEGY (A-VAR) algorithm details of the pivotal token information extraction strategy (A-VAR) are shown in Alg. ig. 8 shows the overview of two main information extraction strategies. Prithm 2 Layer-wise Aligned Outlier Tokens Representation (A-VAR) Input: MHA Outputs Output: A Set of Aligned Token Positions
A.5 The 2. Fi Algo 1: 1 2: 0 3: 1	ALGORITHM DETAILS OF PIVOTAL TOKEN INFORMATION EXTRACTION STRATEGY (A-VAR) algorithm details of the pivotal token information extraction strategy (A-VAR) are shown in Alg. ig. 8 shows the overview of two main information extraction strategies. Drithm 2 Layer-wise Aligned Outlier Tokens Representation (A-VAR) Input: MHA Outputs Output: A Set of Aligned Token Positions Initialize an empty set of aligned token positions $A_{index} \leftarrow \emptyset$
A.5 The 2. Fi Algo 1: 1 2: 0 3: 1 4: 1	ALGORITHM DETAILS OF PIVOTAL TOKEN INFORMATION EXTRACTION STRATEGY (A-VAR) algorithm details of the pivotal token information extraction strategy (A-VAR) are shown in Alg. by Shows the overview of two main information extraction strategies. Prithm 2 Layer-wise Aligned Outlier Tokens Representation (A-VAR) Input: MHA Outputs Output: A Set of Aligned Token Positions Initialize an empty set of aligned token positions $A_{index} \leftarrow \emptyset$ Initialize a dictionary for position occurrence counts $Num \leftarrow \{0, \dots, n-1 \mapsto 0\}$
A.5 The 2. Fi 2. Fi Algo 1: 1 2: 0 3: 1 4: 1 5: 1	ALGORITHM DETAILS OF PIVOTAL TOKEN INFORMATION EXTRACTION STRATEGY (A-VAR) algorithm details of the pivotal token information extraction strategy (A-VAR) are shown in Alg. ig. 8 shows the overview of two main information extraction strategies. Prithm 2 Layer-wise Aligned Outlier Tokens Representation (A-VAR) Input: MHA Outputs Output: A Set of Aligned Token Positions Initialize an empty set of aligned token positions $A_{index} \leftarrow \emptyset$ Initialize a dictionary for position occurrence counts $Num \leftarrow \{0, \ldots, n-1 \mapsto 0\}$ for each layer l in layers do
A.5 The 2. Fi Algo 1: 1 2: 0 3: 1 4: 1 5: 1 6: 7:	ALGORITHM DETAILS OF PIVOTAL TOKEN INFORMATION EXTRACTION STRATEGY (A-VAR) algorithm details of the pivotal token information extraction strategy (A-VAR) are shown in Alg. ig. 8 shows the overview of two main information extraction strategies. Prithm 2 Layer-wise Aligned Outlier Tokens Representation (A-VAR) Input: MHA Outputs Output: A Set of Aligned Token Positions Initialize an empty set of aligned token positions $A_{index} \leftarrow \emptyset$ Initialize a dictionary for position occurrence counts $Num \leftarrow \{0, \dots, n-1 \mapsto 0\}$ for each layer l in $layers$ do Compute and select the top- k token variance indices S_l for each index i in S do
A.5 The 2. Fi Algo 1: 1 2: 0 3: 1 4: 1 5: 1 6: 7: 8:	ALGORITHM DETAILS OF PIVOTAL TOKEN INFORMATION EXTRACTION STRATEGY (A-VAR) algorithm details of the pivotal token information extraction strategy (A-VAR) are shown in Alg. ig. 8 shows the overview of two main information extraction strategies. Prithm 2 Layer-wise Aligned Outlier Tokens Representation (A-VAR) Input: MHA Outputs Output: A Set of Aligned Token Positions Initialize an empty set of aligned token positions $A_{index} \leftarrow \emptyset$ Initialize a dictionary for position occurrence counts $Num \leftarrow \{0, \dots, n-1 \mapsto 0\}$ for each layer <i>l</i> in <i>layers</i> do Compute and select the top- <i>k</i> token variance indices S_l for each index <i>i</i> in S_l do Increment Num[i] by 1
A.5 The 2. Fi 2. Fi 1: 1 2: 0 3: 1 4: 1 5: 1 6: 7: 8: 9:	ALGORITHM DETAILS OF PIVOTAL TOKEN INFORMATION EXTRACTION STRATEGY (A-VAR) algorithm details of the pivotal token information extraction strategy (A-VAR) are shown in Alg. ig. 8 shows the overview of two main information extraction strategies. Prithm 2 Layer-wise Aligned Outlier Tokens Representation (A-VAR) Input: MHA Outputs Output: A Set of Aligned Token Positions Initialize an empty set of aligned token positions $A_{index} \leftarrow \emptyset$ Initialize a dictionary for position occurrence counts $Num \leftarrow \{0, \dots, n-1 \mapsto 0\}$ for each layer l in layers do Compute and select the top-k token variance indices S_l for each index i in S_l do Increment $Num[i]$ by 1 end for
A.5 The 2. Fi 2. Fi 2: 0 3: 1 4: 1 5: 1 6: 7: 8: 9:	ALGORITHM DETAILS OF PIVOTAL TOKEN INFORMATION EXTRACTION STRATEGY (A-VAR) algorithm details of the pivotal token information extraction strategy (A-VAR) are shown in Alg. ig. 8 shows the overview of two main information extraction strategies. Drithm 2 Layer-wise Aligned Outlier Tokens Representation (A-VAR) Input: MHA Outputs Output: A Set of Aligned Token Positions Initialize an empty set of aligned token positions $A_{index} \leftarrow \emptyset$ Initialize a dictionary for position occurrence counts $Num \leftarrow \{0, \ldots, n-1 \mapsto 0\}$ for each layer l in layers do Compute and select the top- k token variance indices S_l for each index i in S_l do Increment $Num[i]$ by 1 end for
A.5 The : 2. Fi Algo 1: 1 2: 0 3: 1 4: 1 5: 1 6: 7: 8: 9: 10: 0	ALGORITHM DETAILS OF PIVOTAL TOKEN INFORMATION EXTRACTION STRATEGY (A-VAR) algorithm details of the pivotal token information extraction strategy (A-VAR) are shown in Alg. ig. 8 shows the overview of two main information extraction strategies. Drithm 2 Layer-wise Aligned Outlier Tokens Representation (A-VAR) Input: MHA Outputs Output: A Set of Aligned Token Positions Initialize an empty set of aligned token positions $A_{index} \leftarrow \emptyset$ Initialize a dictionary for position occurrence counts $Num \leftarrow \{0, \dots, n-1 \mapsto 0\}$ for each layer l in layers do Compute and select the top-k token variance indices S_l for each index i in S_l do Increment $Num[i]$ by 1 end for $A_{index} \leftarrow top-k(\{Num[i]\}_{n=1}^{n-1} k\})$

To show the effectiveness of the designed prompt template, Fig. 9 presents an example of a styletransferring prompt. The figure shows how a text related to database security is accurately transformed into SQL code. An ablation experiment on generation length is shown in Tab. 4 to prove robustness.



Figure 8: Annotations and explanation of the information extraction and token fusion process, showing two main methods.

Continuation Scene	64	128	256	512
Accuracy (%)				
INTER	73.65	72.95	77.82	71.73
A-VAR	72.95	72.96	80.56	78.94
VAR	70.98	77.03	75.15	72.98

Table 4: Ablation experiments on robustness to different generation lengths on GPT-Neo-2.7B. The impact of text length on our proposed method is minimal, and our method is more suitable for real-world scenarios with longer texts.

Besides, the prompt engineering template for the baseline is provided. In this case, due to the lengthy dialog input, words such as "judge" and "charge" are ignored by the LLM. Instead, attention is given to the nearest word "search", leading to an incorrect answer, as illustrated in Fig. 10.



Figure 9: Prompt Design and LLM Output in Style Transfer Scene. (a) Designed prompt input. (b) LLM output.



Figure 10: Prompt Design and LLM Output in Prompt Engineering Baseline. (a) Designed prompt input. (b) LLM output.

A.7 MSE LOSS IN TEXT MIXTURE TASK

In the text mixture task, eight copyrighted sub-dataset classes in the Pile {0, 1, 2, 3, 4, 5, 6, 7} are used for mixture. We generate mixed text by sampling text segments based on predefined proportions (15%, 15%, 70%) for each label. The segments are then randomly combined to create the final mixed text. MSE loss is computed between predicted dataset contribution scores and predefined portions, which is lower than 0.05 in average, as Tab. 5 shows.

864	MSE Loss(1)	G	PT-Neo-2.7	'B	GPT-J-6B			
866	11012 1005(4)	(0.15, 0.15, 0.70)	INTER	VAR	A-VAR	INTER	VAR	A-VAR
867		[0,7,6]	0.00644	0.00245	0.00448	0.00925	0.00433	0.00184
868		[1,0,5]	0.00229	0.01572	0.00501 0.00342	0.00231	0.01814 0.04381	0.01711 0.03847
869		[2,5,0]	0.00773	0.00771	0.00729	0.00941	0.01251	0.00762
870		[3,1,4] [4,1,6]	0.03448 0.01747	0.03803 0.02788	0.01674 0.00652	0.04416	0.03830 0.03489	0.02203 0.03057
871		[4,7,0]	0.00474	0.00295	0.00385	0.00410	0.00723	0.00559
872		[5,6,0] [7,4,6]	0.00440 0.00365	0.00439 0.00887	0.01920 0.00649	0.02089	0.00949	0.00499
873		Mean	0.00939	0.01432	0.00811	0.01331	0.01941	0.01462

Table 5: MSE Loss for Different Models in Text Mixture Task.

SIMPLE BUT EFFECTIVE: WHY DOES INTERVAL SAMPLING OUTPERFORM OUTLIER A.8 METHODS IN MOST CASES?

Tab. 2 shows that in vanilla classification scenes, interval-based sampling outperforms pivotal-token-based (outlier) methods. We try to explain why this happens from the perspective of information theory and take k = 3 in the smallest data scale as an example.

Why interval-based sampling method is effective? We formularize the overall information extraction and classification process based on mutual information (MI). MI quantifies the statistical association between variables, exploring both linear and non-linear dependencies, and can be used to identify relationships in potential causality. A larger MI shows a closer relationship between the two variables.

The process of the proposed framework can be summarized into two parts: first extracting token representations from MHA outputs, then feeding into the LSTM-based contribution analysis framework. The latter can be simplified as a classification. TODOTODOTODO need revise the framework

Problem Formulation We choose k = 3 with the smallest data scale as basic settings in Tab. 2. Therefore, in each layer l with an input training sample X, given MHA outputs O_l across all layers, the extracted three representations are denoted as S_{1l}, S_{2l}, S_{3l} with the extraction policy F_{θ} .

Our initial goal is to maximize the mutual information between extracted representations S and the corresponding data sample Y label. That is, $\max I(Y; S)$, where I(;) denotes mutual information.

Due to the constraints on the complexity of S_l , information must be efficiently compressed. To achieve this, we must minimize the MI between the MHA outputs and the extracted representations, as well as between the extracted representations themselves. Therefore, the final objective is to maximize the information bottleneck as shown in eq. equation 11:

$$\max_{F_{\theta}} \sum_{l} \left(I(Y; S_{l}) - \beta I(O_{l}; S_{l}) - \beta \sum_{1 \le n1 < n2 \le 3} I(S_{n1}; S_{n2}) \right)$$
(11)

where β is set to 0.5, and

$$I(Y;S_l) = I(Y;S_{1l}) + I(Y;S_{2l}) + I(Y;S_{3l})$$
(12)

$$I(O_l; S_l) = I(O_l; S_{1l}) + I(O_l; S_{2l}) + I(O_l; S_{3l})$$
(13)

$$\sum_{1 \le n1 \le n2 \le 3} I(S_{n1}; S_{n2}) = I(S_{1l}; S_{2l}) + I(S_{1l}; S_{3l}) + I(S_{2l}; S_{3l})$$
(14)

We denote the above formulation as Information Bottleneck (IB). The larger this metric is, the more necessary information the process maintains.

Relationship between IB and Classification Accuracy An intuitive idea is that better IB metrics mean higher accuracy. To investigate the relationship between the Information Bottleneck (IB) metric and classification accuracy, we conducted experiments on GPT-Neo-2.7B and GPT-J-6B. As depicted in Tab. 6 and Tab. 7, the IB metric closely correlates with classification accuracy. The interval-based method shows better IB metrics, which corresponds to its higher accuracy. This correlation may explain the superior performance observed with interval-based samples.

In summary, we have designed an IB-based metric that quantitatively explains the differences between extraction methods. This paves the way for research into optimized extraction strategies in theory.

Methods	INTER.	VAR.	A-VAR.
Framework Classification Accuracy (%)	85.5	76.9	83.8
Mutual Information			
$I(O; S_1)(\downarrow)$	-6.86	11.09	10.13
$I(O; S_2)(\downarrow)$	-1.63	12.63	12.46
$I(O; S_3)(\downarrow)$	1.96	12.87	12.88
$I(S_1; S_2)(\downarrow)$	1.52	42.87	28.79
$I(S_1; S_3)(\downarrow)$	1.50	38.55	23.26
$I(S_2; S_3)(\downarrow)$	18.04	42.68	27.59
$I(Y; S_1)(\uparrow)$	36.44	1.17	2.78
$I(Y;S_2)(\uparrow)$	0.17	0.39	0.63
$I(Y;S_3)(\uparrow)$	1.44	0.31	0.65
$I(O;S)(\downarrow)$	-6.53	36.59	35.46
Final: Information Bottleneck (↑)	30.79	-78.47	-53.50

Table 6: MI comparison with interval-based and outlier methods on GPT-Neo-2.7B. The classification accuracy is from Tab. 2, and the calculated information bottleneck shows exactly the same trend with framework accuracy. Experiments are conducted under settings of k = 3 at the smallest data scale.

Methods	INTER.	VAR.	A-VAR.
Framework Classification Accuracy (%)	87.1	83.5	80.4
Mutual Information			
$I(O; S_1)(\downarrow)$	0.95	3.97	12.75
$I(O; S_2)(\downarrow)$	1.52	3.53	16.16
$I(O; S_3)(\downarrow)$	1.54	3.25	17.11
$I(S_1; S_2)(\downarrow)$	1.27	32.01	36.42
$I(S_1; S_3)(\downarrow)$	1.18	28.60	33.02
$I(S_2; S_3)(\downarrow)$	8.43	30.13	43.82
$I(Y; S_1)(\uparrow)$	35.21	2.46	5.05
$I(Y;S_2)(\uparrow)$	4.77	1.32	6.33
$I(Y;S_3)(\uparrow)$	8.76	0.76	5.17
$I(O;S)(\downarrow)$	4.01	10.75	46.01
Final: Information Bottleneck (↑)	41.29	-46.21	-63.09

Table 7: MI comparison with interval-based and outlier methods on GPT-J-6B. The classification accuracy is from Tab. 2, and the calculated information bottleneck shows exactly the same trend with framework accuracy. Experiments are conducted under settings of k = 3 at the smallest data scale.