

# PDSS: A Privacy-Preserving Framework for Step-by-Step Distillation of Large Language Models

Anonymous ACL submission

## Abstract

In the context of real-world applications, leveraging large language models (LLMs) for domain-specific tasks often faces two major challenges: domain-specific knowledge privacy and constrained resources. To address these issues, we propose PDSS, a privacy-preserving framework for step-by-step distillation of LLMs. PDSS works on a server-client architecture, wherein client transmits perturbed prompts to the server’s LLM for rationale generation. The generated rationales are then decoded by the client and used to enrich the training of task-specific small language model(SLM) within a multi-task learning paradigm. PDSS introduces two privacy protection strategies: the *Exponential Mechanism Strategy* and the *Encoder-Decoder Strategy*, balancing prompt privacy and rationale usability. Experiments demonstrate the effectiveness of PDSS in various text generation tasks, enabling the training of task-specific SLM with enhanced performance while prioritizing data privacy protection.

## 1 Introduction

Large Language Models(LLMs), boasting billions of parameters and remarkable text generation abilities, have risen as a revolutionary force in artificial intelligence. Prominent models, such as GPT-4 (OpenAI, 2023), LLaMA(Touvron et al., 2023), and Qwen(Bai et al., 2023), have garnered the attention of researchers and practitioners alike, demonstrating unparalleled proficiency across numerous tasks. Nevertheless, the sheer size of these models presents significant obstacles for real-world deployment, particularly in environments with limited resources. Meanwhile, as LLMs gain escalating popularity and widespread utilization, privacy concerns have moved to the forefront, especially when it comes to user data and model inference. In contrast, Small Language Models(SLMs) often exhibit

superior computational efficiency and faster convergence rates, rendering them perfectly suited for real-time applications or resource-constrained environments. Nonetheless, SLMs also possess certain drawbacks stemming from their performance limitations. The question then arises: *How can we effectively combine the predictive prowess of LLMs with the nimbleness of SLMs, all while adhering to privacy requirements?*

To address these challenges, we introduce PDSS, a privacy-preserving framework for step-by-step distillation of LLMs. In our envisioned setup, there’s a high-powered server capable of deploying an LLM, paired with a client possessing more limited computational resources running SLM. The challenge lies in maintaining the privacy of client data while leveraging the server’s LLM to aid in training the client’s SLM for text generation tasks, thereby elevating its performance. PDSS aims to bridge this gap, enabling secure and efficient knowledge transfer between LLM and SLM, and ultimately enhancing the capabilities of the SLM without compromising privacy.

As illustrated in Figure 1, within our framework, the process works as follows. Initially, the client transmits perturbed prompts to the server’s LLM, which are protected by the PDSS prompt encoder module, thus ensuring privacy protection. Subsequently, the server’s LLM generates perturbed rationales from these prompts through the *Chain of Thought (COT)* approach (Wei et al., 2022) and relays them back to the client. Upon receiving these perturbed rationales, the client’s rationales decoder module reconstructs them into their original, aligned form corresponding to the raw prompt. Ultimately, the client incorporates these rationales as supplementary and enriching information for training its *Task-Specific SLM* within a multi-task learning paradigm (Wei et al., 2022; Hsieh et al., 2023; Zhang and Yang, 2021). These rationales justify the predicted labels and serve as insightful

guidance for training smaller and domain-specific models.

Within the PDSS framework, to achieve a balance between preserving the privacy of user prompts and enhancing the usability of rationales, we introduce two privacy protection strategies incorporated into the the prompt encoder module and the rationales decoder module: the *Exponential Mechanism Strategy* and the *Encoder-Decoder Strategy*. In the *Exponential Mechanism Strategy*, we utilize an exponential mechanism to obfuscate the prompts (Tong et al., 2023), followed by decoding the perturbed rationales through In-Context Learning (ICL) (Dong et al., 2022). In the *Encoder-Decoder strategy*, we utilize an Encoder-Decoder SLM specifically designed to encode raw prompts into perturbed prompts and subsequently decode perturbed rationales back into their original form. To effectively train this unified Encoder-Decoder SLM, we utilize a multi-task learning paradigm (Zhang and Yang, 2021), encompassing both the encoding and decoding training processes.

Our contributions are summarized as follows:

- **Privacy-Preserving Framework for LLM Distillation.** We propose PDSS, a novel framework that facilitates secure and efficient knowledge transfer from LLM to SLM in resource-constrained environments while adhering to privacy requirements. PDSS addresses the challenges posed by the massive size of LLMs for real-world deployment and the privacy concerns surrounding user data. By utilizing perturbed prompts and rationales, PDSS ensures data privacy while leveraging the predictive prowess of LLMs to enhance the performance of SLMs.
- **Innovative Privacy Protection Strategies.** Within PDSS, we introduce two privacy protection strategies: the *Exponential Mechanism Strategy* and the *Encoder-Decoder Strategy*. The former utilizes an exponential mechanism to obfuscate user prompts, while the latter employs a specialized Encoder-Decoder SLM to encode and decode perturbed prompts and rationales. These strategies effectively balance user privacy and the usability of rationales, allowing for secure and enhanced training of the client’s SLM without compromising on privacy concerns.
- **Empirical Evaluation and Enhanced Per-**

**formance of Task-Specific SLM.** Through experiments on various text generation tasks, PDSS demonstrates the effectiveness of its framework in training task-specific SLM with enhanced performance. By harnessing the rationales generated by the server-side LLM, PDSS provides valuable task-specific knowledge to the SLM, enabling them to achieve significant improvements with the support of the LLM while prioritizing data privacy protections.

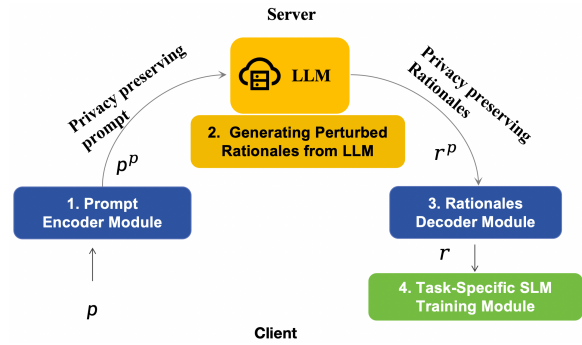


Figure 1: Overview of our proposed PDSS workflow.

**[1. Raw Private Prompt]**  
 Question: A person broke up with me, what do I have now?  
 Answer Choices: (a) annoyance (b) hospitalized (c) schizophrenia (d) broken heart (e) stress

**[2. Prompt Encode]**  
 Question: A Students unhappy up offers me, Whit do I find now?  
 Answer Choices: (a) benefittance (b) admisionized (c) schizophrenia (d) disappointed craft (e) crazy

**[3. Perturbed Rationales Generation]**  
 The question suggests that the student is unhappy, so the appropriate response would be related to their emotional state. Of the given choices, "disappointed craft" (which could imply a disappointed or unfulfilled state) fits the context, as it indicates dissatisfaction. The other options don't directly relate to the student's unhappiness.

**[4. Rationales Decode]**  
 The question suggests that the person broke up with the person, so the appropriate response would be related to the person's emotional state. Of the given choices, "broken heart" (which could imply a broken or heartbroken state) fits the context, as it indicates the person's emotional state. The other options don't directly relate to the person's emotional state.

Figure 2: Privacy-Preserving Rationals Generation Example.

## 2 Related Work

### 2.1 Chain of Thought in Large Language Models

The Chain of Thought (COT) approach has recently garnered significant attention in the realm of LLMs, thanks primarily to its remarkable ability to enhance the reasoning capabilities of these models.

This innovative concept was first introduced by (Wei et al., 2022). Their research demonstrated that by prompting LLMs to produce a sequence of intermediary reasoning steps (rationales), the models’ performance in handling intricate reasoning tasks could be notably boosted. This groundbreaking study opened the door for further explorations into COT. Since the introduction of COT, several studies have delved into its extensions and variations. For example, (Kojima et al., 2022) proposed the use of zero-shot COT, where the model is prompted to generate reasoning steps (rationales) without relying on prior examples. COT has also been applied to various domains, including arithmetic reasoning (Cobbe et al., 2021), commonsense reasoning (Klein and Nabi, 2020).

Nonetheless, despite the impressive feats achieved by LLMs, the adoption of LLMs in domain-specific applications with constrained resources poses a significant challenge (Fan et al., 2023) (Kang et al., 2023). Recent studies by (Hsieh et al., 2023) (Ho et al., 2022) (Li et al., 2023), have capitalized on the generated rationales as a form of insightful supervision to train smaller and domain-specific models. However, previous studies have not addressed the domain-specific data privacy issue that arises when LLMs and domain-specific smaller models are deployed across different parties. In our work, we endeavor to address this significant challenge.

## 2.2 Privacy Preserving LLM Inference

With the escalating popularity and widespread utilization of LLMs, privacy concerns have taken center stage, particularly regarding user data and model inference. Previous research efforts aimed at preserving privacy during LLM inference have predominantly focused on several key techniques, including differential privacy (DP) (Dwork, 2006), fully homomorphic encryption (FHE) (Gentry, 2009), and secure multi-party computation (MPC) (Yao, 1986) protocols.

Numerous studies have delved into the intricacies of LLM inference leveraging DP techniques. Notably, methods like SANTEXT+ (Yue et al., 2021), CUSTEXT+ (Chen et al., 2022), TextObfuscator (Zhou et al., 2023) and InferDPT (Tong et al., 2023) have harnessed differential privacy to sequentially replace sensitive words in the text with semantically similar alternatives from a predefined word adjacency list.

FHE and MPC techniques have also garnered

attention as viable methods for ensuring privacy during LLM inference. For instance, CipherGPT (Hou et al., 2023) proposes a secure matrix multiplication and a novel protocol for securely computing GELU within transformer architecture using FHE and MPC protocols to facilitate secure two-party GPT inference. Likewise, Puma (Dong et al., 2023) has adopted FHE and MPC in its transformer architecture for secure third-party LLM inference. While FHE and MPC can be utilized for privacy-preserving text generation tasks, their practical applications remain limited primarily due to significant computational and communication overheads.

The advancements in privacy-preserving techniques, such as differential privacy, FHE, and MPC, offer promising solutions to mitigate privacy risks associated with LLM inference. However, balancing privacy and efficiency remains a challenge that requires further exploration and refinement.

## 3 The Proposed PDSS Framework

### 3.1 Overview

In this section, we introduce PDSS, an innovative privacy-preserving framework specifically designed for distilling step-by-step LLMs. The PDSS framework can enhance the performance of SLMs while maintaining privacy, leveraging the capabilities of LLM. We illustrate the PDSS in Figure 1 and describe the associated training algorithm in Algorithm 1. The workflow of PDSS is outlined as follows:

1. In the client, **Prompt Encoder Module** perturbs these prompts before sending them to the server-side LLM.
2. In the server, the server-side LLM **generates perturbed rationales** based on these perturbed prompts and sends them back to the client.
3. In the client, **Rationales Decoder Module** decodes the perturbed rationales.
4. In the client, **Task-Specific SLM Training Module** employs both the original label data and the filter rationales data for multi-task learning.

### 3.2 Prompt Encoder Module

In the prompt encoder module, as illustrated in Figure 3, we propose two privacy protection strategies:

247 **1. Exponential Mechanism Encoder Strategy.** In the first strategy, we utilize an exponential mechanism (McSherry and Talwar, 2007)(Tong et al., 2023), which satisfies the criteria for the  $\epsilon - DP$ . This strategy works by replacing each token in the prompt with a semantically similar one sampled from either a predetermined adjacency list or a randomly generated adjacency list, based on exponential mechanism.

257 *The Definition of Exponential Mechanism* (Tong et al., 2023). For a given scoring function  $u : X \times Y \rightarrow R$ , a randomized mechanism  $M(X, u, Y)$  is  $\epsilon - DP$  compliant if it satisfies:

$$P_r[y|x] \propto \exp\left(\frac{\epsilon \cdot u(x, y)}{2 \Delta u}\right) \quad (1)$$

263 where the sensitivity  $\Delta u$  is defined as:

$$\Delta u = \max_{x, x' \in X, y \in Y} |u(x, y) - u(x', y)| \quad (2)$$

265 **2. Encoder-Decoder Encoder Strategy.** The tokens within a prompt differ significantly in terms of their importance and degree of privacy. Applying a uniform privacy budget  $\epsilon$  across all tokens may not lead to the most optimal solution. To further optimize the privacy-utility balance, we propose an Encoder-Decoder strategy. This strategy is built upon the first exponential mechanism. In the Encoder-Decoder strategy, we utilize an Encoder-Decoder SLM specifically designed to encode raw prompts into perturbed prompts and subsequently decode perturbed rationales back into their original form. This strategy involves two training process: encoding training process and decoding training process. In this section, we mainly focus on encoding training process, as illustrated in Figure 3.

283 Initially, an encoding training process is required for the Encoder-Decoder SLM. Formally, let's denote a public dataset as  $P = \{(p_i, p_i^\epsilon)\}_{i=1}^N$ , where  $p_i$  represents raw private prompt,  $p_i^\epsilon$  represents perturbed prompt generated using the first exponential mechanism with a privacy budget of  $\epsilon$ . In the encoding training process, we train the Encoder-Decoder SLM:  $g_\phi(p_i) \rightarrow p_i^\epsilon$ . The details of encoding training process is illustrated in Algorithm 1.

294 The Encoder objective can be formulated as follows: 295

$$\mathcal{L}_{\text{Encoder}}(\phi; \mathcal{P}) = \mathbb{E}_{(p, p^\epsilon) \sim \mathcal{P}} \ell_{\text{CE}}(g_\phi(p), p^\epsilon) \quad (3)$$

297 where  $\ell_{\text{CE}}$  is the cross-entropy loss. 298

299 As illustrated in Figure 2, we can observe an exemplary comparison between the original input and its perturbed input in Step 1 and Step 2. This perturbed prompt serves as the new, privacy-enhanced input for further processing. 300 301 302

303 By incorporating this perturbation mechanism, we ensure that the privacy of the original prompt is preserved. This approach not only satisfies the privacy requirements but also enables effective data utilization for downstream tasks, striking a balance between privacy and utility. 304 305 306 307 308



Figure 3: Prompt Encoder Module.

### 3.3 Generating Perturbed Rationales from LLM 309 310

311 When the server-side LLM receives the perturbed prompt, we leverage the Chain-of-Thought (CoT) prompting technique introduced by (Wei et al., 2022) to generate rationales from the LLM using this perturbed prompt. These generated rationales, which are also perturbed, are then transmitted to the client. For instance, as illustrated in Figure 2, given a perturbed prompt in the Step 2, the LLM generates perturbed rationales in the Step 3. 312 313 314 315 316 317 318 319

### 3.4 Rationales Decoder Module 320

321 Once the client receives the perturbed rationales from the server-side LLM, it must initiate a "decoder" process within the rationales decoder module to decode the rationales. In rationales decoder module, as illustrated in Figure 4, we also propose two strategies correspond to the two protection strategy of the prompt encoder module: 322 323 324 325 326 327

#### 1. Exponential Mechanism Decoder Strategy. 328

329 In the first decoding strategy, which corresponds to Exponential Mechanism Encoder strategy. Here, we utilize In-Context Learning(ICL) (Dong et al., 2022) (Tong et al., 2023) with the SLM to decode the perturbed 330 331 332 333

rationales. we can input a sample  $x_i = (p, p^p, r^p)_i$  into the SLM to prompt the generation of rationales, where  $p$  represents raw private prompt,  $p^p$  represents perturbed prompt and  $r^p$  represents perturbed rationales generated from LLM.  $(p^p, r^p)_i$  can be viewed as an example for SLM in ICL. This allows the SLM to generate rationales  $r_i$  that are aligned with the original, unperturbed prompt.

2. **Encoder-Decoder Decoder Strategy.** In the second decoding strategy, which corresponds to Encoder-Decoder Encoder strategy. The rationales decoder module also use the same the Encoder-Decoder SLM with Section 3.2.

Initially, a decoding training process is required for the Encoder-Decoder SLM. Formally, let's denote a public dataset as  $R = \{(x_i, r_i)\}_{i=1}^N$ , where  $x_i$  represents an input, where  $x_i = (p, p^p, r^p)_i$ ,  $p$  represents raw private prompt,  $p^p$  represents perturbed prompt generated from Encoder-Decoder SLM,  $r^p$  represents perturbed rationales generated from LLM.  $r_i$  represents the raw rationale of raw prompt  $p$  generated from LLM. In the decoding training process, we train the Encoder-Decoder SLM:  $g_\phi(x_i) \rightarrow r_i$ . The details of decoding training process is illustrated in Algorithm 1.

The Decoder objective can be formulated as follows:

$$\mathcal{L}_{\text{Decoder}}(\phi; \mathcal{R}) = \mathbb{E}_{(x,r) \sim \mathcal{R}} \ell_{\text{CE}}(g_\phi(x), r) \quad (4)$$

where  $\mathcal{L}_{\text{Decoder}}$  is the rational decoder loss, and  $\ell_{\text{CE}}$  is the cross-entropy loss.

Subsequently, once the decoding training process of Encoder-Decoder SLM is finished, we can input a sample  $x_i = (p, p^p, r^p)_i$  into the SLM, where  $r^p$  represents perturbed rationales generated from LLM. This allows the SLM to generate rationales  $r_i$  that are aligned with the original, unperturbed prompt.

We approach the training of the Encoder-Decoder SLM as a multi-task learning problem encompassing both the encoding and decoding training processes. The multi-task learning objective can be formulated as follows:

$$\mathcal{L}_1 = \alpha \mathcal{L}_{\text{Encoder}} + (1 - \alpha) \mathcal{L}_{\text{Decoder}} \quad (5)$$

where  $\alpha$  is the hyperparameters that control the weight of encoder and decoder loss.

As illustrated in Figure 2, we can observe an exemplary comparison between the perturbed rationales from LLM and its decoded rationales from SLM in Step 3 and Step 4. It's worth noting that although the SLM has the ability to generate aligned rationales independently, the quality often falls short due to its limited capabilities. By leveraging the perturbed rationales, we effectively transfer the powerful capabilities of the server-side LLM to enhance the Encoder-Decoder SLM, thereby improving the overall quality of the generated rationales.

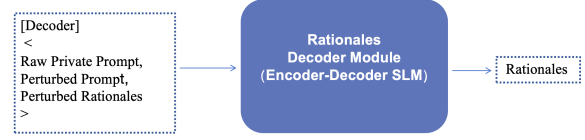


Figure 4: Rationales Decoder Module.

---

#### Algorithm 1 PDSS

---

##### Input:

- 1:  $T$ : total number of rounds;
- 2:  $\mathcal{P}$ : encoding training datasets;
- 3:  $\mathcal{R}$ : decoding training datasets;
- 4:  $\mathcal{D}$ : task-Spec training datasets;
- 5:  $\eta_\phi$ : learning rate of Encoder-Decoder SLM;
- 6:  $\eta_\omega$ : learning rate of Task-Specific SLM.

##### Output: $g_\phi, f_\omega$ .

- 7:  $\triangleright$  Multi-Task Training for Encoder-Decoder SLM based on Public Datasets  $\mathcal{P}$  and  $\mathcal{R}$ .
  - 8: **for** each epoch  $t \in [T]$  **do**
  - 9:      $\phi^{t+1} \leftarrow \phi^t - \eta_\phi \nabla \mathcal{L}_1$ .
  - 10: **end for**
  - 11:  $\triangleright$  Generated  $p^p$  using the updated Encoder.
  - 12:  $p^p = SLM_{\text{Encoder}}(p)$ .
  - 13:  $\triangleright$  Generated perturbed rationales from LLM on the server.
  - 14:  $r^p = LLM(p^p)$ .
  - 15:  $\triangleright$  Decoded perturbed rationales using the updated Encoder-Decoder SLM.
  - 16:  $r = SLM_{\text{Decoder}}(r^p)$ .
  - 17:  $\triangleright$  Multi-Task Training for Task-Specific SLM based on Datasets  $\mathcal{D}$ .
  - 18: **for** each epoch  $t \in [T]$  **do**
  - 19:      $\omega^{t+1} \leftarrow \omega^t - \eta_\omega \nabla \mathcal{L}_2$ .
  - 20: **end for**
-

### 3.5 Task-Specific SLM Training Module

In our work, we undertake the training of the client’s Task-Specific SLM tailored for text generation tasks. Initially, we elaborate on the prevalent framework for learning task-specific models. Leveraging this established framework, we enhance it by integrating rationales produced from the rationales decoder module into the training process. Formally, let’s denote a dataset as  $D = \{(x_i, (y_i, r_i))\}_{i=1}^N$ , where  $x_i$  represents an input,  $y_i$  represents the associated expected output label, and  $r_i$  is the corresponding desired rationale.

We conceptualize learning with rationales as a *multi-task learning* problem, as illustrated in Figure 5. Specifically, we train the model  $f_\omega(x_i) \rightarrow (y_i, r_i)$  to accomplish not just the prediction of task labels but also the generation of the corresponding rationales based on textual inputs. This multi-task training ensures that our model not only produces accurate predictions but also provides insightful justifications for its decisions. By doing so, we enhance the transparency and explainability of the model. The multi-task learning objective can be formulated as follows:

$$\mathcal{L}_2 = \beta \mathcal{L}_{\text{Label}} + (1 - \beta) \mathcal{L}_{\text{Rationale}} \quad (6)$$

where  $\mathcal{L}_{\text{Label}}$  is the label prediction loss:

$$\mathcal{L}_{\text{Label}}(\omega; \mathcal{D}) = \mathbb{E}_{(x,y) \sim \mathcal{D}} \ell_{\text{CE}}(f_\omega(x), y) \quad (7)$$

and  $\mathcal{L}_{\text{Rationale}}$  is the rationale generation loss:

$$\mathcal{L}_{\text{Rationale}}(\omega; \mathcal{D}) = \mathbb{E}_{(x,r) \sim \mathcal{D}} \ell_{\text{CE}}(f_\omega(x), r) \quad (8)$$

where  $\ell_{\text{CE}}$  is the cross-entropy loss,  $f_\omega(\cdot)$  is the Task-Specific SLM model, and  $\beta$  is the hyperparameters that control the weight of label prediction loss and rationale generation loss.



Figure 5: Task-Specific SLM Training Module.

## 4 Experiments

### 4.1 Setup

We have established a scenario to evaluate the performance of the PDSS framework across a range of text generation tasks. This setup involves a

client-server architecture, where the client holds two downstream SLMs :an *Encoder-Decoder SLM*, which specializes in encoder-decoder functionalities and a *Task-Specific SLM*, tailored for specific tasks. On the server-side, we host a LLM for more general and powerful text generation capabilities. Specifically, we have chosen Qwen-14B(Bai et al., 2023) as LLM, while both SLMs are Qwen-0.5B(Bai et al., 2023). Table 1 outlines the detailed configurations of both the LLM and the SLMs.

**Datasets and Evaluation Metrics.** We conduct a comprehensive evaluation of PDSS on 4 QA datasets. Specifically, we include CommonsenseQA(CQA) (Talmor et al., 2018), OpenBookQA(OBQA) (Mihaylov et al., 2018), BoolQ (Clark et al., 2019), ARC-E(Clark et al., 2018). For these datasets, we primarily use **Accuracy** as the evaluation metric.

**Baselines.** Since we incorporate two distinct strategies in the prompt encoder module and rationales decoder module, we denote PDSS method with the Exponential Mechanism Strategy as *PDSS-EM* and PDSS method with the Encoder-Decoder Strategy as *PDSS-ED*. We conduct a comparative analysis to evaluate the performance of our PDSS framework, which comprises both *PDSS-EM* and *PDSS-ED*.

These baselines included:

- FewShot-LLM, which represents the few-shot capabilities of LLM on the server;
- FewShot-SLM, which represents the few-shot performance of SLM on the client;
- Standalone, where the client independently fine-tunes its local model using its own private dataset;
- DSS(Hsieh et al., 2023), where the client fine-tunes its local model by distilling step-by-step LLM method without privacy-preserving.

### 4.2 Overall Performance Evaluation

In this section, we undertake a comprehensive analysis of the task performance of PDSS. We assess both the PDSS-EM and PDSS-ED methods against other baselines on Task-Specific SLM across various privacy budgets, denoted by  $\epsilon$ .

The results, as presented in Table 2, clearly illustrate that both PDSS-EM and PDSS-ED exhibit significantly better performance when compared to FewShot-SLM and Standalone methods. With

Setting	Server	Client	Client
Model Type	LLM	Encoder-Decoder SLM	Task-Specific SLM
Model Name	Qwen-14B	Qwen-0.5B	Qwen-0.5B
Parameters(Billion)	14	0.5	0.5

Table 1: LLM and SLMs Setting of PDSS.

an increase in the privacy budget  $\epsilon$ , both the performance of PDSS-EM and PDSS-ED have risen notably. Furthermore, PDSS-ED demonstrates notably superior performance compared to PDSS-EM under the same privacy budget  $\epsilon$ . Specifically, under a privacy budget of  $\epsilon = 3$ , PDSS-EM surpasses the Standalone method by 3.4% and 17% in the CQA and OBQA datasets, respectively, while PDSS-ED outperforms it by 5.2% and 22.4%. Similarly, when the privacy budget is increased to  $\epsilon = 10$ , PDSS-EM exceeds the Standalone approach by 6.3% and 21.6% within the CQA and OBQA datasets, respectively, and PDSS-ED beats it by 7.2% and 28.6%. Remarkably, across all datasets evaluated, when the privacy budget is set to  $\epsilon = 10$ , PDSS achieves comparable performance to DSS, highlighting its efficacy and versatility in balancing privacy and utility.

Method	CQA	OBQA	BoolQ	ARC-E
FewShot-LLM	80.9	82.8	85.2	80.3
FewShot-SLM	25.7	28.6	59.7	40.7
Standalone	55.7	43.4	78.4	50.3
DSS	59.3	55.1	80.5	57.6
PDSS-EM( $\epsilon = 1$ )	57.7	49.2	80.1	52.3
PDSS-EM( $\epsilon = 3$ )	57.6	50.8	79	52.6
PDSS-EM( $\epsilon = 5$ )	58.8	<b>53.2</b>	80	55.3
PDSS-EM( $\epsilon = 10$ )	<b>59.2</b>	52.8	<b>80.2</b>	<b>56.2</b>
PDSS-ED( $\epsilon = 1$ )	58.2	50.8	80.3	56.4
PDSS-ED( $\epsilon = 3$ )	58.6	53.1	80.2	56.5
PDSS-ED( $\epsilon = 5$ )	58.3	53.4	80.4	56.3
PDSS-ED( $\epsilon = 10$ )	<b>59.7</b>	<b>55.8</b>	<b>80.7</b>	<b>57.9</b>

Table 2: We compare the performance of Task-Specific SLM trained with PDSS-EM and PDSS-ED across different privacy budgets  $\epsilon$  against the Task-Specific SLM trained using baseline methods.

### 4.3 Reducing Training Data Evaluation

In this section, we conduct an in-depth analysis to explore the influence of training data size on model performance. We compare the PDSS method with the Standalone approach, varying the amount of

Task	Method	25%	50%	75%	100%
CQA	PDSS-EM	49	53.5	<b>56.7</b>	57.6
	PDSS-ED	54.2	54.6	<b>56.1</b>	58.6
	Standalone	-	-	-	55.7
OBQA	PDSS-EM	34.8	42.2	<b>45.6</b>	50.8
	PDSS-ED	41.4	43.6	<b>50.6</b>	53.1
	Standalone	-	-	-	44.2
BoolQ	PDSS-EM	63	74	<b>78.7</b>	79
	PDSS-ED	72.8	77.6	<b>79.1</b>	80.2
	Standalone	-	-	-	78.4
ARC-E	PDSS-EM	45.3	<b>52.2</b>	53.1	53.8
	PDSS-ED	48	49.7	<b>55.9</b>	56.5
	Standalone	-	-	-	50.3

Table 3: We compare the performance of Task-Specific SLM trained with PDSS-EM( $\epsilon = 3$ ) and PDSS-ED( $\epsilon = 3$ ) against Standalone, across a range of dataset sizes from 25% to 100%. The '-' indicates a method does not apply to the corresponding dataset sizes.

training data used. Table 3 provides a clear illustration of how PDSS(with  $\epsilon = 3$ ) consistently outperforms the Standalone method.

Remarkably, PDSS achieves superior performance even with significantly fewer training samples compared to Standalone. More specifically, when trained on merely 75% of the complete CQA, OBQA, and BoolQ datasets, both PDSS-EM and PDSS-ED surpasses the performance of Standalone fine-tuning that has been trained on the entirety of these datasets. Likewise, by using only 50% of the full ARC-E dataset, PDSS-EM exceeds the results achieved by Standalone fine-tuning on the complete dataset. Furthermore, PDSS-ED exhibits significantly better performance than PDSS-EM across various dataset sizes (ranging from 25% to 100%). The results indicate that PDSS is capable of extracting more valuable information from smaller datasets, making it a promising approach in data-scarce environments.

### 4.4 Perturbed Rationales Evaluation

In this section, we focus on analyzing the quality of the perturbed rationales( $r^p$ ) generated from the

527 perturbed prompt of LLM based on PDSS-EM and  
 528 PDSS-ED methods and compare them with the  
 529 rationales( $r$ ) generated from raw prompt of the  
 530 LLM. To evaluate the similarity between  $r^p$  and  $r$ ,  
 531 we use *TokenRatio* metric. A higher *TokenRatio*  
 532 indicates a greater degree of similarity between the  
 533 perturbed and original rationales. For more details  
 534 about *TokenRatio*, please refer to Appendix C.

535 As shown in Table 4, with an increase in the  
 536 privacy budget  $\epsilon$  and a corresponding decrease in  
 537 perturbation, both the *TokenRatio* of PDSS-EM  
 538 and PDSS-ED have risen notably. Furthermore,  
 539 in most of tasks, the *TokenRatio* of PDSS-ED is  
 540 higher than that of PDSS-EM in the same level of  
 541 privacy budget  $\epsilon$ . The experimental results confirm  
 542 that the *TokenRatio* observed in the perturbed ratio-  
 543 nales produced by both PDSS-EM and PDSS-ED,  
 544 positively correlate with the privacy budget  $\epsilon$ . This  
 545 suggests that as the privacy constraints are relaxed  
 546 (higher  $\epsilon$  values), the perturbed rationales become  
 547 more similar to the original rationales. This find-  
 548 ing is significant as it demonstrates the trade-off  
 549 between privacy protection and the utility of the  
 550 generated rationales.

Method	CQA	OBQA	BoolQ	ARC-E
PDSS-EM( $\epsilon = 1$ )	19.8	26.2	26.6	24.6
PDSS-EM( $\epsilon = 3$ )	29.2	37.2	35.5	33.9
PDSS-EM( $\epsilon = 5$ )	48.8	59.6	55.2	53.9
PDSS-EM( $\epsilon = 10$ )	69.7	72	74.6	68.2
PDSS-ED( $\epsilon = 1$ )	26.7	33.1	29.7	31
PDSS-ED( $\epsilon = 3$ )	33.1	40.9	40.4	42.9
PDSS-ED( $\epsilon = 5$ )	49.6	61	57.5	63.5
PDSS-ED( $\epsilon = 10$ )	57.2	68.3	68	74.2

Table 4: We conduct a comparative analysis to assess the perturbed rationales produced by PDSS-EM and PDSS-ED methods against the original, unperturbed (raw) rationales that are directly generated from the raw prompt of the LLM.

#### 4.5 Decoded Rationales Evaluation

551 In this section, we delve into the quality analysis of  
 552 the decoded rationales produced by the rationales  
 553 decoder module based on PDSS-EM and PDSS-  
 554 ED methods. We compare these decoded rationales  
 555 against those generated directly from raw prompt  
 556 of the LLM. We utilize the *TokenRatio* metric to  
 557 assess their similarities.

558 As shown in Table 5, in contrast to FewShot-  
 559 SLM, it becomes apparent that the decoded ratio-

561 nales’ quality based on PDSS-EM and PDSS-ED  
 562 methods isn’t solely reliant on the locally decoded  
 563 SLM. The perturbed rationales crafted by the LLM  
 564 indeed fulfill their intended purpose. When juxtapo-  
 565 sed with Table 4, it’s clear that at comparable  $\epsilon$   
 566 levels, the *TokenRatio* for the decoded rationales  
 567 surpass those of the perturbed rationales in the  
 568 PDSS-EM and PDSS-ED methods. This under-  
 569 scores the effectiveness of the rationales decoder  
 570 module in the PDSS-EM and PDSS-ED methods.  
 571 Furthermore, with the increase of the privacy bud-  
 572 get  $\epsilon$ , the *TokenRatio* for the decoded rationales  
 573 generated by both PDSS-EM and PDSS-ED have  
 574 increased significantly. This suggests that as the  
 575 privacy constraints are relaxed (higher  $\epsilon$  values),  
 576 the decoded rationales become more similar to the  
 577 original rationales. For more details about compar-  
 578 ative analysis of perturbed rationales and decoded  
 579 rationales, please refer to Appendix D.

Method	CQA	OBQA	BoolQ	ARC-E
FewShot-SLM	43.3	43.4	51.9	42.6
PDSS-EM( $\epsilon = 1$ )	38.3	37.1	38.4	41.5
PDSS-EM( $\epsilon = 3$ )	41.9	41.3	41.7	45.6
PDSS-EM( $\epsilon = 5$ )	53.1	54	55	58.3
PDSS-EM( $\epsilon = 10$ )	71.1	63	73.6	70.4
PDSS-ED( $\epsilon = 1$ )	57.2	53.4	45.2	57.5
PDSS-ED( $\epsilon = 3$ )	59	55.1	48	59.4
PDSS-ED( $\epsilon = 5$ )	59.8	59.5	55.7	65.5
PDSS-ED( $\epsilon = 10$ )	62	62.3	63.4	70.1

Table 5: We conduct a comparative analysis to assess the decoded rationales produced by PDSS-EM and PDSS-ED methods against the original, unperturbed (raw) rationales that are directly generated from the raw prompt of the LLM.

## 5 Conclusions

580 We introduced PDSS, a privacy-preserving frame-  
 581 work for LLM distillation, addressing domain-  
 582 specific knowledge privacy and resource con-  
 583 straints. PDSS employs a server-client architec-  
 584 ture with prompt encoding, rationale generating,  
 585 rationale decoding, and task-specific SLM training,  
 586 bridging the gap between LLM and SLM while  
 587 maintaining data privacy. Experiments on various  
 588 text generation tasks demonstrate PDSS’s ability  
 589 to enhance SLM performance with LLM support  
 590 while prioritizing data privacy.



## 592 Limitations

593 Our current study faces limitations due to compu-  
594 tational and storage constraints, which hinder our  
595 ability to experiment with larger model sizes. Addi-  
596 tionally, our evaluation of PDSS has been restricted  
597 to the Qwen model architecture, leaving the possi-  
598 bility that PDSS may need to be further explored  
599 in other model architectures. We intend to tackle  
600 these issues in future research endeavors.

## 601 References

602 Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang,  
603 Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei  
604 Huang, et al. 2023. Qwen technical report. *arXiv*  
605 *preprint arXiv:2309.16609*.

606 Huimin Chen, Fengran Mo, Yanhao Wang, Cen Chen,  
607 Jian-Yun Nie, Chengyu Wang, and Jamie Cui. 2022.  
608 A customized text sanitization mechanism with dif-  
609 ferential privacy. *arXiv preprint arXiv:2207.01193*.

610 Christopher Clark, Kenton Lee, Ming-Wei Chang,  
611 Tom Kwiatkowski, Michael Collins, and Kristina  
612 Toutanova. 2019. Boolq: Exploring the surprising  
613 difficulty of natural yes/no questions. *arXiv preprint*  
614 *arXiv:1905.10044*.

615 Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot,  
616 Ashish Sabharwal, Carissa Schoenick, and Oyvind  
617 Tafjord. 2018. Think you have solved question an-  
618 swering? try arc, the ai2 reasoning challenge. *arXiv*  
619 *preprint arXiv:1803.05457*.

620 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian,  
621 Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias  
622 Plappert, Jerry Tworek, Jacob Hilton, Reiichiro  
623 Nakano, et al. 2021. Training verifiers to solve math  
624 word problems. *arXiv preprint arXiv:2110.14168*.

625 Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiy-  
626 ong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and  
627 Zhifang Sui. 2022. A survey on in-context learning.  
628 *arXiv preprint arXiv:2301.00234*.

629 Ye Dong, Wen-jie Lu, Yancheng Zheng, Haoqi Wu,  
630 Derun Zhao, Jin Tan, Zhicong Huang, Cheng Hong,  
631 Tao Wei, and Wenguang Cheng. 2023. Puma: Secure  
632 inference of llama-7b in five minutes. *arXiv preprint*  
633 *arXiv:2307.12533*.

634 Cynthia Dwork. 2006. Differential privacy. In *Inter-  
635 national colloquium on automata, languages, and  
636 programming*, pages 1–12. Springer.

637 Tao Fan, Yan Kang, Guoqiang Ma, Weijing Chen, Wen-  
638 bin Wei, Lixin Fan, and Qiang Yang. 2023. Fate-  
639 llm: A industrial grade federated learning frame-  
640 work for large language models. *arXiv preprint*  
641 *arXiv:2310.10049*.

642 Craig Gentry. 2009. *A fully homomorphic encryption*  
643 *scheme*. Stanford university.

Namgyu Ho, Laura Schmid, and Se-Young Yun. 2022. Large language models are reasoning teachers. *arXiv preprint arXiv:2212.10071*. 644  
645  
646

Xiaoyang Hou, Jian Liu, Jingyu Li, Yuhan Li, Wen-jie Lu, Cheng Hong, and Kui Ren. 2023. Ciphergpt: Secure two-party gpt inference. *Cryptology ePrint Archive*. 647  
648  
649  
650

Cheng-Yu Hsieh, Chun-Liang Li, Chih-Kuan Yeh, Hootan Nakhost, Yasuhisa Fujii, Alexander Ratner, Ranjay Krishna, Chen-Yu Lee, and Tomas Pfister. 2023. Distilling step-by-step! outperforming larger language models with less training data and smaller model sizes. *arXiv preprint arXiv:2305.02301*. 651  
652  
653  
654  
655  
656

Yan Kang, Tao Fan, Hanlin Gu, Lixin Fan, and Qiang Yang. 2023. Grounding foundation models through federated transfer learning: A general framework. *arXiv preprint arXiv:2311.17431*. 657  
658  
659  
660

Tassilo Klein and Moin Nabi. 2020. Contrastive self-supervised learning for commonsense reasoning. *arXiv preprint arXiv:2005.00669*. 661  
662  
663

Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213. 664  
665  
666  
667  
668

Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, Joe Davison, Mario Šaško, Gungjan Chhablani, Bhavitvya Malik, Simon Brandeis, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas Patry, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Clément Delangue, Théo Matussière, Lysandre Debut, Stas Bekman, Pierric Cistac, Thibault Goehringer, Victor Mustar, François Lagunas, Alexander Rush, and Thomas Wolf. 2021. *Datasets: A community library for natural language processing*. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 175–184, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics. 669  
670  
671  
672  
673  
674  
675  
676  
677  
678  
679  
680  
681  
682  
683  
684  
685

Liunian Harold Li, Jack Hessel, Youngjae Yu, Xiang Ren, Kai-Wei Chang, and Yejin Choi. 2023. Symbolic chain-of-thought distillation: Small models can also "think" step-by-step. *arXiv preprint arXiv:2306.14050*. 686  
687  
688  
689  
690

Frank McSherry and Kunal Talwar. 2007. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*, pages 94–103. IEEE. 691  
692  
693  
694

Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. 2018. Can a suit of armor conduct electricity? a new dataset for open book question answering. *arXiv preprint arXiv:1809.02789*. 695  
696  
697  
698

OpenAI. 2023. Gpt-4. 699

700	Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2018. Commonsenseqa: A question answering challenge targeting commonsense knowledge. <i>arXiv preprint arXiv:1811.00937</i> .	753
701		754
702		755
703		756
704	Meng Tong, Kejiang Chen, Yuang Qi, Jie Zhang, Weiming Zhang, and Nenghai Yu. 2023. Privinfer: Privacy-preserving inference for black-box large language model. <i>arXiv preprint arXiv:2310.12214</i> .	757
705		758
706		759
707		760
708	Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. <i>arXiv preprint arXiv:2302.13971</i> .	761
709		762
710		763
711		
712		
713		
714	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. <i>Advances in neural information processing systems</i> , 35:24824–24837.	764
715		765
716		766
717		767
718		768
719	Andrew Chi-Chih Yao. 1986. How to generate and exchange secrets. In <i>27th annual symposium on foundations of computer science (Sfcs 1986)</i> , pages 162–167. IEEE.	769
720		770
721		771
722		772
723	Xiang Yue, Minxin Du, Tianhao Wang, Yaliang Li, Huan Sun, and Sherman SM Chow. 2021. Differential privacy for text analytics via natural text sanitization. <i>arXiv preprint arXiv:2106.01221</i> .	773
724		774
725		
726		
727	Yu Zhang and Qiang Yang. 2021. A survey on multi-task learning. <i>IEEE Transactions on Knowledge and Data Engineering</i> , 34(12):5586–5609.	775
728		776
729		777
730	Xin Zhou, Yi Lu, Ruotian Ma, Tao Gui, Yuran Wang, Yong Ding, Yibo Zhang, Qi Zhang, and Xuan-Jing Huang. 2023. Textobfuscator: Making pre-trained language model a privacy protector via obfuscating word representations. In <i>Findings of the Association for Computational Linguistics: ACL 2023</i> , pages 5459–5473.	778
731		779
732		780
733		781
734		782
735		783
736		
737	<b>A Rationales Generation through COT</b>	
738	We utilize the rationales data generated by server-side LLM through chain-of-thought (CoT)(Wei et al., 2022)(Hsieh et al., 2023) technique to enhance the performance of the client’s task-specific SLM. These rationales justify the predicted labels and serve as insightful guidance for training smaller and domain-specific models. Consider the following example: when asked “Question:A beaver is know for building prowess, their supplies come from where? Answer Choices: (a) british columbia (b) body of water (c) wooded area (d) pay debts (e) zoo”. Utilizing the chain-of-thought (CoT) technique, the LLM can generate intermediate rationales like, "The answer must be the place where beavers get their supplies. Of the above choices,	
739		
740		
741		
742		
743		
744		
745		
746		
747		
748		
749		
750		
751		
752		
	only wooded areas have the supplies that beavers need.” Such rationales bridge the gap between the input and the final answer, often encapsulating valuable task-related knowledge. This knowledge would traditionally require extensive data for smaller and task-specific models to acquire. Therefore, we harness these rationales as enriched training material for small language models, employing a multi-task training paradigm that encompasses both label prediction task and rationale prediction task.	
	<b>B More on Experimental Details</b>	
	<b>B.1 Hyperparameter Settings</b>	
	<b>SLM Parameters.</b> During the training process for both the Encoder-Decoder SLM and the Task-Specific SLM, we specifically configured the parameters. We set the batch size to 32 and employed the AdamW optimizer. The maximum number of training steps ranged from 400 to 1500. Additionally, we assigned the values of 0.5 to both $\alpha$ and $\beta$ . Furthermore, the learning rates for $\eta_\phi$ and $\eta_\omega$ were established at $5e-5$ .	
	<b>B.2 Data Splitting</b>	
	For the datasets CQA/OBQA/BoolQ//ARC-E/, all splits (training, validation, and test) were downloaded from HuggingFace (Lhoest et al., 2021). During the training of the Encoder-Decoder SLM, we randomly divided the training data into two equal parts. One part was designated as the public dataset, while the other part was allocated as the private dataset for the client.	
	<b>B.3 Dataset Licenses</b>	
	For the datasets CQA/OBQA/BoolQ//ARC-E/ were downloaded from HuggingFace(Lhoest et al., 2021) and under Apache License, Version 2.0.	
	<b>B.4 Machine Configuration</b>	
	The experiments were conducted on machines equipped with 4 Nvidia V100 32G.	
	<b>C The Definition of TokenRatio Metric</b>	
	<b>TokenRatio(<math>r', r</math>).</b> This metric calculates the unique words( $u$ ) in $r'$ and counts how many of these words are also present in $r$ , denoted as $i$ . The <i>TokenRatio</i> is then calculated as $i$ divided by the total number of unique words in $r'$ ( $  u  $ ).	

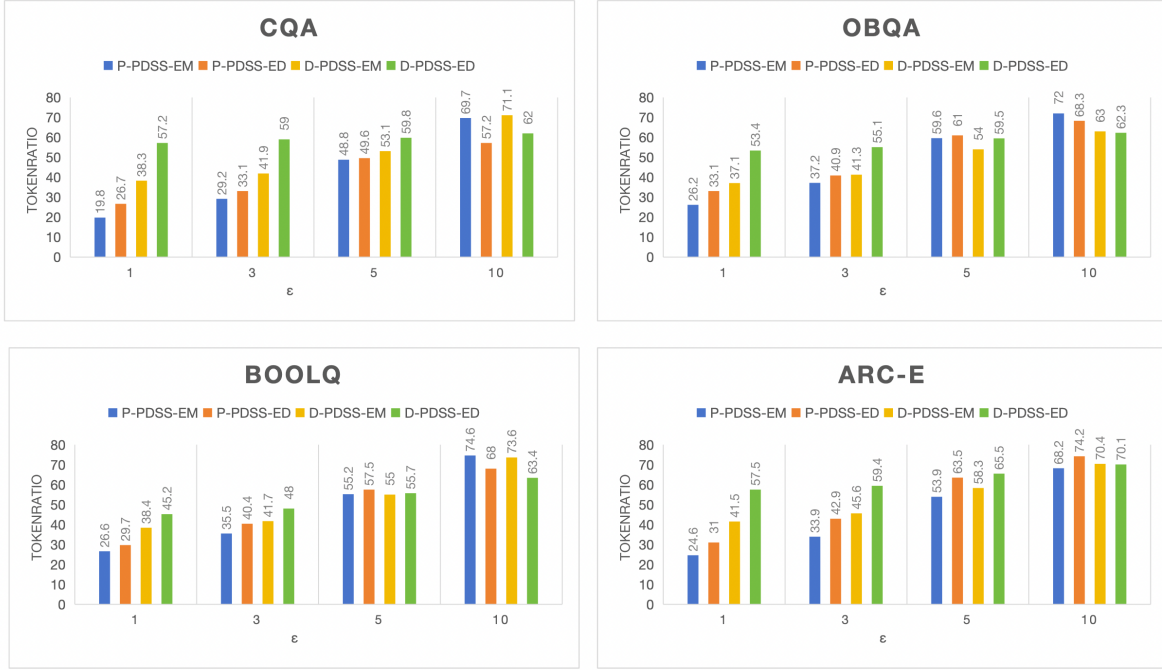


Figure 6: Comparative Analysis of Perturbed Rationales and Decoded Rationales.

## D Comparative Analysis of Perturbed Rationales and Decoded Rationales

As shown in Figure 6, we conduct a comparison of the quality between the perturbed rationales and the decoded rationales, employing both the PDSS-EM and PDSS-ED methods across various privacy budgets denoted by  $\epsilon$ . For clarity, we designate the perturbed rationales generated using the PDSS-EM and PDSS-ED methods as P-PDSS-EM and P-PDSS-ED, respectively. Similarly, the decoded rationales derived from these methods are denoted as D-PDSS-EM and D-PDSS-ED. It's clear that at comparable  $\epsilon$  levels, the *TokenRatio* for decoded rationales consistently surpasses that of perturbed rationales in most tasks, when utilizing the PDSS-EM and PDSS-ED methods. This finding underscores the remarkable effectiveness of the rationales decoder module within both the PDSS-EM and PDSS-ED frameworks.