

---

# Logits are All We Need to Adapt Closed Models

---

Gaurush Hiranandani<sup>1</sup> Haolun Wu<sup>\* 2 3</sup> Subhjoyoti Mukherjee<sup>† 4</sup> Sanmi Koyejo<sup>2</sup>

## Abstract

Many commercial Large Language Models (LLMs) are often closed-source, limiting developers to prompt tuning for aligning content generation with specific applications. While these models currently do not provide access to token logits, we argue that if such access were available, it would enable more powerful adaptation techniques beyond prompt engineering. In this paper, we propose a token-level probability reweighting framework that, given access to logits and a small amount of task-specific data, can effectively steer black-box LLMs toward application-specific content generation. Our approach views next-token prediction through the lens of supervised classification. We show that aligning black-box LLMs with task-specific data can be formulated as a label noise correction problem, leading to *Plugin* model – an autoregressive probability reweighting model that operates solely on logits. We provide theoretical justification for why reweighting logits alone is sufficient for task adaptation. Extensive experiments with multiple datasets, LLMs, and reweighting models demonstrate the effectiveness of our method, advocating for broader access to token logits in closed-source models. We provide our code at [this URL](https://github.com/gaurushhiranandani/plugin).

## 1. Introduction

The rise of Large Language Models (LLMs) has revolutionized generative Artificial Intelligence, yet the most capable models are often closed-source or black-box (Achiam et al., 2023; Bai et al., 2022a). These models generate text based on input prompts but keep their internal weights and training data undisclosed, limiting transparency and customization.

<sup>\*</sup>Work done while visiting Stanford University. <sup>†</sup>Work done while PhD Candidate at UW Madison. <sup>1</sup>Typeface <sup>2</sup>Stanford University <sup>3</sup>Mila - Quebec AI Institute <sup>4</sup>Adobe Research. Correspondence to: Gaurush Hiranandani <gaurush@typeface.ai>, Haolun Wu <haolunwu@stanford.edu>.

*Proceedings of the 42<sup>nd</sup> International Conference on Machine Learning*, Vancouver, Canada. PMLR 267, 2025. Copyright 2025 by the author(s).

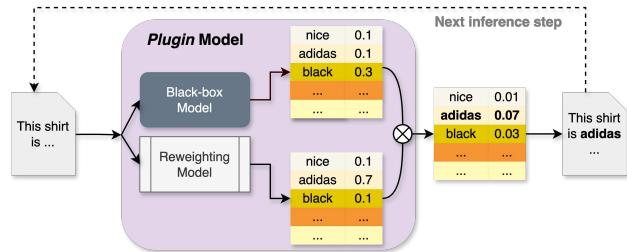


Figure 1. Inference phase of the *Plugin* model. The token probabilities are a product of the probabilities from the black-box model and a reweighting model that denotes label transitioning.

Despite these constraints, closed-source LLMs are widely adopted across applications ranging from travel itinerary generation to tax advice, with developers largely relying on prompt optimization to achieve domain-specific outputs.

However, this reliance on prompt engineering is insufficient for specialized tasks, e.g., those requiring brand-specific tone or style. Consider a content writer aiming to generate product descriptions that reflect a brand’s unique identity. Black-box LLMs, trained on broad datasets, often fail to meet such nuanced requirements. With access limited to generated tokens, developers resort to zero-shot (Kojima et al., 2022) or few-shot (Song et al., 2023) prompting techniques. However, if model weights were accessible, advanced techniques like Parameter-Efficient Fine-Tuning (PEFT) using LoRA (Hu et al., 2021), QLoRA (Dettmers et al., 2024), prefix tuning (Li & Liang, 2021), or adapters (Hu et al., 2023a) could be employed for fine-tuning. Yet, due to intellectual property concerns and the high costs of development, most commercial LLMs remain closed-source, and even with API-based fine-tuning options, concerns over data privacy discourage developers from sharing proprietary data.

In this paper, we propose a middle ground between general-purpose LLM creators and developers seeking application-specific alignment. We argue that providing access to token logits, in addition to generated text, would enable more effective customization for downstream tasks. Viewing next-token prediction as a classification problem, we draw an analogy between LLMs and supervised classification models. Since decoder-only LLMs are trained to predict the next token given preceding tokens, aligning black-box LLMs to domain-specific data can be reframed as a label noise correction problem in supervised classification. In this analogy, the LLM’s broad training data serves as proxy labels, while

application-specific data represents true labels. This can be interpreted as a distribution shift scenario. For example, in “label shift” (Lipton et al., 2018), certain tokens may appear more frequently in application-specific data than in the LLM’s original corpus. In “class-dependent or independent label noise” (Patrini et al., 2017), synonymous expressions or stylistic variations in application data may diverge from those seen during model training.

Inspired by the label noise correction method of Patrini et al. (2017), which estimates a transition matrix to correct class-dependent noise, we adapt this idea to black-box LLM alignment. Unlike prior work that modifies the loss and retrains the model, we lack access to the LLM’s training data and cannot retrain the model. Instead, we estimate an autoregressive transition matrix from application-specific data and use it to reweight token probabilities at inference.

This autoregressive extension is novel, as it accounts for dependencies on previously generated tokens when adjusting logits for the next token. By adapting label noise correction techniques to autoregressive language modeling, we present a practical method to align black-box LLMs using only logits—without requiring access to model weights or original training data.

Our contributions are summarized as follows:

1. We formulate the problem of adapting black-box LLMs for application-specific content generation as a loss correction approach, requiring only token logits at each generation step. This bridges label noise correction in supervised classification with autoregressive language modeling (Sections 2 and 3).
2. We propose an autoregressive probability reweighting framework, enabling token-level probability adjustment during inference. The resulting *Plugin* model dynamically reweights logits to align generation with task-specific data (Section 4).
3. We provide theoretical guarantees, showing that under mild assumptions, the *Plugin* model consistently aligns probability estimates with the target distribution given sufficient application-specific samples. To our knowledge, this is the first work to establish such consistency in an autoregressive label noise setting (Section 5).
4. We conduct extensive experiments across four language generation datasets and three black-box LLMs. Our results, supported by multiple ablations, demonstrate that the *Plugin* model outperforms baselines in adapting black-box LLMs for domain-specific content generation (Section 7). Based on our results, we advocate for publishing token logits alongside outputs in closed-source LLMs.

## 2. Preliminaries

We begin by establishing the notation. The index set is denoted as  $[c] = \{1, \dots, c\}$  for any positive integer  $c$ . Vectors are represented in boldface, for example,  $\mathbf{v}$ , while matrices are denoted using uppercase letters, such as  $V$ . The coordinates of a vector are indicated with subscripts, for instance,  $v_j$ . The all-ones vector is denoted by  $\mathbf{1}$ , with its size being clear from the context. The  $c$ -dimensional simplex is represented as  $\Delta^{c-1} \subset [0, 1]^c$ . Finally, a sequence  $(x_t, x_{t-1}, \dots, x_1)$  of size  $t$  is denoted by  $x_{t:1}$ .

We assume access to language data for the target task, while the black-box LLM, trained on broad world knowledge, is treated as having learned from a noisy version of this data. We seek to adapt the black-box model to align with the task-specific distribution. To formalize this, we extend the label-noise framework from supervised classification (Patrini et al., 2017) to the decoder-only language modeling.

Decoder-only models are trained using a next-token prediction objective. At each step, this setup resembles a supervised classification problem with  $|V|$  classes, where  $V$  is the vocabulary of tokens. Formally, the label space at step  $t$  is  $\mathcal{X}_t = \{\mathbf{e}^i : i \in [|V|]\}$ , where  $\mathbf{e}^i$  denotes the  $i$ -th standard canonical vector in  $\mathbb{R}^{|V|}$ , i.e.,  $\mathbf{e}^i \in \{0, 1\}^{|V|}, \mathbf{1}^T \mathbf{e}^i = 1$ . The task at each step  $t$  is to predict the next token  $\mathbf{x}_t$  (denoted as one-hot vector) given a sequence of tokens  $\mathbf{x}_{t-1:1}$ .

One observes examples  $(\mathbf{x}_t, \mathbf{x}_{t-1:1})$  drawn from an unknown distribution  $p^*(\mathbf{x}_t, \mathbf{x}_{t-1:1}) = p^*(\mathbf{x}_t | \mathbf{x}_{t-1:1})p^*(\mathbf{x}_{t-1:1})$  over  $V \times V^{[t-1]}$ , with expectations denoted by  $E_{\mathbf{x}_t, \mathbf{x}_{t-1:1}}^*$ . Cross-entropy loss is typically used for training over the vocabulary tokens. Assuming access to token logits, and thus the softmax outputs, from the black-box LLM, we interpret the softmax output as a vector approximating the class-conditional probabilities  $p^*(\mathbf{x}_t | \mathbf{x}_{t-1:1})$ , denoted as  $b(\mathbf{x}_t | \mathbf{x}_{t-1:1}) \in \Delta^{|V|-1}$ .

To quantify the discrepancy between the target label  $\mathbf{x}_t = \mathbf{e}^i$  at step  $t$  and the model’s predicted output, we define a loss function  $\ell : |V| \times \Delta^{|V|-1} \rightarrow \mathbb{R}$ . A common choice in next-token prediction tasks is the cross-entropy loss:

$$\begin{aligned} \ell(\mathbf{e}^i, b(\mathbf{x}_t | \mathbf{x}_{t-1:1})) &= -(\mathbf{e}^i)^T \log b(\mathbf{x}_t | \mathbf{x}_{t-1:1}) \\ &= -\log b(\mathbf{x}_t = \mathbf{e}^i | \mathbf{x}_{t-1:1}). \end{aligned} \quad (1)$$

With some abuse of notation, the loss in vector form  $\ell : \Delta^{|V|-1} \rightarrow \mathbb{R}^{|V|}$ , computed on every possible label is  $\ell(b(\mathbf{x}_t | \mathbf{x}_{t-1:1})) = (\ell(\mathbf{e}^1, b(\mathbf{x}_t | \mathbf{x}_{t-1:1})), \dots, \ell(\mathbf{e}^{|V|}, b(\mathbf{x}_t | \mathbf{x}_{t-1:1})))^T$ .

## 3. Loss Robustness

We extend label noise modeling to the autoregressive language setting, focusing on asymmetric or class-conditional

noise. At each step  $t$ , the label  $\mathbf{x}_t$  in the black-box model’s training data is flipped to  $\tilde{\mathbf{x}}_t \in V$  with probability  $p^*(\tilde{\mathbf{x}}_t | \mathbf{x}_t)$ , while preceding tokens  $(\mathbf{x}_{t-1:1})$  remain unchanged. As a result, the black-box model observes samples from a noisy distribution:  $p^*(\tilde{\mathbf{x}}_t, \mathbf{x}_{t-1:1}) = \sum_{\mathbf{x}_t} p^*(\tilde{\mathbf{x}}_t | \mathbf{x}_t) p^*(\mathbf{x}_t | \mathbf{x}_{t-1:1}) p^*(\mathbf{x}_{t-1:1})$ .

We define the noise transition matrix  $T_t \in [0, 1]^{|V| \times |V|}$  at step  $t$ , where each entry  $T_{tij} = p^*(\tilde{\mathbf{x}}_t = \mathbf{e}^j | \mathbf{x}_t = \mathbf{e}^i)$  represents the probability of label flipping. This matrix is row-stochastic but not necessarily symmetric.

To handle asymmetric label noise, we modify the loss  $\ell$  for robustness. Initially, assuming a known  $T_t$ , we apply a loss correction inspired by (Patrini et al., 2017; Sukhbaatar et al., 2015). We then relax this assumption by estimating  $T_t$  directly, forming the basis of our *Plugin* model approach.

We observe that a language model trained with no loss correction would result in a predictor for noisy labels  $b(\tilde{\mathbf{x}}_t | \mathbf{x}_{t-1:1})$ . We can make explicit the dependence on  $T_t$ . For example, with cross-entropy we have:

$$\begin{aligned} \ell(\mathbf{e}^i, b(\tilde{\mathbf{x}}_t | \mathbf{x}_{t-1:1})) &= -\log b(\tilde{\mathbf{x}}_t = \mathbf{e}^i | \mathbf{x}_{t-1:1}) \\ &= -\log \sum_{j=1}^{|V|} p^*(\tilde{\mathbf{x}}_t = \mathbf{e}^i | \mathbf{x}_t = \mathbf{e}^j) b(\mathbf{x}_t = \mathbf{e}^j | \mathbf{x}_{t-1:1}) \\ &= -\log \sum_{j=1}^{|V|} T_{tji} b(\mathbf{x}_t = \mathbf{e}^j | \mathbf{x}_{t-1:1}), \end{aligned} \quad (2)$$

or in matrix form

$$\ell(b(\tilde{\mathbf{x}}_t | \mathbf{x}_{t-1:1})) = -\log T_t^\top b(\mathbf{x}_t | \mathbf{x}_{t-1:1}). \quad (3)$$

This loss compares the noisy label  $\tilde{\mathbf{x}}_t$  to the noisy predictions averaged via the transition matrix  $T_t$  at step  $t$ . Cross-entropy loss, commonly used for next-token prediction, is a *proper composite loss* with the softmax function as its *inverse link function* (Patrini et al., 2017). Consequently, from Theorem 2 of Patrini et al. (2017), the minimizer of the *forwardly-corrected loss* in Equation (3) on noisy data aligns with the minimizer of the true loss on clean data, i.e.,

$$\begin{aligned} \operatorname{argmin}_w E_{\tilde{\mathbf{x}}_t, \mathbf{x}_{t-1:1}}^* [\ell(\mathbf{x}_t, T_t^\top b(\mathbf{x}_t | \mathbf{x}_{t-1:1}))] \\ = \operatorname{argmin}_w E_{\mathbf{x}_t, \mathbf{x}_{t-1:1}}^* [\ell(\mathbf{x}_t, b(\mathbf{x}_t | \mathbf{x}_{t-1:1}))], \end{aligned}$$

where  $w$  are the language model’s weights, implicitly embedded in the softmax output  $b$  from the black-box model. This result suggests that if  $T_t$  were known, we could transform the softmax output  $b(\mathbf{x}_t | \mathbf{x}_{t-1:1})$  using  $T_t^T$ , use the transformed predictions as final outputs, and retrain the model accordingly. However, since  $T_t$  is unknown and training data is inaccessible, estimating  $T_t$  from clean data is essential to our approach.

### 3.1. Estimation of Transition Matrix

We assume access to a small amount of target data for the task. Given that the black-box model is expressive enough to approximate  $p^*(\tilde{\mathbf{x}}_t | \mathbf{x}_{t-1:1})$  (Assumption (2) in Theorem 3 of Patrini et al. (2017)), the transition matrix  $T_t$  can be estimated from this target data. Considering the supervised classification setting at step  $t$ , let  $\mathcal{X}_t^i$  represent all target data samples where  $\mathbf{x}_t = \mathbf{e}^i$  and the preceding tokens are  $(\mathbf{x}_{t-1:1})$ . A naive estimate of the transition matrix is:  $\hat{T}_{tji} = b(\tilde{\mathbf{x}}_t = \mathbf{e}^j | \mathbf{x}_t = \mathbf{e}^i) = \frac{1}{|\mathcal{X}_t^i|} \sum_{x \in \mathcal{X}_t^i} b(\tilde{\mathbf{x}}_t = \mathbf{e}^j | \mathbf{x}_{t-1:1})$ . While this setup works for a single step  $t$ , there are two key challenges in extending it across all steps in the token prediction task:

1. **Limited sample availability:** The number of samples where  $\mathbf{x}_t = \mathbf{e}^i$  and the preceding tokens  $(\mathbf{x}_{t-1}, \dots, \mathbf{x}_1)$  match exactly is limited in the clean data, especially with large vocabulary sizes (e.g.,  $|V| = O(100K)$  for LLaMA (Dubey et al., 2024)). This necessitates modeling the transition matrix as a function of features derived from  $\mathbf{x}_{t-1:1}$ , akin to text-based autoregressive models.
2. **Large parameter space:** With a vocabulary size of  $|V| = O(100K)$ , the transition matrix  $T_t$  has approximately 10 billion parameters. This scale may exceed the size of the closed-source LLM and cannot be effectively learned from limited target data. Thus, structural restrictions must be imposed on  $T_t$  to reduce its complexity.

To address these challenges, we impose the restriction that the transition matrix  $T_t$  is diagonal. While various constraints could be applied to simplify the problem, assuming  $T_t$  is diagonal offers two key advantages. First, it allows the transition matrix—effectively a vector in this case—to be modeled using standard autoregressive language models, such as a *GPT-2 model with  $k$  transformer blocks*, a *LLaMA model with  $d$ -dimensional embeddings*, or a fine-tuned *GPT-2-small* model. These architectures can be adjusted based on the size of the target data. Second, a diagonal transition matrix corresponds to a symmetric or class-independent label noise setup, where  $\mathbf{x}_t = \mathbf{e}^i$  flips to any other class with equal probability in the training data. This assumption, while simplifying, remains realistic within the framework of label noise models.

Enforcing a diagonal structure ensures efficient estimation of the transition matrix while maintaining practical applicability within our framework. Next, we outline our approach for adapting closed-source language models to target data.

## 4. Proposed Method: The Plugin Approach

To estimate the autoregressive transition vector, we train an autoregressive language model on target data, which oper-

ates alongside the black-box model during inference. This model acts as an autoregressive reweighting mechanism, adjusting the token probabilities produced by the black-box model. The combined approach, integrating probabilities from the black-box and reweighting models, is referred to as the *Plugin* model. The term *Plugin* is inspired by classification literature, where plugin methods reweight probabilities to adapt to distribution shifts (Koyejo et al., 2014; Narasimhan et al., 2015; Hiranandani et al., 2021). We now detail the training and inference phases, summarized in Algorithm 1 (Appendix A) and illustrated in Figure 1.

#### 4.1. Training the Plugin Model

During each training iteration, a sequence  $s$  of  $m$  tokens is passed through both the black-box model and the reweighting model to obtain token probabilities  $\{\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_m\}$  and  $\{\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_m\}$ , respectively, where each  $\mathbf{b}_i, \mathbf{r}_i \in \Delta^{|V|-1}$ . The final token probability from the *Plugin* model is computed by normalizing the element-wise product of these probabilities:

$$\mathbf{p}_i = \frac{\mathbf{b}_i \odot \mathbf{r}_i}{\|\mathbf{b}_i \odot \mathbf{r}_i\|_1}. \quad (4)$$

The sequence-level cross-entropy loss is given by:

$$\ell_s = -\frac{1}{m} \sum_{i=1}^m \sum_{j=1}^{|V|} \log(\mathbf{p}_i) \odot \mathbf{e}_j, \quad (5)$$

where the  $j$ -th token appears at the  $i$ -th position in the sequence  $s$ . During backpropagation, only the reweighting model parameters are updated, while the black-box model remains frozen. This formulation extends naturally to batch training, refining  $\mathbf{r}_i$  over iterations to approximate the transition vector governing label shifts in the target data.

#### 4.2. Inference from the Plugin Model

Given a fully trained reweighting model and access to the black-box model, token generation proceeds autoregressively. At the first step, the black-box model produces token probabilities  $\mathbf{b}_1$ , while the reweighting model outputs  $\mathbf{r}_1$ . The *Plugin* model selects the first token as  $\mathbf{x}_1 = \text{argmax}_V(\mathbf{b}_1 \odot \mathbf{r}_1)$ . For subsequent steps, given the previously generated tokens  $\mathbf{x}_{t-1:1}$ , we obtain probabilities  $\mathbf{b}_t$  from the black-box model and  $\mathbf{r}_t$  from the reweighting model. The *Plugin* model then predicts the next token as:  $\mathbf{x}_t = \text{argmax}_V(\mathbf{b}_t \odot \mathbf{r}_t)$ .

The process continues until a stopping criterion is met. Note that, this manuscript focuses on greedy decoding for inference. Other decoding strategies, such as temperature scaling, top- $p$  sampling, or beam search, can be incorporated by normalizing the element-wise product of probabilities and using them as the final token distribution, as in Equation (4).

## 5. Theoretical Analysis

We establish the convergence properties of *Plugin*, showing that after  $t$  tokens, it accurately estimates the autoregressive noise transition matrix. Modeling the matrix as a function of an unknown parameter  $\theta_*$ , we prove that optimizing the autoregressive loss over token sequences enables consistent estimation of  $\theta_*$  with high probability. To our knowledge, this is the first finite-time convergence analysis for transition matrix estimation under autoregressive noisy loss.

Let  $\mathcal{F}^{t-1}$  denote the history of selected tokens up to time  $t-1$ . Let an unknown parameter  $\theta_* \in \Theta \subseteq \mathbb{R}^d$  governs the transition dynamics of label flipping between token pairs. The transition matrix at time  $t$ , denoted as  $T_t(\theta_* | \mathcal{F}^{t-1})$ , depends on  $\theta_*$  and all previously observed tokens. Before proving our main result, we first make a few assumptions.

**Assumption 5.1.** Let  $T_t(\theta_*; x_i, x_j, \mathcal{F}^{t-1})$  denote the  $(i, j)$ -th component of the transition matrix, and let  $f_{I_t}(\theta_*; x_i, x_j, \mathcal{F}^{t-1})$  be the transition function that determines the transition from  $x_i$  to  $x_j$ , where  $I_t$  is the  $x_i$  token selected at time  $t$ . Let  $x_i, x_j \in \mathbb{R}^d$ . We assume that  $\nabla f_{I_t}(\theta_*; x_i, x_j, \mathcal{F}^{t-1}) < \lambda_0$  and  $\nabla^2 f_{I_t}(\theta_*; x_i, x_j, \mathcal{F}^{t-1}) < \lambda_1$  for some constant  $\lambda_0 > 0$ ,  $\lambda_1 > 0$  and for all steps  $t$ .

Assumption 5.1 assumes the transition matrix depends on the history-dependent function  $f_{I_t}(\cdot)$  with bounded gradient and Hessian, similar to assumptions in (Singh et al., 2023; Zhang et al., 2024) for other deep models.

**Assumption 5.2.** We assume the cross-entropy loss (5) is clipped by  $\epsilon > 0$  and upper bounded as  $\ell_t^{\text{clipped}} \leq C|V|^2(Y_t - f_{I_t}(\theta_*; x_i, x_j, \mathcal{F}^{t-1}))^2$  for any time  $t$ , where  $Y_t$  is the predicted token class,  $f_{I_t}$  determines the true class and satisfies Assumption 5.1, and  $C > 0$  is a constant.

Assumption 5.2 ensures that the clipped log loss is upper bounded by a smoother squared loss. For the remaining of this section we refer to this squared loss at time  $t$  as  $\ell_t(\theta)$ . Let the *Plugin* model minimize the loss  $\ell_1(\theta), \ell_2(\theta), \dots, \ell_t(\theta)$  over  $t$  iterations. Let  $\hat{\theta}_t = \text{argmin}_{\theta \in \Theta} \sum_{s=1}^t \ell_s(\theta)$ . At every iteration  $t$ , the *Plugin* algorithm looks into the history  $\mathcal{F}^{t-1}$  and samples a token  $\mathbf{x}_t \sim \mathbf{p}_{\hat{\theta}_t} = \mathbf{b}_t \odot \mathbf{r}_{\hat{\theta}_t}$ .

Let  $\hat{\mathcal{L}}_t(\theta) = \frac{1}{t} \sum_{s=1}^t \ell_s(\theta)$  and its expectation  $\mathcal{L}_t(\theta) = \frac{1}{t} \sum_{s=1}^t \mathbb{E}_{x_s \sim \mathbf{p}_{\hat{\theta}_{s-1}}} [\ell_s(\theta) | \mathcal{F}^{s-1}]$ . We impose regularity and smoothness assumptions on the loss function  $\ell_t(\theta)$  as stated in Assumption B.1 (Appendix B). We are now ready to prove the main theoretical result of the paper.

**Theorem 1.** Suppose  $\ell_1(\theta), \dots, \ell_t(\theta) : \mathbb{R}^{|V|} \rightarrow \mathbb{R}$  are loss functions from a distribution that satisfies Assumptions 5.1, 5.2, and B.1. Define  $\mathcal{L}_t(\theta) = \frac{1}{t} \sum_{s=1}^t \mathbb{E}_{x_s \sim \mathbf{p}_{\hat{\theta}_{s-1}}} [\ell_s(\theta) | \mathcal{F}^{s-1}]$  where,

$\hat{\theta}_t = \operatorname{argmin}_{\theta \in \Theta} \sum_{s=1}^t \ell_s(\theta)$ . If  $t$  is large enough such that  $\frac{\gamma \log(dt)}{t} \leq c' \min \left\{ \frac{1}{C_1 C_2 |V|^4}, \frac{\max_{\theta \in \Theta} (\mathcal{L}_t(\theta) - \mathcal{L}_t(\theta_*))}{C_2} \right\}$  then for a constant  $\gamma \geq 2$ , universal constants  $C_1, C_2, c'$ , we have that

$$(1 - \rho_t) \frac{\sigma_t^2}{t} - \frac{C_1^2}{t^{\gamma/2}} \leq \mathbb{E} \left[ \mathcal{L}_t(\hat{\theta}_t) - \mathcal{L}_t(\theta_*) \right] \leq (1 + \rho_t) \frac{\sigma_t^2}{t} + \frac{\max_{\theta \in \Theta} (\mathcal{L}_t(\theta) - \mathcal{L}_t(\theta_*))}{t^\gamma}$$

where  $\sigma_t^2 := \mathbb{E} \left[ \frac{1}{2} \left\| \nabla \hat{\mathcal{L}}_t(\theta_*) \right\|_{(\nabla^2 \mathcal{L}_t(\theta_*))^{-1}}^2 \right]$ , and  $\rho_t := (C_1 C_2 + 2\eta^2 \lambda_1^2) \sqrt{\frac{\gamma \log(dt)}{t}}$ .

Theorem 1 bounds the difference between the estimated and true average loss functions, showing that this gap diminishes as the number of training tokens increases. Since  $\hat{\theta}_t = \operatorname{argmin}_{\theta \in \Theta} \sum_{s=1}^t \ell_s(\theta)$ , the *Plugin* model progressively refines its estimation of the unknown parameter  $\theta_*$ . As the transition matrix  $T_t(\theta_*; x_i, x_j, \mathcal{F}^{t-1})$  is derived from  $f_{I_t}(\theta_*; x_i, x_j, \mathcal{F}^{t-1})$ , which depends on  $\theta_*$ , training on sufficiently many tokens ensures an accurate estimation of each component of  $T_t(\theta_* | \mathcal{F}^{t-1})$ .

Our proof reformulates the problem as a sequential hypothesis testing setting to estimate the average loss function  $\mathcal{L}_t(\hat{\theta}_t)$  using the sequence of losses  $\ell_1(\theta), \dots, \ell_t(\theta)$  (Naghshvar & Javidi, 2013; Lattimore & Szepesvári, 2020). Unlike prior work (Frostig et al., 2015; Chaudhuri et al., 2015), which assumes i.i.d. losses, the loss at time  $t$  in our setting depends on all previous losses. Additionally, Mukherjee et al. (2022) study a different active regression setting without considering cross-entropy loss or transition noise matrices as in Patrini et al. (2017). We provide a brief overview of the proof technique in Remark B.9 (Appendix B), highlighting key novelties.

## 6. Related Work

**Parameter-Efficient Fine-Tuning (PEFT).** PEFT methods adapt LLMs to downstream tasks while minimizing computational overhead. LoRA (Hu et al., 2021) and QLoRA (Dettmers et al., 2024) introduce low-rank updates and quantization for efficient fine-tuning, while prefix tuning (Li & Liang, 2021), adapters (Hu et al., 2023b), and soft prompting (Lester et al., 2021) modify task-specific representations through trainable layers or embeddings. Torroba-Hennigen et al. (2025) further explore the equivalence between gradient-based transformations and adapter-based tuning. However, these methods require access to model weights, gradients, or architecture details, making them unsuitable for closed-source LLMs and inapplicable as baselines in our setup. In contrast, our approach operates solely

on token logits, enabling adaptation without modifying the underlying model. Thus, we emphasize that the *Plugin* model is not an alternative to fine-tuning, but rather an approach that uniquely stands for adapting black-box LLMs which only provide logit access.

**Steering and Aligning LLMs.** LLM alignment methods primarily use reinforcement learning or instruction tuning. RLHF and DPO (Christiano et al., 2017; Ouyang et al., 2022; Rafailov et al., 2024) optimize model behavior via human preferences, with DPO eliminating reward modeling. Constitutional AI (Bai et al., 2022b) aligns models using self-generated principles, while instruction tuning (Wei et al., 2021; Sanh et al., 2022) adapts them via task-specific demonstrations. Unlike our approach, these methods require model weights and training data, limiting their applicability as baselines in our setup.

**Calibration of LLMs.** LLM calibration methods aim to align model confidence with predictive accuracy and adjust confidence scores but do not alter token predictions (Ulmer et al., 2024; Shen et al., 2024; Huang et al., 2024; Kapoor et al., 2024; Zhu et al., 2023; Zhang et al., 2023). In contrast, our method reweights token probabilities at inference, enabling adaptation of black-box LLMs without modifying the model or requiring fine-tuning.

**Black-box LLMs.** Prior work explores various approaches for adapting black-box LLMs without fine-tuning, though they differ fundamentally from our method. (Gao et al., 2024) infer user preferences through interactive edits but do not adapt models based on past language data. Diffusion-LM (Li et al., 2022) formulates text generation as a non-autoregressive denoising process, whereas our approach reweights token probabilities autoregressively without requiring black-box model weights. Discriminator-based methods (Dathathri et al., 2020; Mireshghallah et al., 2022; Yang & Klein, 2021; Krause et al., 2021) control generation based on predefined attributes, contrasting with our method, which enables free-form text adaptation. DExperts (Liu et al., 2021; 2024) combines expert and anti-expert probabilities; we incorporate a similar probability combining strategy in a modified baseline without a de-expert component. In-context learning (Long et al., 2023; Dong et al., 2024) offers a common adaptation technique for black-box models and serves as a baseline in our setup.

## 7. Experiments

We divide this section into four parts. Section 7.1 evaluates *Plugin* on four text generation datasets across three black-box language models. Since the *Plugin* model is trained on top of black-box models, we refer to black-box models interchangeably as *base models*. Section 7.2 discusses how *Plugin* can be applied on top of any prompt-tuning method as a wrapper, when logits are accessible. Section 7.3 presents ablation studies analyzing the impact of black-box

Table 1. Performance comparison on E2E NLG dataset. We show mean and standard deviation of the metrics over five seeds.

Model	Method	BLEU	Rouge-1	Rouge-2	Rouge-L	METEOR	CIDEr	NIST
GPT2-M	Zeroshot	0.0247	0.3539	0.1003	0.2250	0.3015	0.0156	0.6133
GPT2-M	ICL-1	0.0543 $\pm$ 0.026	0.3431 $\pm$ 0.048	0.1299 $\pm$ 0.033	0.2280 $\pm$ 0.047	0.3434 $\pm$ 0.051	0.0260 $\pm$ 0.042	0.7767 $\pm$ 0.060
GPT2-M	ICL-3	0.0750 $\pm$ 0.035	0.3955 $\pm$ 0.028	0.1676 $\pm$ 0.020	0.2649 $\pm$ 0.052	0.3977 $\pm$ 0.063	0.0252 $\pm$ 0.049	0.8993 $\pm$ 0.076
GPT2-M	NewModel	0.2377 $\pm$ 0.011	0.5049 $\pm$ 0.014	0.2742 $\pm$ 0.013	0.3902 $\pm$ 0.006	0.4521 $\pm$ 0.016	0.3938 $\pm$ 0.019	1.1927 $\pm$ 0.069
GPT2-M	WeightedComb	0.1709 $\pm$ 0.008	0.4817 $\pm$ 0.020	0.2447 $\pm$ 0.011	0.3720 $\pm$ 0.014	0.4071 $\pm$ 0.025	0.3329 $\pm$ 0.027	1.0864 $\pm$ 0.002
GPT2-M	TempNet	0.1036 $\pm$ 0.010	0.3425 $\pm$ 0.016	0.1526 $\pm$ 0.012	0.2735 $\pm$ 0.010	0.2615 $\pm$ 0.016	0.4116 $\pm$ 0.023	0.2826 $\pm$ 0.057
GPT2-M	<b>Plugin (Ours)</b>	0.1863 $\pm$ 0.010	<b>0.5227</b> $\pm$ 0.011	0.2612 $\pm$ 0.013	0.3728 $\pm$ 0.003	<b>0.4857</b> $\pm$ 0.012	0.3544 $\pm$ 0.013	1.1241 $\pm$ 0.009
GPT2-XL	Zeroshot	0.0562	0.4013	0.1636	0.2862	0.3697	0.0187	0.5338
GPT2-XL	ICL-1	0.0686 $\pm$ 0.032	0.4016 $\pm$ 0.042	0.1404 $\pm$ 0.052	0.2745 $\pm$ 0.025	0.3503 $\pm$ 0.019	0.0353 $\pm$ 0.015	0.7944 $\pm$ 0.067
GPT2-XL	ICL-3	0.0980 $\pm$ 0.035	0.4188 $\pm$ 0.040	0.1923 $\pm$ 0.046	0.2912 $\pm$ 0.031	0.3925 $\pm$ 0.027	0.0250 $\pm$ 0.017	0.9390 $\pm$ 0.054
GPT2-XL	NewModel	0.2377 $\pm$ 0.011	0.5049 $\pm$ 0.014	0.2742 $\pm$ 0.013	0.3902 $\pm$ 0.006	0.4521 $\pm$ 0.016	0.3938 $\pm$ 0.019	1.1927 $\pm$ 0.069
GPT2-XL	WeightedComb	0.1184 $\pm$ 0.010	0.4237 $\pm$ 0.016	0.1858 $\pm$ 0.012	0.3004 $\pm$ 0.010	0.3776 $\pm$ 0.016	0.1818 $\pm$ 0.023	1.0261 $\pm$ 0.057
GPT2-XL	TempNet	0.1325 $\pm$ 0.013	0.4642 $\pm$ 0.017	0.2516 $\pm$ 0.016	0.3021 $\pm$ 0.022	0.4126 $\pm$ 0.025	0.3627 $\pm$ 0.033	0.8027 $\pm$ 0.047
GPT2-XL	<b>Plugin (Ours)</b>	<b>0.2470</b> $\pm$ 0.009	<b>0.5536</b> $\pm$ 0.007	<b>0.3084</b> $\pm$ 0.007	<b>0.4213</b> $\pm$ 0.008	<b>0.5057</b> $\pm$ 0.009	<b>0.5455</b> $\pm$ 0.013	<b>1.2736</b> $\pm$ 0.051
LLaMA-3.1-8B	Zeroshot	0.3226	0.6917	0.4050	0.5004	0.6041	0.9764	1.1310
LLaMA-3.1-8B	ICL-1	0.3301 $\pm$ 0.037	0.6914 $\pm$ 0.027	0.4126 $\pm$ 0.026	0.5023 $\pm$ 0.018	0.6037 $\pm$ 0.015	0.9715 $\pm$ 0.057	1.1735 $\pm$ 0.066
LLaMA-3.1-8B	ICL-3	0.3527 $\pm$ 0.033	0.6936 $\pm$ 0.036	0.4217 $\pm$ 0.017	0.5127 $\pm$ 0.017	0.6202 $\pm$ 0.009	0.9927 $\pm$ 0.018	1.1672 $\pm$ 0.047
LLaMA-3.1-8B	NewModel	0.2452 $\pm$ 0.008	0.5347 $\pm$ 0.005	0.2905 $\pm$ 0.006	0.4097 $\pm$ 0.005	0.4812 $\pm$ 0.009	0.4571 $\pm$ 0.021	<b>1.2281</b> $\pm$ 0.041
LLaMA-3.1-8B	WeightedComb	0.3517 $\pm$ 0.004	<b>0.7040</b> $\pm$ 0.004	<b>0.4249</b> $\pm$ 0.004	<b>0.5181</b> $\pm$ 0.003	<b>0.6206</b> $\pm$ 0.002	<b>1.0947</b> $\pm$ 0.010	1.1737 $\pm$ 0.015
LLaMA-3.1-8B	TempNet	0.3502 $\pm$ 0.023	0.6927 $\pm$ 0.006	0.4216 $\pm$ 0.023	0.5027 $\pm$ 0.017	0.6124 $\pm$ 0.019	0.9625 $\pm$ 0.025	1.1713 $\pm$ 0.027
LLaMA-3.1-8B	<b>Plugin (Ours)</b>	<b>0.3691</b> $\pm$ 0.013	<b>0.7113</b> $\pm$ 0.002	<b>0.4374</b> $\pm$ 0.004	<b>0.5247</b> $\pm$ 0.002	<b>0.6392</b> $\pm$ 0.009	<b>1.1441</b> $\pm$ 0.030	1.1749 $\pm$ 0.034

Table 2. Performance comparison on Web NLG dataset. We show mean and standard deviation of the metrics over five seeds.

Model	Method	BLEU	Rouge-1	Rouge-2	Rouge-L	METEOR	CIDEr	NIST
GPT2-M	Zeroshot	0.0213	0.2765	0.1014	0.1872	0.2111	0.0479	0.2340
GPT2-M	ICL-1	0.0317 $\pm$ 0.013	0.3388 $\pm$ 0.021	0.1318 $\pm$ 0.013	0.2346 $\pm$ 0.019	0.2876 $\pm$ 0.042	0.0732 $\pm$ 0.053	0.2715 $\pm$ 0.042
GPT2-M	ICL-3	0.0461 $\pm$ 0.014	0.3388 $\pm$ 0.018	0.1378 $\pm$ 0.016	0.2291 $\pm$ 0.010	0.3408 $\pm$ 0.027	0.0748 $\pm$ 0.031	<b>0.3283</b> $\pm$ 0.037
GPT2-M	NewModel	0.1071 $\pm$ 0.005	0.3260 $\pm$ 0.010	0.1496 $\pm$ 0.014	0.2724 $\pm$ 0.013	0.2642 $\pm$ 0.008	0.4327 $\pm$ 0.023	0.2916 $\pm$ 0.031
GPT2-M	WeightedComb	0.0692 $\pm$ 0.007	<b>0.3593</b> $\pm$ 0.010	<b>0.1568</b> $\pm$ 0.008	<b>0.2834</b> $\pm$ 0.015	0.2379 $\pm$ 0.030	0.1916 $\pm$ 0.028	0.2996 $\pm$ 0.037
GPT2-M	TempNet	0.1045 $\pm$ 0.012	0.3526 $\pm$ 0.014	0.1526 $\pm$ 0.014	0.2731 $\pm$ 0.018	0.3326 $\pm$ 0.026	0.4237 $\pm$ 0.033	0.3002 $\pm$ 0.048
GPT2-M	<b>Plugin (Ours)</b>	<b>0.1280</b> $\pm$ 0.007	<b>0.4590</b> $\pm$ 0.005	<b>0.2226</b> $\pm$ 0.005	<b>0.3515</b> $\pm$ 0.006	<b>0.3832</b> $\pm$ 0.010	<b>0.7280</b> $\pm$ 0.039	<b>0.3060</b> $\pm$ 0.017
GPT2-XL	Zeroshot	0.0317	0.2992	0.1321	0.2417	0.1969	0.0491	0.1826
GPT2-XL	ICL-1	0.0510 $\pm$ 0.024	0.3223 $\pm$ 0.026	0.1526 $\pm$ 0.016	0.2562 $\pm$ 0.031	0.2591 $\pm$ 0.009	0.1336 $\pm$ 0.029	0.2235 $\pm$ 0.033
GPT2-XL	ICL-3	0.0744 $\pm$ 0.016	0.3383 $\pm$ 0.036	<b>0.1682</b> $\pm$ 0.016	0.2651 $\pm$ 0.028	<b>0.3071</b> $\pm$ 0.014	0.1675 $\pm$ 0.024	0.2550 $\pm$ 0.021
GPT2-XL	NewModel	0.1071 $\pm$ 0.005	0.3260 $\pm$ 0.010	0.1496 $\pm$ 0.014	0.2724 $\pm$ 0.013	0.2642 $\pm$ 0.008	0.4327 $\pm$ 0.023	0.2916 $\pm$ 0.031
GPT2-XL	WeightedComb	0.0636 $\pm$ 0.006	<b>0.3453</b> $\pm$ 0.007	0.1666 $\pm$ 0.003	<b>0.2782</b> $\pm$ 0.005	0.2871 $\pm$ 0.006	0.2460 $\pm$ 0.005	<b>0.2981</b> $\pm$ 0.018
GPT2-XL	TempNet	0.0925 $\pm$ 0.008	0.3357 $\pm$ 0.009	0.1663 $\pm$ 0.014	0.2764 $\pm$ 0.011	0.3025 $\pm$ 0.009	0.4226 $\pm$ 0.013	0.2837 $\pm$ 0.027
GPT2-XL	<b>Plugin (Ours)</b>	<b>0.1673</b> $\pm$ 0.004	<b>0.4616</b> $\pm$ 0.007	<b>0.2527</b> $\pm$ 0.007	<b>0.3757</b> $\pm$ 0.008	<b>0.3895</b> $\pm$ 0.007	<b>0.8987</b> $\pm$ 0.013	0.2646 $\pm$ 0.003
LLaMA-3.1-8B	Zeroshot	0.1453	0.5278	0.3030	0.3982	0.4314	0.6991	<b>0.2684</b>
LLaMA-3.1-8B	ICL-1	<b>0.2166</b> $\pm$ 0.031	0.5944 $\pm$ 0.027	0.3706 $\pm$ 0.025	<b>0.4667</b> $\pm$ 0.013	0.5651 $\pm$ 0.045	<b>1.5719</b> $\pm$ 0.024	0.2462 $\pm$ 0.038
LLaMA-3.1-8B	ICL-3	0.2031 $\pm$ 0.027	0.5937 $\pm$ 0.019	<b>0.3821</b> $\pm$ 0.015	0.4653 $\pm$ 0.024	<b>0.5682</b> $\pm$ 0.046	1.3826 $\pm$ 0.051	0.2469 $\pm$ 0.045
LLaMA-3.1-8B	NewModel	0.1284 $\pm$ 0.005	0.3506 $\pm$ 0.009	0.1673 $\pm$ 0.007	0.2879 $\pm$ 0.009	0.2921 $\pm$ 0.008	0.4999 $\pm$ 0.030	<b>0.2973</b> $\pm$ 0.008
LLaMA-3.1-8B	WeightedComb	0.1922 $\pm$ 0.012	<b>0.5986</b> $\pm$ 0.019	0.3612 $\pm$ 0.012	0.4659 $\pm$ 0.008	0.4470 $\pm$ 0.030	1.1855 $\pm$ 0.075	0.2575 $\pm$ 0.020
LLaMA-3.1-8B	TempNet	0.2315 $\pm$ 0.010	0.5916 $\pm$ 0.015	0.3794 $\pm$ 0.012	0.4620 $\pm$ 0.010	0.5581 $\pm$ 0.036	1.4826 $\pm$ 0.043	0.2513 $\pm$ 0.020
LLaMA-3.1-8B	<b>Plugin (Ours)</b>	<b>0.2542</b> $\pm$ 0.004	<b>0.6375</b> $\pm$ 0.005	<b>0.3873</b> $\pm$ 0.005	<b>0.4869</b> $\pm$ 0.007	<b>0.5724</b> $\pm$ 0.004	<b>1.5911</b> $\pm$ 0.046	0.2590 $\pm$ 0.003

model quality, *Plugin*'s complexity, and architecture choices. Section 7.4 shows qualitative analysis and case studies.

We evaluate *Plugin* on four text generation benchmarks: (a) E2E NLG (Dušek et al., 2020), (b) Web NLG (Gardent et al., 2017), (c) CommonGen (Lin et al., 2020), and (d) the Adidas product description dataset (adi, 2023). For the first three datasets, we use the train-validation-test splits from the Transformers library (Wolf, 2020). To introduce distribution shifts, we filter Web NLG's training data to include only *infrastructure* descriptions, while validation and test sets retain *person* descriptions. Similarly, CommonGen's training set is restricted to samples having *man*, while validation and test sets remain unchanged. Details of this setup are in Section 7.4. The Adidas dataset is split into validation and test sets. Data statistics are provided in Table 6, Appendix C.1.

## 7.1. Text Generation Performance Comparison

We evaluate *Plugin* on the text generation task using only the validation and test splits of all four datasets, reserving the train split for ablation studies (Section 7.3). *Plugin* and baseline models are trained on the small validation set, with performance measured on the test set. Additionally, we allocate 40% of the validation data as *hyper-validation* for cross-validation of hyperparameters.

Performance is reported using seven standard natural language generation metrics: (a) BLEU (Papineni et al., 2002), (b) ROUGE-1 (Lin, 2004), (c) ROUGE-2 (Lin, 2004), (d) ROUGE-L (Lin & Och, 2004), (e) METEOR (Banerjee & Lavie, 2005), (f) CIDEr (Vedantam et al., 2015), and (g) NIST (Doddington, 2002). All experiments are repeated

Table 3. Performance comparison on CommonGen dataset. We show mean and standard deviation of the metrics over five seeds.

Model	Method	BLEU	Rouge-1	Rouge-2	Rouge-L	METEOR	CIDEr	NIST
GPT2-M	Zeroshot	0.0153	0.2216	0.0409	0.1527	0.2848	0.0001	0.3686
GPT2-M	ICL-1	0.0157 $\pm$ 0.013	0.2580 $\pm$ 0.024	0.0362 $\pm$ 0.096	0.1388 $\pm$ 0.102	0.2871 $\pm$ 0.107	0.0222 $\pm$ 0.076	0.3704 $\pm$ 0.101
GPT2-M	ICL-3	0.0552 $\pm$ 0.010	0.3610 $\pm$ 0.019	0.1248 $\pm$ 0.045	0.2680 $\pm$ 0.089	0.4079 $\pm$ 0.133	0.1366 $\pm$ 0.125	0.5340 $\pm$ 0.087
GPT2-M	NewModel	0.1260 $\pm$ 0.007	0.4106 $\pm$ 0.016	0.1683 $\pm$ 0.013	0.3740 $\pm$ 0.009	0.3600 $\pm$ 0.024	0.4570 $\pm$ 0.058	0.7113 $\pm$ 0.025
GPT2-M	WeightedComb	0.0567 $\pm$ 0.005	0.3918 $\pm$ 0.010	0.1353 $\pm$ 0.005	0.3280 $\pm$ 0.010	0.2929 $\pm$ 0.016	0.2623 $\pm$ 0.042	0.4353 $\pm$ 0.028
GPT2-M	TempNet	0.1248 $\pm$ 0.015	0.4048 $\pm$ 0.014	0.1528 $\pm$ 0.015	0.3526 $\pm$ 0.014	0.3883 $\pm$ 0.017	0.4492 $\pm$ 0.023	0.4037 $\pm$ 0.058
GPT2-M	Plugin (Ours)	0.1366 $\pm$ 0.003	0.4533 $\pm$ 0.007	0.1878 $\pm$ 0.003	0.3934 $\pm$ 0.006	0.4095 $\pm$ 0.011	0.5572 $\pm$ 0.022	0.6395 $\pm$ 0.061
GPT2-XL	Zeroshot	0.0317	0.2992	0.1321	0.2417	0.1969	0.0491	0.1826
GPT2-XL	ICL-1	0.0508 $\pm$ 0.023	0.3201 $\pm$ 0.035	0.1526 $\pm$ 0.097	0.2562 $\pm$ 0.103	0.2591 $\pm$ 0.089	0.1336 $\pm$ 0.092	0.2235 $\pm$ 0.069
GPT2-XL	ICL-3	0.0744 $\pm$ 0.011	0.3383 $\pm$ 0.014	0.1682 $\pm$ 0.030	0.2651 $\pm$ 0.072	0.3071 $\pm$ 0.073	0.1675 $\pm$ 0.066	0.2550 $\pm$ 0.047
GPT2-XL	NewModel	0.1260 $\pm$ 0.007	0.4106 $\pm$ 0.016	0.1683 $\pm$ 0.013	0.3740 $\pm$ 0.009	0.3600 $\pm$ 0.024	0.4570 $\pm$ 0.058	0.7113 $\pm$ 0.025
GPT2-XL	WeightedComb	0.0614 $\pm$ 0.020	0.3364 $\pm$ 0.024	0.1347 $\pm$ 0.009	0.2969 $\pm$ 0.019	0.2921 $\pm$ 0.018	0.2763 $\pm$ 0.010	0.3352 $\pm$ 0.051
GPT2-XL	TempNet	0.1154 $\pm$ 0.020	0.3937 $\pm$ 0.026	0.1482 $\pm$ 0.017	0.3625 $\pm$ 0.013	0.3389 $\pm$ 0.019	0.4376 $\pm$ 0.018	0.5927 $\pm$ 0.047
GPT2-XL	Plugin (Ours)	0.1791 $\pm$ 0.014	0.4932 $\pm$ 0.007	0.2288 $\pm$ 0.004	0.4347 $\pm$ 0.007	0.4702 $\pm$ 0.006	0.7283 $\pm$ 0.012	0.6554 $\pm$ 0.038
LLaMA-3.1-8B	Zeroshot	0.0643	0.2776	0.1181	0.2488	0.3857	0.3155	0.3347
LLaMA-3.1-8B	ICL-1	0.0615 $\pm$ 0.027	0.2697 $\pm$ 0.033	0.1158 $\pm$ 0.062	0.2469 $\pm$ 0.087	0.3822 $\pm$ 0.069	0.3005 $\pm$ 0.072	0.3059 $\pm$ 0.094
LLaMA-3.1-8B	ICL-3	0.0635 $\pm$ 0.016	0.2748 $\pm$ 0.024	0.1225 $\pm$ 0.018	0.3120 $\pm$ 0.047	0.4012 $\pm$ 0.029	0.3250 $\pm$ 0.022	0.3794 $\pm$ 0.034
LLaMA-3.1-8B	NewModel	0.0753 $\pm$ 0.004	0.3716 $\pm$ 0.005	0.1122 $\pm$ 0.003	0.3404 $\pm$ 0.004	0.2665 $\pm$ 0.006	0.1919 $\pm$ 0.015	0.6900 $\pm$ 0.046
LLaMA-3.1-8B	WeightedComb	0.1789 $\pm$ 0.005	0.3485 $\pm$ 0.012	0.1797 $\pm$ 0.008	0.2981 $\pm$ 0.012	0.3637 $\pm$ 0.011	0.5503 $\pm$ 0.046	0.5450 $\pm$ 0.020
LLaMA-3.1-8B	TempNet	0.1524 $\pm$ 0.008	0.3372 $\pm$ 0.015	0.1524 $\pm$ 0.010	0.3298 $\pm$ 0.017	0.3676 $\pm$ 0.015	0.3986 $\pm$ 0.033	0.5286 $\pm$ 0.023
LLaMA-3.1-8B	Plugin (Ours)	0.2665 $\pm$ 0.010	0.5800 $\pm$ 0.002	0.3139 $\pm$ 0.005	0.5037 $\pm$ 0.004	0.5829 $\pm$ 0.003	1.0876 $\pm$ 0.020	0.7031 $\pm$ 0.007

Table 4. Performance comparison on Adidas dataset. We show mean and standard deviation of the metrics over five seeds.

Model	Method	BLEU	Rouge-1	Rouge-2	Rouge-L	METEOR	CIDEr	NIST
GPT2-M	Zeroshot	0.0046	0.2488	0.0189	0.1353	0.1653	0.0312	0.6860
GPT2-M	ICL-1	0.0088 $\pm$ 0.054	0.2667 $\pm$ 0.047	0.0247 $\pm$ 0.066	0.1358 $\pm$ 0.041	0.1762 $\pm$ 0.028	0.0464 $\pm$ 0.089	0.6793 $\pm$ 0.078
GPT2-M	ICL-3	0.0121 $\pm$ 0.047	0.2693 $\pm$ 0.028	0.0262 $\pm$ 0.054	0.1470 $\pm$ 0.020	0.1806 $\pm$ 0.030	0.0415 $\pm$ 0.104	0.7037 $\pm$ 0.081
GPT2-M	NewModel	0.0515 $\pm$ 0.016	0.2690 $\pm$ 0.014	0.0637 $\pm$ 0.014	0.1697 $\pm$ 0.008	0.1918 $\pm$ 0.013	0.0550 $\pm$ 0.086	0.6682 $\pm$ 0.047
GPT2-M	WeightedComb	0.0565 $\pm$ 0.014	0.2630 $\pm$ 0.028	0.0495 $\pm$ 0.018	0.1565 $\pm$ 0.015	0.1938 $\pm$ 0.019	0.0585 $\pm$ 0.088	0.6456 $\pm$ 0.156
GPT2-M	TempNet	0.0442 $\pm$ 0.017	0.2672 $\pm$ 0.019	0.0482 $\pm$ 0.022	0.1582 $\pm$ 0.020	0.1902 $\pm$ 0.017	0.0525 $\pm$ 0.031	0.6533 $\pm$ 0.098
GPT2-M	Plugin (Ours)	0.0486 $\pm$ 0.006	0.2766 $\pm$ 0.002	0.0515 $\pm$ 0.007	0.1684 $\pm$ 0.005	0.1994 $\pm$ 0.004	0.0626 $\pm$ 0.017	0.7919 $\pm$ 0.024
GPT2-XL	Zeroshot	0.0075	0.2309	0.0278	0.1438	0.1487	0.0184	0.4956
GPT2-XL	ICL-1	0.0109 $\pm$ 0.039	0.2567 $\pm$ 0.082	0.0265 $\pm$ 0.054	0.1519 $\pm$ 0.038	0.1649 $\pm$ 0.052	0.0318 $\pm$ 0.171	0.5133 $\pm$ 0.162
GPT2-XL	ICL-3	0.0295 $\pm$ 0.037	0.2509 $\pm$ 0.071	0.0395 $\pm$ 0.043	0.1536 $\pm$ 0.039	0.1658 $\pm$ 0.041	0.0321 $\pm$ 0.109	0.5176 $\pm$ 0.116
GPT2-XL	NewModel	0.0515 $\pm$ 0.016	0.2690 $\pm$ 0.014	0.0637 $\pm$ 0.014	0.1697 $\pm$ 0.008	0.1918 $\pm$ 0.013	0.0550 $\pm$ 0.086	0.6682 $\pm$ 0.047
GPT2-XL	WeightedComb	0.0567 $\pm$ 0.016	0.2210 $\pm$ 0.027	0.0714 $\pm$ 0.015	0.1550 $\pm$ 0.024	0.1674 $\pm$ 0.017	0.0183 $\pm$ 0.117	0.4105 $\pm$ 0.109
GPT2-XL	TempNet	0.0539 $\pm$ 0.018	0.2598 $\pm$ 0.026	0.0686 $\pm$ 0.014	0.1562 $\pm$ 0.019	0.1863 $\pm$ 0.029	0.0462 $\pm$ 0.120	0.5263 $\pm$ 0.117
GPT2-XL	Plugin (Ours)	0.0600 $\pm$ 0.017	0.2710 $\pm$ 0.025	0.0722 $\pm$ 0.018	0.1725 $\pm$ 0.017	0.1995 $\pm$ 0.018	0.1195 $\pm$ 0.138	0.6375 $\pm$ 0.120
LLaMA-3.1-8B	Zeroshot	0.0120	0.2470	0.0318	0.1493	0.1526	0.0424	0.5285
LLaMA-3.1-8B	ICL-1	0.0220 $\pm$ 0.044	0.2472 $\pm$ 0.072	0.0405 $\pm$ 0.068	0.1434 $\pm$ 0.057	0.1686 $\pm$ 0.041	0.0555 $\pm$ 0.133	0.5078 $\pm$ 0.142
LLaMA-3.1-8B	ICL-3	0.0177 $\pm$ 0.041	0.2385 $\pm$ 0.065	0.0364 $\pm$ 0.071	0.1408 $\pm$ 0.030	0.1712 $\pm$ 0.029	0.0587 $\pm$ 0.102	0.5775 $\pm$ 0.145
LLaMA-3.1-8B	NewModel	0.0506 $\pm$ 0.011	0.2700 $\pm$ 0.011	0.0634 $\pm$ 0.006	0.1749 $\pm$ 0.006	0.1995 $\pm$ 0.009	0.0575 $\pm$ 0.051	0.6570 $\pm$ 0.072
LLaMA-3.1-8B	WeightedComb	0.0357 $\pm$ 0.017	0.2583 $\pm$ 0.014	0.0661 $\pm$ 0.015	0.1560 $\pm$ 0.011	0.1706 $\pm$ 0.016	0.0745 $\pm$ 0.086	0.5927 $\pm$ 0.077
LLaMA-3.1-8B	TempNet	0.0472 $\pm$ 0.016	0.2647 $\pm$ 0.022	0.0625 $\pm$ 0.012	0.1625 $\pm$ 0.020	0.1857 $\pm$ 0.013	0.0586 $\pm$ 0.103	0.5926 $\pm$ 0.137
LLaMA-3.1-8B	Plugin (Ours)	0.0611 $\pm$ 0.018	0.2714 $\pm$ 0.029	0.0742 $\pm$ 0.020	0.1759 $\pm$ 0.019	0.1990 $\pm$ 0.020	0.1293 $\pm$ 0.152	0.6361 $\pm$ 0.134

over five random seeds, and we report the mean and standard deviation for each metric.

We compare *Plugin* with the following baselines: (a) **Zeroshot**: The black-box model directly performs text generation without additional adaptation. (b) **ICL-1** (Long et al., 2023): One randomly selected validation sample is used as an in-context example. (c) **ICL-3** (Long et al., 2023): Three randomly selected validation samples are used as in-context examples. (d) **NewModel**: A new language model is trained using the validation data. (e) **WeightedComb** (Liu et al., 2021): A new model is trained alongside the black-box model, with token probabilities computed as  $\alpha n + (1 - \alpha)b$ , where  $n$  represents the probabilities from the new model and  $\alpha$  is cross-validated in  $\{0.25, 0.50, 0.75\}$ . (f) **TempNet** (Qiu et al., 2024), a recent logit-scaling approach that

learns a global temperature per input and uniformly scales logits during generation. Since the black-box model weights are inaccessible, fine-tuning-based approaches are not applicable in our setting. Nonetheless, we include a comparison with LoRA in Appendix C.4 for completeness. This highlights *Plugin*'s competitiveness despite operating under stricter access constraints than required for PEFT.

All methods use the same prompts where applicable (Appendix C.2) and employ greedy decoding. The base (black-box) models used are GPT2-M (Radford et al., 2019), GPT2-XL (Radford et al., 2019), and LLaMA-3.1-8B (Dubey et al., 2024). *NewModel*, *WeightedComb*, and the reweighting model in *Plugin* share the same architecture. For GPT-based models, these use a Transformer encoder with one hidden layer and default configurations. For LLaMA-based mod-

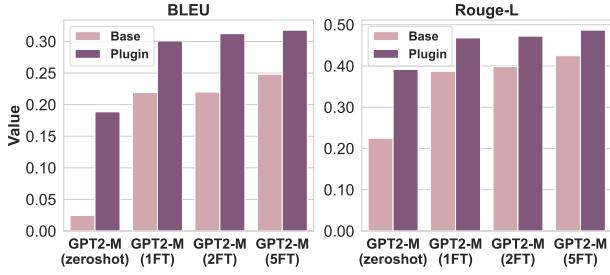


Figure 2. *Plugin* with increasingly fine-tuned GPT2-M models on the E2E NLG dataset. Results demonstrate that as the quality of the base model improves, the performance of the *Plugin* improves.

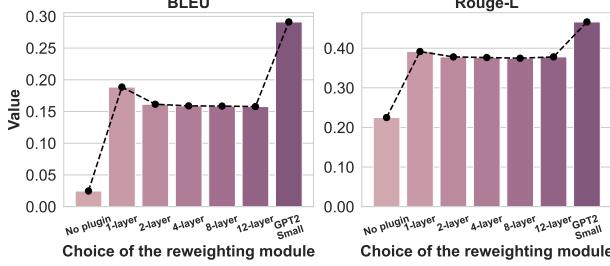


Figure 3. Performance of GPT2-M with varying reweighting model complexities on E2E NLG (BLEU, ROUGE-L). A single-layer reweighting model yields significant gains, while additional layers degrade performance due to overfitting. Initializing with GPT2-Small as the reweighting model improves performance, demonstrating the benefits of leveraging small pretrained models.

els, the architecture consists of a Transformer encoder with one hidden layer, 256 hidden size, 1024 intermediate size, and one attention head. Learning rate and weight decay are cross-validated over  $\{1e-5, 5e-5, 1e-4, 5e-4, 1e-3, 5e-3\}$  and  $\{0.01, 0.1, 1, 10\}$ , respectively. Models are trained using AdamW with warmup followed by linear decay, and early stopping is applied if the *hyper-validation* loss does not decrease for five consecutive epochs.

As shown in Tables 1–4 (the best is bold, the second best is underlined), *Plugin* outperforms baselines across nearly all datasets, black-box models, and evaluation metrics. *New-Model* occasionally achieves higher NIST scores due to increased repetition of less-frequent input tokens, but this comes at the cost of coherence, as reflected by other metrics. *WeightedComb* does not perform well, indicating one combination for all tokens is not a good modeling choice. *TempNet*, which learns a single temperature per input and uniformly scales logits during generation, also underperforms. In contrast, *Plugin* reweights logits at each timestep, enabling finer, context-sensitive adjustments.

We note that the absolute numbers may not appear competitive with state-of-the-art results, because (a) we restrict to greedy decoding (Section 4.2), and (b) Web NLG and CommonGen use distribution-shifted subsets.

We also conduct a human evaluation on 100 Adidas dataset samples, where three subjects compare outputs from *Plugin*

Table 5. Performance comparison of BDPL and BDPL + Plugin.

Dataset	Method	BLEU	Rouge-L	METEOR	CIDEr	NIST
E2E NLG	BDPL	0.2287	0.3922	0.4628	0.4216	0.8625
	<b>BDPL + Plugin</b>	<b>0.4527</b>	<b>0.6027</b>	<b>0.6214</b>	<b>0.7002</b>	<b>2.0817</b>
WEB NLG	BDPL	0.1024	0.3017	0.3527	0.4321	0.2631
	<b>BDPL + Plugin</b>	<b>0.2137</b>	<b>0.5928</b>	<b>0.5766</b>	<b>1.0826</b>	<b>0.6142</b>
CommonGen	BDPL	0.1023	0.2936	0.3362	0.2517	0.4226
	<b>BDPL + Plugin</b>	<b>0.2614</b>	<b>0.5241</b>	<b>0.5016</b>	<b>0.8251</b>	<b>0.9261</b>
Adidas	BDPL	0.0417	0.1710	0.1826	0.0861	0.6034
	<b>BDPL + Plugin</b>	<b>0.0623</b>	<b>0.1759</b>	<b>0.2148</b>	<b>0.1325</b>	<b>0.7024</b>

and *ICL-3* using LLaMA-3.1 as the base model. Evaluators select the prediction closest to the ground truth, with *Plugin* preferred in 81% of cases. Details are in Appendix C.7.

## 7.2. Plugin as a Wrapper

If logit access is available, *Plugin* can be applied on top of any prompt-based method using its best-found prompt. For example, our *Zeroshot* prompt is reused across methods. We also apply *Plugin* to the Black-box Discrete Prompt Learning (BDPL) approach from Diao et al. (2022), following their recommended 75 API call budget. Table 5 shows results on all datasets with GPT2-XL as the base model. *Plugin* outperforms BDPL (see Tables 1–4), and their combination yields further gains, underscoring the utility of logit-level access in strengthening prompt-based methods.

## 7.3. Ablation Study

We now show ablation studies that reflect various aspects of the *Plugin* model. We display the results using GPT2-M as base model on the E2E NLG dataset. The observation is similar on other base models and datasets (Appendix C.5).

**Impact of Base Model Quality.** We fine-tune GPT2-M for varying epochs, denoted as 1FT (one epoch), 2FT (two epochs), and 5FT (five epochs), and train a *Plugin* model for each. Figure 2 shows that as the base model’s task-specific quality improves, the *Plugin*’s performance improves.

**Complexity of the Reweighting Model in *Plugin*.** We train *Plugin* models with reweighting architectures varying from 1 to 12 transformer layers while keeping other configurations unchanged. Additionally, we train a variant where the reweighting model is initialized with GPT2-Small. As shown in Figure 3, a single-layer reweighting model yields significant improvements over the base GPT2-M model, while additional layers (e.g., 2, 4, 8, 12) offer diminishing returns and slight performance decline due to overfitting on the small validation set of E2E NLG. This suggests that more data is required for learning complex reweighting models. Notably, initializing with a pretrained GPT2-Small substantially improves performance, underscoring the advantage of using small pretrained models for reweighting due to their inherent autoregressive properties.

## 7.4. Qualitative Analysis and Case Study

**Plugin adapting to distribution shift.** We evaluate *Plugin* on distribution-shifted Web NLG and CommonGen us-

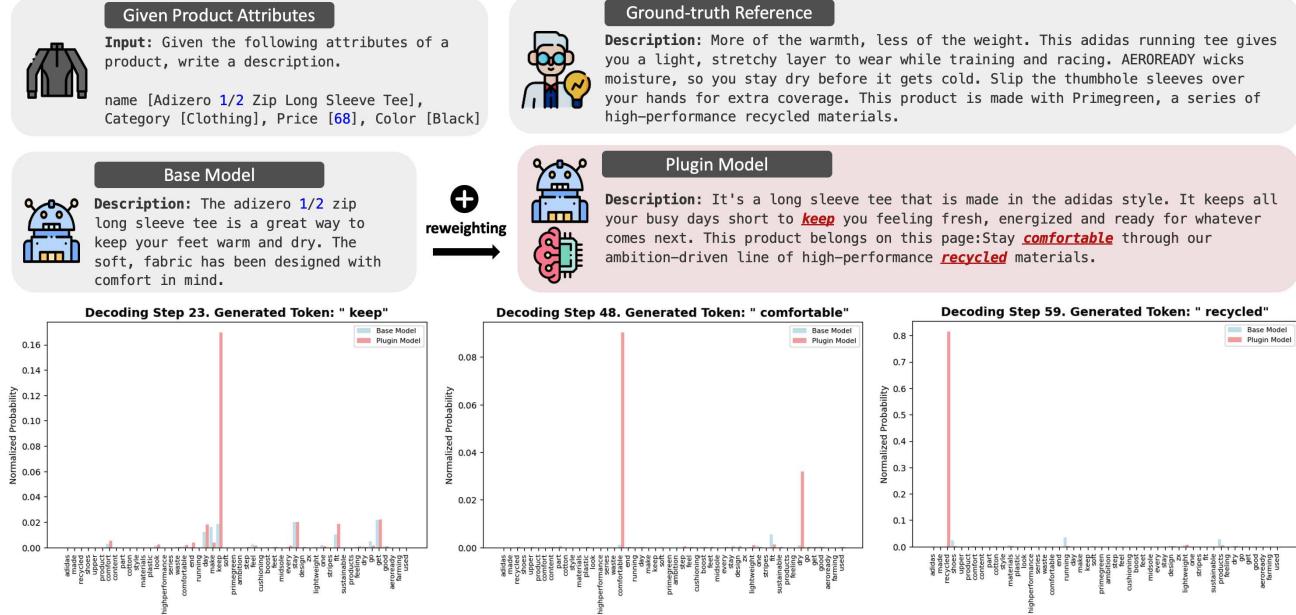


Figure 4. Comparison of the adaptation ability between the base model and *Plugin* on Adidas dataset. *Plugin*, enhanced with a reweighting model, generates text that better aligns with the “Adidas domain”. The bottom row illustrates token probabilities for key Adidas-related words at different decoding steps, showing how the reweighting model influences token selection.

ing LLaMA-3.1-8B as the base model. Web NLG training data contains only *Infrastructure* concepts, while validation and test sets include *Person* concepts. Similarly, CommonGen training data features *man*, whereas validation and test sets contain both *man* and *woman*. The base model is fine-tuned on the training data, and *Plugin* is trained on validation data using the fine-tuned model as the base. These settings reflect different degrees of domain shift, even *adversarial* to some extent as the training distributions induce biases (e.g., overemphasis on infrastructure or male-related concepts), and *Plugin* is tasked with correcting them during inference.

Using GPT-4o (Hurst et al., 2024) as an evaluator, the fine-tuned Web NLG model generates only 17.99% *Person*-related sentences, while *Plugin* increases this to 71.34%. On CommonGen, the fine-tuned model generates 10.37% *Woman*-related sentences, whereas *Plugin* improves this to 31.92%. These results highlight *Plugin*’s ability to adapt under distribution shift and mitigate biases in the base model.

**Case study: *Plugin* adapting to domain (extreme distribution shifts).** We examine token probabilities during inference for LLaMA-3.1-8B and *Plugin* to assess domain adaptation in the Adidas dataset, which features product-centric language and a brand-specific tone that diverge significantly from the general pretraining distribution of the black-box LLMs. This experimental setup can also be viewed as *extreme distribution shift*. Removing stopwords, we extract the top-50 most frequent words, defining the “Adidas domain”. Figure 4 illustrates this adaptation: the first row shows product attributes and ground-truth references; the second row compares outputs from the base model (left)

and *Plugin* (right); the third row visualizes model probabilities for “Adidas domain” words at three decoding steps.

As seen in Figure 4, *Plugin* dynamically reweights probabilities to align with domain-specific language. At step 23, “keep” is significantly upweighted. At step 48, “comfortable” and “dry” gain prominence over “fit,” which the base model favors. At step 59, “recycled” is preferred by *Plugin*, aligning with the ground truth, while the base model favors “running” and “products”. This demonstrates that *Plugin* effectively steers generation toward domain-specific terminology, whereas the base model, trained on broad corpora, lacks inherent domain preference.

Unlike methods that prune or suppress tokens, *Plugin* softly reweights token probabilities without eliminating any vocabulary candidates. This preserves full coverage while amplifying domain-specific terms. To quantify this, we measure the total occurrences of the top-50 “Adidas domain” words in generated outputs: *Plugin* includes 25.6% of these terms compared to 13.8% in the base model, indicating substantially improved alignment with domain language.

## 8. Conclusion

We propose *Plugin*, a token-level probability reweighting framework that adapts black-box LLMs using only logits and small task-specific data. Framing next-token prediction as a label noise correction problem, we demonstrate both theoretical guarantees and empirical effectiveness across multiple datasets and models. Our findings highlight the potential of logit-based adaptation and advocate for broader access to token logits in closed-source LLMs.

## Acknowledgements

HW acknowledges support by Fonds de recherche du Québec – Nature et technologies (FRQNT) and Borealis AI. SK acknowledges support by NSF 2046795 and 2205329, IES R305C240046, the MacArthur Foundation, Stanford HAI, OpenAI, and Google.

## Impact Statement

This work introduces a powerful middle ground between fully black-box APIs and fully white-box access to large language models (LLMs), addressing a critical constraint faced by developers: the inability to adapt models when weights and architecture are inaccessible. By leveraging token-level logits—without requiring access to model weights or architecture—our approach enables meaningful adaptation of closed-source LLMs for domain-specific tasks. This has far-reaching implications for both research and industry: it empowers developers to customize models within privacy-preserving, IP-sensitive environments while ensuring greater control, transparency, and safety. Our findings advocate for broader logit access as a scalable, secure, and effective interface—bridging the gap between usability and protection of proprietary models—and open new possibilities for equitable, context-aware language generation in real-world applications.

While *Plugin* effectively adapts black-box LLMs, it has some limitations, too. Since it only reweights token probabilities without modifying internal representations or embeddings, it may struggle with tasks requiring deep structural adaptations, such as executing complex reasoning. Further research on this aspect is needed. Additionally, although *Plugin* avoids full fine-tuning, training a separate reweighting model introduces computational overhead compared to prompt tuning or in-context learning, with efficiency depending on the complexity of the reweighting model and the availability of task-specific data.

## References

Adidas us retail products dataset. *Kaggle*, 2023. URL <https://www.kaggle.com/datasets/whenamancodes/adidas-us-retail-products-dataset>.

Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.

Bai, Y., Jones, A., Ndousse, K., Askell, A., Chen, A., Das-Sarma, N., Drain, D., Fort, S., Ganguli, D., Henighan, T., et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022a.

Bai, Y., Kadavath, S., Kundu, S., Askell, A., Kernion, J., Jones, A., Chen, A., Goldie, A., Mirhoseini, A., McKinnon, C., et al. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*, 2022b.

Banerjee, S. and Lavie, A. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *ACL workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pp. 65–72, 2005.

Chaudhuri, K., Kakade, S. M., Netrapalli, P., and Sanghavi, S. Convergence rates of active learning for maximum likelihood estimation. *Advances in Neural Information Processing Systems*, 28, 2015.

Christiano, P. F., Leike, J., Brown, T., Martic, M., Legg, S., and Amodei, D. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30, 2017.

Dathathri, S., Madotto, A., Lan, J., Hung, J., Frank, E., Molino, P., Yosinski, J., and Liu, R. Plug and play language models: A simple approach to controlled text generation. In *International Conference on Learning Representations*, 2020.

Dettmers, T., Pagnoni, A., Holtzman, A., and Zettlemoyer, L. Qlora: Efficient finetuning of quantized llms. *Advances in Neural Information Processing Systems*, 36, 2024.

Diao, S., Huang, Z., Xu, R., Li, X., Lin, Y., Zhou, X., and Zhang, T. Black-box prompt learning for pre-trained language models. *arXiv preprint arXiv:2201.08531*, 2022.

Doddington, G. Automatic evaluation of machine translation quality using n-gram co-occurrence statistics. In *International conference on Human Language Technology Research*, pp. 138–145, 2002.

Dong, Q., Li, L., Dai, D., Zheng, C., Ma, J., Li, R., Xia, H., Xu, J., Wu, Z., Chang, B., et al. A survey on in-context learning. In *Conference on Empirical Methods in Natural Language Processing*, pp. 1107–1128, 2024.

Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A., Yang, A., Fan, A., et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.

Dušek, O., Novikova, J., and Rieser, V. Evaluating the state-of-the-art of end-to-end natural language generation: The e2e nlg challenge. *Computer Speech & Language*, 59: 123–156, 2020.

Frostig, R., Ge, R., Kakade, S. M., and Sidford, A. Competing with the empirical risk minimizer in a single pass. In *Conference on learning theory*, pp. 728–763. PMLR, 2015.

Gao, G., Taymanov, A., Salinas, E., Mineiro, P., and Misra, D. Aligning llm agents by learning latent preference from user edits. *arXiv preprint arXiv:2404.15269*, 2024.

Gardent, C., Shimorina, A., Narayan, S., and Perez-Beltrachini, L. Creating training corpora for nlg micro-planning. In *Annual Meeting of the Association for Computational Linguistics*, pp. 179–188, 2017.

Hiranandani, G., Mathur, J., Narasimhan, H., Fard, M. M., and Koyejo, S. Optimizing black-box metrics with iterative example weighting. In *International Conference on Machine Learning*, pp. 4239–4249. PMLR, 2021.

Hsu, D., Kakade, S., Zhang, T., et al. A tail inequality for quadratic forms of subgaussian random vectors. *Electronic Communications in Probability*, 17, 2012.

Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., and Chen, W. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.

Hu, Z., Wang, L., Lan, Y., Xu, W., Lim, E.-P., Bing, L., Xu, X., Poria, S., and Lee, R. LLM-adapters: An adapter family for parameter-efficient fine-tuning of large language models. In *Conference on Empirical Methods in Natural Language Processing*, pp. 5254–5276, 2023a.

Hu, Z., Wang, L., Lan, Y., Xu, W., Lim, E.-P., Bing, L., Xu, X., Poria, S., and Lee, R. Llm-adapters: An adapter family for parameter-efficient fine-tuning of large language models. In *Conference on Empirical Methods in Natural Language Processing*, pp. 5254–5276, 2023b.

Huang, Y., Liu, Y., Thirukovalluru, R., Cohan, A., and Dhingra, B. Calibrating long-form generations from large language models. *arXiv preprint arXiv:2402.06544*, 2024.

Hurst, A., Lerer, A., Goucher, A. P., Perelman, A., Ramesh, A., Clark, A., Ostrow, A., Welihinda, A., Hayes, A., Radford, A., et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024.

Kapoor, S., Gruver, N., Roberts, M., Pal, A., Dooley, S., Goldblum, M., and Wilson, A. Calibration-tuning: Teaching large language models to know what they don’t know. In *Workshop on Uncertainty-Aware NLP*, pp. 1–14, 2024.

Kojima, T., Gu, S. S., Reid, M., Matsuo, Y., and Iwasawa, Y. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35, 2022.

Koyejo, O. O., Natarajan, N., Ravikumar, P. K., and Dhillon, I. S. Consistent binary classification with generalized performance metrics. *Advances in neural information processing systems*, 27, 2014.

Krause, B., Gotmare, A. D., McCann, B., Keskar, N. S., Joty, S., Socher, R., and Rajani, N. F. Gedi: Generative discriminator guided sequence generation. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pp. 4929–4952, 2021.

Lattimore, T. and Szepesvári, C. *Bandit algorithms*. Cambridge University Press, 2020.

Lester, B., Al-Rfou, R., and Constant, N. The power of scale for parameter-efficient prompt tuning. In *Conference on Empirical Methods in Natural Language Processing*, pp. 3045–3059, 2021.

Li, X., Thickstun, J., Gulrajani, I., Liang, P. S., and Hashimoto, T. B. Diffusion-lm improves controllable text generation. *Advances in Neural Information Processing Systems*, 35:4328–4343, 2022.

Li, X. L. and Liang, P. Prefix-tuning: Optimizing continuous prompts for generation. In *Annual Meeting of the Association for Computational Linguistics*, pp. 4582–4597, 2021.

Lin, B. Y., Zhou, W., Shen, M., Zhou, P., Bhagavatula, C., Choi, Y., and Ren, X. Commongen: A constrained text generation challenge for generative commonsense reasoning. In *Findings of the Association for Computational Linguistics*, pp. 1823–1840, 2020.

Lin, C.-Y. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pp. 74–81, 2004.

Lin, C.-Y. and Och, F. J. Automatic evaluation of machine translation quality using longest common subsequence and skip-bigram statistics. In *Annual meeting of the association for computational linguistics*, pp. 605–612, 2004.

Lipton, Z., Wang, Y.-X., and Smola, A. Detecting and correcting for label shift with black box predictors. In *International conference on machine learning*, pp. 3122–3130. PMLR, 2018.

Liu, A., Sap, M., Lu, X., Swayamdipta, S., Bhagavatula, C., Smith, N. A., and Choi, Y. Dexperts: Decoding-time controlled text generation with experts and anti-experts. In *Annual Meeting of the Association for Computational Linguistics*, 2021.

Liu, A., Han, X., Wang, Y., Tsvetkov, Y., Choi, Y., and Smith, N. A. Tuning language models by proxy. *arXiv preprint arXiv:2401.08565*, 2024.

Long, Q., Wang, W., and Pan, S. Adapt in contexts: Retrieval-augmented domain adaptation via in-context learning. In *Conference on Empirical Methods in Natural Language Processing*, pp. 6525–6542, 2023.

Mireshghallah, F., Goyal, K., and Berg-Kirkpatrick, T. Mix and match: Learning-free controllable text generation using energy language models. In *Annual Meeting of the Association for Computational Linguistics*, pp. 401–415, 2022.

Mukherjee, S., Tripathy, A. S., and Nowak, R. Chernoff sampling for active testing and extension to active regression. In *International Conference on Artificial Intelligence and Statistics*, pp. 7384–7432. PMLR, 2022.

Naghshvar, M. and Javidi, T. Active sequential hypothesis testing. 2013.

Narasimhan, H., Ramaswamy, H., Saha, A., and Agarwal, S. Consistent multiclass algorithms for complex performance measures. In *International Conference on Machine Learning*, pp. 2398–2407. PMLR, 2015.

Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022.

Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. Bleu: a method for automatic evaluation of machine translation. In *Annual meeting of the Association for Computational Linguistics*, pp. 311–318, 2002.

Patrini, G., Rozza, A., Krishna Menon, A., Nock, R., and Qu, L. Making deep neural networks robust to label noise: A loss correction approach. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1944–1952, 2017.

Qiu, Z.-H., Guo, S., Xu, M., Zhao, T., Zhang, L., and Yang, T. To cool or not to cool? temperature network meets large foundation models via dro. *arXiv preprint arXiv:2404.04575*, 2024.

Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.

Rafailov, R., Sharma, A., Mitchell, E., Manning, C. D., Ermon, S., and Finn, C. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36, 2024.

Sanh, V., Webson, A., Raffel, C., Bach, S., Sutawika, L., Alyafeai, Z., Chaffin, A., Stiegler, A., Raja, A., Dey, M., et al. Multitask prompted training enables zero-shot task generalization. In *International Conference on Learning Representations*, 2022.

Shen, M., Das, S., Greenwald, K., Sattigeri, P., Wornell, G. W., and Ghosh, S. Thermometer: Towards universal calibration for large language models. In *International Conference on Machine Learning*, 2024.

Singh, S. P., Hofmann, T., and Schölkopf, B. The hessian perspective into the nature of convolutional neural networks. *arXiv preprint arXiv:2305.09088*, 2023.

Song, C. H., Wu, J., Washington, C., Sadler, B. M., Chao, W.-L., and Su, Y. Llm-planner: Few-shot grounded planning for embodied agents with large language models. In *IEEE/CVF International Conference on Computer Vision*, pp. 2998–3009, 2023.

Sukhbaatar, S., Bruna, J., Paluri, M., Bourdev, L., and Fergus, R. Training convolutional networks with noisy labels. In *International Conference on Learning Representations*, 2015.

Torroba-Hennigen, L., Lang, H., Guo, H., and Kim, Y. On the duality between gradient transformations and adapters. *arXiv preprint arXiv:2502.13811*, 2025.

Ulmer, D., Gubri, M., Lee, H., Yun, S., and Oh, S. J. Calibrating large language models using their generations only. *arXiv preprint arXiv:2403.05973*, 2024.

Vedantam, R., Lawrence Zitnick, C., and Parikh, D. Cider: Consensus-based image description evaluation. In *IEEE conference on computer vision and pattern recognition*, pp. 4566–4575, 2015.

Wei, J., Bosma, M., Zhao, V., Guu, K., Yu, A. W., Lester, B., Du, N., Dai, A. M., and Le, Q. V. Finetuned language models are zero-shot learners. In *International Conference on Learning Representations*, 2021.

Wolf, T. Transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*, 2020.

Yang, K. and Klein, D. Fudge: Controlled text generation with future discriminators. In *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 3511–3535, 2021.

Zhang, H., Zhang, Y.-F., Yu, Y., Madeka, D., Foster, D., Xing, E., Lakkaraju, H., and Kakade, S. A study on the calibration of in-context learning. *arXiv preprint arXiv:2312.04021*, 2023.

Zhang, Y., Chen, C., Ding, T., Li, Z., Sun, R., and Luo, Z.-Q. Why transformers need adam: A hessian perspective. *arXiv preprint arXiv:2402.16788*, 2024.

Zhu, C., Xu, B., Wang, Q., Zhang, Y., and Mao, Z. On the calibration of large language models and alignment. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 9778–9795, 2023.

## A. Algorithm Details

We provide summarized form of the training and inference algorithm for the *Plugin* model below.

---

### Algorithm 1 Training and Inference for the Plugin Model

---

**Input:** Black-box model  $B$ , reweighting model  $R$ , clean training data  $\mathcal{D}$ , vocabulary  $V$

**Output:** Plugin model predictions  $\mathbf{x}_{1:T}$  for a given sequence

```

1: Training Phase:
2: for each sequence  $s \in \mathcal{D}$  do
3:   Compute token probabilities  $\{b_1, b_2, \dots, b_m\}$  using  $B$ .
4:   Compute token probabilities  $\{r_1, r_2, \dots, r_m\}$  using  $R$ .
5:   Combine probabilities:  $\mathbf{p}_i = \frac{b_i \odot r_i}{\|b_i \odot r_i\|_1}$  for  $i \in [m]$ .
6:   Compute sequence loss  $\ell_s = -\frac{1}{m} \sum_{i=1}^m \sum_{j=1}^{|V|} \log(\mathbf{p}_i) \odot \mathbf{e}_j$ .
7:   Update parameters of  $R$  using back-propagation. Freeze  $B$ .
8: end for
9: Inference Phase:
10: Initialize sequence  $\mathbf{x}_{1:T} = \{\}$ .
11: for each token position  $t = 1$  to  $T$  do
12:   Compute token probabilities  $b_t$  using  $B$ .
13:   Compute token probabilities  $r_t$  using  $R$ .
14:   Combine probabilities:  $\mathbf{p}_t = \frac{b_t \odot r_t}{\|b_t \odot r_t\|_1}$ .
15:   Predict token:  $\mathbf{x}_t = \text{argmax}_V(\mathbf{p}_t)$ .
16:   Append  $\mathbf{x}_t$  to  $\mathbf{x}_{1:T}$ .
17: end for
18: Return:  $\mathbf{x}_{1:T}$ 

```

---

## B. Proof of Main Convergence Theorem

We define the following assumption on the smoothness and regularity of the loss function.

**Assumption B.1.** We assume the following assumptions hold with probability 1:

1. **(Convexity of  $\ell_s$ ):** The loss function  $\ell_s$  is convex for all time  $s \in [t]$ .
2. **(Smoothness of  $\ell_s$ ):** The  $\ell_s$  is smooth such that the first, second, and third derivatives exist at all interior points in  $\Theta$ .
3. **(Regularity Conditions):**
  - (a)  $\Theta$  is compact and  $\ell_s(\theta)$  is bounded for all  $\theta \in \Theta$  and for all  $s \in [t]$ .
  - (b)  $\theta_*$  is an interior point in  $\Theta$ .
  - (c)  $\nabla^2 \ell_s(\theta_*)$  is positive definite, for all  $s \in [t]$ .
  - (d) There exists a neighborhood  $\mathcal{B}$  of  $\theta_*$  and a constant  $C_1$ , such that  $\nabla^2 \ell_s(\theta)$  is  $C_1$ -Lipschitz. Hence, we have that  $\|\nabla^2 \ell_s(\theta) - \nabla^2 \ell_s(\theta')\|_* \leq C_1 \|\theta - \theta'\|_{\nabla^2 \ell_s(\theta_*)}$ , for  $\theta, \theta'$  in this neighborhood.
4. **(Concentration at  $\theta_*$ ):** We further assume that  $\|\nabla \ell_s(\theta_*)\|_{(\nabla^2 \ell_s(\theta_*))^{-1}} \leq C_2$  hold with probability one.

**Lemma B.2. (Proposition 2 of (Hsu et al., 2012))** Let  $\mathbf{u}_1, \dots, \mathbf{u}_n$  be a martingale difference vector sequence (i.e.,  $\mathbb{E} [\mathbf{u}_i | \mathbf{u}_1, \dots, \mathbf{u}_{i-1}] = 0$  for all  $i = 1, \dots, n$ ) such that

$$\sum_{i=1}^n \mathbb{E} [\|\mathbf{u}_i\|^2 | \mathbf{u}_1, \dots, \mathbf{u}_{i-1}] \leq v \quad \text{and} \quad \|\mathbf{u}_i\| \leq b$$

for all  $i = 1, \dots, n$ , almost surely. For all  $t > 0$

$$\Pr \left[ \left\| \sum_{i=1}^n \mathbf{u}_i \right\| > \sqrt{v} + \sqrt{8vt} + (4/3)bt \right] \leq e^{-t}$$

**Lemma B.3.** The probability that  $\|\nabla \widehat{\mathcal{L}}_t(\boldsymbol{\theta}_*)\|_{(\nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*))^{-1}}$  crosses the threshold  $\sqrt{\frac{c\gamma \log(dt)}{t}} > 0$  is bounded by

$$\mathbb{P} \left( \|\nabla \widehat{\mathcal{L}}_t(\boldsymbol{\theta}_*)\|_{(\nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*))^{-1}} \geq C_2 \sqrt{\frac{c\gamma \log(dt)}{t}} \right) \leq \frac{1}{t^{c\gamma}}.$$

*Proof.* Define  $\mathbf{u}_s := \nabla(Y_s - f_{I_s}(\boldsymbol{\theta}_*; x_i, x_j, \mathcal{F}^{s-1}))^2$ . Then we have  $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_t$  as random vectors such that

$$\mathbb{E} \left[ \left\| \sum_{s=1}^t \mathbf{u}_s \right\|_{(\nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*))^{-1}}^2 \mid \mathbf{u}_1, \dots, \mathbf{u}_{s-1} \right] = \mathbb{E} \left[ \sum_{s=1}^t \mathbf{u}_s^\top (\nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*))^{-1} \mathbf{u}_s \mid \mathbf{u}_1, \dots, \mathbf{u}_{s-1} \right] \leq t C_2^2$$

Also we have that  $\|\mathbf{u}_s\| \leq C_2$ . Finally we have that

$$\mathbb{E}[\nabla_{\boldsymbol{\theta}=\boldsymbol{\theta}_*} \mathbf{u}_s] = -2 \sum_{s=1}^t p_{\widehat{\boldsymbol{\theta}}_{s-1}}(f_{I_s}(\boldsymbol{\theta}_*; x_i, x_j, \mathcal{F}^{s-1}) - f_{I_s}(\boldsymbol{\theta}_*; x_i, x_j, \mathcal{F}^{s-1})) \nabla_{\boldsymbol{\theta}=\boldsymbol{\theta}_*} f_{I_s}(\boldsymbol{\theta}_*; x_i, x_j, \mathcal{F}^{s-1}) = 0.$$

Then following Lemma B.2 and by setting  $\epsilon = c\gamma \log(dt)$  we can show that

$$\begin{aligned} & \mathbb{P} \left( \left\| \frac{1}{t} \sum_{s=1}^t \mathbf{u}_s \right\|_{(\nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*))^{-1}}^2 - \mathbb{E} \left[ \left\| \frac{1}{t} \sum_{s=1}^t \mathbf{u}_s \right\|_{(\nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*))^{-1}}^2 \right] > \frac{1}{t} \sqrt{8tC_2^2\epsilon} + \frac{4C_2}{3\epsilon} \right) \\ &= \mathbb{P} \left( \left\| \frac{1}{t} \sum_{s=1}^t \mathbf{u}_s \right\|_{(\nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*))^{-1}}^2 > C_1^2 + C_2 \sqrt{\frac{8\epsilon}{t}} + \frac{4C_2}{3\epsilon} \right) \\ &\leq \mathbb{P} \left( \left\| \sum_{s=1}^t \mathbf{u}_s \right\|_{(\nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*))^{-1}}^2 > C_2 \sqrt{\frac{8\epsilon}{t}} \right) = \mathbb{P} \left( \left\| \sum_{s=1}^t \mathbf{u}_s \right\|_{(\nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*))^{-1}}^2 > 4C_2 \sqrt{\frac{c\gamma \log(dt)}{t}} \right) \\ &\leq \exp(-c\gamma \log(dt)) = \left( \frac{1}{dt} \right)^{c\gamma} \leq \frac{1}{t^{c\gamma}} \end{aligned}$$

The claim of the lemma follows.  $\square$

**Lemma B.4.** Let the  $j$ -th row and  $k$ -th column entry in the Hessian matrix  $\nabla_{\boldsymbol{\theta}=\boldsymbol{\theta}'}^2(\ell_s(\boldsymbol{\theta}))$  be denoted as  $[\nabla_{\boldsymbol{\theta}=\boldsymbol{\theta}'}^2(\ell_s(\boldsymbol{\theta}))]_{jk}$ . Then we have that

$$[\nabla_{\boldsymbol{\theta}=\boldsymbol{\theta}'}^2(\ell_s(\boldsymbol{\theta}))]_{jk} = 2 \frac{\partial f_{I_s}(\boldsymbol{\theta}; x_i, x_j, \mathcal{F}^{s-1})}{\partial \boldsymbol{\theta}_j} \frac{\partial f_{I_s}(\boldsymbol{\theta}; x_i, x_j, \mathcal{F}^{s-1})}{\partial \boldsymbol{\theta}_k} + 2 \left( f_{I_s}(\boldsymbol{\theta}; x_i, x_j, \mathcal{F}^{s-1}) - Y_s \right) \frac{\partial^2 f_{I_s}(\boldsymbol{\theta}; x_i, x_j, \mathcal{F}^{s-1})}{\partial \boldsymbol{\theta}_j \partial \boldsymbol{\theta}_k}.$$

*Proof.* This lemma follows from Frostig et al. (2015); Mukherjee et al. (2022) adapted to our setting for the squared loss, and transition function  $f_{I_s}(\boldsymbol{\theta}_*; x_i, x_j, \mathcal{F}^{s-1})$ . We want to evaluate the Hessian  $\nabla_{\boldsymbol{\theta}=\boldsymbol{\theta}'}^2(\ell_s(\boldsymbol{\theta}))$  at any  $\boldsymbol{\theta}' \in \boldsymbol{\Theta}$ . We denote the  $j$ -th row and  $k$ -th column entry in the Hessian matrix as  $[\nabla_{\boldsymbol{\theta}=\boldsymbol{\theta}'}^2(\ell_s(\boldsymbol{\theta}))]_{jk}$ . Then we can show that

$$\begin{aligned} & [\nabla_{\boldsymbol{\theta}=\boldsymbol{\theta}'}^2(\ell_s(\boldsymbol{\theta}))]_{jk} := \frac{\partial}{\partial \boldsymbol{\theta}_j} \left[ \frac{\partial(f_{I_s}(\boldsymbol{\theta}; x_i, x_j, \mathcal{F}^{s-1}) - Y_s)^2}{\partial \boldsymbol{\theta}_k} \right] = \frac{\partial}{\partial \boldsymbol{\theta}_j} \left[ 2(f_{I_s}(\boldsymbol{\theta}; x_i, x_j, \mathcal{F}^{s-1}) - Y_s) \frac{\partial f_{I_s}(\boldsymbol{\theta}; x_i, x_j, \mathcal{F}^{s-1})}{\partial \boldsymbol{\theta}_k} \right] \\ &= \frac{\partial}{\partial \boldsymbol{\theta}_j} \left[ 2f_{I_s}(\boldsymbol{\theta}; x_i, x_j, \mathcal{F}^{s-1}) \frac{\partial f_{I_s}(\boldsymbol{\theta}; x_i, x_j, \mathcal{F}^{s-1})}{\partial \boldsymbol{\theta}_k} - 2Y_s \frac{\partial f_{I_s}(\boldsymbol{\theta}; x_i, x_j, \mathcal{F}^{s-1})}{\partial \boldsymbol{\theta}_k} \right] \\ &= 2 \frac{\partial f_{I_s}(\boldsymbol{\theta}; x_i, x_j, \mathcal{F}^{s-1})}{\partial \boldsymbol{\theta}_j} \frac{\partial f_{I_s}(\boldsymbol{\theta}; x_i, x_j, \mathcal{F}^{s-1})}{\partial \boldsymbol{\theta}_k} + 2f_{I_s}(\boldsymbol{\theta}; x_i, x_j, \mathcal{F}^{s-1}) \frac{\partial^2 f_{I_s}(\boldsymbol{\theta}; x_i, x_j, \mathcal{F}^{s-1})}{\partial \boldsymbol{\theta}_j \partial \boldsymbol{\theta}_k} \\ &\quad - 2Y_s \frac{\partial^2 f_{I_s}(\boldsymbol{\theta}; x_i, x_j, \mathcal{F}^{s-1})}{\partial \boldsymbol{\theta}_j \partial \boldsymbol{\theta}_k} - 2 \frac{\partial f_{I_s}(\boldsymbol{\theta}; x_i, x_j, \mathcal{F}^{s-1})}{\partial \boldsymbol{\theta}_j} \frac{\partial Y_s}{\partial \boldsymbol{\theta}_k} \\ &= 2 \frac{\partial f_{I_s}(\boldsymbol{\theta}; x_i, x_j, \mathcal{F}^{s-1})}{\partial \boldsymbol{\theta}_j} \frac{\partial f_{I_s}(\boldsymbol{\theta}; x_i, x_j, \mathcal{F}^{s-1})}{\partial \boldsymbol{\theta}_k} + 2 \left( f_{I_s}(\boldsymbol{\theta}; x_i, x_j, \mathcal{F}^{s-1}) - Y_s \right) \frac{\partial^2 f_{I_s}(\boldsymbol{\theta}; x_i, x_j, \mathcal{F}^{s-1})}{\partial \boldsymbol{\theta}_j \partial \boldsymbol{\theta}_k} \end{aligned}$$

The claim of the lemma follows.  $\square$

**Lemma B.5.** Let the  $j$ -th row and  $k$ -th column entry in the Hessian matrix  $\nabla_{\theta=\theta'}^2(\mathbb{E}[\ell_s(\theta)|\mathcal{F}^{s-1}])$  be denoted as  $[\nabla_{\theta=\theta'}^2(\mathbb{E}[\ell_s(\theta)|\mathcal{F}^{s-1}])]_{jk}$ . Then we have that

$$\begin{aligned} \left[ \nabla_{\theta=\theta'}^2 \mathbb{E}[\ell_s(\theta)|\mathcal{F}^{s-1}] \right]_{jk} &= 2 \sum_{i=1}^{|V|} p_{\hat{\theta}_{s-1}}(i) \left( \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_j} \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_k} \right. \\ &\quad \left. + 2 \left( f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1}) - f_{I_s}(\theta_*; x_i, x_j, \mathcal{F}^{s-1}) \right) \frac{\partial^2 f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_j \partial \theta_k} \right). \end{aligned}$$

*Proof.* This lemma follows from [Frostig et al. \(2015\)](#); [Mukherjee et al. \(2022\)](#) adapted to our setting for the squared loss, transition function  $f_{I_s}(\theta_*; x_i, x_j, \mathcal{F}^{s-1})$ , and the sampling distribution  $\mathbf{p}_{\hat{\theta}_{s-1}}$ . We show it here for completeness. Now we want to evaluate the Hessian  $\nabla_{\theta=\theta'}^2(\mathbb{E}[\ell_s(\theta)|\mathcal{F}^{s-1}])$  at any  $\theta' \in \Theta$ . We denote the  $j$ -th row and  $k$ -th column entry in the Hessian matrix as  $[\nabla_{\theta=\theta'}^2(\mathbb{E}[\ell_s(\theta)|\mathcal{F}^{s-1}])]_{jk}$ . Then we can show that

$$\begin{aligned} \nabla_{\theta=\theta'}^2 \mathbb{E}[\ell_s(\theta)|\mathcal{F}^{s-1}] &= \nabla_{\theta=\theta'}^2 \left( f_{I_s}^2(\theta; x_i, x_j, \mathcal{F}^{s-1}) + \mathbb{E}[Y_s^2|\mathcal{F}^{s-1}] - 2\mathbb{E}[Y_s|\mathcal{F}^{s-1}]f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1}) \right) \\ &= \nabla_{\theta=\theta'}^2 \sum_{i=1}^{|V|} p_{\hat{\theta}_{s-1}}(i) \left( f_i^2(\theta; x_i, x_j, \mathcal{F}^{s-1}) + f_i^2(\theta'; x_i, x_j, \mathcal{F}^{s-1}) + \frac{1}{2} - 2f_{I_s}(\theta_*; x_i, x_j, \mathcal{F}^{s-1})f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1}) \right) \\ &= \nabla_{\theta=\theta'}^2 \sum_{i=1}^{|V|} p_{\hat{\theta}_{s-1}}(i) \left( \left( f_{I_s}(\theta_*; x_i, x_j, \mathcal{F}^{s-1}) - f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1}) \right)^2 + \frac{1}{2} \right) \\ &= \nabla_{\theta=\theta'}^2 \sum_{i=1}^{|V|} p_{\hat{\theta}_{s-1}}(i) \left( \left( f_{I_s}(\theta_*; x_i, x_j, \mathcal{F}^{s-1}) - f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1}) \right)^2 \right) \end{aligned} \tag{6}$$

We now denote the  $j$ -th row and  $k$ -th column entry of the Hessian Matrix  $\nabla_{\theta=\theta'}^2((f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1}) - f_i(\theta_*; x_i, x_j, \mathcal{F}^{s-1}))^2)$  as  $[\nabla_{\theta=\theta'}^2((f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1}) - f_{I_s}(\theta_*; x_i, x_j, \mathcal{F}^{s-1}))^2)]_{jk}$ . Then we can show that

$$\begin{aligned} &[\nabla_{\theta=\theta_*}^2((f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1}) - f_{I_s}(\theta_*; x_i, x_j, \mathcal{F}^{s-1}))^2)]_{jk} \\ &:= \frac{\partial}{\partial \theta_j} \left[ \frac{\partial(f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1}) - f_{I_s}(\theta_*; x_i, x_j, \mathcal{F}^{s-1}))^2}{\partial \theta_k} \right] \\ &= \frac{\partial}{\partial \theta_j} \left[ 2(f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1}) - f_{I_s}(\theta_*; x_i, x_j, \mathcal{F}^{s-1})) \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_k} \right] \\ &= \frac{\partial}{\partial \theta_j} \left[ 2f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1}) \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_k} - 2f_i(\theta_*) \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_k} \right] \\ &= 2 \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_j} \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_k} + 2f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1}) \frac{\partial^2 f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_j \partial \theta_k} \\ &\quad - 2f_{I_s}(\theta_*; x_i, x_j, \mathcal{F}^{s-1}) \frac{\partial^2 f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_j \partial \theta_k} - 2 \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_j} \frac{\partial f_{I_s}(\theta_*; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_k} \\ &= 2 \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_j} \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_k} + 2 \left( f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1}) - f_{I_s}(\theta_*; x_i, x_j, \mathcal{F}^{s-1}) \right) \frac{\partial^2 f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_j \partial \theta_k} \end{aligned}$$

Plugging this back in Equation (6) we get that

$$\begin{aligned} \left[ \nabla_{\theta=\theta'}^2 \mathbb{E}[\ell_s(\theta)|\mathcal{F}^{s-1}] \right]_{jk} &= 2 \sum_{i=1}^{|V|} p_{\hat{\theta}_{s-1}}(i) \left( \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_j} \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_k} \right. \\ &\quad \left. + 2 \left( f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1}) - f_{I_s}(\theta_*; x_i, x_j, \mathcal{F}^{s-1}) \right) \frac{\partial^2 f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_j \partial \theta_k} \right). \end{aligned}$$

□

**Lemma B.6.** *The sum of the difference of the Hessians  $\sum_{s=1}^t \nabla_{\theta=\theta'}^2 \ell_s(\theta) - \mathbb{E} [\nabla_{\theta=\theta'}^2 \ell_s(\theta) \mid \mathcal{F}^{s-1}]$  is given by*

$$\begin{aligned} \sum_{s=1}^t \nabla_{\theta=\theta'}^2 \ell_s(\theta) - \mathbb{E} [\nabla_{\theta=\theta'}^2 \ell_s(\theta) \mid \mathcal{F}^{s-1}] &= \sum_{s=1}^t \left( -2(Y_s - f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})) \frac{\partial^2 f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_j \partial \theta_k} \right. \\ &\quad + 2 \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_j} \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_k} \\ &\quad \left. - 2 \sum_{i=1}^{|V|} p_{\hat{\theta}_{s-1}}(i) \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_j} \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_k} \right). \end{aligned}$$

*Proof.* This lemma directly follows from Lemma B.4 and Lemma B.5. First note that the difference  $\nabla_{\theta=\theta'}^2 \ell_s(\theta) - \mathbb{E} [\nabla_{\theta=\theta'}^2 \ell_s(\theta) \mid \mathcal{F}^{s-1}]_{jk}$  is given by

$$\begin{aligned} &\nabla_{\theta=\theta'}^2 \ell_s(\theta) - \mathbb{E} [\nabla_{\theta=\theta'}^2 \ell_s(\theta) \mid \mathcal{F}^{s-1}] \\ &\stackrel{(a)}{=} 2 \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_j} \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_k} + 2 \left( f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1}) - Y_s \right) \frac{\partial^2 f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_j \partial \theta_k} \\ &\quad - 2 \sum_{i=1}^{|V|} p_{\hat{\theta}_{s-1}}(i) \left( \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_j} \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_k} - \left( f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1}) - f_{I_s}(\theta_*; x_i, x_j, \mathcal{F}^{s-1}) \right) \cdot \right. \\ &\quad \left. \frac{\partial^2 f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_j \partial \theta_k} \right) \\ &= -2(Y_s - f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})) \frac{\partial^2 f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_j \partial \theta_k} + 2 \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_j} \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_k} \\ &\quad - 2 \sum_{i=1}^{|V|} p_{\hat{\theta}_{s-1}}(i) \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_j} \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_k} \tag{7} \end{aligned}$$

where, (a) follows from Lemma B.4 and Lemma B.5. Plugging this equality in Equation (7) below we get

$$\begin{aligned} &\sum_{s=1}^t \nabla_{\theta=\theta'}^2 \ell_s(\theta) - \mathbb{E} [\nabla_{\theta=\theta'}^2 \ell_s(\theta) \mid \mathcal{F}^{s-1}] \\ &= \sum_{s=1}^t \left( -2(Y_s - f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})) \frac{\partial^2 f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_j \partial \theta_k} + 2 \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_j} \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_k} \right. \\ &\quad - 2 \sum_{i=1}^{|V|} p_{\hat{\theta}_{s-1}}(i) \left( \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_j} \frac{\partial f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_k} - 2 \left( f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1}) - f_{I_s}(\theta_*; x_i, x_j, \mathcal{F}^{s-1}) \right) \cdot \right. \\ &\quad \left. \left. \frac{\partial^2 f_{I_s}(\theta; x_i, x_j, \mathcal{F}^{s-1})}{\partial \theta_j \partial \theta_k} \right) \right). \end{aligned}$$

The claim of the lemma follows. □

**Lemma B.7.** *Let  $\widehat{\mathcal{L}}_t(\theta_*) = \frac{1}{t} \sum_{s=1}^t \ell_s(\theta_*)$  and  $\nabla^2 \mathcal{L}_t(\theta_*) = \frac{1}{t} \sum_{s=1}^t \nabla^2 \mathbb{E} [\ell_s(\theta_*) \mid \mathcal{F}^{s-1}]$ . Then we can bound the*

$$\mathbb{P} \left( \lambda_{\max}(\nabla^2 \widehat{\mathcal{L}}_t(\theta_*) - \nabla^2 \mathcal{L}_t(\theta_*)) > \sqrt{\frac{8C|V|^2 n^2 \lambda_1^2 c \gamma \log(dt)}{t}} \right) \leq \frac{2}{(dt)^\gamma},$$

where  $c > 0$  is a constant.

*Proof.* This lemma is different than Frostig et al. (2015); Mukherjee et al. (2022) as it requires a different concentration bound to take into account the squared loss Assumption 5.2 and the vocabulary size. Recall that  $\widehat{\mathcal{L}}_t(\boldsymbol{\theta}_*) = \frac{1}{t} \sum_{s=1}^t \ell_s(\boldsymbol{\theta}_*)$  and  $\nabla^2 \mathcal{L}_s(\boldsymbol{\theta}^*) = \nabla^2 \mathbb{E}[\ell_s(\boldsymbol{\theta}_*) | \mathcal{F}^{s-1}]$ . We define  $\nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*) = \frac{1}{t} \sum_{s=1}^t \nabla^2 \mathbb{E}[\ell_s(\boldsymbol{\theta}_*) | \mathcal{F}^{s-1}]$ . Denote,  $\mathbf{V}_s = 2 \nabla_{\boldsymbol{\theta}=\boldsymbol{\theta}_*} f_{I_s}(\boldsymbol{\theta}; x_i, x_j, \mathcal{F}^{s-1}) \nabla_{\boldsymbol{\theta}=\boldsymbol{\theta}_*} f_{I_s}(\boldsymbol{\theta}; x_i, x_j, \mathcal{F}^{s-1})^\top - 2 \sum_{i=1}^{|V|} p_{\widehat{\boldsymbol{\theta}}_{s-1}}(i) \nabla_{\boldsymbol{\theta}=\boldsymbol{\theta}_*} f_{I_s}(\boldsymbol{\theta}; x_i, x_j, \mathcal{F}^{s-1}) \nabla_{\boldsymbol{\theta}=\boldsymbol{\theta}_*} f_{I_s}(\boldsymbol{\theta}; x_i, x_j, \mathcal{F}^{s-1})^\top$ . Then we can show that,

$$\begin{aligned}
 & \mathbb{P} \left( \lambda_{\max}(\nabla^2 \widehat{\mathcal{L}}_t(\boldsymbol{\theta}_*) - \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}^*)) > \sqrt{\frac{8C^2|V|^4\eta^2\lambda_1^2c\gamma\log(dt)}{t}} \right) \\
 &= \mathbb{P} \left( \lambda_{\max} \left( \nabla_{\boldsymbol{\theta}=\boldsymbol{\theta}_*}^2 \frac{1}{t} \sum_{s=1}^t \ell_s(\boldsymbol{\theta}) - \frac{1}{t} \sum_{s=1}^t \nabla_{\boldsymbol{\theta}=\boldsymbol{\theta}_*}^2 \mathbb{E}[\ell_s(\boldsymbol{\theta}) | \mathcal{F}^{s-1}] \right) > \sqrt{\frac{8C^2|V|^4\eta^2\lambda_1^2c\gamma\log(dt)}{t}} \right) \\
 &= \mathbb{P} \left( \lambda_{\max} \left( \nabla_{\boldsymbol{\theta}=\boldsymbol{\theta}_*}^2 \frac{1}{t} \sum_{s=1}^t \left( \ell_s(\boldsymbol{\theta}) - \nabla_{\boldsymbol{\theta}=\boldsymbol{\theta}_*}^2 \mathbb{E}[\ell_s(\boldsymbol{\theta}) | \mathcal{F}^{s-1}] \right) \right) > \sqrt{\frac{8C^2|V|^4\eta^2\lambda_1^2c\gamma\log(dt)}{t}} \right) \\
 &\stackrel{(a)}{\leq} \mathbb{P} \left( \lambda_{\max} \left( \frac{C|V|^2}{t} \sum_{s=1}^t \left( Y_s - f_{I_s}(\boldsymbol{\theta}_*; x_i, x_j, \mathcal{F}^{s-1}) \right) \nabla_{\boldsymbol{\theta}=\boldsymbol{\theta}_*}^2 f_{I_s}(\boldsymbol{\theta}_*; x_i, x_j, \mathcal{F}^{s-1}) \right. \right. \\
 &\quad \left. \left. + \frac{C|V|^2}{t} \sum_{s=1}^t \mathbf{V}_s \right) > \sqrt{\frac{8C^2|V|^4\eta^2\lambda_1^2c\gamma\log(dt)}{t}} \right) \\
 &\leq \mathbb{P} \left( \lambda_{\max} \left( \frac{1}{t} \sum_{s=1}^t -2 \left( Y_s - f_{I_s}(\boldsymbol{\theta}_*; x_i, x_j, \mathcal{F}^{s-1}) \right) \nabla_{\boldsymbol{\theta}=\boldsymbol{\theta}_*}^2 f_{I_s}(\boldsymbol{\theta}_*; x_i, x_j, \mathcal{F}^{s-1}) \right) > \frac{1}{2} \sqrt{\frac{8\eta^2\lambda_1^2c\gamma\log(dt)}{t}} \right) \\
 &\quad + \mathbb{P} \left( \lambda_{\max} \left( \frac{1}{t} \sum_{s=1}^t \mathbf{V}_s \right) > \frac{1}{2} \sqrt{\frac{8\eta^2\lambda_1^2c\gamma\log(dt)}{t}} \right) \\
 &\stackrel{(b)}{\leq} \mathbb{P} \left( \frac{1}{t} \sum_{s=1}^t -2 \left( Y_s - f_{I_s}(\boldsymbol{\theta}_*; x_i, x_j, \mathcal{F}^{s-1}) \right) \lambda_{\max} \left( \nabla_{\boldsymbol{\theta}=\boldsymbol{\theta}_*}^2 f_{I_s}(\boldsymbol{\theta}_*; x_i, x_j, \mathcal{F}^{s-1}) \right) > \frac{1}{2} \sqrt{\frac{8\eta^2\lambda_1^2c\gamma\log(dt)}{t}} \right) \\
 &\quad + \mathbb{P} \left( \frac{1}{t} \sum_{s=1}^t \lambda_{\max}(\mathbf{V}_s) > \frac{1}{2} \sqrt{\frac{8\eta^2\lambda_1^2c\gamma\log(dt)}{t}} \right) \\
 &\stackrel{(c)}{\leq} 2 \exp \left( -\frac{t^2 8\eta^2\lambda_1^2c\gamma\log(dt)}{4t} \cdot \frac{1}{2tc\eta^2\lambda_1^2} \right) \stackrel{(d)}{\leq} 2 \left( \frac{1}{dt} \right)^\gamma. \tag{8}
 \end{aligned}$$

where, (a) follows from substituting the value of  $\nabla_{\boldsymbol{\theta}=\boldsymbol{\theta}_*}^2 \ell_s(\boldsymbol{\theta}) - \nabla_{\boldsymbol{\theta}=\boldsymbol{\theta}_*}^2 \mathbb{E}[\ell_s(\boldsymbol{\theta}) | \mathcal{F}^{s-1}]$  from Lemma B.6, and (b) follows by triangle inequality, (c) follows by using two concentration inequalities stated below, and (d) follows by simplifying the equations.

Denote  $Q_s = -2 \left( Y_s - f_{I_s}(\boldsymbol{\theta}_*; x_i, x_j, \mathcal{F}^{s-1}) \right) \lambda_{\max} \left( \nabla_{\boldsymbol{\theta}=\boldsymbol{\theta}_*}^2 f_{I_s}(\boldsymbol{\theta}_*; x_i, x_j, \mathcal{F}^{s-1}) \right)$ . Also note that

$\lambda_{\max} \left( \nabla_{\theta=\theta_*}^2 f_{I_s}(\theta_*; x_i, x_j, \mathcal{F}^{s-1}) \right) \leq \lambda_1$  for all time  $s$  using Assumption B.1.

$$\begin{aligned}
 & \mathbb{P} \left( \sum_{s=1}^t -2 \left( Y_s - f_{I_s}(\theta_*; x_i, x_j, \mathcal{F}^{s-1}) \right) \lambda_{\max} \left( \nabla_{\theta=\theta_*}^2 f_{I_s}(\theta_*; x_i, x_j, \mathcal{F}^{s-1}) \right) \geq \epsilon \right) = \mathbb{P} \left( - \sum_{s=1}^t Q_s \geq \epsilon \right) \\
 &= \mathbb{P} \left( e^{-\lambda \sum_{s=1}^t Q_s} \geq e^{\lambda \epsilon} \right) \stackrel{(a)}{\leq} e^{-\lambda \epsilon} \mathbb{E} \left[ e^{-\lambda \sum_{s=1}^t Q_s} \right] = e^{-\lambda \epsilon} \mathbb{E} \left[ \mathbb{E} \left[ e^{-\lambda \sum_{s=1}^t Q_s} \mid \hat{\theta}_{t-1} \right] \right] \\
 &\stackrel{(b)}{=} e^{-\lambda \epsilon} \mathbb{E} \left[ \mathbb{E} \left[ e^{-\lambda Q_t} \mid \hat{\theta}_{t-1} \right] \mathbb{E} \left[ e^{-\lambda \sum_{s=1}^{t-1} Q_s} \mid \hat{\theta}_{t-1} \right] \right] \\
 &\leq e^{-\lambda \epsilon} \mathbb{E} \left[ \exp \left( 2\lambda^2 \lambda_1^2 \eta^2 \right) \mathbb{E} \left[ e^{-\lambda \sum_{s=1}^{t-1} Q_s} \mid \hat{\theta}_{t-1} \right] \right] \\
 &= e^{-\lambda \epsilon} e^{2\lambda^2 \eta^2 \lambda_1^2} \mathbb{E} \left[ e^{-\lambda \sum_{s=1}^{t-1} Q_s} \right] \\
 &\vdots \\
 &\stackrel{(c)}{\leq} e^{-\lambda \epsilon} e^{2\lambda^2 t \eta^2 \lambda_1^2} \stackrel{(d)}{\leq} \exp \left( -\frac{2\epsilon^2}{t \eta^2 \lambda_1^2} \right).
 \end{aligned}$$

where (a) follows by Markov's inequality, (b) follows as  $Q_s$  is conditionally independent given  $\hat{\theta}_{s-1}$ , (c) follows by unpacking the term for  $t$  times and (d) follows by taking  $\lambda = \epsilon/4t\lambda_1^2\eta^2$  where  $\lambda_1$  is defined in Assumption 5.1. Next we bound the second term of (8) below.

$$\begin{aligned}
 \mathbb{P} \left( \sum_{s=1}^t \lambda_{\max}(\mathbf{V}_s) \geq \epsilon \right) &= \mathbb{P} \left( \lambda \sum_{s=1}^t \lambda_{\max}(\mathbf{V}_s) \geq \lambda \epsilon \right) = \mathbb{P} \left( e^{\lambda \sum_{s=1}^t \lambda_{\max}(\mathbf{V}_s)} \geq e^{\lambda \epsilon} \right) \stackrel{(a)}{\leq} e^{-\lambda \epsilon} \mathbb{E} \left[ e^{\lambda \sum_{s=1}^t \lambda_{\max}(\mathbf{V}_s)} \right] \\
 &= e^{-\lambda \epsilon} \mathbb{E} \left[ \mathbb{E} \left[ e^{\lambda \sum_{s=1}^t \lambda_{\max}(\mathbf{V}_s)} \mid \hat{\theta}_{t-1} \right] \right] \\
 &\stackrel{(b)}{=} e^{-\lambda \epsilon} \mathbb{E} \left[ \mathbb{E} \left[ e^{\lambda \lambda_{\max}(\mathbf{V}_t)} \mid \hat{\theta}_{t-1} \right] \mathbb{E} \left[ e^{\lambda \sum_{s=1}^{t-1} \lambda_{\max}(\mathbf{V}_s)} \mid \hat{\theta}_{t-1} \right] \right] \\
 &\stackrel{(c)}{\leq} e^{-\lambda \epsilon} \mathbb{E} \left[ \exp \left( 2c\lambda^2 \lambda_1^2 \eta^2 \right) \mathbb{E} \left[ e^{\lambda \sum_{s=1}^{t-1} \lambda_{\max}(\mathbf{V}_s)} \mid \hat{\theta}_{t-1} \right] \right] \\
 &= e^{-\lambda \epsilon} e^{2c\lambda^2 \eta^2 \lambda_1^2} \mathbb{E} \left[ e^{\lambda \sum_{s=1}^{t-1} \lambda_{\max}(\mathbf{V}_s)} \right] \\
 &\vdots \\
 &\stackrel{(d)}{\leq} e^{-\lambda \epsilon} e^{2c\lambda^2 t \eta^2 \lambda_1^2} \stackrel{(e)}{\leq} \exp \left( -\frac{2\epsilon^2}{t c \eta^2 \lambda_1^2} \right)
 \end{aligned}$$

where (a) follows by Markov's inequality, (b) follows as  $\lambda_{\max}(\mathbf{V}_s)$  is conditionally independent given  $\hat{\theta}_{s-1}$ . In the inequality (c) using the always valid upper bound of  $2\lambda_1$ , we have that  $\mathbb{E}[\lambda_{\max}(\mathbf{V}_t)] \leq 2\lambda_1$ . So the term in inequality (c) will become  $e^{-\lambda \epsilon} e^{2\lambda^2 t \eta^2 \lambda_1^2 + 4t\lambda \lambda_1}$ . Hence, we can upper bound the inequality (c) by a constant  $c > 0$  such that we have  $\mathbb{E}[e^{\lambda \lambda_{\max}(\mathbf{V}_t)} \mid \hat{\theta}_{t-1}] \leq e^{2\lambda^2 \lambda_1^2 \eta^2} e^{2\lambda \times 2\lambda_1} = \exp(2\lambda^2 \lambda_1^2 \eta^2 + 4\lambda \lambda_1) \leq \exp(2c\lambda^2 \lambda_1^2 \eta^2)$ . The inequality (d) follows by unpacking the term for  $t$  times and (e) follows by taking  $\lambda = \epsilon/4t\lambda_1^2\eta^2$  and  $\lambda_1$  defined in Assumption 5.1.  $\square$

**Lemma B.8.** Let  $\hat{\theta}_t - \theta_* = \left( \nabla^2 \hat{\mathcal{L}}_t(\tilde{\theta}_t) \right)^{-1} \nabla \hat{\mathcal{L}}_t(\theta_*)$  where  $\tilde{\theta}_t$  is between  $\hat{\theta}_t$  and  $\theta_*$ . Then we can show that

$$\left\| \hat{\theta}_t - \theta_* \right\|_{\nabla^2 \mathcal{L}_t(\theta_*)} \leq \left\| \left( \nabla^2 \mathcal{L}_t(\theta_*) \right)^{1/2} \left( \nabla^2 \hat{\mathcal{L}}_t(\tilde{\theta}_t) \right)^{-1} \left( \nabla^2 \mathcal{L}_t(\theta_*) \right)^{1/2} \right\| \left\| \nabla \hat{\mathcal{L}}_t(\theta_*) \right\|_{\left( \nabla^2 \mathcal{L}_t(\theta_*) \right)^{-1}}.$$

*Proof.* We begin with the definition of  $\|\hat{\theta}_t - \theta_*\|_{\nabla^2 \mathcal{L}_t(\theta_*)}$  as follows:

$$\begin{aligned}
 \|\hat{\theta}_t - \theta_*\|_{\nabla^2 \mathcal{L}_t(\theta_*)} &\stackrel{(a)}{=} \sqrt{(\hat{\theta}_t - \theta_*)^T \nabla^2 \mathcal{L}_t(\theta_*) (\hat{\theta}_t - \theta_*)} \\
 &\stackrel{(b)}{=} \sqrt{\left( (\nabla^2 \hat{\mathcal{L}}_t(\tilde{\theta}_t))^{-1} \nabla \hat{\mathcal{L}}_t(\theta_*) \right)^T \nabla^2 \mathcal{L}_t(\theta_*) \left( (\nabla^2 \hat{\mathcal{L}}_t(\tilde{\theta}_t))^{-1} \nabla \hat{\mathcal{L}}_t(\theta_*) \right)} \\
 &\stackrel{(c)}{\leq} \left\| \nabla^2 \mathcal{L}_t(\theta_*)^{1/2} \left( \nabla^2 \hat{\mathcal{L}}_t(\tilde{\theta}_t) \right)^{-1} \nabla^2 \mathcal{L}_t(\theta_*)^{1/2} \right\| \sqrt{\left( \nabla \hat{\mathcal{L}}_t(\theta_*)^T (\nabla^2 \mathcal{L}_t(\theta_*))^{-1} \nabla \hat{\mathcal{L}}_t(\theta_*) \right)} \\
 &= \left\| \left( \nabla^2 \mathcal{L}_t(\theta_*) \right)^{1/2} \left( \nabla^2 \hat{\mathcal{L}}_t(\tilde{\theta}_t) \right)^{-1} \left( \nabla^2 \mathcal{L}_t(\theta_*) \right)^{1/2} \right\| \left\| \nabla \hat{\mathcal{L}}_t(\theta_*) \right\|_{(\nabla^2 \mathcal{L}_t(\theta_*))^{-1}}.
 \end{aligned}$$

where, (a) follows as  $\|x\|_M = \sqrt{x^T M x}$ , (b) follows as  $\|\hat{\theta}_t - \theta_*\|_{\nabla^2 \mathcal{L}_t(\theta_*)} = \left( \nabla^2 \hat{\mathcal{L}}_t(\tilde{\theta}_t) \right)^{-1} \nabla \hat{\mathcal{L}}_t(\theta_*)$ , and (c) follows from Cauchy Schwarz inequality. The claim of the lemma follows.  $\square$

*Remark B.9.* The proof of Theorem 1 consists of several steps. In the first step we relate  $\nabla^2 \hat{\mathcal{L}}_t(\theta)$  to  $\nabla^2 \mathcal{L}_t(\theta_*)$  for any  $\theta$  in a ball  $\mathcal{B}$  around  $\theta_*$ . The ball  $\mathcal{B}$  is assumed in Assumption B.1 to be a neighborhood where  $\nabla^2 \ell_s(\theta)$  satisfies a Lipschitz property. Assumption B.1 in Appendix B are standard and have also been made by Frostig et al. (2015); Chaudhuri et al. (2015); Mukherjee et al. (2022). Using Assumption 5.1 and Assumption B.1, we can show that for a large enough sequences of tokens  $t$  stated in Theorem 1 we have the following: (1)  $\nabla^2 \mathcal{L}_t(\theta_*)$  lies between in the positive semidefinite order by scaled multiples of  $\nabla^2 \hat{\mathcal{L}}_t(\theta)$  for any  $\theta \in \mathcal{B}$ , and (2) the empirical error minimizing  $\hat{\theta}_t$  is in the ball  $\mathcal{B}$  with probability  $1 - 1/t^\gamma$ , which is the good event  $\mathcal{E}$ . Then we use a Taylor series expansion around  $\hat{\theta}_t$  and using the fact that  $\nabla \hat{\mathcal{L}}_t(\hat{\theta}_t) = 0$  along with the relation between  $\nabla^2 \hat{\mathcal{L}}_t(\theta)$  and  $\nabla^2 \mathcal{L}_t(\theta_*)$ , we can obtain an upper bound to  $\|\hat{\theta}_t - \theta_*\|_{\nabla^2 \mathcal{L}_t(\theta_*)}$  in terms of  $\|\nabla \hat{\mathcal{L}}_t(\theta_*)\|_{(\nabla^2 \mathcal{L}_t(\theta_*))^{-1}}$  that can be shown to be decreasing with  $t$ . Further,  $\|\hat{\theta}_t - \theta_*\|_{\nabla^2 \mathcal{L}_t(\theta_*)}$  can also be used to obtain an upper bound to  $\mathcal{L}_t(\hat{\theta}_t) - \mathcal{L}_t(\theta_*)$  using a Taylor series expansion. Finally we can bound  $\mathbb{E}[\mathcal{L}_t(\hat{\theta}_t) - \mathcal{L}_t(\theta^*)] = \mathbb{E}[(\mathcal{L}_t(\hat{\theta}_t) - \mathcal{L}_t(\theta^*))I(\mathcal{E})] + \mathbb{E}[(\mathcal{L}_t(\hat{\theta}_t) - \mathcal{L}_t(\theta^*))I(\mathcal{E}^C)]$  where  $I(\cdot)$  is the indicator. Since  $\mathbb{P}(\mathcal{E}^C) \leq 1/t^\gamma$ , the second term can be bounded as  $\max_{\theta \in \Theta} (\mathcal{L}_t(\theta) - \mathcal{L}_t(\theta^*)) / t^\gamma$ , while the first term simplifies to  $(1 + \rho_t)\sigma_t^2/t$ .

**Theorem 1. (Restatement of main theorem)** Suppose  $\ell_1(\theta), \ell_2(\theta), \dots, \ell_t(\theta) : \mathbb{R}^{|V|} \rightarrow \mathbb{R}$  are loss functions from a distribution that satisfies Assumptions 5.1, 5.2, and B.1. Define  $\mathcal{L}_t(\theta) = \frac{1}{t} \sum_{s=1}^t \mathbb{E}_{x_s \sim \mathbf{p}_{\hat{\theta}_{s-1}}} [\ell_s(\theta) | \mathcal{F}^{s-1}]$  where,

$\hat{\theta}_t = \operatorname{argmin}_{\theta \in \Theta} \sum_{s=1}^t \ell_s(\theta)$ . If  $t$  is large enough such that  $\frac{\gamma \log(dt)}{t} \leq c' \min \left\{ \frac{1}{C_1 C_2 |V|^4}, \frac{\max_{\theta \in \Theta} (\mathcal{L}_t(\theta) - \mathcal{L}_t(\theta_*))}{C_2} \right\}$  then for a constant  $\gamma \geq 2$ , universal constants  $C_1, C_2, c'$ , we can show that

$$(1 - \rho_t) \frac{\sigma_t^2}{t} - \frac{C_1^2}{t^{\gamma/2}} \leq \mathbb{E} [\mathcal{L}_t(\hat{\theta}_t) - \mathcal{L}_t(\theta_*)] \leq (1 + \rho_t) \frac{\sigma_t^2}{t} + \frac{\max_{\theta \in \Theta} (\mathcal{L}_t(\theta) - \mathcal{L}_t(\theta_*))}{t^\gamma},$$

where  $\sigma_t^2 := \mathbb{E} \left[ \frac{1}{2} \left\| \nabla \hat{\mathcal{L}}_t(\theta_*) \right\|_{(\nabla^2 \mathcal{L}_t(\theta_*))^{-1}}^2 \right]$ , and  $\rho_t := (C_1 C_2 + 2\eta^2 \lambda_1^2) \sqrt{\frac{\gamma \log(dt)}{t}}$ .

*Proof. Step 1:* We first bound the  $\left\| \nabla^2 \hat{\mathcal{L}}_t(\theta) - \nabla^2 \mathcal{L}_t(\theta_*) \right\|_*$  as follows

$$\begin{aligned}
 \left\| \nabla^2 \hat{\mathcal{L}}_t(\theta) - \nabla^2 \mathcal{L}_t(\theta_*) \right\|_* &\stackrel{(a)}{\leq} \left\| \nabla^2 \hat{\mathcal{L}}_t(\theta) - \nabla^2 \hat{\mathcal{L}}_t(\theta_*) \right\|_* + \left\| \nabla^2 \hat{\mathcal{L}}_t(\theta_*) - \nabla^2 \mathcal{L}_t(\theta_*) \right\|_* \\
 &\stackrel{(b)}{\leq} C_1 \|\theta - \theta_*\|_{\nabla^2 \mathcal{L}_t(\theta_*)} + \sqrt{\frac{8C^2 |V|^4 \eta^2 \lambda_1^2 c \gamma \log(dt)}{t}}
 \end{aligned} \tag{9}$$

where, (a) follows from triangle inequality, and (b) is due to Assumption B.1.3.d and Lemma B.7.

**Step 2 (Approximation of  $\nabla^2 \mathcal{L}_t(\theta_*)$ ):** By choosing a sufficiently smaller ball  $\mathcal{B}_1$  of radius of  $\min \left\{ 1/(10C_1), \text{diameter}(\mathcal{B}) \right\}$ , the first term in (9) can be made small for  $\theta \in \mathcal{B}_1$ . Also, for sufficiently large  $t$ , the second term in (9) can be made

arbitrarily small (smaller than  $1/10$ ), which occurs if  $\sqrt{\frac{\gamma \log(dt)}{t}} \leq \frac{c'}{\sqrt{2C^2|V|^4\eta^2\lambda_1^2}}$ . Hence for large  $t$  and  $\theta \in \mathcal{B}_1$  we have

$$\frac{1}{2}\nabla^2\widehat{\mathcal{L}}_t(\theta) \preceq \nabla^2\mathcal{L}_t(\theta_*) \preceq 2\nabla^2\widehat{\mathcal{L}}_t(\theta) \quad (10)$$

**Step 3 (Show  $\widehat{\theta}_t$  in  $\mathcal{B}_1$ ):** Fix a  $\widetilde{\theta}$  between  $\theta$  and  $\theta_*$  in  $\mathcal{B}_1$ . Apply Taylor's series approximation

$$\widehat{\mathcal{L}}_t(\theta) = \widehat{\mathcal{L}}_t(\theta_*) + \nabla\widehat{\mathcal{L}}_t(\theta_*)^\top(\theta - \theta_*) + \frac{1}{2}(\theta - \theta_*)^\top\nabla^2\widehat{\mathcal{L}}_t(\widetilde{\theta})(\theta - \theta_*)$$

We can further reduce this as follows:

$$\begin{aligned} \widehat{\mathcal{L}}_t(\theta) - \widehat{\mathcal{L}}_t(\theta_*) &\stackrel{(a)}{=} \nabla\widehat{\mathcal{L}}_t(\theta_*)^\top(\theta - \theta_*) + \frac{1}{2}\|\theta - \theta_*\|_{\nabla^2\widehat{\mathcal{L}}_t(\widetilde{\theta})}^2 \\ &\stackrel{(b)}{\geq} \nabla\widehat{\mathcal{L}}_t(\theta_*)^\top(\theta - \theta_*) + \frac{1}{4}\|\theta - \theta_*\|_{\nabla^2\mathcal{L}_t(\theta_*)}^2 \\ &\geq -\|\theta - \theta_*\|_{\nabla^2\mathcal{L}_t(\theta_*)}\left\|\nabla\widehat{\mathcal{L}}_t(\theta_*)\right\|_{(\nabla^2\mathcal{L}_t(\theta_*))^{-1}} + \frac{1}{4}\left(\|\theta - \theta_*\|_{\nabla^2\mathcal{L}_t(\theta_*)}\right)^\top\left(\|\theta - \theta_*\|_{\nabla^2\mathcal{L}_t(\theta_*)}\right) \\ &= \|\theta - \theta_*\|_{\nabla^2\mathcal{L}_t(\theta_*)}\left(-\left\|\nabla\widehat{\mathcal{L}}_t(\theta_*)\right\|_{(\nabla^2\mathcal{L}_t(\theta_*))^{-1}} + \frac{1}{4}\|\theta - \theta_*\|_{\nabla^2\mathcal{L}_t(\theta_*)}\right) \end{aligned} \quad (11)$$

where, (a) follows as  $\|\theta - \theta_*\|_{\nabla^2\widehat{\mathcal{L}}_t(\widetilde{\theta})}^2 := (\theta - \theta_*)^\top\nabla^2\widehat{\mathcal{L}}_t(\widetilde{\theta})(\theta - \theta_*)$ , and (b) follows as  $\widetilde{\theta}$  is in between  $\theta$  and  $\theta_*$  and then using (10). Note that in (11) if the right hand side is positive for some  $\theta \in \mathcal{B}_1$ , then  $\theta$  is not a local minimum. Also, since  $\left\|\nabla\widehat{\mathcal{L}}_t(\theta_*)\right\| \rightarrow 0$ , for a sufficiently small value of  $\left\|\nabla\widehat{\mathcal{L}}_t(\theta_*)\right\|$ , all points on the boundary of  $\mathcal{B}_1$  will have values greater than that of  $\theta_*$ . Hence, we must have a local minimum of  $\widehat{\mathcal{L}}_t(\theta)$  that is strictly inside  $\mathcal{B}_1$  (for  $t$  large enough). We can ensure this local minimum condition is achieved by choosing an  $t$  large enough so that  $\sqrt{\frac{\gamma \log(dt)}{t}} \leq c' \min\left\{\frac{1}{C_1C_2}, \frac{\text{diameter}(\mathcal{B})}{C_2}\right\}$ , using Lemma B.3 (and our bound on the diameter of  $\mathcal{B}_1$ ). By convexity, we have that this is the global minimum,  $\widehat{\theta}_t$ , and so  $\widehat{\theta}_t \in \mathcal{B}_1$  for  $t$  large enough. We will assume now that  $t$  is this large from here on.

**Step 4 (Bound  $\|\widehat{\theta}_t - \theta_*\|_{\nabla^2\mathcal{L}_t(\theta_*)}$ ):** For the  $\widehat{\theta}(t)$  that minimizes the sum of squared errors,  $0 = \nabla\widehat{\mathcal{L}}_t(\widehat{\theta}_t)$ . Again, using Taylor's theorem if  $\widehat{\theta}_t$  is an interior point, we have:

$$0 = \nabla\widehat{\mathcal{L}}_t(\widehat{\theta}_t) = \nabla\widehat{\mathcal{L}}_t(\theta_*) + \nabla^2\widehat{\mathcal{L}}_t(\widetilde{\theta}_t)(\widehat{\theta}_t - \theta_*) \quad (12)$$

for some  $\widetilde{\theta}_t$  between  $\theta_*$  and  $\widehat{\theta}_t$ . Now observe that  $\widetilde{\theta}_t$  is in  $\mathcal{B}_1$  (since, for  $t$  large enough,  $\widehat{\theta}_t \in \mathcal{B}_1$ ). Thus it follows from (12) that,

$$\widehat{\theta}_t - \theta_* = \left(\nabla^2\widehat{\mathcal{L}}_t(\widetilde{\theta}_t)\right)^{-1}\nabla\widehat{\mathcal{L}}_t(\theta_*) \quad (13)$$

where the invertibility is guaranteed by (10) and the positive definiteness of  $\nabla^2\mathcal{L}_t(\theta_*)$  (by Assumption B.1 (3c)). We finally derive the upper bound to  $\|\widehat{\theta}_t - \theta_*\|_{\nabla^2\mathcal{L}_t(\theta_*)}$  as follows

$$\begin{aligned} \|\widehat{\theta}_t - \theta_*\|_{\nabla^2\mathcal{L}_t(\theta_*)} &\stackrel{(a)}{\leq} \left\|\left(\nabla^2\mathcal{L}_t(\theta_*)\right)^{1/2}\left(\nabla^2\widehat{\mathcal{L}}_t(\widetilde{\theta}_t)\right)^{-1}\left(\nabla^2\mathcal{L}_t(\theta_*)\right)^{1/2}\right\|\left\|\nabla\widehat{\mathcal{L}}_t(\theta_*)\right\|_{(\nabla^2\mathcal{L}_t(\theta_*))^{-1}} \\ &\stackrel{(b)}{\leq} cC_2\sqrt{\frac{\gamma \log(dt)}{t}} \end{aligned} \quad (14)$$

where (a) follows from Lemma B.8, and (b) from Lemma B.3, (11), and  $c$  is some universal constant.

**Step 5 (Introducing  $\widetilde{z}_t$ ):** Fix a  $\widetilde{z}_t$  between  $\theta_*$  and  $\widehat{\theta}_t$ . Apply Taylor's series

$$\mathcal{L}_t(\widehat{\theta}_t) - \mathcal{L}_t(\theta_*) = \frac{1}{2}\left(\widehat{\theta}_t - \theta_*\right)^\top\nabla^2\mathcal{L}_t(\widetilde{z}_t)(\widehat{\theta}_t - \theta_*) \quad (15)$$

Now note that both  $\tilde{\boldsymbol{\theta}}_t$  and  $\tilde{\mathbf{z}}_t$  are between  $\hat{\boldsymbol{\theta}}_t$  and  $\boldsymbol{\theta}_*$ , which implies  $\tilde{\boldsymbol{\theta}}_t \rightarrow \boldsymbol{\theta}_*$  and  $\tilde{\mathbf{z}}_t \rightarrow \boldsymbol{\theta}_*$  since  $\hat{\boldsymbol{\theta}}_t \rightarrow \boldsymbol{\theta}_*$ . By (9) and (14) and applying the concentration inequalities give us

$$\left\| \nabla^2 \hat{\mathcal{L}}_t(\tilde{\boldsymbol{\theta}}_t) - \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*) \right\|_* \leq \rho_t \quad (16)$$

$$\left\| \nabla^2 \mathcal{L}_t(\tilde{\mathbf{z}}_t) - \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*) \right\|_* \leq C_1 \|\tilde{\mathbf{z}}_t - \boldsymbol{\theta}_*\|_{\nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*)} \leq \rho_t \quad (17)$$

where  $\rho_t = c(C_1 C_2 + 2\eta^2 \lambda_1^2) \sqrt{\frac{\gamma \log(dt)}{t}}$ .

**Step 6 (Define  $\mathbf{M}_{1,t}$  and  $\mathbf{M}_{2,t}$ ):** It follows from the inequality (16) that

$$\begin{aligned} \nabla^2 \hat{\mathcal{L}}_t(\tilde{\boldsymbol{\theta}}_t) &\preceq (1 + \rho_t) \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*) \implies \nabla^2 \hat{\mathcal{L}}_t(\tilde{\boldsymbol{\theta}}_t) - \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*) \preceq \rho_t \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*) \\ &\implies \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*)^{-1/2} (\hat{\mathcal{L}}_t(\tilde{\boldsymbol{\theta}}_t) - \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*)) \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*)^{-1/2} \preceq \rho_t I \\ &\implies \|\nabla^2 \hat{\mathcal{L}}_t(\tilde{\boldsymbol{\theta}}_t) - \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*)\|_* \leq \rho_t. \end{aligned}$$

Then we can use the inequalities (16) and (17) to show that

$$\begin{aligned} (1 - \rho_t) \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*) &\preceq \nabla^2 \hat{\mathcal{L}}_t(\tilde{\boldsymbol{\theta}}_t) \preceq (1 + \rho_t) \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*) \\ (1 - \rho_t) \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*) &\preceq \nabla^2 \mathcal{L}_t(\tilde{\mathbf{z}}_t) \preceq (1 + \rho_t) \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*). \end{aligned}$$

Now we define the two quantities  $\mathbf{M}_{1,t}$  and  $\mathbf{M}_{2,t}$  as follows:

$$\begin{aligned} \mathbf{M}_{1,t} &:= \left( \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*) \right)^{1/2} \left( \nabla^2 \hat{\mathcal{L}}_t(\tilde{\boldsymbol{\theta}}_t) \right)^{-1} \left( \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*) \right)^{1/2} \\ \mathbf{M}_{2,t} &:= \left( \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*) \right)^{-1/2} \nabla^2 \mathcal{L}_t(\tilde{\mathbf{z}}_t) \left( \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*) \right)^{-1/2}. \end{aligned}$$

**Step 7 (Lower bound  $\mathcal{L}_t(\hat{\boldsymbol{\theta}}_t) - \mathcal{L}_t(\boldsymbol{\theta}_*)$ ):** Now for the lower bound it follows from Equation (15) that

$$\begin{aligned} \mathcal{L}_t(\hat{\boldsymbol{\theta}}_t) - \mathcal{L}_t(\boldsymbol{\theta}_*) &= \frac{1}{2} \left( \hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta}_* \right)^\top \nabla^2 \mathcal{L}_t(\tilde{\mathbf{z}}_t) \left( \hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta}_* \right) \\ &= \frac{1}{2} \left( \hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta}_* \right)^\top \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*)^{\frac{1}{2}} \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*)^{-\frac{1}{2}} \nabla^2 \mathcal{L}_t(\tilde{\mathbf{z}}_t) \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*)^{-\frac{1}{2}} \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*)^{\frac{1}{2}} \left( \hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta}_* \right) \\ &\stackrel{(a)}{=} \frac{1}{2} \mathbf{u}^T \mathbf{M}_{2,t} \mathbf{u} \end{aligned}$$

where, in (a) we define the vector  $\mathbf{u} := \left( \hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta}_* \right)^\top \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*)^{\frac{1}{2}}$ . Now observe from the definition of and then using the min-max theorem we can show that

$$\begin{aligned} \mathcal{L}_t(\hat{\boldsymbol{\theta}}_t) - \mathcal{L}_t(\boldsymbol{\theta}_*) &\geq \frac{1}{2} \lambda_{\min}(\mathbf{M}_{2,t}) \mathbf{u}^T \mathbf{u} \\ &= \frac{1}{2} \lambda_{\min}(\mathbf{M}_{2,t}) \left\| \hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta}_* \right\|_{\nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*)}^2 \\ &= \frac{1}{2} \lambda_{\min}(\mathbf{M}_{2,t}) \left\| \nabla^2 \hat{\mathcal{L}}_t(\tilde{\boldsymbol{\theta}}_t) \left( \hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta}_* \right) \right\|_{\left( \nabla^2 \hat{\mathcal{L}}_t(\tilde{\boldsymbol{\theta}}_t) \right)^{-1} \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*) \left( \nabla^2 \hat{\mathcal{L}}_t(\tilde{\boldsymbol{\theta}}_t) \right)^{-1}}^2 \\ &\geq \frac{1}{2} \left( \lambda_{\min}(\mathbf{M}_{1,t}) \right)^2 \lambda_{\min}(\mathbf{M}_{2,t}) \left\| \nabla^2 \hat{\mathcal{L}}_t(\tilde{\boldsymbol{\theta}}_t) \left( \hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta}_* \right) \right\|_{\left( \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*) \right)^{-1}}^2 \\ &\stackrel{(a)}{=} \frac{1}{2} \left( \lambda_{\min}(\mathbf{M}_{1,t}) \right)^2 \lambda_{\min}(\mathbf{M}_{2,t}) \left\| \nabla \hat{\mathcal{L}}_t(\boldsymbol{\theta}_*) \right\|_{\left( \nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*) \right)^{-1}}^2 \end{aligned}$$

where, in (a) we use the Equation (13).

**Step 8:** Define  $I(\mathcal{E})$  as the indicator that the desired previous events hold, which we can ensure with probability greater than  $1 - 2 \left( \frac{1}{dt} \right)^\gamma$ . Then we can show that:

$$\begin{aligned}
 \mathbb{E} \left[ \mathcal{L}_t(\hat{\theta}_t) - \mathcal{L}_t(\theta_*) \right] &\geq \mathbb{E} \left[ \left( \mathcal{L}_t(\hat{\theta}_t) - \mathcal{L}_t(\theta_*) \right) I(\mathcal{E}) \right] \\
 &\geq \frac{1}{2} \mathbb{E} \left[ \left( \lambda_{\min}(\mathbf{M}_{1,t}) \right)^2 \lambda_{\min}(\mathbf{M}_{2,t}) \left\| \nabla \hat{\mathcal{L}}_t(\theta_*) \right\|_{(\nabla^2 \mathcal{L}_t(\theta_*))^{-1}}^2 I(\mathcal{E}) \right] \\
 &\geq (1 - c' \rho_t) \frac{1}{2} \mathbb{E} \left[ \left\| \nabla \hat{\mathcal{L}}_t(\theta_*) \right\|_{(\nabla^2 \mathcal{L}_t(\theta_*))^{-1}}^2 I(\mathcal{E}) \right] \\
 &= (1 - c' \rho_t) \frac{1}{2} \mathbb{E} \left[ \left\| \nabla \hat{\mathcal{L}}_t(\theta_*) \right\|_{(\nabla^2 \mathcal{L}_t(\theta_*))^{-1}}^2 (1 - I(\text{not } \mathcal{E})) \right] \\
 &\stackrel{(a)}{=} (1 - c' \rho_t) \left( \sigma_t^2 - \frac{1}{2} \mathbb{E} \left[ \left\| \nabla \hat{\mathcal{L}}_t(\theta_*) \right\|_{(\nabla^2 \mathcal{L}_t(\theta_*))^{-1}}^2 I(\text{not } \mathcal{E}) \right] \right) \\
 &\geq (1 - c' \rho_t) \sigma_t^2 - \mathbb{E} \left[ \left\| \nabla \hat{\mathcal{L}}_t(\theta_*) \right\|_{(\nabla^2 \mathcal{L}_t(\theta_*))^{-1}}^2 I(\text{not } \mathcal{E}) \right]
 \end{aligned}$$

where, in (a) we have  $\sigma_t^2 := \left\| \nabla \hat{\mathcal{L}}_t(\theta_*) \right\|_{(\nabla^2 \mathcal{L}_t(\theta_*))^{-1}}^2$ , and  $c'$  is an universal constant.

**Step 9:** Define the random variable  $Z = \left\| \nabla \hat{\mathcal{L}}_t(\theta_*) \right\|_{(\nabla^2 \mathcal{L}_t(\theta_*))^{-1}}$ . With a failure event probability of less than  $2 \left( \frac{1}{dt} \right)^\gamma$  for any  $z_0$ , we have:

$$\begin{aligned}
 \mathbb{E} \left[ Z^2 I(\text{not } \mathcal{E}) \right] &= \mathbb{E} \left[ Z^2 I(\text{not } \mathcal{E}) I(Z^2 < z_0) \right] + \mathbb{E} \left[ Z^2 I(\text{not } \mathcal{E}) I(Z^2 \geq z_0) \right] \\
 &\leq z_0 \mathbb{E}[I(\text{not } \mathcal{E})] + \mathbb{E} \left[ Z^2 I(Z^2 \geq z_0) \right] \\
 &\leq \frac{z_0}{2t^\gamma} + \mathbb{E} \left[ Z^2 \frac{Z^2}{z_0} \right] \\
 &\leq \frac{z_0}{2t^\gamma} + \frac{\mathbb{E}[Z^4]}{z_0} \\
 &\leq \frac{\sqrt{\mathbb{E}[Z^4]}}{t^{\gamma/2}}
 \end{aligned}$$

where  $z_0 = t^{\gamma/2} \sqrt{\mathbb{E}[Z^4]}$ .

**Step 10 (Upper Bound):** For an upper bound we have that:

$$\begin{aligned}
 \mathbb{E} \left[ \mathcal{L}_t(\hat{\theta}_t) - \mathcal{L}_t(\theta_*) \right] &= \mathbb{E} \left[ \left( \mathcal{L}_t(\hat{\theta}_t) - \mathcal{L}_t(\theta_*) \right) I(\mathcal{E}) \right] + \mathbb{E} \left[ \left( \mathcal{L}_t(\hat{\theta}_t) - \mathcal{L}_t(\theta_*) \right) I(\text{not } \mathcal{E}) \right] \\
 &\leq \mathbb{E} \left[ \left( \mathcal{L}_t(\hat{\theta}_t) - \mathcal{L}_t(\theta_*) \right) I(\mathcal{E}) \right] + \frac{\max_{\theta \in \Theta} (\mathcal{L}_t(\theta) - \mathcal{L}_t(\theta_*))}{t^\gamma}
 \end{aligned}$$

since the probability of not  $\mathcal{E}$  is less than  $\frac{1}{t^\gamma}$ . Now for an upper bound of the first term, observe that

$$\begin{aligned}
 \mathbb{E} \left[ \left( \mathcal{L}_t(\hat{\boldsymbol{\theta}}_t) - \mathcal{L}_t(\boldsymbol{\theta}_*) \right) I(\mathcal{E}) \right] &\leq \frac{1}{2} \mathbb{E} \left[ \left( \lambda_{\max}(\mathbf{M}_{1,t}) \right)^2 \lambda_{\max}(\mathbf{M}_{2,t}) \left\| \nabla \hat{\mathcal{L}}_t(\boldsymbol{\theta}_*) \right\|_{(\nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*))^{-1}}^2 I(\mathcal{E}) \right] \\
 &\leq (1 + c' \rho_t) \frac{1}{2} \mathbb{E} \left[ \left\| \nabla \hat{\mathcal{L}}_t(\boldsymbol{\theta}_*) \right\|_{(\nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*))^{-1}}^2 I(\mathcal{E}) \right] \\
 &\leq (1 + c' \rho_t) \frac{1}{2} \mathbb{E} \left[ \left\| \nabla \hat{\mathcal{L}}_t(\boldsymbol{\theta}_*) \right\|_{(\nabla^2 \mathcal{L}_t(\boldsymbol{\theta}_*))^{-1}}^2 \right] \\
 &= (1 + c' \rho_t) \frac{\sigma_t^2}{t}
 \end{aligned}$$

where,  $c'$  is another universal constant.  $\square$

## C. Experimental Details

### C.1. Dataset Statistics

We provide the processed data statistics in Table 6. We highlight that due to the black-box assumption of the base model, the training set is used for ablation and qualitative analysis in Section 7.3 and Section 7.4.

Table 6. Processed Dataset Statistics. Training set is only used for ablation and qualitative analysis due to the black-box model assumption.

Dataset	Train	Validation	Test
E2E NLG	33,525	4,299	4,693
Web NLG	2,732 (filtered by categories)	844	720
CommonGen	1,476 (filtered for “man”)	2,026	1,992
Adidas	—	745	100

### C.2. Prompts

We now describe the prompts we used for the four datasets and three models.

#### E2E NLG Dataset

- For the **GPT2-M** model, we use the prompt:

Given the following aspects of a restaurant, [attributes], a natural language sentence describing the restaurant is:

- For the **GPT2-XL** model, the prompt is:

Imagine you are writing a one-sentence description for a restaurant, given the following aspects: [attributes], a human-readable natural language sentence describing the restaurant is:

- For the **LLaMA-3.1-8B** model, we use:

Please convert the following attributes into a coherent sentence. Do not provide an explanation.

#### Web NLG Dataset

- For the **GPT2-M** model, we use the prompt:

Convert the following facts into a coherent sentence: Facts: [facts]  
Sentence:

- For the **GPT2-XL** model, the prompt is:

You are given the following facts. Facts: [facts] A short, coherent sentence summarizing the facts is:

- For the **LLaMA-3.1-8B** model, we use:

Do not provide an explanation or follow-up. Just convert the following facts of an entity into a coherent sentence. Facts: [facts] Sentence:

### CommonGen Dataset

- For the **GPT2-M** and **GPT2-XL** models, we use the same prompt:

One coherent sentence that uses all the following concepts: [concepts], is:

- For the **LLaMA-3.1-8B** model, we use:

Please write a coherent sentence that uses all the following concepts. Concepts: [concepts] Sentence:

### Adidas Dataset

- For the **GPT2-M** and **GPT2-XL** models, we use the same prompt:

Given the following attributes of a product, write a description.  
Attributes: [attributes] Description:

- For the **LLaMA-3.1-8B** model, we use:

Please write a description of this product given the following attributes.  
Attributes: [attributes] Description:

For **in-context learning**, we simply add a sentence at the beginning of the prompt before adding the samples in the prompt:  
Below are a list of demonstrations:

For the qualitative analysis on the distribution shift in Section 7.4, we ask GPT-4o with the following prompt:  
For Web NLG dataset:

Focus on all the samples, how much percentage is related to ‘‘Person’’?

For CommonGen dataset:

Focus on those samples whose target is related to gender, how much percentage is related to ‘‘woman’’?

### C.3. Metrics

We report performance using seven standard metrics often used in the natural language generation tasks. These are: (a) BLEU (Papineni et al., 2002) (measures n-gram overlap between the generated and reference texts, emphasizing precision), (b) ROUGE-1 (Lin, 2004) (computes unigram recall to measure the overlap between generated and reference texts), (c) ROUGE-2 (Lin, 2004) (extends ROUGE-1 to bigrams, measuring the recall of two-word sequences), (d) ROUGE-L (Lin & Och, 2004) (uses the longest common subsequence to evaluate recall), (e) METEOR (Banerjee & Lavie, 2005) (combines unigram precision, recall, and semantic matching to assess similarity), (f) CIDEr (Vedantam et al., 2015) (measures consensus in n-gram usage across multiple references, with tf-idf weighting), and (g) NIST (Doddington, 2002) (similar to BLEU but weights n-grams by their informativeness, favoring less frequent and meaningful phrases).

### C.4. Performance and Efficiency Comparison with Parameter-Efficient Fine-Tuning

While our work focuses on black-box LLM adaptation where model weights are inaccessible, we include a controlled comparison with Parameter-Efficient Fine-Tuning (PEFT) methods. Specifically, we implement LoRA (Hu et al., 2021) with

Table 7. Comparison between *Plugin* and PEFT (LoRA,  $r=8$ ) on four datasets using GPT2-XL and LLaMA-3.1-8B as base models. We show mean and standard deviation of the metrics over five seeds.

Model	Method	BLEU	Rouge-1	Rouge-2	Rouge-L	METEOR	CIDEr	NIST
<b>E2E NLG</b>								
GPT2-XL	Zeroshot	0.0562	0.4013	0.1636	0.2862	0.3697	0.0187	0.5338
GPT2-XL	LoRA ( $r=8$ )	$0.2517 \pm 0.012$	$0.5712 \pm 0.010$	$0.3079 \pm 0.013$	$0.4317 \pm 0.011$	$0.5162 \pm 0.014$	$0.5225 \pm 0.012$	$1.2172 \pm 0.011$
GPT2-XL	Plugin (Ours)	$0.2470 \pm 0.009$	$0.5536 \pm 0.007$	$0.3084 \pm 0.007$	$0.4213 \pm 0.008$	$0.5057 \pm 0.009$	$0.5455 \pm 0.013$	$1.2736 \pm 0.051$
LLaMA-3.1-8B	Zeroshot	0.3226	0.6917	0.4050	0.5004	0.6041	0.9764	1.1310
LLaMA-3.1-8B	LoRA ( $r=8$ )	$0.3702 \pm 0.016$	$0.7125 \pm 0.010$	$0.4236 \pm 0.014$	$0.5345 \pm 0.012$	$0.6413 \pm 0.017$	$1.1028 \pm 0.033$	$1.1827 \pm 0.035$
LLaMA-3.1-8B	Plugin (Ours)	$0.3691 \pm 0.013$	$0.7113 \pm 0.002$	$0.4374 \pm 0.004$	$0.5247 \pm 0.002$	$0.6392 \pm 0.009$	$1.1441 \pm 0.030$	$1.1749 \pm 0.034$
<b>Web NLG</b>								
GPT2-XL	Zeroshot	0.0317	0.2992	0.1321	0.2417	0.1969	0.0491	0.1826
GPT2-XL	LoRA ( $r=8$ )	$0.1723 \pm 0.007$	$0.4604 \pm 0.010$	$0.2618 \pm 0.011$	$0.3628 \pm 0.015$	$0.4012 \pm 0.017$	$0.9018 \pm 0.028$	$0.2736 \pm 0.014$
GPT2-XL	Plugin (Ours)	$0.1673 \pm 0.004$	$0.4616 \pm 0.007$	$0.2527 \pm 0.007$	$0.3757 \pm 0.008$	$0.3895 \pm 0.007$	$0.8987 \pm 0.013$	$0.2646 \pm 0.003$
LLaMA-3.1-8B	Zeroshot	0.1453	0.5278	0.3030	0.3982	0.4314	0.6991	0.2684
LLaMA-3.1-8B	LoRA ( $r=8$ )	$0.2638 \pm 0.008$	$0.6238 \pm 0.010$	$0.3927 \pm 0.009$	$0.4726 \pm 0.009$	$0.5927 \pm 0.013$	$1.6421 \pm 0.028$	$0.2379 \pm 0.008$
LLaMA-3.1-8B	Plugin (Ours)	$0.2542 \pm 0.004$	$0.6375 \pm 0.005$	$0.3873 \pm 0.005$	$0.4869 \pm 0.007$	$0.5724 \pm 0.004$	$1.5911 \pm 0.046$	$0.2590 \pm 0.003$
<b>CommonGen</b>								
GPT2-XL	Zeroshot	0.0317	0.2992	0.1321	0.2417	0.1969	0.0491	0.1826
GPT2-XL	LoRA ( $r=8$ )	$0.1826 \pm 0.027$	$0.5027 \pm 0.010$	$0.2137 \pm 0.014$	$0.4447 \pm 0.016$	$0.4726 \pm 0.009$	$0.7182 \pm 0.027$	$0.6725 \pm 0.043$
GPT2-XL	Plugin (Ours)	$0.1791 \pm 0.014$	$0.4932 \pm 0.007$	$0.2288 \pm 0.004$	$0.4347 \pm 0.007$	$0.4702 \pm 0.006$	$0.7283 \pm 0.012$	$0.6554 \pm 0.038$
LLaMA-3.1-8B	Zeroshot	0.0643	0.2776	0.1181	0.2488	0.3857	0.3155	0.3347
LLaMA-3.1-8B	LoRA ( $r=8$ )	$0.2736 \pm 0.018$	$0.5829 \pm 0.009$	$0.3206 \pm 0.009$	$0.5026 \pm 0.012$	$0.5927 \pm 0.016$	$1.1121 \pm 0.034$	$0.7926 \pm 0.028$
LLaMA-3.1-8B	Plugin (Ours)	$0.2665 \pm 0.010$	$0.5800 \pm 0.002$	$0.3139 \pm 0.005$	$0.5037 \pm 0.004$	$0.5829 \pm 0.003$	$1.0876 \pm 0.020$	$0.7031 \pm 0.007$
<b>Adidas</b>								
GPT2-XL	Zeroshot	0.0075	0.2309	0.0278	0.1438	0.1487	0.0184	0.4956
GPT2-XL	LoRA ( $r=8$ )	$0.0629 \pm 0.028$	$0.2816 \pm 0.030$	$0.0719 \pm 0.029$	$0.1816 \pm 0.038$	$0.2037 \pm 0.018$	$0.1231 \pm 0.126$	$0.6576 \pm 0.134$
GPT2-XL	Plugin (Ours)	$0.0600 \pm 0.017$	$0.2710 \pm 0.025$	$0.0722 \pm 0.018$	$0.1725 \pm 0.017$	$0.1995 \pm 0.018$	$0.1195 \pm 0.138$	$0.6375 \pm 0.120$
LLaMA-3.1-8B	Zeroshot	0.0120	0.2470	0.0318	0.1493	0.1526	0.0424	0.5285
LLaMA-3.1-8B	LoRA ( $r=8$ )	$0.0721 \pm 0.020$	$0.2697 \pm 0.031$	$0.0756 \pm 0.028$	$0.1821 \pm 0.020$	$0.2023 \pm 0.038$	$0.1302 \pm 0.178$	$0.6137 \pm 0.172$
LLaMA-3.1-8B	Plugin (Ours)	$0.0611 \pm 0.018$	$0.2714 \pm 0.029$	$0.0742 \pm 0.020$	$0.1759 \pm 0.019$	$0.1990 \pm 0.020$	$0.1293 \pm 0.152$	$0.6361 \pm 0.134$

rank-8 matrices on the `query` and `value` projections of GPT2-XL and LLaMA-3.1-8B, and fine-tune the base models using the same task-specific data.

The performance results are shown in Table 7. Consider GPT2-XL as a reference example, *Plugin* adds a 1-layer autoregressive Transformer with 30.72M parameters, while LoRA ( $r=8$ ) introduces only 2.46M trainable parameters. However, *Plugin* requires no modification of the base model and can be deployed post hoc. Despite the access advantage of LoRA, the performance gap is minimal. As for computational efficiency, *Plugin* requires 196.2B FLOPs (up to 64 decoding steps), while LoRA uses 188.8B FLOPs—a difference of less than 5%. The gap narrows or inverts depending on model configuration. These results suggest that *Plugin* offers a competitive adaptation solution even under white-box conditions, while maintaining broader applicability in black-box settings.

### C.5. Further Quantitative Analysis and Ablation

Following Section 7.3, we present the same ablation analysis using GPT2-M on the remaining three datasets. As shown in Figure 5, the trends mirror those in Figure 2: the *Plugin* model consistently improves performance as the base model becomes stronger with additional fine-tuning, underscoring the robustness and versatility of our approach. Similarly, Figure 6 confirms the pattern observed in Figure 3: a single-layer reweighting model yields optimal performance, while deeper configurations tend to overfit and degrade quality. Across all datasets, initializing the reweighting model with a pretrained GPT2-Small consistently boosts effectiveness.

### C.6. Influence of the architecture of the reweighting model in *Plugin*

We vary the choice of the reweighting model architecture. We find that a causal transformer layer identical to those used in the base model performs best, as it can leverage the base model’s logits and aggregate contextual information from prior tokens to better adapt the base model to the new data distribution. This conclusion is reinforced by Figure 7, where the transformer architecture consistently outperforms both the MLP (two layer with ReLU activation) and linear layers across

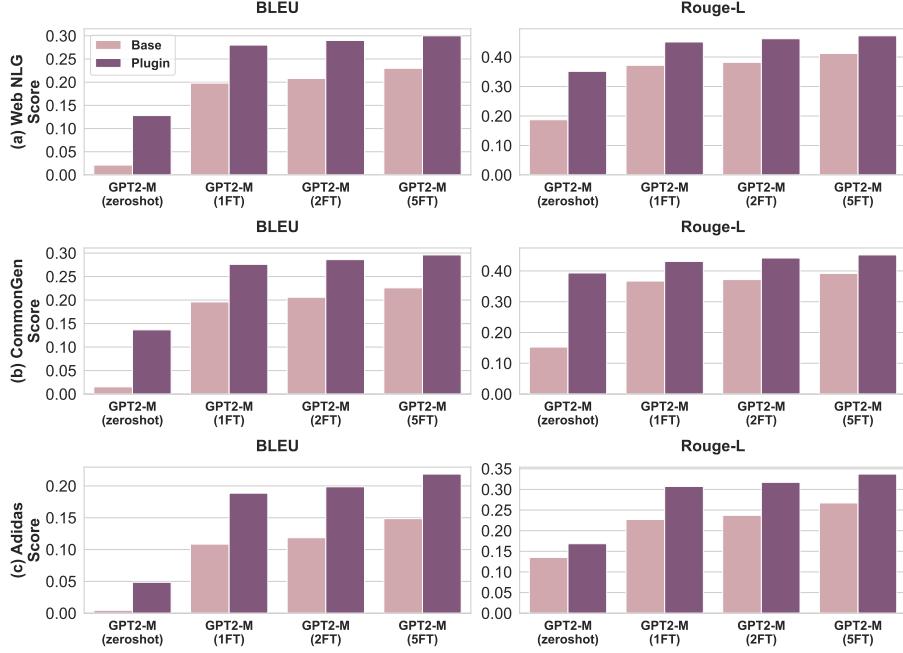


Figure 5. Performance of applying a single-layer reweighting model across increasingly fine-tuned GPT2-M models on the three datasets. Results demonstrate consistent improvements introduced by our method regardless of the strength of the base model.

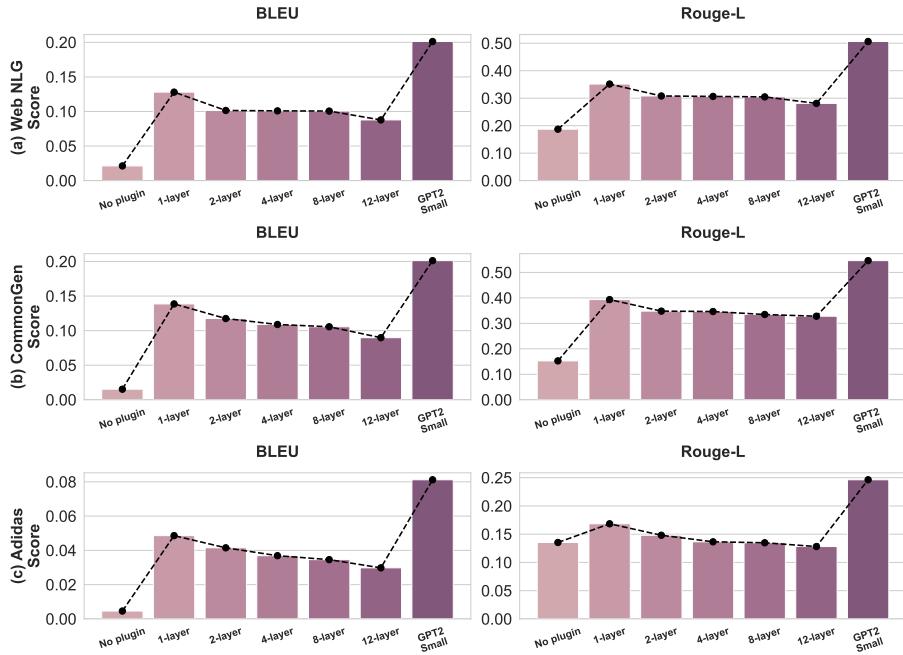


Figure 6. Performance of GPT2-M with varying reweighting model complexities on the three datasets, measured by BLEU and Rouge-L. Results demonstrate that a single reweighting layer achieves significant improvements, while increasing the number of layers beyond this leads to performance degradation, likely due to overfitting. Using a pretrained GPT2-Small as the reweighting model largely boosts the performance, highlighting the benefits of leveraging pretrained models.

all metrics, as indicated by higher means and narrower standard deviation bands. These results highlight the importance of leveraging the architectural capacity of transformers to effectively adapt the logits of the base black-box model.

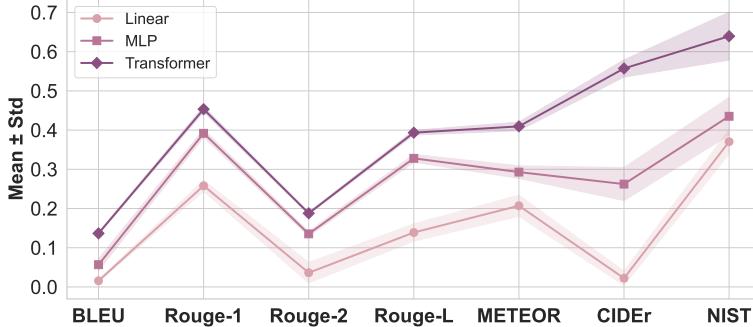


Figure 7. Performance comparison of the weighting model architecture in *Plugin*. The transformer layer achieves the best performance with consistently higher means and narrower standard deviations. Shaded bands represent the standard deviation around the mean.

### C.7. Details for Adidas Qualitative Studies

**Human Evaluation.** We conduct a human evaluation on 100 test passages from the Adidas product dataset, comparing outputs generated with and without applying the reweighting model, using LLaMA-3.1-8B as the base model. Three human evaluators are presented with a ground-truth Adidas product description and two randomly ordered descriptions: one generated with the reweighting layer and one without (i.e., we use the base model with ICL-3 as a much stronger baseline due to the low quality of the zero-shot). Evaluators are prompted to select the prediction closest to the ground truth. Results show that the output generated with the reweighting model is preferred on an average of 80.7 out of all 100 cases. The output descriptions from the base model without the reweighting are generally short and general. This demonstrates that our approach effectively adapts a closed model to the unique style of the given dataset.

In this section, we display some details for the qualitative analysis on the Adidas product description dataset.

**Details of Extracting Adidas Style Words.** We discuss the details on extracting the most frequent 50 words in the Adidas product description dataset as the “Adidas style” words. We argue that there does not exist a gold-standard way to define the “style” words for a dataset. We extract these style words through a minimal preprocessing pipeline: converting text to lowercase, removing special characters and numbers, and filtering out common English stopwords. We deliberately preserve the original word forms without lemmatization or stemming to maintain distinct style markers (e.g., keeping “comfortable” distinct from “comfort”, “running” distinct from “run”). After tokenization using NLTK’s word tokenizer, we count word frequencies across all product descriptions and select the top 50 most frequent words. This approach captures the exact vocabulary used in Adidas’ product descriptions, including specific product features.

A statistics of the frequency of these top-50 words is shown in Figure 8.

**Case Studies for Human Side-by-side Evaluation.** Below we provide more case studies on the Adidas product description datasets. We display the given product attributes, base model’s zero-shot outputs, base model with ICL-3’s outputs, *Plugin* model’s outputs, and the ground-truth references for each case.

#### ♣ Case 1

##### Given Product Attributes:

- name [Cushioned Angle Stripe Low-Cut Socks 6 Pairs]
- category [Accessories]
- price [13]
- color [Multicolor]

**Base Model (Zero-shot):** Cushioned Angle Stripe Low-Cut Socks 6 Pairs, Multicolor.

**Base Model (ICL-3):** These adidas socks are made with a cushioned sole for extra comfort and support. The low-cut

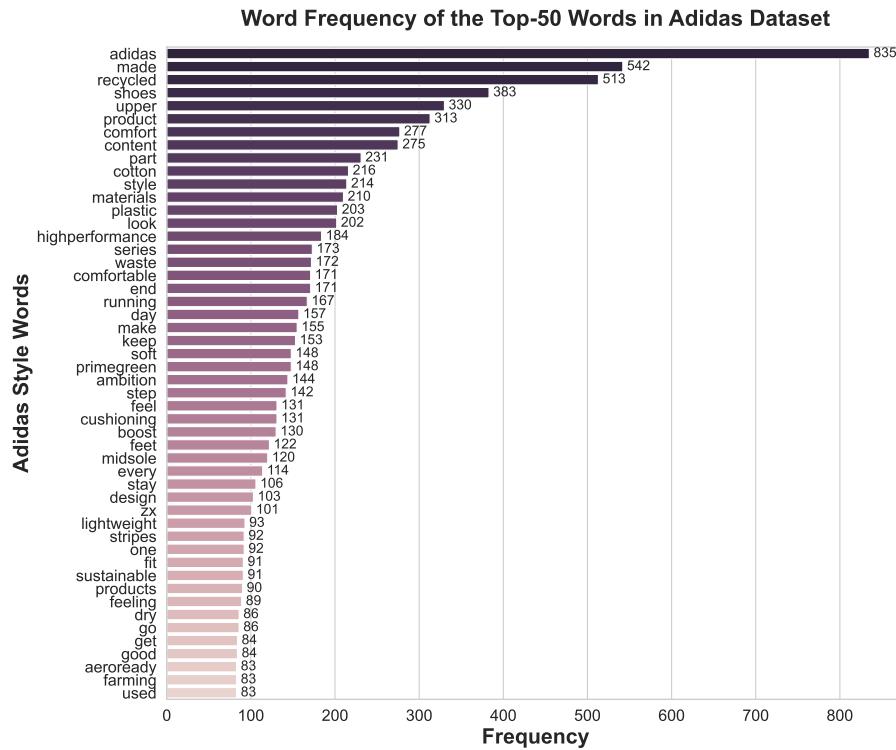


Figure 8. Word Frequency of the Top-50 Words in Adidas Dataset.

design is perfect for everyday wear, whether you're running errands or hitting the gym.

**Plugin Model:** Cushioned Angle Stripe Low-Cut Socks 6 Pairs are the perfect blend of style and comfort. The low-cut design gives you a sleek, modern look while still providing plenty of support for your feet. Made with high-quality materials, these socks will keep your feet feeling comfortable and looking chic all day long.

**Ground-truth Reference:** With lots of cushioning in the heel and toe and a soft, stretchy feel, these kids' socks are big on comfort. Whether they're bouncing around between classes, practice or after-school activities, their little feet will stay dry thanks to moisture-wicking yarn and a mesh upper. An adidas Badge of Sport and 3-Stripes give the socks a classic adidas feel.

## ♠ Case 2

**Given Product Attributes:**

- name [Badge of Sport Tee]
- category [Clothing]
- price [21]
- color [Grey]

**Base Model (Zero-shot):** A white t-shirt with a badge of the sport you love on it.

**Base Model (ICL-3):** The adidas Badge of Sport tee is a classic. It's made with soft cotton and features the iconic 3-Stripes on the shoulders. The crew neck has a ribbed hem for a comfortable fit.

**Plugin Model:** There's nothing more classic than a white tee with your favorite team's logo. This adidas Badge of Sport Tee is an everyday essential that can be worn to the gym or on game days. It's made with recycled materials, so you're making a sustainable choice. The soft cotton jersey is lightweight, breathable, and comfortable.

**Ground-truth Reference:** Put your brand pride on full display. A simple yet versatile look, this long sleeve t-shirt

keeps you comfortable while on the move or lounging at home. A contrast adidas Badge of Sport stands out on the chest. This product is made with Primegreen, a series of high-performance recycled materials.

#### ◆ Case 3

##### Given Product Attributes:

- name [Stretch Cotton Trunk Briefs 3 Pairs]
- category [Clothing]
- price [21]
- color [Grey]

**Base Model (Zero-shot):** Stretch Cotton Trunk Briefs 3 Pairs, Grey.

**Base Model (ICL-3):** These adidas men's trunks are made from a soft cotton blend that feels great against the skin. The stretchy fabric moves with you, while the flatlock seams prevent chafing and irritation. A drawcord waistband allows for an adjustable fit.

**Plugin Model:** Stretch Cotton Trunk Briefs 3 Pairs. The stretch cotton fabric is comfortable to wear. These briefs are made with a classic design in mind, so you can feel confident and comfortable. They have been designed with a low waist, which gives them a flattering fit and makes them easy to get on.

**Ground-truth Reference:** Just look at these adidas boxer briefs. They're made with soft, quick-dry fabric that keeps you feeling fresh and a stay-put design that keeps you covered from your first rep to your last. Get comfortable and give it your all. Our cotton products support sustainable cotton farming. This is part of our ambition to end plastic waste.

#### ♥ Case 4

##### Given Product Attributes:

- name [adidas x Zoe Saldana Long Sleeve Tee]
- category [Clothing]
- price [28]
- color [Black]

**Base Model (Zero-shot):** This is a long sleeve tee with the adidas logo on the front and Zoe Saldana's name on the back. It has a crew neck, short sleeves, and a straight hem. The fabric is made of 100% cotton and is machine washable.

**Base Model (ICL-3):** This adidas x Zoe Saldana long sleeve tee is made with 100% recycled polyester. The soft, lightweight fabric has a smooth finish and features a crew neck and short sleeves. It's finished with an adidas logo on the chest.

**Plugin Model:** This is not your average training t-shirt. A modern, slim fit and the softest fabric make it one of our most comfortable tops to date. The lightweight feel makes this tee perfect for all-day wear. Pair it with a pair of leggings or joggers for that ultimate athleisure look you're always striving for ambition and comfort.

**Ground-truth Reference:** Your growth, your goals and the challenges you face along the way are all part of your journey. This adidas tee was created in collaboration with Zoe Saldana to inspire you to move with confidence, no matter where you are on your path. Smooth material lays softly on your body so you can feel comfortable in your own skin. A solid color offers styling versatility. This product is made with recycled content as part of our ambition to end plastic waste.