

DoRAN: STABILIZING WEIGHT-DECOMPOSED LOW-RANK ADAPTATION VIA NOISE INJECTION AND AUXILIARY NETWORKS

Anonymous authors

Paper under double-blind review

ABSTRACT

Parameter-efficient fine-tuning (PEFT) methods have become the standard paradigm for adapting large-scale models. Among these techniques, Weight-Decomposed Low-Rank Adaptation (DoRA) has been shown to improve both the learning capacity and training stability of the vanilla Low-Rank Adaptation (LoRA) method by explicitly decomposing pre-trained weights into magnitude and directional components. In this work, we propose **DoRAN**, a new variant of DoRA designed to further stabilize training and boost the sample efficiency of DoRA. Our approach includes two key stages: (i) injecting noise into the denominator of DoRA’s weight decomposition, which serves as an adaptive regularizer to mitigate instabilities; and (ii) replacing static low-rank matrices with auxiliary networks that generate them dynamically, enabling parameter coupling across layers and yielding better sample efficiency in both theory and practice. Comprehensive experiments on vision and language benchmarks show that DoRAN consistently outperforms LoRA, DoRA, and other PEFT baselines. These results underscore the effectiveness of combining stabilization through noise-based regularization with network-based parameter generation, offering a promising direction for robust and efficient fine-tuning of foundation models.

1 INTRODUCTION

The rapid growth of large-scale pre-trained models has reshaped modern machine learning, enabling state-of-the-art performance across a wide range of vision and language tasks (Bommasani et al., 2022; Wei et al., 2022). The prohibitive cost of fully fine-tuning these models has spurred significant interest in parameter-efficient fine-tuning (PEFT) methods, which achieve strong downstream performance while modifying only a small fraction of parameters. A variety of PEFT methods have been developed to adapt large-scale pre-trained models without incurring the cost of full fine-tuning. Early approaches, such as adapter layers, insert small trainable modules between frozen transformer blocks, while prefix- and prompt-tuning prepend learnable vectors to the model’s input space. More recently, Low-Rank Adaptation (LoRA) has become a foundational technique for efficient fine-tuning of large-scale models across multiple AI domains. Originally proposed by Hu et al. (2022), LoRA introduces low-rank trainable adapters into existing weight matrices, reducing computational overhead while maintaining high task performance. Beyond its success in natural language processing (NLP), LoRA has demonstrated strong adaptability in computer vision (Yang et al., 2024), speech processing, and reinforcement learning (Huan & Shun, 2025). LoRA also plays a pivotal role in federated and distributed learning, enabling efficient personalization while preserving data privacy (Sun et al., 2025; Xu et al., 2025).

Despite its effectiveness in providing a flexible low-rank adaptation, LoRA’s rigid low-rank formulation can limit both representational capacity and training stability. Several extensions have sought to address these shortcomings. Notably, DoRA (Liu et al., 2024) introduces a directional-based normalization approach: instead of learning flexible low-rank updates directly, it decomposes pre-trained weights into magnitude and directional components. This design improves representational power relative to vanilla LoRA and mitigates some instability.

While DoRA represents a significant advance, it still inherits two challenges. First, DoRA relies on strict normalization, which makes it sensitive to optimization instabilities: when the adapted weight

norm approaches zero, gradients can explode, destabilizing training. Second, DoRA still employs static low-rank matrices that are optimized independently at each layer, underutilizing cross-layer correlations and often leading to redundant or poorly coordinated updates. These challenges motivate the following central question of our work:

*(Q) Can we unify the strengths of **direction-based normalization** and **flexible low-rank adaptation** into a single framework that is both stable and expressive?*

To address this question, we propose **Weight-Decomposed Low-Rank Adaptation with Noise injection (DoRAN)**, a PEFT approach that simultaneously **stabilizes optimization** and **enriches low-rank representations**. DoRAN introduces two key components. **(i) Stabilization via learnable noise:** We augment DoRA’s normalization with a learnable offset in the denominator, which functions as an adaptive regularizer. This simple yet powerful modification eliminates singularities at small norms, ensures gradients remain well-conditioned, and allows the model to smoothly interpolate between two regimes: purely directional updates (as in DoRA) and linear scaling (as in LoRA). In this way, DoRAN adaptively preserves magnitude information of the gradient signal instead of discarding it entirely like DoRA, as we will reveal in Section 3.2, leading to more expressive and stable updates. **(ii) Dynamic generation of low-rank factors:** Instead of directly learning a separate low-rank adapter at each layer, we dynamically generate the adapters through a hypernetwork with a shared backbone and head-specific output heads **between query and value matrices for each layer**. This coupling encourages **information sharing between query and value across layers**, enforces structural consistency across the model, and improves sample efficiency by leveraging shared latent structure. **This design is motivated by recent studies which demonstrated that hypernetworks enhances generalization and stability by enabling adaptive weight generation conditioned on task representations (Ha et al., 2017).** Besides, this design also provides a richer inductive bias that regularizes the adaptation space and mitigates redundancy across layers.

Crucially, as we will show in our experiments, both components contribute significantly to the performance improvements. Moreover, these two components are synergistic: the learnable noise term ensures that gradient signals remain stable and informative, while the hypernetwork ensures that these stabilized signals are used to generate structured low-rank updates. Together, they provide a unified framework that bridges the gap between direction-based normalization and flexible low-rank adaptation, offering both robustness and expressiveness for fine-tuning large-scale models.

We validate DoRAN through extensive experiments on vision (VTAB-1K, FGVC) and language (commonsense reasoning with LLaMA-7B/13B) benchmarks. DoRAN consistently outperforms LoRA, DoRA, and other PEFT baselines, while introducing negligible additional overhead. We also provide theoretical and empirical evidence that DoRAN significantly improves sample efficiency by leveraging shared structure **between query and value matrices through applying specific-hypernetworks** across layers.

Contributions. In summary, our contributions are threefold:

- We propose a stabilized weight decomposition with learnable noise that removes the singularity of DoRA and adaptively preserves magnitude information, and a hypernetwork-based reparameterization of low-rank factors, enabling **parameter coupling between query and value matrices across layers** and improved sample efficiency.
- We provide a theoretical analysis showing that this design significantly enhances the sample efficiency from an *exponential order* to a *polynomial order*.
- We demonstrate through various experimental settings that DoRAN achieves state-of-the-art parameter efficiency, consistently outperforming LoRA, DoRA, and other PEFT baselines on both vision and language tasks.

Together, we show that combining noise-based stabilization with hypernetwork-driven parameter generation offers a robust and generalizable path forward for fine-tuning foundation models.

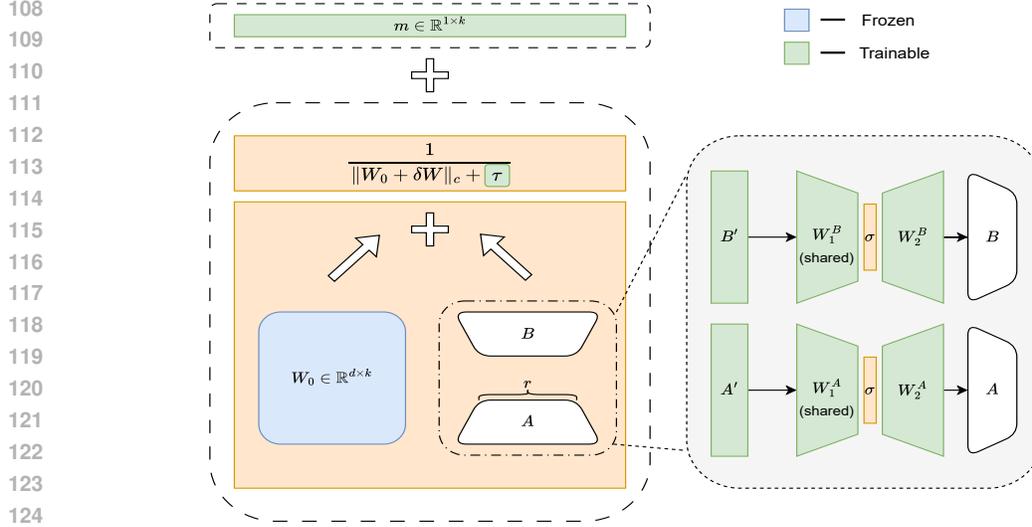


Figure 1: Illustration of DoRAN. The matrices W_1^A and W_1^B are shared across query and value projection layers, as well as across all attention heads within the same Transformer block. Each layer maintains its own independent hypernetwork.

2 PRELIMINARIES

Notation. We denote $[n] = \{1, 2, \dots, n\}$ for an integer n . For a vector $u \in \mathbb{R}^d$, we will use both notations $u = (u^{(1)}, u^{(2)}, \dots, u^{(d)})$ and $u = (u_1, u_2, \dots, u_d)$ interchangeably. Given a multi-index vector $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_d) \in \mathbb{N}^d$, we write $u^\alpha = u_1^{\alpha_1} u_2^{\alpha_2} \dots u_d^{\alpha_d}$, $|\alpha| = \alpha_1 + \alpha_2 + \dots + \alpha_d$, and $\alpha! = \alpha_1! \alpha_2! \dots \alpha_d!$. $\|u\|$ denotes the Euclidean norm of u , while $|S|$ represents the cardinality of a set S . For two positive sequences $(a_n)_{n \geq 1}$ and $(b_n)_{n \geq 1}$, we write $a_n = \mathcal{O}(b_n)$ or $a_n \lesssim b_n$ if there exists a constant $C > 0$ such that $a_n \leq C b_n$ for all n . We denote $a_n = \mathcal{O}_P(b_n)$ if a_n/b_n is bounded in probability. Next, we use the notation $a_n = \tilde{\mathcal{O}}_P(b_n)$ when $a_n = \mathcal{O}_P(b_n \log^c(b_n))$ for some $c > 0$. Finally, we define the inner product of two matrices as $\langle A, B \rangle = \text{Trace}(A^\top B)$.

MSA: Multi-head Self-attention. Self-attention is the central mechanism of the Transformer architecture (Vaswani et al., 2017; Dosovitskiy, 2021a). Given an input sequence $X \in \mathbb{R}^{n \times d}$ with n tokens of d dimension, the model projects X into query, key, and value matrices via learned linear transformations $Q = XW_Q$, $K = XW_K$ and $V = XW_V$ where $W_Q, W_K, W_V \in \mathbb{R}^{d \times d_h}$. The attention weight is defined as

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^\top}{\sqrt{d_k}}\right)V.$$

Multi-head Self-attention (MSA) extends this mechanism by introducing multiple parallel attention operations or "heads", where each attention head processes the input sequence from a different representational subspace. Formally, for h heads, the output of an MSA layer is calculated as

$$\begin{aligned} \text{MSA}(X_Q, X_K, X_V) &= \text{Concat}(\text{head}_1, \dots, \text{head}_h)W_O, \\ \text{head}_i &= \text{Attention}(XW_Q^{(i)}, XW_K^{(i)}, XW_V^{(i)}). \end{aligned}$$

LoRA and DoRA. Low-rank Adaptation (LoRA) (Hu et al., 2022) is one of the most widely used PEFT approaches. The key idea is to decompose the weight updates into low-rank matrices and fine-tune these low-rank matrices while keeping the original weights frozen. Formally, given a pre-trained weight matrix $W \in \mathbb{R}^{d \times k}$, LoRA freezes W and models the weight update $\Delta W \in \mathbb{R}^{d \times k}$ as two low-rank matrices B and A :

$$W = W_0 + \Delta W = W_0 + BA$$

where $B \in \mathbb{R}^{d \times r}$ and $A \in \mathbb{R}^{r \times k}$, with rank $r \ll \min(d, k)$. On performing weight decomposition on fine-tuned weight matrices, Liu et al. (2024) found that LoRA and full fine-tuning show different learning patterns. To resolve this difference, weight-decomposed low-rank adaptation (DoRA)

(Liu et al., 2024) explicitly decomposes the pre-trained weight into its magnitude and directional component and fine-tunes both components. Then, DoRA performs updates as follows:

$$\mathbf{W} = m \frac{\mathbf{V} + \Delta \mathbf{V}}{\|\mathbf{V} + \Delta \mathbf{V}\|_c} = m \frac{\mathbf{W}_0 + \mathbf{B}\mathbf{A}}{\|\mathbf{W}_0 + \mathbf{B}\mathbf{A}\|_c},$$

where \mathbf{B} , \mathbf{A} are trainable low-rank matrices, $m \in \mathbb{R}^{1 \times k}$ is a learnable vector, and $\|\cdot\|_c$ is the vector-wise norm of a matrix across each column. In this work, we aim to improve DoRA by introducing noise-stabilized normalization to prevent instability and by utilizing auxiliary networks to generate low-rank updates, thereby enhancing both robustness and parameter efficiency.

3 METHODOLOGY

Motivations. DoRA enhances stability by decomposing pre-trained weights into magnitude and directional components. Nevertheless, adaptation can still suffer from instability. In particular, when the normalization denominator becomes excessively small, DoRA may encounter gradient explosions, as we will discuss in Section 3.2. Moreover, standard DoRA employs static, **distinct low-rank matrices between query and value**. When each adapter (\mathbf{A} , \mathbf{B}) is optimized independently, they risk diverging toward inconsistent or redundant solutions. To address these limitations, we propose **Stabilized Weight-Decomposed Low-Rank Adaptation via Noise Injection and Auxiliary Networks (DoRAN)**, which integrates two complementary components. **First**, DoRAN introduces a learnable stabilization noise term in the normalization step, which provides more stable and expressive updates. **Second**, instead of directly optimizing the low-rank matrices, DoRAN employs hypernetworks to dynamically generate low-rank matrices.

3.1 DoRAN: DoRA WITH NOISE INJECTION AND AUXILIARY NETWORKS

Noise Injection. We first enhance stability by introducing a trainable stabilization noise term τ :

$$\mathbf{W} = m \frac{\mathbf{W}_0 + \mathbf{B}\mathbf{A}}{\|\mathbf{W}_0 + \mathbf{B}\mathbf{A}\|_c + \tau},$$

where $\tau \in \mathbb{R}^+$ acts as an adaptive regularizer, ensuring stable normalization. We design τ as a *learnable scalar* rather than a fixed constant for a key reasons: As we will reveal in the gradient analysis (Section 3.2), the value of τ determines the balance between two adaptation regimes: when $\tau \rightarrow 0$, the update reduces to DoRA-style direction learning, while large τ yields LoRA-like scaling. By making τ learnable, the model can automatically interpolate between these extremes, adaptively finding the most effective balance between directional learning and stable norm control.

Auxiliary Networks. In addition to the term τ , DoRAN further enhances flexibility by implementing the low-rank matrices \mathbf{A} and \mathbf{B} through auxiliary networks rather than optimizing them directly:

$$\mathbf{B} = g_B(\mathbf{B}') := \mathbf{W}_2^B \sigma(\mathbf{W}_1^B \mathbf{B}'); \quad \mathbf{A} = g_A(\mathbf{A}') := \mathbf{W}_2^A \sigma(\mathbf{W}_1^A \mathbf{A}'),$$

where \mathbf{A}' , \mathbf{B}' are learnable embedding, and g_B, g_A are two-layer feedforward networks with the activation σ . In our experiments, the inputs \mathbf{B}' , \mathbf{A}' and the first layers \mathbf{W}_1^A , \mathbf{W}_1^B are shared between query and value projections in each layer, while \mathbf{W}_2^A , \mathbf{W}_2^B are distinct to emit specific adaptation matrices for the query and value projections. To ensure parameter efficiency and facilitate knowledge sharing across attention heads, this hypernetwork is shared among attention heads in each layer. Bringing these components together, we have the following DoRAN updates, whose demonstration can be found in Figure 1:

$$\mathbf{W} = m \frac{\mathbf{W}_0 + g_B(\mathbf{B}')g_A(\mathbf{A}')}{\|\mathbf{W}_0 + g_B(\mathbf{B}')g_A(\mathbf{A}')\|_c + \tau}.$$

Stabilization–Hypernetwork Synergy in DoRAN. We emphasize that the stabilization noise and the hypernetworks play complementary roles. As we will discuss in Section 3.2, the learnable noise term τ stabilizes optimization by ensuring that the gradients passed to the low-rank factors remain well-conditioned and preserve both orthogonal and parallel components of the signal. On the other hand, the hypernetwork then ensures that these gradients are not used to produce arbitrary low-rank updates, but rather to generate structured and consistent factors. In particular, the parameter sharing

mechanism acts as an implicit regularizer: rather than producing arbitrary and potentially redundant low-rank solutions, the hypernetwork encourages consistency and reuse of common structure [between query and value matrices](#) across layers, which has been shown in prior work to improve generalization of learned weights (Ha et al., 2017; Bertinetto et al., 2016).

3.2 GRADIENT ANALYSIS

In this section, we examine the gradient properties of DoRAN to further highlight its theoretical benefits. Let $\mathbf{W}' = \mathbf{W}_0 + \mathbf{B}\mathbf{A}$. Let $\mathcal{L}(\mathbf{W})$ denote the loss and define $\mathbf{G} = \frac{\partial \mathcal{L}}{\partial \mathbf{W}}$ as the upstream gradient. By the chain rule, the total gradient with respect to \mathbf{W}' can be given by:

$$\frac{\partial \mathcal{L}}{\partial \mathbf{W}'} = m \left[\frac{\mathbf{G}}{\|\mathbf{W}'\|_c + \tau} - \frac{\langle \mathbf{G}, \mathbf{W}' \rangle}{(\|\mathbf{W}'\|_c + \tau)^2} \frac{\mathbf{W}'}{\|\mathbf{W}'\|_c} \right].$$

Writing \mathbf{G} as the sum of its component orthogonal to \mathbf{W}' (\mathbf{G}_\perp) and its projection onto \mathbf{W}' ($\text{proj}_{\mathbf{W}'}(\mathbf{G})$), we obtain the equivalent decomposition:

$$\frac{\partial \mathcal{L}}{\partial \mathbf{W}'} = \frac{m}{\|\mathbf{W}'\|_c + \tau} \mathbf{G}_\perp + \frac{m\tau}{(\|\mathbf{W}'\|_c + \tau)^2} \text{proj}_{\mathbf{W}'}(\mathbf{G}).$$

This formulation reveals that the *orthogonal component* of the gradient, which drives updates in the direction subspace, is scaled by a factor of $m/(\|\mathbf{W}'\|_c + \tau)$. On the other hand, the *parallel component* (which controls norm changes) is preserved with a damping factor $\tau/(\|\mathbf{W}'\|_c + \tau)$. In other words, the parallel gradient $\text{proj}_{\mathbf{W}'}(\mathbf{G})$ is not completely discarded (as in DoRA), hence the magnitude information is adaptively preserved. Therefore, in the limit $\tau \rightarrow 0$, the parallel component vanishes, recovering the DoRA update rule where only the direction of \mathbf{W}' is optimized. Conversely, as $\tau \rightarrow \infty$, the update reduces to a rescaled linear adaptation $\mathbf{W} \approx (m/\tau)\mathbf{W}' = (m/\tau)(\mathbf{W}_0 + \mathbf{A}\mathbf{B})$, which resembles LoRA. The gradients with respect to m and τ are given by:

$$\frac{\partial \mathcal{L}}{\partial m} = \left\langle \mathbf{G}, \frac{\mathbf{W}'}{\|\mathbf{W}'\|_c + \tau} \right\rangle, \quad \frac{\partial \mathcal{L}}{\partial \tau} = -m \left\langle \mathbf{G}, \frac{\mathbf{W}'}{(\|\mathbf{W}'\|_c + \tau)^2} \right\rangle.$$

This equation reveals that m governs global scaling, while τ adaptively interpolates between strict normalization and linear scaling. To sum up, relative to DoRA and LoRA, the adaptive τ of DoRAN provides the following effects:

- **Stability:** The denominator $\|\mathbf{W}'\|_c + \tau$ avoids the singularity at small norms, providing bounded gradients, which leads to more stable updates.
- **Magnitude learning via \mathbf{W}' :** Unlike DoRA, the parallel component is partially retained, allowing \mathbf{A} , \mathbf{B} to co-adapt both direction and norm.
- **Continuum of behaviors:** By adjusting τ , the method interpolates smoothly between DoRA ($\tau \rightarrow 0$) and LoRA-like scaling ($\tau \gg \|\mathbf{W}'\|_c$). Therefore, τ enables the model to discover the most effective balance between direction learning and stable norm control.

4 THEORETICAL ANALYSIS

This section presents the theoretical explanation of the advantages of the reparameterization technique implemented in DoRAN, through its connection to Mixture of Experts (MoE).

DoRA meets MoE. Recent works have shown that single-headed self-attention can be reinterpreted as a Mixture of Experts (MoE) model (Truong et al., 2025a), where each head or projection serves as an expert and the attention mechanism acts as the gating function. Building on this view, applying DoRA to an attention head can also be reformulated as an MoE model. Specifically, the attention head before and after applying DoRA can be expressed as follows:

$$\begin{aligned} \text{head}_{\text{pre}} &= \text{softmax} \left(\mathbf{X}\mathbf{W}_Q\mathbf{W}_K^\top\mathbf{X}^\top / \sqrt{d_h} \right) \mathbf{X}\mathbf{W}_V \in \mathbb{R}^{N \times d_h}, \\ \text{head}_{\text{post}} &= \text{softmax} \left(\mathbf{X}m_Q \frac{\mathbf{W}_Q + \mathbf{B}_Q\mathbf{A}_Q}{\|\mathbf{W}_Q + \mathbf{B}_Q\mathbf{A}_Q\| \sqrt{d_h}} \mathbf{W}_K^\top\mathbf{X}^\top \right) \mathbf{X}m_V \frac{\mathbf{W}_V + \mathbf{B}_V\mathbf{A}_V}{\|\mathbf{W}_V + \mathbf{B}_V\mathbf{A}_V\|}, \end{aligned}$$

where the softmax is applied to each row of the matrix inside. Therefore, the i -th row of the $\text{head}_{\text{post}}$ can be written as,

$$\text{head}_{\text{post},i} = \sum_{j=1}^N \underbrace{\text{softmax} \left(\left\{ \frac{\widetilde{\mathbf{X}} \mathbf{E}_i^\top m_Q (\mathbf{W}_Q + \mathbf{B}_Q \mathbf{A}_Q)}{\|\mathbf{W}_Q + \mathbf{B}_Q \mathbf{A}_Q\| \sqrt{d_h}} \mathbf{W}_K^\top \mathbf{E}_k \widetilde{\mathbf{X}} \right\}_{k=1}^N \right)}_{\text{gating score for } j\text{-th expert}} \underbrace{m_V \frac{(\mathbf{W}_V + \mathbf{B}_V \mathbf{A}_V)^\top}{\|\mathbf{W}_V + \mathbf{B}_V \mathbf{A}_V\|} \mathbf{E}_j \widetilde{\mathbf{X}}}_{j\text{-th expert}}, \quad (1)$$

where we define $\widetilde{\mathbf{X}} := \text{Vec}(\mathbf{X}) \in \mathbb{R}^{Nd}$ and the matrix $\mathbf{E}_j \in \mathbb{R}^{d \times Nd}$ that extracts the j -th token of \mathbf{X} , i.e., $\mathbf{E}_j \widetilde{\mathbf{X}} = \mathbf{x}_j$. Thus, the attention head can be expressed as an MoE with N experts and we will leverage this connection to analyze the properties of DoRA and DoRAN.

Problem setup: Let $(\mathbf{X}_1, \mathbf{Y}_1), \dots, (\mathbf{X}_n, \mathbf{Y}_n) \in \mathbb{R}^d \times \mathbb{R}^d$ be i.i.d. samples generated from the following regression model:

$$\mathbf{Y}_i = f_{G_*}(\mathbf{X}_i) + \varepsilon_i, \quad i = 1, 2, \dots, n. \quad (2)$$

We assume that $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n$ are i.i.d. samples from some probability distribution μ with bounded support; $\varepsilon_1, \varepsilon_2, \dots, \varepsilon_n \in \mathbb{R}^d$ are independent Gaussian noise with $\mathbb{E}(\varepsilon_i | \mathbf{X}_i) = 0$ and $\text{Var}(\varepsilon_i | \mathbf{X}_i) = \sigma^2 I_d$ for all i . The regression function f_{G_*} takes the form of an MoE model with L unknown experts,

$$f_{G_*}(\mathbf{X}) := \sum_{j=1}^L \text{softmax} \left(\left\{ \mathbf{X}^\top m_{Q,k} \frac{\mathbf{C}_Q + \mathbf{B}_{Q,k}^* \mathbf{A}_{Q,k}^*}{\|\mathbf{C}_Q + \mathbf{B}_{Q,k}^* \mathbf{A}_{Q,k}^*\|} \mathbf{C}_K \mathbf{X} + c_k^* \right\}_{k=1}^L \right) \left(m_{V,j} \frac{\mathbf{C}_V + \mathbf{B}_{V,j}^* \mathbf{A}_{V,j}^*}{\|\mathbf{C}_V + \mathbf{B}_{V,j}^* \mathbf{A}_{V,j}^*\|} \right) \mathbf{X}, \quad (3)$$

and $G_* = \sum_{j=1}^L \exp(c_j^*) \delta_{(\mathbf{B}_{Q,j}^*, \mathbf{A}_{Q,j}^*, \mathbf{B}_{V,j}^*, \mathbf{A}_{V,j}^*)}$ represents the mixing measure, which is a combination of Dirac measures δ , associated with the true unknown parameters $(c_j^*, \mathbf{B}_{Q,j}^*, \mathbf{A}_{Q,j}^*, \mathbf{B}_{V,j}^*, \mathbf{A}_{V,j}^*)_{j=1}^L$ in the compact parameter space $\Theta \subset \mathbb{R} \times \mathbb{R}^{d \times r} \times \mathbb{R}^{r \times d} \times \mathbb{R}^{d \times r} \times \mathbb{R}^{r \times d}$. The pre-trained matrices $\mathbf{C}_Q, \mathbf{C}_K, \mathbf{C}_V \in \mathbb{R}^{d \times d}$ are fixed and given (we changed the notation from matrix \mathbf{W} to \mathbf{C} to avoid confusion in subsequent analysis). To facilitate the analysis, we also assume that the magnitude parameters $\{m_{Q,j}, m_{V,j}\}_j$ are positive and known.

We note that the regression function in Eq. (3) resembles the adaptation strategy of DoRA, with different low-rank matrices added to the query and value matrices of the pre-trained model. We prove that this mechanism causes suboptimal sample complexity in estimating the unknown low-rank matrices $(\mathbf{B}_{Q,j}^*, \mathbf{A}_{Q,j}^*, \mathbf{B}_{V,j}^*, \mathbf{A}_{V,j}^*)$ (see Section 4.1). We refer to this as the non-shared structure. On the other hand, we employ the sharing technique in DoRAN, with added noise to the ℓ_2 norms of the query and value matrices. We prove that these adjustments improve sample efficiency compared to the original technique (see Section 4.2).

4.1 NON-SHARED STRUCTURE CAUSES SUBOPTIMAL SAMPLE COMPLEXITY

We proceed to analyze the convergence rate of the low-rank matrices under the non-shared structure of the regression function in Eq. (3). To study the convergence behavior, it is natural to approach the problem from the perspective of estimating the ground-truth mixing measure G_* . For this purpose, we adopt the least-squares method (van de Geer, 2000) as follows:

$$\widehat{G}_n \in \arg \min_{G \in \mathcal{G}_{L'}(\Theta)} \sum_{i=1}^n \left(\mathbf{Y}_i - f_G(\mathbf{X}_i) \right)^2, \quad (4)$$

where we denote by $\mathcal{G}_{L'}(\Theta) := \{G = \sum_{j=1}^{\ell} \exp(c_j) \delta_{\mathbf{B}_{Q,j}, \mathbf{A}_{Q,j}, \mathbf{B}_{V,j}, \mathbf{A}_{V,j}} : 1 \leq \ell \leq L'\}$ the set of all mixing measures with at most L' atoms. Since the true number of experts L is typically

unknown, we assume that the number of fitted experts L' chosen such that it is larger than L . To determine the convergence rates of the above estimator \widehat{G}_n , we use the following loss function, which is constructed based on the Voronoi cells concept from Manole & Ho (2022).

Voronoi loss: For a mixing measure G with $L' > L$ atoms, the Voronoi cell set $\{\mathcal{V}_j \equiv \mathcal{V}_j(G) : j \in [L]\}$ is generated by the atoms of G_* as $\mathcal{V}_j := \{i \in [L'] : \|\mathbf{H}_i - \mathbf{H}_j^*\| \leq \|\mathbf{H}_i - \mathbf{H}_\ell^*\|, \forall \ell \neq j\}$, where $\mathbf{H} := (\mathbf{B}_Q, \mathbf{A}_Q, \mathbf{B}_V, \mathbf{A}_V)$. Then, we define the Voronoi loss for the non-shared structure,

$$\begin{aligned} \mathcal{D}_{1,r}(G, G_*) &:= \sum_{j=1}^L \left| \sum_{i \in \mathcal{V}_j} \exp(c_i) - \exp(c_j^*) \right| \\ &+ \sum_{j=1}^L \sum_{i \in \mathcal{V}_j} \exp(c_i) (\|\Delta \mathbf{B}_{Q,ij}\|^r + \|\Delta \mathbf{A}_{Q,ij}\|^r + \|\Delta \mathbf{B}_{V,ij}\|^r + \|\Delta \mathbf{A}_{V,ij}\|^r), \end{aligned}$$

where we use the notation $\Delta \mathbf{A}_{ij} := \mathbf{A}_i - \mathbf{A}_j^*$ for a matrix \mathbf{A} to quantify the estimation error. It is evident that once the order of the loss value is determined, the convergence rate of the estimators can be readily inferred. The following theorem presents the minimax rate of the Voronoi loss.

Theorem 1. *Under the setting of non-shared structure defined in Eq. (2) and Eq. (3), the following minimax lower bound of estimating G_* using \widehat{G}_n defined in Eq. (4) holds for any $r \in \mathbb{N}$:*

$$\inf_{\widehat{G}_n \in \mathcal{G}_{L'}} \sup_{G \in \mathcal{G}_{L'}(\Theta) \setminus \mathcal{G}_{L-1}(\Theta)} \mathbb{E}_{f_G} [\mathcal{D}_{1,r}(\widehat{G}_n, G)] \gtrsim \frac{1}{\sqrt{n}},$$

where \mathbb{E}_{f_G} denotes the expectation taken w.r.t. the product measure f_G^n .

The proof of Theorem 1 is in Appendix D.1. Since the minimax lower bound above holds for any natural number r , which is the order of the estimation error in the Voronoi loss, it follows that the convergence rates of the low-rank matrix estimators $\{\mathbf{B}_{Q,j}^n, \mathbf{A}_{Q,j}^n, \mathbf{B}_{V,j}^n, \mathbf{A}_{V,j}^n\}$ would be slower than any polynomial rates $\mathcal{O}_P(n^{-1/2r})$. Hence, the convergence rate could become as slow as $\mathcal{O}_P(1/\log^a(n))$ for some constant $a > 0$. As a consequence, we need an exponential sample size of order $\mathcal{O}(\exp(\epsilon^{-1/a}))$ to obtain the estimation error ϵ . This observation underscores the limited sample efficiency of the non-sharing structure. At a high level, the main technical obstacle to the above convergence rate arises from the fact that normalizing the adapted matrix induces an interaction among the parameters through the following partial differential equation (PDE):

$$\left\langle \mathbf{C}_Q + \mathbf{B}\mathbf{A}, \frac{\partial}{\partial(\mathbf{B}\mathbf{A})} \exp \left(\mathbf{X}^\top m_Q \frac{\mathbf{C}_Q + \mathbf{B}\mathbf{A}}{\|\mathbf{C}_Q + \mathbf{B}\mathbf{A}\|} \mathbf{C}_K \mathbf{X} \right) \right\rangle = 0. \quad (5)$$

This PDE causes the linear dependence of the terms in the Taylor expansion of the density discrepancy $f_{\widehat{G}_n}(\mathbf{X}) - f_{G_*}(\mathbf{X})$ that affect the convergence of expert estimation.

4.2 SHARED-STRUCTURE WITH ADDED NOISE

Now, we consider the following non-linear reparameterization that shares the parameters within the low-rank adapting matrices:

$$\mathbf{A}_Q = \mathbf{A}_V = \sigma_1(\mathbf{W}_1 \mathbf{A}), \quad \mathbf{B}_Q = \mathbf{B}_V = \sigma_2(\mathbf{W}_2 \mathbf{B}), \quad (6)$$

where $\mathbf{A} \in \mathbb{R}^{m \times d}$, $\mathbf{B} \in \mathbb{R}^{m' \times r}$, $\mathbf{W}_1 \in \mathbb{R}^{r \times m}$, $\mathbf{W}_2 \in \mathbb{R}^{d \times m'}$ are learnable matrices with given dimension m, m' ; σ_1, σ_2 are some non-linear activation functions. Moreover, we perturb the ℓ_2 -norm components of the query and value matrices by adding a small noise term in the denominator. In addition to the aforementioned benefits of the added noise, we note that it also helps avoid the interaction in Eq. (5). With these adjustments, the ground-truth regression function is:

$$\begin{aligned} f_{\widetilde{G}_*}(\mathbf{X}) &:= \sum_{j=1}^L \text{softmax} \left(\left\{ \mathbf{X}^\top m_{Q,k} \frac{\mathbf{C}_Q + \sigma_2(\mathbf{W}_{2,k}^* \mathbf{B}_k^*) \sigma_1(\mathbf{W}_{1,k}^* \mathbf{A}_k^*)}{\|\mathbf{C}_Q + \sigma_2(\mathbf{W}_{2,k}^* \mathbf{B}_k^*) \sigma_1(\mathbf{W}_{1,k}^* \mathbf{A}_k^*)\| + \tau_Q} \mathbf{C}_K \mathbf{X} + c_k^* \right\}_{k=1}^L \right)_j \\ &\times \left(m_{V,j} \frac{\mathbf{C}_V + \sigma_2(\mathbf{W}_{2,j}^* \mathbf{B}_j^*) \sigma_1(\mathbf{W}_{1,j}^* \mathbf{A}_j^*)}{\|\mathbf{C}_V + \sigma_2(\mathbf{W}_{2,j}^* \mathbf{B}_j^*) \sigma_1(\mathbf{W}_{1,j}^* \mathbf{A}_j^*)\| + \tau_V} \right) \mathbf{X}, \quad (7) \end{aligned}$$

here the ground-truth mixing measure is $\tilde{G}_* := \sum_{j=1}^L \exp(c_j^*) \delta_{(\mathbf{W}_{2,j}^* \mathbf{B}_j^*, \mathbf{W}_{1,j}^* \mathbf{A}_j^*)}$. Similarly, we consider the least-square method for our estimator, i.e., $\tilde{G}_n \in \arg \min_{\tilde{G} \in \tilde{\mathcal{G}}_{L'}(\Theta)} \sum_{i=1}^n \left(\mathbf{Y}_i - f_{\tilde{G}}(\mathbf{X}_i) \right)^2$, where $\tilde{\mathcal{G}}_{L'}(\Theta) := \{G = \sum_{i=1}^{\ell} \exp(c_i) \delta_{(\mathbf{W}_{2,i} \mathbf{B}_i, \mathbf{W}_{1,i} \mathbf{A}_i)} : 1 \leq \ell \leq L'\}$.

The Voronoi loss function in this setting is constructed as

$$\begin{aligned} \mathcal{D}_2(\tilde{G}, \tilde{G}_*) &:= \sum_{j=1}^L \left| \sum_{i \in \mathcal{V}_j} \exp(c_i) - \exp(c_j^*) \right| + \sum_{j \in [L]: |\mathcal{V}_j|=1} \sum_{i \in \mathcal{V}_j} \exp(c_i) (\|\Delta(\mathbf{W}_2 \mathbf{B})_{ij}\| + \|\Delta(\mathbf{W}_1 \mathbf{A})_{ij}\|) \\ &+ \sum_{j \in [L]: |\mathcal{V}_j| > 1} \sum_{i \in \mathcal{V}_j} \exp(c_i) (\|\Delta(\mathbf{W}_2 \mathbf{B})_{ij}\|^2 + \|\Delta(\mathbf{W}_1 \mathbf{A})_{ij}\|^2), \end{aligned}$$

where we denote, for example, $\Delta(\mathbf{W}_2 \mathbf{B})_{ij} := \mathbf{W}_{2,i} \mathbf{B}_i - \mathbf{W}_{2,j}^* \mathbf{B}_j^*$. Before presenting the main theorem, we introduce some mild assumptions on the activation functions σ_1 and σ_2 . Due to the space constraint, we defer these assumptions until the proof of Theorem 2 in Appendix D.2.

Theorem 2. *Assume that the activation functions σ_1 and σ_2 meet the assumptions specified in Appendix D.2. Then, the estimator \tilde{G}_n converges to the true mixing measure \tilde{G}_* at the following rate: $\mathcal{D}_2(\tilde{G}_n, \tilde{G}_*) = \mathcal{O}_P(\sqrt{\log(n)/n})$.*

The proof of Theorem 2 is Appendix D.2. Based on the construction of the loss $\mathcal{D}_2(\tilde{G}, \tilde{G}_*)$, Theorem 2 suggests that the convergence rate of the estimators $\mathbf{W}_{2,j} \mathbf{B}_j$ and $\mathbf{W}_{1,j} \mathbf{A}_j$ to the true matrices $\mathbf{W}_{2,j}^* \mathbf{B}_j^*$ and $\mathbf{W}_{1,j}^* \mathbf{A}_j^*$ are ranging from order $\mathcal{O}_P([\log(n)/n]^{\frac{1}{2}})$, if the j -th Voronoi cell \mathcal{V}_j has one element, to order $\mathcal{O}_P([\log(n)/n]^{\frac{1}{4}})$, if the j -th Voronoi cell \mathcal{V}_j has more than one element. In both cases, we only need the sample size with the order at most $\mathcal{O}(\epsilon^{-4})$ to achieve the desired estimation error ϵ , compared with the required sample size of order $\mathcal{O}(\exp(\epsilon^{-1/a}))$ of the non-shared structure analyzed in the previous section.

5 EXPERIMENTS

Experimental Settings. This section presents experimental results evaluating the effectiveness of our proposed methods: τ -DoRA, which refers to DoRA with our stabilization term τ , and DoRAN. We consider two tasks: Image Classification on VTAB-1K and FGVC, and Commonsense Reasoning. We compare our method against several PEFT methods, including LoRA (Hu et al., 2022), *Prefix Tuning* (Li & Liang, 2021), *Parallel Adapter* (He et al., 2022), PiSSA (Meng et al., 2024), and DoRA Liu et al. (2024). We also conduct a Sample Efficiency task comparing DoRAN and DoRA on the Commonsense Reasoning task. **Baseline results for the vision tasks are directly adapted from Xin et al. (2024), while those for the language tasks are taken from Liu et al. (2024). For a fair comparison, DoRAN is evaluated under the same experimental settings as DoRA and LoRA, including using the same ranks and identical data augmentation strategies.** For consistency with the theoretical setting, we fine-tune the query and value projection matrices of each attention layer. Additionally, we present an extended variant of DoRAN that applies low-rank matrices to the proj.up and proj.down matrices in both LLaMA-7B and LLaMA-13B, as described in Ablation B.5. Full hyperparameter details are in Appendix B.1.

Image Classification. We first consider the vision domain, where we aim to fine-tune the ViT-B/16 variant of the Vision Transformer architecture (Dosovitskiy, 2021b), which was fine-tuned on the ImageNet-21K (Deng et al., 2009) dataset.

This experiment consists of two benchmarks. The first is the Visual Task Adaptation Benchmark (VTAB-1K), a transfer-learning benchmark that evaluates vision models on 19 image classification tasks across three domains—Natural, Specialized, and Structured—using only 1000 labeled examples per task. Per-domain results are reported in Table 1, with detailed results in Appendix B.3. When incorporating stabilization into DoRA (τ -DoRA), as described in Section 3.1, the resulting τ -DoRA improves overall performance by 0.5% compared to the original DoRA with nearly no additional computational overhead, suggesting the practical benefits of including this stabilization term. Moreover, with hypernetwork-based parameter generation, our method DoRAN significantly outperforms all baselines, notably improving over DoRA by 1.1% on Natural, 1.6% on Specialized,

and 1.8% on Structured tasks. While hypernetworks inevitably introduce additional parameters, sharing half of the architecture across attention heads and query/value projections greatly reduces this overhead. As a result, DoRAN requires only 0.09% more trainable parameters relative to total parameters of ViT compared to DoRA, yet achieves significant performance gains.

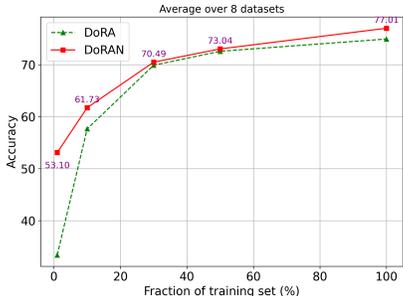


Figure 2: Average sample efficiency on the commonsense reasoning datasets.

The second benchmark is Fine-Grained Visual Classification (FGVC), a family of datasets designed to distinguish highly similar categories within a domain. These datasets evaluate a model’s ability to capture subtle, local visual cues such as textures, parts, and patterns, and are widely used to assess recognition beyond coarse object categories. As reported in Table 2, τ -DoRA significantly improves performance over standard DoRA across all datasets, with an average gain of 1.6%. Furthermore, DoRAN, which combines added noise with hypernetwork-based reparameterization, achieves the best results on nearly all datasets (except Stanford Cars), with notable improvements—+2.3% over DoRA and +4.7% over LoRA—while introducing only 0.09% additional trainable parameters relative to the total parameters of ViT. These results highlight the effectiveness of integrating added noise with the reparameterization mechanism.

Table 2: Image Classification results on the FGVC datasets

Method	#Params (%)	CUB-200 -2011	NABirds	Oxford Flowers	Stanford Dogs	Stanford Cars	Average
FFT	-	87.3	82.7	98.8	89.4	84.5	88.54
LoRA	0.55	84.6	78.2	98.9	85.1	77.1	84.78
DoRA	0.57	87.3	80	99.1	87.6	81.9	87.18
PiSSA	0.55	87.1	80.5	99.1	84.8	82.3	86.66
Adapter	0.47	87.1	84.3	98.5	89.8	68.6	85.66
Prefix	0.42	87.5	82	98	74.2	90.2	86.38
τ -DoRA	0.59	88.3	83.4	99.2	90.2	82.9	88.8
DoRAN	0.66	88.5	85.3	99.2	90.8	83.7	89.5

Commonsense Reasoning. Having shown the effectiveness of both τ -DoRA and DoRAN on vision tasks, we now evaluate them on language tasks using the commonsense reasoning benchmark, which includes eight sub-tasks. Following Hu et al. (2023), we merge all datasets into a 150k training set and test on LLaMA-7B and 13B (Touvron et al., 2023). As shown in Table 3, τ -DoRA outperforms most datasets on LLaMA-7B and improves DoRA by +1.9% (7B) and +0.2% (13B), demonstrating the benefit of added noise. DoRAN, combining noise with hypernetwork-based reparameterization, achieves the best average accuracy on both scales: 77.01% on 7B (+2.0% over DoRA, \approx +3.0% over LoRA) and 80.76% on 13B (+0.7% over DoRA). On LLaMA-7B, it performs especially well on HellaSwag (83.45%), ARC-c (65.02%), and OBQA (79.6%, best in block). Overall, DoRAN consistently improves commonsense reasoning, with gains persisting on harder benchmarks (e.g., ARC-c), while adding only negligible parameters compared to DoRA and LoRA.

Sample Efficiency. In Section 4, we outlined the theoretical advantages of incorporating a shared structure with added noise to improve sample efficiency. Here, we empirically validate this claim by comparing DoRAN with DoRA on the commonsense reasoning task using the LLaMA-7B setting. Following d’Ascoli et al. (2021), we subsample each class at fractions $f = \{1\%, 10\%, 30\%, 50\%, 100\%\}$ and scale the number of training epochs by $1/f$ to keep the total number of training examples constant. The results, presented in Figure 2 and detailed in Appendix B.2, show that DoRAN consistently outperforms DoRA across all fractions. The improvement is especially pronounced in the low-data regime, exceeding 20% when only 1% of the dataset is used, thereby demonstrating the superior sample efficiency of DoRAN over vanilla DoRA.

Table 3: Performance on the Commonsense Reasoning task

Model	Method	#Params (%)	BoolQ	PIQA	SIQA	HellaSwag	WinoGrande	ARC-e	ARC-c	OBQA	Average
LLaMA-7B	Prefix	0.11	64.3	76.8	79.3	42.1	72.1	72.9	54	60.6	65.26
	LoRA	0.25	67.2	79.4	76.6	78.3	78.4	77.1	61.5	74.2	74.09
	DoRA	0.25	67.22	79.98	76.66	80.66	79.72	79.5	61.01	74.8	74.94
	Adapter	0.99	63	79.2	76.3	67.9	75.7	74.5	57.1	72.4	70.76
	τ -DoRA	0.25	69.45	81.39	77.18	83.76	79.56	80.26	64.59	78.6	76.85
	DoRAN	0.26	69.82	81.01	77.89	83.45	79.56	79.76	65.02	79.6	77.01
LLaMA-13B	Prefix	0.03	65.3	75.4	72.1	55.2	68.6	79.5	62.9	68	68.38
	LoRA	0.2	71.7	82.4	79.6	90.4	83.6	83.1	68.5	82.1	80.18
	DoRA	0.2	72.2	83.19	80.81	88.92	81.93	82.95	69.37	81	80.05
	Adapter	0.8	71.8	83	79.2	88.1	82.4	82.5	67.3	81.8	79.51
	τ -DoRA	0.2	71.01	84.39	80.96	89.65	83.74	83	67.24	82.2	80.27
	DoRAN	0.21	71.8	84.5	80.6	89.79	83.19	83.29	68.69	84.2	80.76

6 CONCLUSION

We propose **DoRAN**, a stable yet efficient variant of DoRA. By introducing a noise-stabilized normalization and low-rank updates generated by auxiliary networks, DoRAN addresses key limitations of standard DoRA. Our theoretical analysis established improved convergence guarantees, and empirical results confirmed gains in both stability and sample efficiency.

Limitations and future works: The auxiliary network design introduces additional hyperparameters and structural choices that may complicate practical deployment. Moreover, DoRAN inherits the computational overhead associated with DoRA’s normalization term, as computing the matrix norm $\|\mathbf{W}_0 + \mathbf{B}\mathbf{A}\|$ requires materializing the full matrix, which can be inefficient in certain cases and applications. It is important to note, however, that DoRAN is designed to address an orthogonal limitation of DoRA: its instability and sample inefficiency during training. Improving computational efficiency is therefore complementary to our contribution. Future work may incorporate techniques that mitigate DoRA’s computational cost into the DoRAN framework, yielding a more efficient and robust fine-tuning method. Additionally, we will explore improved hypernetwork architectures to extend DoRAN to broader model families and application domains.

REFERENCES

- Luca Bertinetto, João F Henriques, Jack Valmadre, Philip HS Torr, and Andrea Vedaldi. Learning feed-forward one-shot learners. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 29, 2016.
- Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, Shyamal Buch, Dallas Card, Rodrigo Castellon, Niladri Chatterji, Annie Chen, Kathleen Creel, Jared Quincy Davis, Dora Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kawin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren Gillespie, Karan Goel, Noah Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, Omar Khattab, Pang Wei Koh, Mark Krass, Ranjay Krishna, Rohith Kudipudi, Ananya Kumar, Faisal Ladhak, Mina Lee, Tony Lee, Jure Leskovec, Isabelle Levent, Xiang Lisa Li, Xuechen Li, Tengyu Ma, Ali Malik, Christopher D. Manning, Suvir Mirchandani, Eric Mitchell, Zanele Munyikwa, Suraj Nair, Avaniika Narayan, Deepak Narayanan, Ben Newman, Allen Nie, Juan Carlos Niebles, Hamed Nilforoshan, Julian Nyarko, Giray Ogut, Laurel Orr, Isabel Papadimitriou, Joon Sung Park, Chris Piech, Eva Portelance, Christopher Potts, Aditi Raghunathan, Rob Reich, Hongyu Ren, Frieda Rong, Yusuf Roohani, Camilo Ruiz, Jack Ryan, Christopher Ré, Dorsa Sadigh, Shiori Sagawa, Keshav Santhanam, Andy Shih, Krishnan Srinivasan, Alex Tamkin, Rohan Taori, Armin W. Thomas, Florian Tramèr, Rose E. Wang, William Wang, Bohan Wu, Jiajun Wu, Yuhuai Wu, Sang Michael Xie, Michihiro Yasunaga, Jiaxuan You, Matei Zaharia, Michael Zhang, Tianyi Zhang, Xikun Zhang, Yuhui Zhang, Lucia Zheng, Kaitlyn Zhou, and Percy Liang. On the opportunities and risks of foundation models, 2022. URL <https://arxiv.org/abs/2108.07258>.
- Róbert Csordás, Piotr Piekos, Kazuki Irie, and Jurgen Schmidhuber. Switchhead: Accelerating transformers with mixture-of-experts attention. In *The Thirty-eighth Annual Conference on Neu-*

- 540 *ral Information Processing Systems*, 2024. URL [https://openreview.net/forum?id=](https://openreview.net/forum?id=80SS169GAz)
541 [80SS169GAz](https://openreview.net/forum?id=80SS169GAz).
- 542 Damai Dai, Li Dong, Shuming Ma, Bo Zheng, Zhifang Sui, Baobao Chang, and Furu Wei. StableMoE: Stable routing strategy for mixture of experts. In Smaranda Muresan, Preslav Nakov,
543 and Aline Villavicencio (eds.), *Proceedings of the 60th Annual Meeting of the Association for*
544 *Computational Linguistics (Volume 1: Long Papers)*, pp. 7085–7095, Dublin, Ireland, May
545 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.acl-long.489. URL
546 <https://aclanthology.org/2022.acl-long.489/>.
- 547 <https://aclanthology.org/2022.acl-long.489/>.
- 548 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hier-
549 archical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*,
550 pp. 248–255, 2009. doi: 10.1109/CVPR.2009.5206848.
- 551 Nghiem T Diep, Huy Nguyen, Chau Nguyen, Minh Le, Duy MH Nguyen, Daniel Sonntag, Mathias
552 Niepert, and Nhat Ho. On zero-initialized attention: Optimal prompt and gating factor estimation.
553 *International Conference on Machine Learning (ICML)*, 2025.
- 554 Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale.
555 In *International Conference on Learning Representations (ICLR)*, 2021a.
- 556 Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale.
557 In *International Conference on Learning Representations (ICLR)*, 2021b.
- 558 Nan Du, Yanping Huang, Andrew M Dai, Simon Tong, Dmitry Lepikhin, Yuanzhong Xu, Maxim
559 Krikun, Yanqi Zhou, Adams Wei Yu, Orhan Firat, et al. Glam: Efficient scaling of language
560 models with mixture-of-experts. In *International Conference on Machine Learning*, pp. 5547–
561 5569. PMLR, 2022.
- 562 Stéphane d’Ascoli, Hugo Touvron, Matthew L Leavitt, Ari S Morcos, Giulio Biroli, and Levent
563 Sagun. Convit: Improving vision transformers with soft convolutional inductive biases. In *Inter-*
564 *national conference on machine learning*, pp. 2286–2296. PMLR, 2021.
- 565 David Eigen, Marc’Aurelio Ranzato, and Ilya Sutskever. Learning factored representations in a deep
566 mixture of experts. In *ICLR Workshops*, 2014.
- 567 William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: scaling to trillion parameter
568 models with simple and efficient sparsity. *J. Mach. Learn. Res.*, 23(1), January 2022. ISSN
569 1532-4435.
- 570 David Ha, Andrew M Dai, and Quoc V Le. Hypernetworks. In *International Conference on Learn-*
571 *ing Representations (ICLR)*, 2017.
- 572 Jiaao He, Jiezhong Qiu, Aohan Zeng, Zhilin Yang, Jidong Zhai, and Jie Tang. Fastmoe: A fast
573 mixture-of-expert training system, 2021. URL <https://arxiv.org/abs/2103.13262>.
- 574 Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. Towards
575 a unified view of parameter-efficient transfer learning. In *International Conference on Learning*
576 *Representations (ICLR)*, 2022.
- 577 Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, An-
578 drea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp.
579 In *International conference on machine learning*, pp. 2790–2799. PMLR, 2019.
- 580 Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,
581 and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *ICLR. Open-*
582 *Review.net*, 2022. URL [http://dblp.uni-trier.de/db/conf/iclr/iclr2022.](http://dblp.uni-trier.de/db/conf/iclr/iclr2022.html#HuSWALWWC22)
583 [html#HuSWALWWC22](http://dblp.uni-trier.de/db/conf/iclr/iclr2022.html#HuSWALWWC22).
- 584 Zhiqiang Hu, Lei Wang, Yihui Lan, Wanyu Xu, Ee-Peng Lim, Lidong Bing, Xing Xu, Soujanya
585 Poria, and Roy Lee. LLM-adapters: An adapter family for parameter-efficient fine-tuning of large
586 language models. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023*
587 *Conference on Empirical Methods in Natural Language Processing*, pp. 5254–5276, Singapore,
588 December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.
589 319. URL <https://aclanthology.org/2023.emnlp-main.319/>.

- 594 Muchen Huan and Jianhong Shun. Fine-tuning transformers efficiently: A survey on lora and its
595 impact. *Preprints*, February 2025. doi: 10.20944/preprints202502.1637.v1. URL [https://](https://doi.org/10.20944/preprints202502.1637.v1)
596 doi.org/10.20944/preprints202502.1637.v1.
597
- 598 Robert A. Jacobs, Michael I. Jordan, Steven J. Nowlan, and Geoffrey E. Hinton. Adaptive mixtures
599 of local experts. *Neural Computation*, 3, 1991.
600
- 601 Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and
602 Ser-Nam Lim. Visual prompt tuning. In *Computer Vision – ECCV 2022: 17th European Confer-*
603 *ence, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXIII*, pp. 709–727, Berlin, Hei-
604 delberg, 2022. Springer-Verlag. ISBN 978-3-031-19826-7. doi: 10.1007/978-3-031-19827-4_41.
605 URL https://doi.org/10.1007/978-3-031-19827-4_41.
- 606 Peng Jin, Bo Zhu, Li Yuan, and Shuicheng Yan. Moh: Multi-head attention as mixture-of-head
607 attention, 2025. URL <https://arxiv.org/abs/2410.11842>.
608
- 609 Michael I Jordan and Robert A Jacobs. Hierarchical mixtures of experts and the em algorithm.
610 *Neural computation*, 6(2):181–214, 1994.
611
- 612 Dawid J. Kopiczko, Tijmen Blankevoort, and Yuki M. Asano. Vera: Vector-based random matrix
613 adaptation. In *International Conference on Learning Representations (ICLR)*, 2024.
- 614 Minh Le, Anh Nguyen, Huy Nguyen, Chau Nguyen, Anh Tran, and Nhat Ho. On the expressiveness
615 of visual prompt experts. *arxiv preprint arxiv 2501.18936*, 2025a.
616
- 617 Minh Le, Chau Nguyen, Huy Nguyen, Quyen Tran, Trung Le, and Nhat Ho. Revisiting prefix-
618 tuning: Statistical benefits of reparameterization among prompts. In *The Thirteenth International*
619 *Conference on Learning Representations*, 2025b.
620
- 621 Dmitry Lepikhin, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang,
622 Maxim Krikun, Noam Shazeer, and Zhifeng Chen. GShard: Scaling giant models with conditional
623 computation and automatic sharding. In *International Conference on Learning Representations*,
624 2021. URL <https://openreview.net/forum?id=qrwe7XHTmYb>.
- 625 Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient
626 prompt tuning. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-
627 tau Yih (eds.), *Proceedings of the 2021 Conference on Empirical Methods in Natural Lan-*
628 *guage Processing*, pp. 3045–3059, Online and Punta Cana, Dominican Republic, November
629 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.243. URL
630 <https://aclanthology.org/2021.emnlp-main.243/>.
- 631 Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation.
632 In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (eds.), *Proceedings of the 59th*
633 *Annual Meeting of the Association for Computational Linguistics and the 11th International Joint*
634 *Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 4582–4597, Online,
635 August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.353.
636 URL <https://aclanthology.org/2021.acl-long.353/>.
637
- 638 Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Frank Wang, Kwang-
639 Ting Cheng, and Min-Hung Chen. Dora: weight-decomposed low-rank adaptation. In *Proceed-*
640 *ings of the 41st International Conference on Machine Learning, ICML’24*. JMLR.org, 2024.
- 641 Rabeeh Karimi Mahabadi, Sebastian Ruder, Mostafa Dehghani, and James Henderson. Parameter-
642 efficient multi-task fine-tuning for transformers via shared hypernetworks, 2021. URL [https://](https://arxiv.org/abs/2106.04489)
643 arxiv.org/abs/2106.04489.
644
- 645 Tudor Manole and Nhat Ho. Refined convergence rates for maximum likelihood estimation under
646 finite mixture models. In *Proceedings of the 39th International Conference on Machine Learning*,
647 volume 162 of *Proceedings of Machine Learning Research*, pp. 14979–15006. PMLR, 17–23 Jul
2022.

- 648 Fanxu Meng, Zhaohui Wang, and Muhan Zhang. Pissa: Principal singular values and singular
649 vectors adaptation of large language models. In A. Globerson, L. Mackey, D. Belgrave,
650 A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing
651 Systems*, volume 37, pp. 121038–121072. Curran Associates, Inc., 2024. doi: 10.52202/
652 079017-3846. URL [https://proceedings.neurips.cc/paper_files/paper/
653 2024/file/db36f4d603cc9e3a2a5e10b93e6428f2-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2024/file/db36f4d603cc9e3a2a5e10b93e6428f2-Paper-Conference.pdf).
- 654 Huy Nguyen, TrungTin Nguyen, and Nhat Ho. Demystifying softmax gating function in Gaussian
655 mixture of experts. In *Advances in Neural Information Processing Systems*, 2023.
- 657 Huy Nguyen, Pedram Akbarian, Fanqi Yan, and Nhat Ho. Statistical perspective of top-k sparse
658 softmax gating mixture of experts. In *International Conference on Learning Representations*,
659 2024a.
- 661 Huy Nguyen, Nhat Ho, and Alessandro Rinaldo. Sigmoid gating is more sample efficient than
662 softmax gating in mixture of experts. In *Advances in Neural Information Processing Systems*,
663 2024b.
- 664 Huy Nguyen, Nhat Ho, and Alessandro Rinaldo. On least square estimation in softmax gating
665 mixture of experts. In *Proceedings of the ICML*, 2024c.
- 667 Ngoc-Quan Pham, Tuan Truong, Quyen Tran, Tan Minh Nguyen, Dinh Phung, and Trung Le. Pro-
668 moting ensemble diversity with interactive bayesian distributional robustness for fine-tuning foun-
669 dation models. In *Forty-second International Conference on Machine Learning*, 2025. URL
670 <https://openreview.net/forum?id=yTWqL3XHCC>.
- 671 Samyam Rajbhandari, Conglong Li, Zhewei Yao, Minjia Zhang, Reza Yazdani Aminabadi, Am-
672 mar Ahmad Awan, Jeff Rasley, and Yuxiong He. Deepspeed-moe: Advancing mixture-of-experts
673 inference and training to power next-generation ai scale, 2022. URL [https://arxiv.org/
674 abs/2201.05596](https://arxiv.org/abs/2201.05596).
- 676 Anastasiia Razdaibiedina, Yuning Mao, Madian Khabza, Mike Lewis, Rui Hou, Jimmy Ba, and
677 Amjad Almahairi. Residual prompt tuning: improving prompt tuning with residual reparameter-
678 ization. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Findings of the
679 Association for Computational Linguistics: ACL 2023*, pp. 6740–6757, Toronto, Canada, July
680 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-acl.421. URL
681 <https://aclanthology.org/2023.findings-acl.421/>.
- 682 Stephen Roller, Sainbayar Sukhbaatar, Arthur Szlam, and Jason E Weston. Hash layers for large
683 sparse models. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (eds.), *Ad-
684 vances in Neural Information Processing Systems*, 2021. URL [https://openreview.net/
685 forum?id=1MgDDWb1ULW](https://openreview.net/forum?id=1MgDDWb1ULW).
- 687 Noam Shazeer, *Azalia Mirhoseini, *Krzysztof Maziarsz, Andy Davis, Quoc Le, Geoffrey Hinton,
688 and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-
689 experts layer. In *International Conference on Learning Representations*, 2017. URL [https://openreview.net/
690 forum?id=BlckMDqlg](https://openreview.net/forum?id=BlckMDqlg).
- 691 Haofeng Sun, Hui Tian, Wanli Ni, Jingheng Zheng, Dusit Niyato, and Ping Zhang. Federated
692 low-rank adaptation for large models fine-tuning over wireless networks. *IEEE Transactions on
693 Wireless Communications*, 24(1):659–675, 2025. doi: 10.1109/TWC.2024.3497998.
- 695 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée
696 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and
697 efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- 699 Tuan Truong, Chau Nguyen, Huy Nguyen, Minh Le, Trung Le, and Nhat Ho. Replora: Reparam-
700 eterizing low-rank adaptation via the perspective of mixture of experts. In *Proceedings of the
701 42st International Conference on Machine Learning (ICML)*, Proceedings of Machine Learning
Research, Vancouver, Canada, 2025a. PMLR.

- 702 Tuan Truong, Quyen Tran, Quan Pham-Ngoc, Nhat Ho, Dinh Phung, and Trung Le. Improving
703 generalization with flat hilbert bayesian inference. In *Proceedings of the 42st International Con-*
704 *ference on Machine Learning (ICML)*, Proceedings of Machine Learning Research, Vancouver,
705 Canada, 2025b. PMLR.
- 706
707 Mojtaba Valipour, Mehdi Rezagholizadeh, Ivan Kobyzev, and Ali Ghodsi. Dylora: Parameter ef-
708 ficient tuning of pre-trained models using dynamic search-free low-rank adaptation, 2023. URL
709 <https://arxiv.org/abs/2210.07558>.
- 710 Sara van de Geer. *Empirical processes in M-estimation*. Cambridge University Press, 2000.
- 711
712 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,
713 Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Infor-*
714 *mation Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- 715 Johannes von Oswald, Christian Henning, Benjamin F. Grewe, and João Sacramento. Continual
716 learning with hypernetworks, 2022. URL <https://arxiv.org/abs/1906.00695>.
- 717
718 Junjie Wang, Guangjing Yang, Wentao Chen, Huahui Yi, Xiaohu Wu, Zhouchen Lin, and Qicheng
719 Lao. Mlae: Masked lora experts for visual parameter-efficient fine-tuning. *arxiv preprint arxiv*
720 *2405.18897*, 2024.
- 721 Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani
722 Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto,
723 Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. Emergent abilities of large lan-
724 guage models. *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856. URL
725 <https://openreview.net/forum?id=yzkSU5zdwD>. Survey Certification.
- 726
727 Yi Xin, Siqi Luo, Xuyang Liu, Yuntao Du, Haodi Zhou, Xinyu Cheng, Christina Luoluo Lee, Jun-
728 long Du, Haozhe Wang, MingCai Chen, et al. V-petl bench: A unified visual parameter-efficient
729 transfer learning benchmark. In *The Thirty-eight Conference on Neural Information Processing*
730 *Systems Datasets and Benchmarks Track*, 2024.
- 731
732 Jie Xu, Karthikeyan Saravanan, Rogier van Dalen, Haaris Mehmood, David Tuckey, and Mete Ozay.
733 Dp-dylora: Fine-tuning transformer-based models on-device under differentially private federated
734 learning using dynamic low-rank adaptation, 2025. URL <https://arxiv.org/abs/2405.06368>.
- 735
736 Menglin Yang, Jialin Chen, Yifei Zhang, Jiahong Liu, Jiasheng Zhang, Qiyao Ma, Harshit Verma,
737 Qianru Zhang, Min Zhou, Irwin King, and Rex Ying. Low-rank adaptation for foundation models:
A comprehensive review, 2024. URL <https://arxiv.org/abs/2501.00365>.
- 738
739 Bin Yu. Assouad, Fano, and Le Cam. *Festschrift for Lucien Le Cam*, pp. 423–435, 1997.
- 740
741 Renrui Zhang, Jiaming Han, Chris Liu, Aojun Zhou, Pan Lu, Yu Qiao, Hongsheng Li, and Peng Gao.
742 Llama-adapter: Efficient fine-tuning of large language models with zero-initialized attention. In
The Twelfth International Conference on Learning Representations, 2024a.
- 743
744 Ruiyi Zhang, Rushi Qiang, Sai Ashish Somayajula, and Pengtao Xie. Autolora: Automatically
745 tuning matrix ranks in low-rank adaptation based on meta learning, 2024b. URL [https://](https://arxiv.org/abs/2403.09113)
746 arxiv.org/abs/2403.09113.
- 747
748
749
750
751
752
753
754
755

Supplement to “DoRAN: Stabilizing Weight-Decomposed Low-Rank Adaptation via Noise Injection and Auxiliary Networks”

In this supplementary material, firstly, we provide the deferred related works and notations in Appendix A. Next, we present the additional experiment details in Appendix B.1. Next, detailed proofs for Theorems 1 and 2 are provided in Appendix D.1 and D.2, respectively. Lastly, we discuss the use of large language models in this paper in Appendix E.

A RELATED WORKS

Parameter-Efficient Fine-Tuning. Fine-tuning large pre-trained models has become increasingly expensive with the growing size of these models. Thus, it has motivated a wide range of parameter-efficient fine-tuning (PEFT) methods that only update a relatively small number of parameters. Existing PEFT methods can be divided into three categories. The first category is referred to as *adapter-based* methods, which insert lightweight modules into the frozen backbone. For example, (Houlsby et al., 2019) proposed adding linear modules to the existing layers in sequence, while (He et al., 2022) proposed integrating these modules in parallel to improve performance.

The second category is *prompt-based* methods, which add learnable soft tokens (prompts) to the initial input. As proposed by (Jia et al., 2022; Le et al., 2025a; Lester et al., 2021; Razdaibiedina et al., 2023; Zhang et al., 2024a; Diep et al., 2025), the methods of this paradigm optimize only the additional prompt embeddings while keeping the backbone parameters frozen. This design enables efficient transfer learning with minimal parameter updates. However, despite their efficiency, prompt-based PEFT approaches introduce additional context tokens during inference, thereby increasing computational cost and latency compared to the original frozen backbone models—a limitation that constrains their practical deployment in latency-sensitive applications.

The third category of PEFT is based on low-rank adaptation, which is pioneered by LoRA (Hu et al., 2022) and its variants. LoRA uses low-rank matrices to approximate the weight updates while keeping the pre-trained weights frozen. A key advantage of LoRA is that these low-rank updates can be merged into the original weights before inference, so no extra inference latency is added compared to the original models. Later works have aimed to improve the stability, efficiency, and performance of LoRA. For example, DyLoRA and AutoLoRA (Valipour et al., 2023; Zhang et al., 2024b) improve the performance on downstream tasks by finding the optimal rank of LoRA matrices using various methods. In the parameter efficiency aspect, VeRA (Kopiczko et al., 2024) employs a single pair of shared low-rank matrices across all layers, which further reduces the number of parameters while still maintaining competitive performance. Liu et al. (2024) found that LoRA’s magnitude and direction updates differ significantly from full fine-tuning, which might limit its learning capacity. Thus, DoRA (Liu et al., 2024) was proposed, which decomposes pre-trained weights into magnitude and direction components and uses LoRA to update the direction component, to better approximate full fine-tuning. In parallel, Truong et al. (2025b) and Pham et al. (2025) explore the integration of Bayesian inference and sharpness-aware techniques into LoRA-based fine-tuning frameworks to enhance robustness and generalization. By introducing uncertainty-aware adaptation mechanisms and probabilistic regularization, these methods provide a principled approach to mitigating overfitting and improving the stability of parameter-efficient model tuning. In the parameter efficiency aspect, VeRA Kopiczko et al. (2024) employs a single pair of shared low-rank matrices across all layers, which further reduces the number of parameters while still maintaining competitive performance.

Hypernetworks. The HyperNetwork framework Ha et al. (2017) introduced a paradigm in which the parameters of a target model are dynamically generated by an auxiliary neural network, termed a HyperNetwork, rather than being directly learned. Initially explored in the context of recurrent neural networks, HyperNetworks demonstrated improved adaptability and generalization by producing context-dependent weight updates (Ha et al., 2017). Subsequent research extended this idea to continual learning, where task-specific weight generation via HyperNetworks helped alleviate catastrophic forgetting (von Oswald et al., 2022). In the realm of parameter-efficient fine-tuning (PEFT), HyperNetworks have proven particularly useful for enabling cross-task adaptation and reducing parameter redundancy. For instance, Mahabadi et al. (2021) proposed generating task-specific adapter weights through a shared HyperNetwork, achieving significant parameter savings while maintain-

ing competitive performance. Similarly, Li & Liang (2021) leveraged HyperNetworks to enhance prompt tuning, replacing direct parameter optimization with a meta-network that predicts prompt parameters. More recent works, such as (Le et al., 2025b) and (Truong et al., 2025a), have analyzed the theoretical benefits of HyperNetwork-driven PEFT, showing improved sample efficiency and generalization, thereby motivating novel fine-tuning techniques.

Mixture of Experts. The Mixture of Experts (MoE) framework (Jacobs et al., 1991; Jordan & Jacobs, 1994) has been widely adopted to scale model capacity without proportional increases in computation (Eigen et al., 2014; Shazeer et al., 2017). Modern implementations such as sparsely-gated MoE layers activate only a small set of experts per token, enabling trillion-parameter models with tractable training cost (Shazeer et al., 2017; Lepikhin et al., 2021; Fedus et al., 2022). Efficient routing and load-balancing are central challenges, addressed via auxiliary losses, stochastic routing, and expert capacity constraints (Lepikhin et al., 2021; Fedus et al., 2022; Dai et al., 2022). From a theoretical side, the convergence behavior of MoE and the sample complexity of estimating parameters and experts in MoE have been extensively explored in (Nguyen et al., 2023; 2024a;c;b). MoE has also been linked to the transformer attention mechanism, where attention heads can be viewed as experts and the softmax query-key interactions as gating distributions (Jin et al., 2025; Csordás et al., 2024). RepLoRA (Truong et al., 2025a) formalizes multi-head self-attention as an MoE and shows that LoRA fine-tunes these embedded experts via low-rank updates. Systems-oriented work has improved scalability with hierarchical routing (Du et al., 2022), hashing (Roller et al., 2021), and specialized runtimes (Rajbhandari et al., 2022; He et al., 2021).

B ADDITIONAL EXPERIMENTAL DETAILS

B.1 IMPLEMENTATION DETAILS

For vision tasks, we experiment with ViT-B/16 (Vaswani et al., 2017) for 100 epochs using 100 warmup steps, a batch size of 64, a Low-Rank Matrix rank of 8, and $\alpha = 8$. Optimization is performed with AdamW and a cosine learning rate scheduler. Learning rate and weight decay are tuned via grid search over $\{0.001, 0.005, 0.01, 0.05, 0.1\}$ and $\{0.0001, 0.0005, 0.001, 0.01, 0.1\}$, respectively. For the hypernetwork in a low-rank matrix B , we set the input dimension to 64, the hidden dimension to 16, and use leaky-ReLU activation.

For commonsense reasoning tasks, we use LLaMA-7B (32 layers) and LLaMA-13B (40 layers) Touvron et al. (2023). Training runs for 3 epochs on a single A100 GPU with 100 warmup steps, batch size of 32, learning rate $1e-4$, dropout 0.05, rank 32, and $\alpha = 64$. We optimize the models using the AdamW optimizer with a linear learning rate scheduler. In LLaMA-7B, the hypernetwork for the low-rank matrix B has an input dimension of 64, a hidden dimension of 32, and leaky-ReLU activation; in LLaMA-13B, the hidden dimension is 40 with the same activation.

B.2 DETAIL OF SAMPLE EFFICIENCY

We provide in Figure 3 the details of the sample efficiency problem in each commonsense reasoning dataset with the LLaMA-7B setting.

B.3 DETAIL OF RESULTS ON VTAB-1K DATASETS

In Table 4, we provide the results of DoRAN in detail for each dataset in the VTAB-1K domain. Compared to DoRA, DoRAN delivers consistent improvements across all datasets and achieves state-of-the-art results on nearly all of them—except for CIFAR100 and sNORB-ele—while requiring only a modest increase in parameters. This highlights the benefit of [sharing query and value matrices and using noise, as shown in Section 4](#).

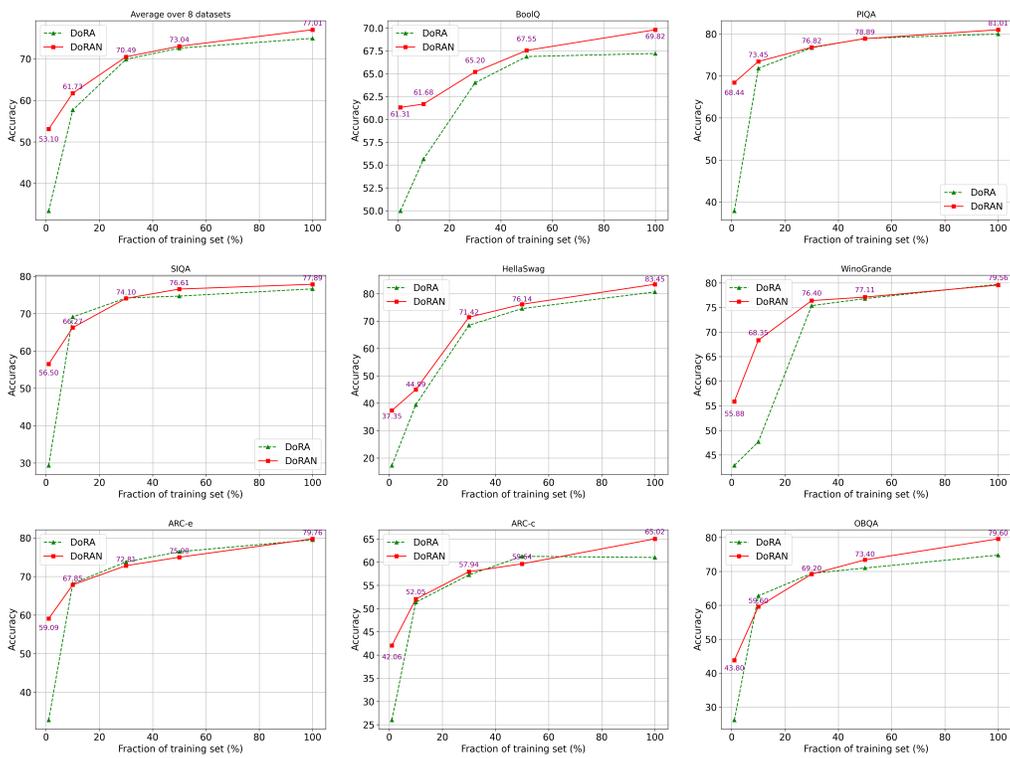


Figure 3: The detail of sample efficiency on each commonsense reasoning dataset with LLaMA-7B settings.

Table 4: Image Classification results on the VTAB-1K dataset

Method	Natural							Specialized				Structured						AVG		
	CIFAR100	Catech101	DTD	Flower102	Pets	SVHN	Sun397	Camelyon	EuroSAT	Resisc45	Retinopathy	Clevr-Count	Clevr-Dist	DMLab	KITTI	dSpre-Loc	dSpre-ori		sNORB-Azim	sNORB-Ele
FFT	68.9	87.7	64.3	97.2	86.9	87.4	38.8	79.7	95.7	84.2	73.9	56.3	58.6	41.7	65.5	57.5	46.7	25.7	29.1	65.6
LoRA	67.1	91.4	69.4	98.2	90.4	85.3	54	84.9	95.3	84.4	73.6	82.9	69.2	49.8	78.5	75.7	47.1	31	44	72.2
DoRA	67.9	90.4	70.6	99	90.2	89.6	54.6	83.9	95.5	85.3	75.9	80.8	69.8	50.5	80.9	79.1	47.7	32.5	39.6	72.8
BSSA	68.9	90.7	71.2	98.7	90.2	87.9	54.7	84.2	95.3	85.5	76	81.4	69.7	51	80.5	78.9	47.2	31.6	41.3	72.9
Adapter	69.2	90.1	68	98.8	89.9	82.8	54.3	84	94.9	81.9	75.5	80.9	65.3	48.6	78.3	74.8	48.5	29.9	41.6	71.4
Prefix	75.5	90.7	65.4	96.6	86	78.5	46.7	79.5	95.1	80.6	74	69.9	58.2	40.9	69.5	72.4	46.8	23.9	34.4	67.6
τ -DoRA	69.4	91.2	71.6	99	90.3	89.6	55.1	84.9	95.5	85.8	76.1	81.9	69.6	52.5	81.0	79.4	47.7	32.8	39.5	73.3
DoRAN	70.2	92.4	71.9	99.1	91.3	90.3	55	86.9	96.4	87.4	76.4	84	70	53.2	81.9	79.9	49.2	35.6	42	74.4

B.4 ABLATION ON THE ROLE OF NOISE

In this ablation study, we assess the role of τ to the performance of DoRAN. To do so, we consider experiments on the FGVC dataset on the following settings: (1) DoRA without hypernetwork and with learnable τ , which corresponds to DoRA- τ in the main experiment; (2) DoRA with the hypernetwork and Gaussian τ , (3) DoRA with the hypernetwork and no noise, and (4) DoRAN. The results are reported in Table 5.

B.5 ABLATION ON FINE-TUNING MODULES

Beyond applying low-rank matrices to the query and value matrices in each layer, we also explore whether our DoRA-based design can generalize when extended to additional modules. We assess the performance of DoRAN on two settings: when it is applied to the query, value, proj_up and proj_down matrices, and when it is applied to every projection matrices while maintaining our pro-

Table 5: Performance with different types of τ on the FGVC datasets

Method	CUB-200-2011	NABirds	Oxford Flowers	Stanford Dogs	Stanford Cars	AVG
Vanilla DoRA	87.3	80	99.1	87.6	81.9	87.18
DoRAN w/o hypernetwork	88.3	83.4	99.2	90.2	82.9	88.8
DoRAN + Gaussian τ	88.3	85.2	99.0	90.5	83.3	89.3
DoRAN w/o τ	87.9	85.4	99.2	90.5	83.5	89.3
DoRAN	88.5	85.3	99.2	90.8	83.7	89.5

posed design for the query and value matrices and using the standard low-rank formulation for the proj_up and proj_down. As in Table 6 and Table 7, in both settings DoRAN achieves the best performance in both LLaMA-7B and LLaMA-13B. These results indicate that our method, originally designed for query and value matrices in multi-head attention, remains effective even when low-rank matrices are applied more broadly across the model.

Table 6: Ablation Study on Low-Rank Matrices in Query, Value, Up, and Down Weights.

Model	Method	#Params (%)	BoolQ	PIQA	SIQA	HellaSwag	WinoGrande	ARC-e	ARC-c	OBQA	Average
LLaMA-7B	LoRA	0.7	69.82	83.51	78.66	85.77	74.27	81.69	66.81	77.8	77.29
	DoRA	0.71	68.69	81.66	78.3	84.24	80.51	81.36	65.44	79.2	77.43
	DoRAN	0.72	70.46	82.1	78.81	84.78	81.61	82.28	66.55	80.6	78.4
LLaMA-13B	LoRA	0.57	72.11	83.73	80.5	90.5	83.74	82.11	68.09	82.4	80.4
	DoRA	0.58	72.42	84.98	81.17	91.81	84.61	84.22	69.88	82.8	81.49
	DoRAN	0.58	72.17	85.91	80.19	92.77	85.71	85.82	71.93	84	82.31

Table 7: Ablation Study on Low-Rank Matrices in Query, Key, Value, Up, and Down Weights.

Model	Method	#Params (%)	BoolQ	PIQA	SIQA	HellaSwag	WinoGrande	ARC-e	ARC-c	OBQA	Average
LLaMA-7B	DoRA	0.84	69.7	82.3	79.5	86.3	81	81.9	66.2	79.2	78.26
	DoRAN	0.84	70.2	83.8	78.6	85	81.1	82	66.7	80.4	78.5
LLaMA-13B	DoRA	0.68	72.4	85.1	80.1	91.9	84.3	84.8	70.9	82.8	81.54
	DoRAN	0.68	72.2	86.1	80.4	92.9	84.9	84.6	71.8	83.8	82.1

B.6 EXPERIMENTS ON RECENT MODEL

To further evaluate the effectiveness of DoRAN on recent and competitive large language models, we conducted an additional experiment using LLaMA3-8B. As shown in Table 8, DoRAN continues to provide measurable benefits at this scale, yielding an average improvement of +0.6% over vanilla DoRA.

B.7 ANALYSIS ON THE VALUES OF τ

To analyze the roles and dynamics of τ , we conducted an experiment on the CIFAR100 dataset and visualize the violin plots of the learned τ values across the 12 ViT layers—both after the first 10 epochs and at the end of training.

At epoch 10, as shown in Figure 4 except the last layer, the distribution of τ values is relatively uniform across layers: both the mean and variance remain relatively similar throughout the depth of the network. This indicates that, early in training, the model has not yet differentiated the roles of different layers in terms of how strongly they should rely on direction learning versus norm control.

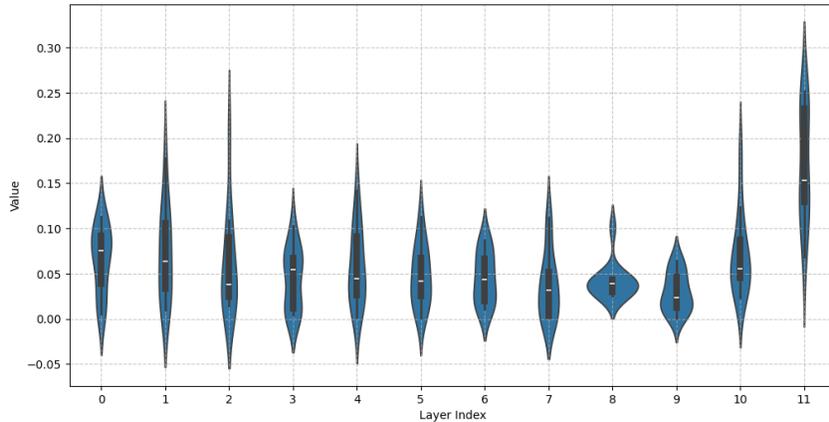
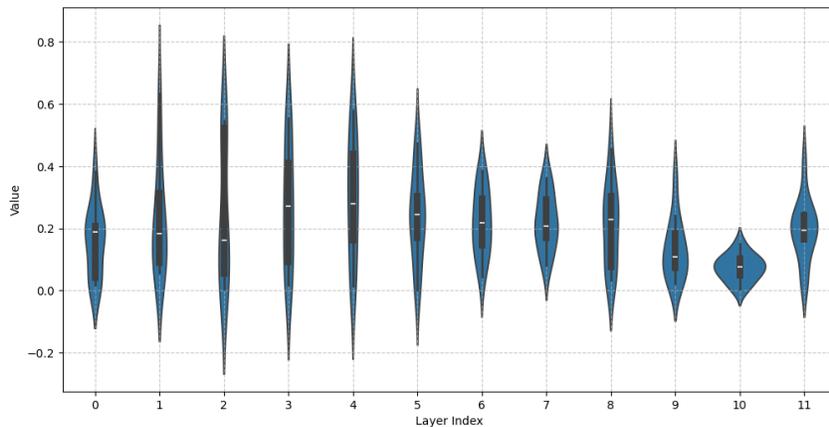
By contrast, in the final epoch, as shown in Figure 5 clear structural patterns emerge. The early layers exhibit larger variance in τ , suggesting that these layers learn more diverse or specialized transformations and therefore benefit from a wider range of direction-vs-norm tradeoffs. The later layers, in comparison, converge to much tighter τ distributions, suggesting that they operate in a more homogeneous regime.

Moreover, the mean τ values increase steadily over training. This trend shows that the model gradually shifts toward stronger direction learning rather than relying solely on norm control, differentiating DoRAN from DoRA and highlighting the additional flexibility introduced by the τ parameter.

Table 8: Performance on the commonsense reasoning task with LLaMA3-8B

Model	Method	BoolQ	PIQA	SIQA	HellaSwag	WinoGrande	ARC-e	ARC-c	OBQA	Average
LLaMA3-8B	DoRA	73.5	87.9	79	93.4	83.7	89.8	78.9	84.4	83.8
	DoRAN	72.9	88.7	80.5	94.7	84.1	90.3	78	85.6	84.4

Overall, these results support our claim that τ adapts meaningfully across depth and provides structural benefits.

Figure 4: Distribution of τ after 10 epochs across 12 layers of the ViTFigure 5: Distribution of final τ after 100 epochs across 12 layers of the ViT

C ADDITIONAL EXPERIMENTS

In this appendix, we perform additional experiments to illustrate how our proposed method DoRAN compares with a strong baseline in the vision domain known as MLAE (Wang et al., 2024).

The configurations for comparing with MLAE: To ensure a fair comparison among low-rank adaptation methods reported in our paper, we emphasize that we apply low-rank fine-tuning to the query and value projection weights for MLAE in both vision and language tasks. This differs from the original implementations of the MLAE paper, which apply low-rank updates to all the projection matrices including the query, key, and value weights.

Because Wang et al. (2024) do not provide hyperparameters and optimization configurations on the FGVC benchmark, we adopt a dropout rate of 0.5 and set the coefficient initialization value to 1.0

for MLAE on this benchmark. All other training settings—including the learning rate, weight decay, and other optimization configurations—are kept consistent with those used in our proposed method. For the VTAB-1K experiments, we follow the hyperparameters and optimizer settings reported in the original paper.

For the Commonsense Reasoning task, we reuse the FGVC hyperparameters for reproducing MLAE: a dropout rate of 0.5 and a coefficient initialization value of 1.0.

Comparison with MLAE: The results are summarized in the tables below. These results show that although MLAE achieves a 0.5% performance gain over DoRAN on FGVC, DoRAN surpasses MLAE by 1.7% on the VTAB-1K benchmark and further outperforms it on the language tasks by 0.91% for LLaMA-7B and 1.04% for LLaMA-13B. This demonstrates that DoRAN is competitive with state-of-the-art approaches including MLAE, while also exhibiting strong flexibility and applicability across multiple domains.

Table 9: Image Classification results on the FGVC datasets

Method	CUB-200 -2011	NABirds	Oxford Flowers	Stanford Dogs	Stanford Cars	Average
LoRA	84.6	78.2	98.9	85.1	77.1	84.78
DoRA	87.3	80	99.1	87.6	81.9	87.18
MLAE	89.5	85.5	99.2	91.2	84.4	90
DoRAN	88.5	85.3	99.2	90.8	83.7	89.5

Table 10: Image Classification results on the VTAB-1K dataset

Method	Natural							Specialized				Structured							AVG	
	CIFAR100	Caltech101	DTD	Flower102	Pets	SVHN	Sun397	Camelyon	EuroSAT	Resisc45	Retinopathy	Clevr-Count	Clevr-Dist	DMLab	KITTI	dSpr-Loc	dSpr-ori	sNORB-Azim		sNORB-Ele
LoRA	67.1	91.4	69.4	98.2	90.4	85.3	54	84.9	95.3	84.4	73.6	82.9	69.2	49.8	78.5	75.7	47.1	31	44	72.2
DoRA	67.9	90.4	70.6	99	90.2	89.6	54.6	83.9	95.5	85.3	75.9	80.8	69.8	50.5	80.9	79.1	47.7	32.5	39.6	72.8
MLAE	71.1	91.4	72.7	99.1	91.8	87.8	55.7	83	95.7	83.7	75.8	81.3	69.1	50.6	79.9	77.5	44.9	28.7	41.2	72.7
DoRAN	70.2	92.4	71.9	99.1	91.3	90.3	55	86.9	96.4	87.4	76.4	84	70	53.2	81.9	79.9	49.2	35.6	42	74.4

Table 11: Performance on the Commonsense Reasoning task

Model	Method	#Params (%)	BoolQ	PIQA	SIQA	HellaSwag	WinoGrande	ARC-e	ARC-c	OBQA	Average
LLaMA-7B	LoRA	0.25	67.2	79.4	76.6	78.3	78.4	77.1	61.5	74.2	74.09
	DoRA	0.25	67.22	79.98	76.66	80.66	79.72	79.5	61.01	74.8	74.94
	MLAE	0.25	69.02	81.18	77.02	78.8	80.51	80.6	63.05	78.6	76.1
	DoRAN	0.26	69.82	81.01	77.89	83.45	79.56	79.76	65.02	79.6	77.01
LLaMA-13B	LoRA	0.2	71.7	82.4	79.6	90.4	83.6	83.1	68.5	82.1	80.18
	DoRA	0.2	72.2	83.19	80.81	88.92	81.93	82.95	69.37	81	80.05
	MLAE	0.2	70.12	83.13	79.22	89.04	81.77	83.25	67.83	83.4	79.72
	DoRAN	0.21	71.8	84.5	80.6	89.79	83.19	83.29	68.69	84.2	80.76

Future direction. To further highlight the flexibility of DoRAN, we note that DoRAN can be naturally combined with MLAE. In MLAE, the stochastic masking mechanism is implemented as

$$\Delta W = M \odot \Lambda \odot \varepsilon$$

Equivalently, for each layer $l \in [L]$, the forward computation can be written as

$$Y = Y_{\text{pretrained}} + Y_{\text{MLAE}}, Y_{\text{MLAE}} = xA^T \text{diag}(\text{Dropout}_p(\lambda))B$$

Importantly, MLAE only introduces a masking mechanism — it does not modify the architecture of the underlying low-rank matrices. This makes it straightforward to incorporate MLAE into DoRAN. A natural combination strategy is to first generate the low-rank matrices A and B using the DoRAN hypernetworks, and then apply MLAE’s masked update within DoRAN’s normalized forward pass:

$$Y_{\text{DoRAN+MLAE}} = x \frac{W_0 + W'}{\|W_0 + W'\| + \tau}, W' = A^\top \text{diag}(\text{Dropout}_p(\lambda))B$$

We believe this hybrid approach could leverage the strengths of both DoRAN and MLAE. Exploring this direction is promising and we leave a thorough investigation of its benefits for future work.

D PROOFS

D.1 PROOF OF THEOREM 1

Firstly, we demonstrate that the following limit holds for any $r \geq 1$:

$$\lim_{\varepsilon \rightarrow 0} \inf_{G \in \mathcal{G}_{L'}(\Theta): \mathcal{D}_{1,r}(G, G_*) \leq \varepsilon} \frac{\|f_G - f_{G_*}\|_{L^2(\mu)}}{\mathcal{D}_{1,r}(G, G_*)} = 0. \quad (8)$$

To prove this, we construct a mixing measure sequence $(G_n)_n$ that satisfies both $\mathcal{D}_{1,r}(G_n, G_*) \rightarrow 0$ and $\|f_{G_n} - f_{G_*}\|_{L^2(\mu)}/\mathcal{D}_{1,r}(G_n, G_*) \rightarrow 0$, as $n \rightarrow \infty$. Specifically, we consider the sequence of mixing measure with $L + 1$ atoms, $G_n = \sum_{i=1}^{L+1} \exp(c_i^n) \delta_{(B_{Q,i}^n, A_{Q,i}^n, B_{V,i}^n, A_{V,i}^n)}$, where,

- $\exp(c_1^n) = \exp(c_2^n) = \frac{1}{2} \exp(c_1^*) + \frac{1}{2n^{r+1}}$ and $\exp(c_i^n) = \exp(c_{i-1}^*)$ for any $3 \leq i \leq L+1$;
- $B_{Q,1}^n = (1 + \frac{1}{n}) B_{Q,1}^* + \frac{1}{\sqrt{n}} \mathbf{S}$, $B_{Q,2}^n = (1 + \frac{1}{n}) B_{Q,1}^* - \frac{1}{\sqrt{n}} \mathbf{S}$ and $B_{Q,i}^n = B_{Q,i-1}^*$ for any $3 \leq i \leq L+1$;
- $A_{Q,1}^n = A_{Q,1}^* + \frac{1}{\sqrt{n}} \mathbf{T}$, $A_{Q,2}^n = A_{Q,1}^* - \frac{1}{\sqrt{n}} \mathbf{T}$ and $A_{Q,i}^n = A_{Q,i-1}^*$ for any $3 \leq i \leq L+1$;
- $B_{V,1}^n = B_{V,2}^n = B_{V,1}^*$ and $B_{V,i}^n = B_{V,i-1}^*$ for any $3 \leq i \leq L+1$,
- $A_{V,1}^n = A_{V,2}^n = A_{V,1}^*$ and $A_{V,i}^n = A_{V,i-1}^*$ for any $3 \leq i \leq L+1$,

where $\mathbf{S} \in \mathbb{R}^{d \times r}$ and $\mathbf{T} \in \mathbb{R}^{r \times d}$ are chosen to satisfy $\mathbf{S}\mathbf{T} = \mathbf{C}_Q$ if rank of \mathbf{C}_Q is less than or equal r . We note that choosing \mathbf{S} and \mathbf{T} so that $\mathbf{S}\mathbf{T} = \mathbf{C}_Q$ can be done easily using the singular vectors and singular values of \mathbf{C}_Q . Then, if we have $\mathbf{S}\mathbf{T} = \mathbf{C}_Q$, it is clear that,

$$B_{Q,1}^n A_{Q,1}^n + B_{Q,2}^n A_{Q,2}^n = 2 \left(1 + \frac{1}{n}\right) B_{Q,1}^* A_{Q,1}^* + \frac{2}{n} \mathbf{C}_Q.$$

This identity is necessary for the later part of our proof. Now, if rank of \mathbf{C}_Q is larger than r , we can write \mathbf{C}_Q as the sum of multiple matrices, let's say $\mathbf{C}_{Q,1}$ and $\mathbf{C}_{Q,2}$, each with rank less than or equal to r and build the mixing measure similarly with more atoms, so that we can recover the sum $\mathbf{C}_{Q,1} + \mathbf{C}_{Q,2} = \mathbf{C}_Q$. Just for this example, the goal is to have the following,

$$B_{Q,1}^n A_{Q,1}^n + B_{Q,2}^n A_{Q,2}^n = 2 \left(1 + \frac{1}{2n}\right) B_{Q,1}^* A_{Q,1}^* + \frac{2}{n} \mathbf{C}_{Q,1},$$

$$B_{Q,3}^n A_{Q,3}^n + B_{Q,4}^n A_{Q,4}^n = 2 \left(1 + \frac{1}{2n}\right) B_{Q,1}^* A_{Q,1}^* + \frac{2}{n} \mathbf{C}_{Q,2}.$$

We go back to the original chosen values above. Computing the loss function $\mathcal{D}_{1,r}(G_n, G_*)$ yields

$$\mathcal{D}_{1,r}(G_n, G_*) = \frac{1}{n^{r+1}} + \left[\exp(c_1^*) + \frac{1}{n^{r+1}} \right] \cdot \frac{1}{n^r} \|B_{Q,1}^*\|^r = \mathcal{O}(n^{-r}). \quad (9)$$

It can be seen that $\mathcal{D}_{1,r}(G_n, G_*) \rightarrow 0$ as $n \rightarrow \infty$.

Now we show that $\|f_{G_n} - f_{G_*}\|_{L^2(\mu)}/\mathcal{D}_{1,r}(G_n, G_*) \rightarrow 0$. Indeed, consider the below quantity,

$$T_n(\mathbf{X}) := \left(\sum_{j=1}^L \exp \left(\mathbf{X}^\top m_{Q,j} \frac{\mathbf{C}_Q + B_{Q,j}^* A_{Q,j}^*}{\|\mathbf{C}_Q + B_{Q,j}^* A_{Q,j}^*\|} \mathbf{C}_K \mathbf{X} + c_j^* \right) \right) \cdot [f_{G_n}(\mathbf{X}) - f_{G_*}(\mathbf{X})],$$

which can be decomposed as follows:

$$\begin{aligned}
T_n(\mathbf{X}) &= \sum_{j=1}^L \sum_{i \in \mathcal{V}_j} \exp(c_i^n) \left[\exp \left(\mathbf{X}^\top m_{Q,j} \frac{\mathbf{C}_Q + \mathbf{B}_{Q,j}^n \mathbf{A}_{Q,j}^n}{\|\mathbf{C}_Q + \mathbf{B}_{Q,j}^n \mathbf{A}_{Q,j}^n\|} \mathbf{C}_K \mathbf{X} \right) \left(m_{V,j} \frac{\mathbf{C}_{V,j} + \mathbf{B}_{V,j}^n \mathbf{A}_{V,j}^n}{\|\mathbf{C}_{V,j} + \mathbf{B}_{V,j}^n \mathbf{A}_{V,j}^n\|} \right) \mathbf{X} \right. \\
&\quad \left. - \exp \left(\mathbf{X}^\top m_{Q,j}^* \frac{\mathbf{C}_Q + \mathbf{B}_{Q,j}^* \mathbf{A}_{Q,j}^*}{\|\mathbf{C}_Q + \mathbf{B}_{Q,j}^* \mathbf{A}_{Q,j}^*\|} \mathbf{C}_K \mathbf{X} \right) \left(m_{V,j} \frac{\mathbf{C}_V + \mathbf{B}_{V,j}^* \mathbf{A}_{V,j}^*}{\|\mathbf{C}_V + \mathbf{B}_{V,j}^* \mathbf{A}_{V,j}^*\|} \right) \mathbf{X} \right] \\
&\quad - \sum_{j=1}^L \sum_{i \in \mathcal{V}_j} \exp(c_i^n) \left[\exp \left(\mathbf{X}^\top m_{Q,j} \frac{\mathbf{C}_Q + \mathbf{B}_{Q,j}^n \mathbf{A}_{Q,j}^n}{\|\mathbf{C}_Q + \mathbf{B}_{Q,j}^n \mathbf{A}_{Q,j}^n\|} \mathbf{C}_K \mathbf{X} \right) \right. \\
&\quad \left. - \exp \left(\mathbf{X}^\top m_{Q,j} \frac{\mathbf{C}_Q + \mathbf{B}_{Q,j}^* \mathbf{A}_{Q,j}^*}{\|\mathbf{C}_Q + \mathbf{B}_{Q,j}^* \mathbf{A}_{Q,j}^*\|} \mathbf{C}_K \mathbf{X} \right) \right] f_{G_n}(\mathbb{X}) \\
&\quad + \sum_{j=1}^L \left(\sum_{i \in \mathcal{V}_j} \exp(c_i^n) - \exp(c_j^*) \right) \exp \left(\mathbf{X}^\top m_{Q,j} \frac{\mathbf{C}_Q + \mathbf{B}_{Q,j}^* \mathbf{A}_{Q,j}^*}{\|\mathbf{C}_Q + \mathbf{B}_{Q,j}^* \mathbf{A}_{Q,j}^*\|} \mathbf{C}_K \mathbf{X} \right) \left[m_{V,j} \frac{\mathbf{C}_V + \mathbf{B}_{V,j}^* \mathbf{A}_{V,j}^*}{\|\mathbf{C}_V + \mathbf{B}_{V,j}^* \mathbf{A}_{V,j}^*\|} \mathbf{X} - f_{G_n}(\mathbf{X}) \right] \\
&:= A_n(\mathbf{X}) - B_n(\mathbf{X}) + C_n(\mathbf{X}).
\end{aligned}$$

It follows from the choices of $\mathbf{B}_{Q,i}^n, \mathbf{A}_{Q,i}^n, \mathbf{B}_{V,i}^n, \mathbf{A}_{V,i}^n$ and c_i^n that

$$\begin{aligned}
A_n(\mathbf{X}) &= \sum_{k=1}^2 \frac{1}{2} \left[\exp(c_1^*) + \frac{1}{n^{r+1}} \right] \left(m_{V,1} \frac{\mathbf{C}_V + \mathbf{B}_{V,1}^* \mathbf{A}_{V,1}^*}{\|\mathbf{C}_V + \mathbf{B}_{V,1}^* \mathbf{A}_{V,1}^*\|} \right) \mathbf{X} \\
&\quad \times \left(\exp \left(\mathbf{X}^\top m_{Q,k} \frac{\mathbf{C}_Q + \mathbf{B}_{Q,k}^n \mathbf{A}_{Q,k}^n}{\|\mathbf{C}_Q + \mathbf{B}_{Q,k}^n \mathbf{A}_{Q,k}^n\|} \mathbf{C}_K \mathbf{X} \right) - \exp \left(\mathbf{X}^\top m_{Q,k} \frac{\mathbf{C}_Q + \mathbf{B}_{Q,k}^* \mathbf{A}_{Q,k}^*}{\|\mathbf{C}_Q + \mathbf{B}_{Q,k}^* \mathbf{A}_{Q,k}^*\|} \mathbf{C}_K \mathbf{X} \right) \right)
\end{aligned}$$

We denote $L(\mathbf{Z}) := \exp \left(\mathbf{X}^\top m_Q \frac{\mathbf{C}_Q + \mathbf{Z}}{\|\mathbf{C}_Q + \mathbf{Z}\|} \mathbf{C}_K \mathbf{X} \right)$. Then, by using first-order Taylor expansion, we have that,

$$L(\mathbf{B}_{Q,1}^n \mathbf{A}_{Q,1}^n) - L(\mathbf{B}_{Q,1}^* \mathbf{A}_{Q,1}^*) = \langle \mathbf{B}_{Q,1}^n \mathbf{A}_{Q,1}^n - \mathbf{B}_{Q,1}^* \mathbf{A}_{Q,1}^*, \frac{\partial L}{\partial \mathbf{B} \mathbf{A}}(\mathbf{B}_{Q,1}^* \mathbf{A}_{Q,1}^*) \rangle + R_1(\mathbf{X}),$$

$$L(\mathbf{B}_{Q,2}^n \mathbf{A}_{Q,2}^n) - L(\mathbf{B}_{Q,1}^* \mathbf{A}_{Q,1}^*) = \langle \mathbf{B}_{Q,2}^n \mathbf{A}_{Q,2}^n - \mathbf{B}_{Q,1}^* \mathbf{A}_{Q,1}^*, \frac{\partial L}{\partial \mathbf{B} \mathbf{A}}(\mathbf{B}_{Q,1}^* \mathbf{A}_{Q,1}^*) \rangle + R_2(\mathbf{X}),$$

where $R_1(\mathbf{X})$ and $R_2(\mathbf{X})$ are Taylor remainder such that $R_1/\mathcal{D}_1(G_n, G_*) \rightarrow 0$ and $R_2/\mathcal{D}_1(G_n, G_*) \rightarrow 0$ as $n \rightarrow \infty$. It follows based on the chosen $\mathbf{B}_{Q,1}^n, \mathbf{A}_{Q,1}^n, \mathbf{B}_{Q,2}^n, \mathbf{A}_{Q,2}^n$,

$$\begin{aligned}
&L(\mathbf{B}_{Q,1}^n \mathbf{A}_{Q,1}^n) - L(\mathbf{B}_{Q,1}^* \mathbf{A}_{Q,1}^*) + L(\mathbf{B}_{Q,2}^n \mathbf{A}_{Q,2}^n) - L(\mathbf{B}_{Q,1}^* \mathbf{A}_{Q,1}^*) \\
&= \frac{1}{n} \langle \mathbf{B}_{Q,1}^* \mathbf{A}_{Q,1}^* + \mathbf{C}_Q, \frac{\partial L}{\partial \mathbf{B} \mathbf{A}}(\mathbf{B}_{Q,1}^* \mathbf{A}_{Q,1}^*) \rangle + R_1 + R_2, \\
&= R_1 + R_2,
\end{aligned}$$

where the inner product is 0 due to Lemma Therefore, we have that $A_n(\mathbf{X})/\mathcal{D}_1(G_n, G_*) \rightarrow 0$ as $n \rightarrow \infty$.

Additionally, it is similar to verify that $B_n(\mathbf{X})/\mathcal{D}_1(G_n, G_*) \rightarrow 0$, and $C_n(\mathbf{X}) = \mathcal{O}(n^{-(r+1)})$. Thus, $T_n(\mathbf{X})/\mathcal{D}_{1,r}(G_n, G_*) \rightarrow 0$ as $n \rightarrow \infty$ for almost every \mathbf{X} .

Since the term $\sum_{k=1}^L \exp \left(\mathbf{X}^\top m_{Q,k} \frac{\mathbf{C}_Q + \mathbf{B}_{Q,k}^* \mathbf{A}_{Q,k}^*}{\|\mathbf{C}_Q + \mathbf{B}_{Q,k}^* \mathbf{A}_{Q,k}^*\|} \mathbf{C}_K \mathbf{X} + c_k^* \right)$ is bounded, we have $[f_{G_n}(\mathbf{X}) - f_{G_*}(\mathbf{X})]/\mathcal{D}_{1,r}(G_n, G_*) \rightarrow 0$ for almost every \mathbb{X} . Therefore, we have

$$\|f_{G_n} - f_{G_*}\|_{L^2(\mu)}/\mathcal{D}_{1,r}(G_n, G_*) \rightarrow 0$$

as $n \rightarrow \infty$. We obtain the claim in equation (8).

Step 2. Now, we are ready to prove the desired result, i.e.,

$$\inf_{G_n \in \mathcal{G}_{L'}(\Theta)} \sup_{G \in \mathcal{G}_{L'}(\Theta) \setminus \mathcal{G}_{L-1}(\Theta)} \mathbb{E}_{f_G} [\mathcal{D}_{1,r}(G_n, G)] \gtrsim n^{-1/2}. \quad (10)$$

Since the noise variables ε_i follow from the Gaussian distribution, we have that $Y_i | \mathbf{X}_i \sim \mathcal{N}(f_{G_*}(\mathbf{X}_i), \sigma^2)$ for all $i \in [n]$. For sufficiently small $\varepsilon > 0$ and a fixed constant $C_1 > 0$ which we will choose later, there exists a mixing measure $G'_* \in \mathcal{G}_{L'}(\Theta)$ such that $\mathcal{D}_{1,r}(G'_*, G_*) = 2\varepsilon$ and $\|f_{G'_*} - f_{G_*}\|_{L^2(\mu)} \leq C_1\varepsilon$ thanks to the result in equation (8). Due to the Le Cam's lemma Yu (1997) and the fact that the Voronoi loss function $\mathcal{D}_{1,r}$ satisfies the weak triangle inequality, we can derive that

$$\begin{aligned} & \inf_{G_n \in \mathcal{G}_{L'}(\Theta)} \sup_{G \in \mathcal{G}_{L'}(\Theta) \setminus \mathcal{G}_{L-1}(\Theta)} \mathbb{E}_{f_G}[\mathcal{D}_{1,r}(G_n, G)] \\ & \gtrsim \frac{\mathcal{D}_{1,r}(G'_*, G_*)}{8} \exp(-n \mathbb{E}_{\mathbf{X} \sim \mu}[\text{KL}(\mathcal{N}(f_{G'_*}(\mathbf{X}), \sigma^2), \mathcal{N}(f_{G_*}(\mathbf{X}), \sigma^2))]) \\ & \gtrsim \varepsilon \cdot \exp(-n \|f_{G'_*} - f_{G_*}\|_{L^2(\mu)}^2) \\ & \gtrsim \varepsilon \cdot \exp(-C_1 n \varepsilon^2), \end{aligned} \quad (11)$$

where the second inequality follows from the KL distance of two multivariate Gaussians, i.e.,

$$\text{KL}(\mathcal{N}(f_{G'_*}(\mathbf{X}), \sigma^2), \mathcal{N}(f_{G_*}(\mathbf{X}), \sigma^2)) = \frac{(f_{G'_*}(\mathbf{X}) - f_{G_*}(\mathbf{X}))^2}{2\sigma^2}.$$

Let $\varepsilon = n^{-1/2}$, then we get that the RHS $\varepsilon \cdot \exp(-C_1 n \varepsilon^2) = n^{-1/2} \exp(-C_1)$. Thus, we achieve the desired minimax lower bound in equation (10).

Lemma 1. Let $L(\mathbf{Z}) := \exp\left(\mathbf{X}^\top m_Q \frac{\mathbf{C}_Q + \mathbf{Z}}{\|\mathbf{C}_Q + \mathbf{Z}\|} \mathbf{C}_K \mathbf{X}\right)$. We have that,

$$\langle \mathbf{C}_Q + \mathbf{Z}, \frac{\partial L}{\partial \mathbf{Z}} \rangle = 0.$$

Proof of Lemma 1: It follows from direct calculation,

$$\begin{aligned} \frac{\partial L}{\partial \mathbf{Z}} &= m_Q \frac{\mathbf{X}(\mathbf{C}_K \mathbf{X})^\top}{\|\mathbf{C}_Q + \mathbf{Z}\|} - m_Q (\mathbf{X}^\top (\mathbf{C}_Q + \mathbf{Z}) \mathbf{C}_K \mathbf{X}) \frac{\mathbf{C}_Q + \mathbf{Z}}{\|\mathbf{C}_Q + \mathbf{Z}\|^3}. \\ \langle \mathbf{C}_Q + \mathbf{Z}, \frac{\partial L}{\partial \mathbf{Z}} \rangle &= \frac{m_Q}{\|\mathbf{C}_Q + \mathbf{Z}\|} \text{Trace}((\mathbf{C}_Q + \mathbf{Z})^\top \mathbf{X} \mathbf{X}^\top \mathbf{C}_K^\top - \mathbf{X}^\top (\mathbf{C}_Q + \mathbf{Z}) \mathbf{C}_K \mathbf{X}) = 0. \end{aligned}$$

D.2 PROOF OF THEOREM 2

Before going to the proof, without loss of generality, we can assume $\mathbf{W}_{1,j}, \mathbf{W}_{2,j}$ are identity matrices for each j . In particular, we may denote $\sigma_1(\mathbf{W}_1 \mathbf{A})$ as $\sigma_1(\mathbf{A})$ for an input matrix \mathbf{A} . Throughout this proof, we also assume without loss of generality that $\mathbf{C}_{K,j} = I_d$ with a note that our techniques can be extended to the general setting of that matrix.

We first start with the following result regarding the convergence rate of the regression function estimation $f_{\tilde{G}_n}$ to the true regression function $f_{\tilde{G}_*}$:

Proposition 1. Given the least square estimator \tilde{G}_n , the convergence rate of the regression function estimator $f_{\tilde{G}_n}$ to the true regression function $f_{\tilde{G}_*}$ under the $L^2(\mu)$ norm is,

$$\|f_{\tilde{G}_n} - f_{\tilde{G}_*}\|_{L^2(\mu)} = \mathcal{O}_P(\sqrt{\log(n)/n}). \quad (12)$$

Our goal now is to demonstrate the following inequality:

$$\inf_{\tilde{G} \in \tilde{\mathcal{G}}_{L'}(\Theta)} \|f_{\tilde{G}} - f_{\tilde{G}_*}\|_{L^2(\mu)} / \mathcal{D}_2(\tilde{G}, \tilde{G}_*) > 0,$$

and since from Proposition 1, the rate of $\|f_{\tilde{G}} - f_{\tilde{G}_*}\|_{L^2(\mu)}$ is $\mathcal{O}_P(\sqrt{\log(n)/n})$, we deduce that $\mathcal{D}_2(\tilde{G}, \tilde{G}_*)$ is also $\mathcal{O}_P(\sqrt{\log(n)/n})$. We will prove this inequality by considering separately the local part when the Voronoi loss is small and the global part when the Voronoi loss is large. Before delving into the proof, we first state some essential assumptions on the activation function σ_1 and σ_2 .

Assumptions. We impose the following assumptions on the activation functions σ_1 and σ_2 :

(A.1) (*Algebraic Independence*) For any pair of matrices $(\mathbf{B}_1, \mathbf{A}_1)$ and $(\mathbf{B}_2, \mathbf{A}_2)$ such that $\sigma_2(\mathbf{B}_1)\sigma_1(\mathbf{A}_1) = \sigma_2(\mathbf{B}_2)\sigma_1(\mathbf{A}_2)$, then it follows that $\mathbf{B}_1 = \mathbf{B}_2$ and $\mathbf{A}_1 = \mathbf{A}_2$.

(A.2) (*Uniform Lipschitz*) Denote

$$\mathbf{F}(\mathbf{X}, \mathbf{A}, \mathbf{B}) := \exp\left(\mathbf{X}^\top m_{Q,k} \frac{\mathbf{C}_Q + \sigma_2(\mathbf{B})\sigma_1(\mathbf{A})}{\|\mathbf{C}_Q + \sigma_2(\mathbf{B})\sigma_1(\mathbf{A})\| + \tau_Q} \mathbf{X}\right) \left(m_{V,j}^* \frac{\mathbf{C}_V + \sigma_2(\mathbf{B})\sigma_1(\mathbf{A})}{\|\mathbf{C}_V + \sigma_2(\mathbf{B})\sigma_1(\mathbf{A})\| + \tau_V}\right) \mathbf{X},$$

then for any $\tau \in \{1, 2\}$,

$$\sum_{|\alpha|=\tau} \left\| \left(\frac{\partial^{|\alpha|} \mathbf{F}}{\partial \mathbf{A}^{\alpha_1} \partial \mathbf{B}^{\alpha_2}}(\mathbb{X}, \mathbf{A}_1, \mathbf{B}_1) - \frac{\partial^{|\alpha|} \mathbf{F}}{\partial \mathbf{A}^{\alpha_1} \partial \mathbf{B}^{\alpha_2}}(\mathbb{X}, \mathbf{A}_2, \mathbf{B}_2) \right) \gamma^\alpha \right\| \leq C \|(\mathbf{A}_1, \mathbf{B}_1) - (\mathbf{A}_2, \mathbf{B}_2)\|^a \|\gamma\|^a,$$

for any vector $\gamma \in \mathbb{R}^{2dr}$ and for some positive constant a and C independent of input \mathbb{X} and parameter $\mathbf{A}_1, \mathbf{B}_1, \mathbf{A}_2, \mathbf{B}_2$. We denote the index $\alpha = (\alpha_1, \alpha_2) \in \mathbb{N}^{r \times d} \times \mathbb{N}^{d \times r}$ and \mathbf{A}^{α_1} will return the entries of \mathbf{A} at the position that the element of α_1 is non-zero.

(A.3) (*Strong identifiability*) We denote,

$$\begin{aligned} \mathbf{M} &:= \frac{\mathbf{C}_Q + \sigma_2(\mathbf{B})\sigma_1(\mathbf{A})}{\|\mathbf{C}_Q + \sigma_2(\mathbf{B})\sigma_1(\mathbf{A})\| + \tau_Q}, \\ \mathbf{N} &:= \frac{\mathbf{C}_V + \sigma_2(\mathbf{B})\sigma_1(\mathbf{A})}{\|\mathbf{C}_V + \sigma_2(\mathbf{B})\sigma_1(\mathbf{A})\| + \tau_V}, \end{aligned}$$

Then, for any non-positive integer ℓ and distinct matrices $\{(\mathbf{B}_j, \mathbf{A}_j)\}_{j \in [\ell]}$, the functions in the set below are linear independent for almost sure \mathbf{X} :

$$\left\{ \begin{aligned} &\left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{A}^{(u)}} \mathbf{X} \right) \mathbf{N}_j^* \mathbf{X}, \quad \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{B}^{(u)}} \mathbf{X} \right) \mathbf{N}_j^* \mathbf{X}, \quad \frac{\partial \mathbf{N}_j^*}{\partial \mathbf{A}^{(u)}} \mathbf{X}, \quad \frac{\partial \mathbf{N}_j^*}{\partial \mathbf{B}^{(u)}} \mathbf{X}, \\ &\left(\mathbf{X}^\top \frac{\partial^2 \mathbf{M}_j^*}{\partial \mathbf{A}^{(u)} \partial \mathbf{A}^{(v)}} \mathbf{X} \right) \mathbf{N}_j^* \mathbf{X}, \quad \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{A}^{(u)}} \mathbf{X} \mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{A}^{(v)}} \mathbf{X} \right) \mathbf{N}_j^* \mathbf{X}, \\ &\left(\mathbf{X}^\top \frac{\partial^2 \mathbf{M}_j^*}{\partial \mathbf{B}^{(u)} \partial \mathbf{B}^{(v)}} \mathbf{X} \right) \mathbf{N}_j^* \mathbf{X}, \quad \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{B}^{(u)}} \mathbf{X} \mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{B}^{(v)}} \mathbf{X} \right) \mathbf{N}_j^* \mathbf{X}, \\ &\left(\mathbf{X}^\top \frac{\partial^2 \mathbf{M}_j^*}{\partial \mathbf{A}^{(u)} \partial \mathbf{B}^{(v)}} \mathbf{X} \right) \mathbf{N}_j^* \mathbf{X}, \quad \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{A}^{(u)}} \mathbf{X} \mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{B}^{(v)}} \mathbf{X} \right) \mathbf{N}_j^* \mathbf{X}, \\ &\frac{\partial^2 \mathbf{N}_j^*}{\partial \mathbf{A}^{(u)} \partial \mathbf{A}^{(v)}} \mathbf{X}, \quad \frac{\partial^2 \mathbf{N}_j^*}{\partial \mathbf{B}^{(u)} \partial \mathbf{B}^{(v)}}, \quad \frac{\partial^2 \mathbf{N}_j^*}{\partial \mathbf{A}^{(u)} \partial \mathbf{B}^{(v)}}, \\ &\left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{A}^{(u)}} \mathbf{X} \right) \frac{\partial \mathbf{N}_j^*}{\partial \mathbf{A}^{(v)}} \mathbf{X}, \quad \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{B}^{(u)}} \mathbf{X} \right) \frac{\partial \mathbf{N}_j^*}{\partial \mathbf{B}^{(v)}} \mathbf{X}, \\ &\left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{A}^{(u)}} \mathbf{X} \right) \frac{\partial \mathbf{N}_j^*}{\partial \mathbf{B}^{(v)}} \mathbf{X}, \quad \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{B}^{(u)}} \mathbf{X} \right) \frac{\partial \mathbf{N}_j^*}{\partial \mathbf{A}^{(v)}} \mathbf{X} : j \in [\ell] \end{aligned} \right\}$$

D.2.1 LOCAL PART

For the local part, we will prove that

$$\lim_{\varepsilon \rightarrow 0} \inf_{\tilde{G} \in \mathcal{G}_{L'}(\Theta) : \mathcal{D}_2(\tilde{G}, \tilde{G}_*) \leq \varepsilon} \|f_{\tilde{G}} - f_{\tilde{G}_*}\|_{L^2(\mu)} / \mathcal{D}_2(\tilde{G}, \tilde{G}_*) > 0.$$

We will prove this by contradiction: assume that the above claim does not hold. This means we can find a sequence of mixing measures $\tilde{G}_n := \sum_{j=1}^{L'} \exp(c_{n,j}) \delta_{\mathbf{B}_{n,j} \mathbf{A}_{n,j}}$ in $\tilde{\mathcal{G}}_{L'}(\Theta)$ such that

$$\left\{ \begin{aligned} &\mathcal{D}_2(\tilde{G}_n, \tilde{G}_*) \rightarrow 0, \\ &\|f_{\tilde{G}_n} - f_{\tilde{G}_*}\|_{L^2(\mu)} / \mathcal{D}_2(\tilde{G}_n, \tilde{G}_*) \rightarrow 0. \end{aligned} \right.$$

as $n \rightarrow \infty$. We denote $\mathcal{V}_j^n := \mathcal{V}_j(\tilde{G}_n)$ as a Voronoi cell of \tilde{G}_n generated by the j -th components of \tilde{G}_* . Without loss of generality, we may assume that those Voronoi cells do not depend on the sample size, i.e., $\mathcal{V}_j = \mathcal{V}_j^n$. Therefore, the Voronoi loss can be rewritten as follows:

$$\begin{aligned} \mathcal{D}_2(\tilde{G}_n, \tilde{G}_*) &= \sum_{j=1}^L \left| \sum_{i \in \mathcal{V}_j} \exp(c_{n,i}) - \exp(c_{j'}^*) \right| + \sum_{j \in [L]: |\mathcal{V}_j|=1} \sum_{i \in \mathcal{V}_j} \exp(c_{n,i}) (\|\Delta \mathbf{B}_{n,ij}\| + \|\Delta \mathbf{A}_{n,ij}\|) \\ &\quad + \sum_{j \in [L]: |\mathcal{V}_j|>1} \sum_{i \in \mathcal{V}_j} \exp(c_{n,i}) (\|\Delta \mathbf{B}_{n,ij}\|^2 + \|\Delta \mathbf{A}_{n,ij}\|^2), \end{aligned}$$

where $\Delta \mathbf{B}_{n,ij} := \mathbf{B}_{n,i} - \mathbf{B}_{j'}^*$ and $\Delta \mathbf{A}_{n,ij} := \mathbf{A}_{n,i} - \mathbf{A}_{j'}^*$ for all $i \in \mathcal{V}_{j'}$ and $j \in [L]$.

Since $\mathcal{D}_2(\tilde{G}_n, \tilde{G}_*) \rightarrow 0$, we have $\sum_{i \in \mathcal{V}_j} \exp(c_{n,i}) \rightarrow \exp(c_j^*)$, $\mathbf{B}_{n,i} \rightarrow \mathbf{B}_{j'}^*$, and $\mathbf{A}_{n,i} \rightarrow \mathbf{A}_{j'}^*$ for any $i \in \mathcal{V}_j$ and $j \in [L]$. Now, the proof of the local part can be divided into three main steps as follows:

Step 1 - Decompose the difference between regression functions.

Recall the notation,

$$\begin{aligned} \mathbf{M} &:= \frac{\mathbf{C}_Q + \sigma_2(\mathbf{B})\sigma_1(\mathbf{A})}{\|\mathbf{C}_Q + \sigma_2(\mathbf{B})\sigma_1(\mathbf{A})\| + \tau_Q}, \\ \mathbf{N} &:= \frac{\mathbf{C}_V + \sigma_2(\mathbf{B})\sigma_1(\mathbf{A})}{\|\mathbf{C}_V + \sigma_2(\mathbf{B})\sigma_1(\mathbf{A})\| + \tau_V}, \end{aligned}$$

and their subscripts will follow the subscripts of the matrices \mathbf{A} , \mathbf{B} .

First, we define

$$T_n(\mathbf{X}) := \left(\sum_{k=1}^L \exp \left(\mathbf{X}^\top m_{Q,k} \underbrace{\frac{\mathbf{C}_Q + \sigma_2(\mathbf{B}_k^*)\sigma_1(\mathbf{A}_k^*)}{\|\mathbf{C}_Q + \sigma_2(\mathbf{B}_k^*)\sigma_1(\mathbf{A}_k^*)\| + \tau_Q}}_{\mathbf{M}_k^*} \mathbf{X} + c_k^* \right) \right) \cdot [f_{\tilde{G}_n}(\mathbf{X}) - f_{\tilde{G}_*}(\mathbf{X})].$$

Then, we can decompose the function $T_n(\mathbb{X})$ as follows:

$$\begin{aligned} T_n(\mathbf{X}) &= \sum_{j=1}^L \sum_{i \in \mathcal{V}_j} \exp(c_{n,i}) \left[\exp(\mathbf{X}^\top m_{Q,j} \mathbf{M}_{n,j} \mathbf{X}) (m_{V,j} \mathbf{N}_{n,j}) \mathbf{X} - \exp(\mathbb{X}^\top m_{Q,j} \mathbf{M}_j^* \mathbb{X}) (m_{V,j} \mathbf{N}_j^*) \mathbb{X} \right] \\ &\quad - \sum_{j=1}^L \sum_{i \in \mathcal{V}_j} \exp(c_{n,i}) \left[\exp(\mathbf{X}^\top m_{Q,j} \mathbf{M}_{n,j} \mathbf{X}) - \exp(\mathbf{X}^\top m_{Q,j} \mathbf{M}_j^* \mathbf{X}) \right] f_{\tilde{G}_n}(\mathbf{X}) \\ &\quad + \sum_{j=1}^L \left(\sum_{i \in \mathcal{V}_j} \exp(c_{n,i}) - \exp(c_j^*) \right) \exp(\mathbf{X}^\top m_{Q,j} \mathbf{M}_j^* \mathbf{X}) \left[(m_{V,j} \mathbf{N}_j^*) \mathbf{X} - f_{\tilde{G}_n}(\mathbf{X}) \right] \\ &:= \tilde{A}_n(\mathbf{X}) - \tilde{B}_n(\mathbf{X}) + \tilde{C}_n(\mathbf{X}). \end{aligned} \tag{13}$$

Decompose $\tilde{A}_n(\mathbf{X})$. We denote,

$$\begin{aligned} \tilde{U}(\mathbf{X}; \mathbf{B}, \mathbf{A}) &:= \exp \left(\mathbf{X}^\top m_Q \frac{\mathbf{C}_Q + \sigma_2(\mathbf{B})\sigma_1(\mathbf{A})}{\|\mathbf{C}_Q + \sigma_2(\mathbf{B})\sigma_1(\mathbf{A})\| + \tau_Q} \mathbf{X} \right) \\ \tilde{V}(\mathbf{X}; \mathbf{B}, \mathbf{A}) &:= \left(m_V \frac{\mathbf{C}_V + \sigma_2(\mathbf{B})\sigma_1(\mathbf{A})}{\|\mathbf{C}_V + \sigma_2(\mathbf{B})\sigma_1(\mathbf{A})\| + \tau_V} \right) \mathbf{X} \\ \tilde{F}(\mathbf{X}; \mathbf{B}, \mathbf{A}) &:= \tilde{U}(\mathbf{X}; \mathbf{B}, \mathbf{A}) \tilde{V}(\mathbf{X}; \mathbf{B}, \mathbf{A}), \end{aligned}$$

and for brevity of notation, we will use \tilde{U} instead of $\tilde{U}(\mathbf{X}; \mathbf{B}, \mathbf{A})$ with its subscript follows the subscripts of $\{\mathbf{B}, \mathbf{A}\}$, similarly for \tilde{V} and \tilde{F} .

We decompose $\tilde{A}_n(\mathbf{X})$ based on the number of element in the Voronoi cells as follows:

$$\begin{aligned}\tilde{A}_n(\mathbf{X}) &= \sum_{j:|\mathcal{V}_j|=1} \sum_{i \in \mathcal{V}_j} \exp(c_{n,i}) [\tilde{F}_{n,i} - \tilde{F}_j^*] + \sum_{j:|\mathcal{V}_j|>1} \sum_{i \in \mathcal{V}_j} \exp(c_{n,i}) [\tilde{F}_{n,i} - \tilde{F}_j^*] \\ &:= \tilde{A}_{n,1}(\mathbf{X}) + \tilde{A}_{n,2}(\mathbf{X}).\end{aligned}$$

We apply the first-order Taylor expansion to \tilde{U} and \tilde{V} ,

$$\tilde{U}_{n,i} = \tilde{U}_j^* + \sum_{|\alpha|=1} (\Delta \mathbf{A}_{n,ij})^{\alpha_1} (\Delta \mathbf{B}_{n,ij})^{\alpha_2} \frac{\partial^{|\alpha|} \tilde{U}}{\partial \mathbf{A}^{\alpha_1} \partial \mathbf{B}^{\alpha_2}}(\mathbf{X}; \mathbf{B}_j^*, \mathbf{A}_j^*) + \tilde{R}_{ij,1}(\mathbf{X}),$$

$$\tilde{V}_{n,i} = \tilde{V}_j^* + \sum_{|\alpha|=1} (\Delta \mathbf{A}_{n,ij})^{\alpha_1} (\Delta \mathbf{B}_{n,ij})^{\alpha_2} \frac{\partial^{|\alpha|} \tilde{V}}{\partial \mathbf{A}^{\alpha_1} \partial \mathbf{B}^{\alpha_2}}(\mathbf{X}; \mathbf{B}_j^*, \mathbf{A}_j^*) + \tilde{R}_{ij,2}(\mathbf{X}),$$

for any i and j such that $i \in \mathcal{V}_j$ and $|\mathcal{V}_j| = 1$. Here, the functions $\tilde{R}_{ij,1}(\mathbb{X})$ and $\tilde{R}_{ij,2}(\mathbb{X})$ denote the Taylor remainders. Plugging the above identities into $\tilde{A}_{n,1}$ leads to

$$\begin{aligned}\tilde{A}_{n,1}(\mathbf{X}) &= \sum_{j:|\mathcal{V}_j|=1} \sum_{i \in \mathcal{V}_j} \frac{\exp(c_{n,i})}{\alpha!} \sum_{|\alpha|=1} \left\{ (\Delta \mathbf{A}_{n,ij})^{\alpha_1} (\Delta \mathbf{B}_{n,ij})^{\alpha_2} \frac{\partial^{|\alpha|} \tilde{U}}{\partial \mathbf{A}^{\alpha_1} \partial \mathbf{B}^{\alpha_2}}(\mathbf{X}; \mathbf{B}_j^*, \mathbf{A}_j^*) \tilde{V}_j^* \right. \\ &\quad \left. + (\Delta \mathbf{A}_{n,ij})^{\alpha_1} (\Delta \mathbf{B}_{n,ij})^{\alpha_2} \frac{\partial^{|\alpha|} \tilde{V}}{\partial \mathbf{A}^{\alpha_1} \partial \mathbf{B}^{\alpha_2}}(\mathbf{X}; \mathbf{B}_j^*, \mathbf{A}_j^*) \tilde{U}_j^* \right\} + \tilde{R}_{n,1}(\mathbf{X}) \\ &= \sum_{j:|\mathcal{V}_j|=1} \sum_{|\alpha|=1} \left\{ \tilde{M}_{n,j,\alpha_1,\alpha_2} \frac{\partial^{|\alpha|} \tilde{U}}{\partial \mathbf{A}^{\alpha_1} \partial \mathbf{B}^{\alpha_2}}(\mathbf{X}; \mathbf{B}_j^*, \mathbf{A}_j^*) \tilde{V}_j^* + \tilde{M}_{n,j,\alpha_1,\alpha_2} \frac{\partial^{|\alpha|} \tilde{V}}{\partial \mathbf{A}^{\alpha_1} \partial \mathbf{B}^{\alpha_2}}(\mathbf{X}; \mathbf{B}_j^*, \mathbf{A}_j^*) \tilde{U}_j^* \right\} \\ &\quad + \tilde{R}_{n,1}(\mathbb{X})\end{aligned}$$

where $\alpha = (\alpha_1, \alpha_2) \in (\mathbb{N}^{r \times d}, \mathbb{N}^{d \times r})$ and the function $\tilde{R}_{n,1}(\mathbb{X})$ satisfies that $\tilde{R}_{n,1}(\mathbb{X})/\mathcal{D}_2(\tilde{G}_n, G_*) \rightarrow 0$. This fact is due to the uniform Lipschitz assumption of the function \tilde{F} . The coefficients $\tilde{M}_{n,j,\alpha_1,\alpha_2}$ are given by:

$$\tilde{M}_{n,j,\alpha_1,\alpha_2} = \sum_{i \in \mathcal{V}_j} \frac{\exp(c_{n,i})}{\alpha!} (\Delta \mathbf{A}_{n,ij})^{\alpha_1} (\Delta \mathbf{B}_{n,ij})^{\alpha_2},$$

for any $|\alpha| = 1$.

Moving to the function $\tilde{A}_{n,2}(\mathbf{X})$, we perform the Taylor expansion up to the second order of \tilde{U} and \tilde{V} then plugging in $\tilde{A}_{n,2}$ yields,

$$\begin{aligned}\tilde{A}_{n,2}(\mathbb{X}) &= \sum_{j:|\mathcal{V}_j|>1} \sum_{1 \leq |\alpha| \leq 2} \left\{ \tilde{M}_{n,j,\alpha_1,\alpha_2} \frac{\partial^{|\alpha|} \tilde{U}}{\partial \mathbf{A}^{\alpha_1} \partial \mathbf{B}^{\alpha_2}}(\mathbf{X}; \mathbf{B}_j^*, \mathbf{A}_j^*) \tilde{V}_j^* + \tilde{M}_{n,j,\alpha_1,\alpha_2} \frac{\partial^{|\alpha|} \tilde{V}}{\partial \mathbf{A}^{\alpha_1} \partial \mathbf{B}^{\alpha_2}}(\mathbf{X}; \mathbf{B}_j^*, \mathbf{A}_j^*) \tilde{U}_j^* \right\} \\ &\quad + \sum_{|\alpha|=1, |\beta|=1} \tilde{M}_{n,j,\alpha_1,\alpha_2,\beta_1,\beta_2} \frac{\partial^{|\alpha|} \tilde{U}}{\partial \mathbf{A}^{\alpha_1} \partial \mathbf{B}^{\alpha_2}}(\mathbf{X}; \mathbf{B}_j^*, \mathbf{A}_j^*) \frac{\partial^{|\beta|} \tilde{V}}{\partial \mathbf{A}^{\beta_1} \partial \mathbf{B}^{\beta_2}}(\mathbf{X}; \mathbf{B}_j^*, \mathbf{A}_j^*) \\ &\quad + \tilde{R}_{n,2}(\mathbb{X})\end{aligned}$$

where the remainder $\tilde{R}_{n,2}(\mathbb{X})$ satisfies that $\tilde{R}_{n,2}(\mathbb{X})/\mathcal{D}_2(\tilde{G}_n, G_*) \rightarrow 0$. The coefficients $\tilde{M}_{n,j,\alpha_1,\alpha_2}$ and $\tilde{M}_{n,j,\alpha_1,\alpha_2,\beta_1,\beta_2}$ take the following forms:

$$\tilde{M}_{n,j,\alpha_1,\alpha_2} = \sum_{i \in \mathcal{V}_j} \frac{\exp(c_{n,i})}{\alpha!} (\Delta \mathbf{A}_{n,ij})^{\alpha_1} (\Delta \mathbf{B}_{n,ij})^{\alpha_2},$$

for any $|\alpha| = 2$ and

$$\tilde{M}_{n,j,\alpha_1,\alpha_2,\beta_1,\beta_2} = \sum_{i \in \mathcal{V}_j} \frac{\exp(c_{n,i})}{\alpha! \beta!} (\Delta \mathbf{A}_{n,ij})^{\alpha_1 + \alpha_2} (\Delta \mathbf{B}_{n,ij})^{\beta_1 + \beta_2},$$

1404 for any $|\alpha| = |\beta| = 1$.

1405
1406 The partial derivatives of \tilde{U} and \tilde{V} can be calculated as following (recall the M , N term appear in
1407 \tilde{U} and \tilde{V}):

$$\begin{aligned}
1408 & \frac{\partial \tilde{U}}{\partial \mathbf{A}^{(u)}} = \exp(\mathbf{X}^\top m_Q \mathbf{M} \mathbf{X}) \mathbf{X}^\top m_Q \frac{\partial \mathbf{M}}{\partial (\mathbf{A}^{(u)})} \mathbf{X}, \quad \frac{\partial \tilde{U}}{\partial \mathbf{B}^{(u)}} = \exp(\mathbf{X}^\top m_Q \mathbf{M} \mathbf{X}) \mathbf{X}^\top m_Q \frac{\partial \mathbf{M}}{\partial (\mathbf{B}^{(u)})} \mathbf{X}, \\
1409 & \\
1410 & \frac{\partial^2 \tilde{U}}{\partial \mathbf{A}^{(u)} \partial \mathbf{A}^{(v)}} = \exp(\mathbf{X}^\top m_Q \mathbf{M} \mathbf{X}) m_Q^2 \left[\left(\mathbf{X}^\top \frac{\partial^2 \mathbf{M}}{\partial (\mathbf{A}^{(u)}) \partial (\mathbf{A}^{(v)})} \mathbf{X} \right) + \left(\mathbf{X}^\top \frac{\partial \mathbf{M}}{\partial (\mathbf{A}^{(u)})} \mathbf{X} \right) \left(\mathbf{X}^\top \frac{\partial \mathbf{M}}{\partial (\mathbf{A}^{(v)})} \mathbf{X} \right) \right], \\
1411 & \\
1412 & \frac{\partial^2 \tilde{U}}{\partial \mathbf{B}^{(u)} \partial \mathbf{B}^{(v)}} = \exp(\mathbf{X}^\top m_Q \mathbf{M} \mathbf{X}) m_Q^2 \left[\left(\mathbf{X}^\top \frac{\partial^2 \mathbf{M}}{\partial (\mathbf{B}^{(u)}) \partial (\mathbf{B}^{(v)})} \mathbf{X} \right) + \left(\mathbf{X}^\top \frac{\partial \mathbf{M}}{\partial (\mathbf{B}^{(u)})} \mathbf{X} \right) \left(\mathbf{X}^\top \frac{\partial \mathbf{M}}{\partial (\mathbf{B}^{(v)})} \mathbf{X} \right) \right], \\
1413 & \\
1414 & \frac{\partial^2 \tilde{U}}{\partial \mathbf{A}^{(u)} \partial \mathbf{B}^{(v)}} = \exp(\mathbf{X}^\top m_Q \mathbf{M} \mathbf{X}) m_Q^2 \left[\left(\mathbf{X}^\top \frac{\partial^2 \mathbf{M}}{\partial (\mathbf{A}^{(u)}) \partial (\mathbf{B}^{(v)})} \mathbf{X} \right) + \left(\mathbf{X}^\top \frac{\partial \mathbf{M}}{\partial (\mathbf{A}^{(u)})} \mathbf{X} \right) \left(\mathbf{X}^\top \frac{\partial \mathbf{M}}{\partial (\mathbf{B}^{(v)})} \mathbf{X} \right) \right], \\
1415 & \\
1416 & \frac{\partial \tilde{V}}{\partial \mathbf{A}^{(u)}} = m_V \frac{\partial \mathbf{N}}{\partial \mathbf{A}^{(u)}} \mathbf{X}, \quad \frac{\partial \tilde{V}}{\partial \mathbf{B}^{(u)}} = m_V \frac{\partial \mathbf{N}}{\partial \mathbf{B}^{(u)}} \mathbf{X}, \\
1417 & \\
1418 & \frac{\partial^2 \tilde{V}}{\partial \mathbf{A}^{(u)} \partial \mathbf{A}^{(v)}} = m_V \frac{\partial^2 \mathbf{N}}{\partial \mathbf{A}^{(u)} \partial \mathbf{A}^{(v)}} \mathbf{X}, \quad \frac{\partial^2 \tilde{V}}{\partial \mathbf{B}^{(u)} \partial \mathbf{B}^{(v)}} = m_V \frac{\partial^2 \mathbf{N}}{\partial \mathbf{B}^{(u)} \partial \mathbf{B}^{(v)}} \mathbf{X}, \quad \frac{\partial^2 \tilde{V}}{\partial \mathbf{A}^{(u)} \partial \mathbf{B}^{(v)}} = m_V \frac{\partial^2 \mathbf{N}}{\partial \mathbf{A}^{(u)} \partial \mathbf{B}^{(v)}} \mathbf{X}. \\
1419 & \\
1420 & \\
1421 & \\
1422 & \\
1423 & \\
1424 & \\
1425 & \\
1426 & \\
1427 & \\
1428 & \\
1429 & \\
1430 & \\
1431 & \\
1432 & \\
1433 & \\
1434 & \\
1435 & \\
1436 & \\
1437 & \\
1438 & \\
1439 & \\
1440 & \\
1441 & \\
1442 & \\
1443 & \\
1444 & \\
1445 & \\
1446 & \\
1447 & \\
1448 & \\
1449 & \\
1450 & \\
1451 & \\
1452 & \\
1453 & \\
1454 & \\
1455 & \\
1456 & \\
1457 &
\end{aligned}$$

Plugging these terms into $\tilde{A}_{n,1}$ and $\tilde{A}_{n,2}$, we obtain that,

$$\begin{aligned}
\tilde{A}_{n,1}(\mathbf{X}) &= \sum_{j:|\mathcal{V}_j|=1} \left\{ \sum_{|\alpha|=1} \left\{ \tilde{M}_{n,j,\alpha_1,\alpha_2} \exp(m_{Q,j} \mathbf{X}^\top M_j^* \mathbf{X}) m_{Q,j} \mathbf{X}^\top \frac{\partial M_j^*}{\partial \mathbf{A}^{\alpha_1} \partial \mathbf{B}^{\alpha_2}} \mathbf{X} \tilde{V}_j^* \right. \right. \\
&+ \left. \left. \tilde{M}_{n,j,\alpha_1,\alpha_2} m_{V,j} \frac{\partial N_j^*}{\partial \mathbf{A}^{\alpha_1} \mathbf{B}^{\alpha_2}} \mathbf{X} \tilde{U}_j \right\} + \tilde{R}_{n,1}(\mathbf{X}) \right. \\
&= \sum_{j:|\mathcal{V}_j|=1} \exp(m_{Q,j} \mathbf{X}^\top M_j^* \mathbf{X}) \left[\sum_{u=(1,1)}^{(d,d)} \tilde{M}_{n,j,u,0} \left(m_{Q,j} \mathbf{X}^\top \frac{\partial M_j^*}{\partial \mathbf{A}^u} \mathbf{X} \right) m_{V,j} N_j^* \mathbf{X} \right. \\
&+ \left. \sum_{u=(1,1)}^{(d,d)} \tilde{M}_{n,j,0,u} \left(m_{Q,j} \mathbf{X}^\top \frac{\partial M_j^*}{\partial \mathbf{B}^u} \mathbf{X} \right) m_{V,j} N_j^* \mathbf{X} + \sum_{u=(1,1)}^{(d,d)} \tilde{M}_{n,j,u,0} m_{V,j} \frac{\partial N_j^*}{\partial \mathbf{A}^u} \mathbf{X} + \tilde{M}_{n,j,0,u} m_{V,j} \frac{\partial N_j^*}{\partial \mathbf{B}^u} \mathbf{X} \right] + \tilde{R}_{n,1}(\mathbf{X}) \\
&= \sum_{j:|\mathcal{V}_j|=1} \exp(m_{Q,j} \mathbf{X}^\top M_j^* \mathbf{X}) \left[\sum_{u=(1,1)}^{(d,d)} \left(\mathbf{X}^\top \frac{\partial M_j^*}{\partial \mathbf{A}^u} \mathbf{X} \right) N_j^* \mathbf{X} \times \tilde{M}_{n,j,u,0} m_{Q,j} m_{V,j} \right. \\
&+ \left. \sum_{u=(1,1)}^{(d,d)} \left(\mathbf{X}^\top \frac{\partial M_j^*}{\partial \mathbf{B}^u} \mathbf{X} \right) N_j^* \mathbf{X} \times \tilde{M}_{n,j,0,u} m_{Q,j} m_{V,j} + \sum_{u=(1,1)}^{(d,d)} \frac{\partial N_j^*}{\partial \mathbf{A}^u} \mathbf{X} \times \tilde{M}_{n,j,u,0} m_{V,j} + \frac{\partial N_j^*}{\partial \mathbf{B}^u} \mathbf{X} \times \tilde{M}_{n,j,0,u} m_{V,j} \right] \\
&+ \tilde{R}_{n,1}(\mathbf{X}), \\
\tilde{A}_{n,2}(\mathbf{X}) &= \sum_{j:|\mathcal{V}_j|>1} \exp(m_{Q,j} \mathbf{X}^\top M_j^* \mathbf{X}) \left[\right. \\
&\sum_u \left(\mathbf{X}^\top \frac{\partial M_j^*}{\partial \mathbf{A}^{(u)}} \mathbf{X} \right) N_j^* \mathbf{X} \tilde{M}_{n,j,u,0} m_{Q,j} m_{V,j} + \sum_u \left(\mathbf{X}^\top \frac{\partial M_j^*}{\partial \mathbf{B}^{(u)}} \mathbf{X} \right) N_j^* \mathbf{X} \tilde{M}_{n,j,0,u} m_{Q,j} m_{V,j} \\
&+ \sum_u \frac{\partial N_j^*}{\partial \mathbf{A}^{(u)}} \mathbf{X} \tilde{M}_{n,j,u,0} m_{V,j} + \sum_u \frac{\partial N_j^*}{\partial \mathbf{B}^{(u)}} \mathbf{X} \tilde{M}_{n,j,0,u} m_{V,j} \\
&+ \sum_{u,v} \left(\mathbf{X}^\top \frac{\partial^2 M_j^*}{\partial \mathbf{A}^{(u)} \partial \mathbf{A}^{(v)}} \mathbf{X} \right) N_j^* \mathbf{X} \tilde{M}_{n,j,u+v,0} m_{Q,j} m_{V,j} \\
&+ \sum_{u,v} \left(\mathbf{X}^\top \frac{\partial M_j^*}{\partial \mathbf{A}^{(u)}} \mathbf{X} \mathbf{X}^\top \frac{\partial M_j^*}{\partial \mathbf{A}^{(v)}} \mathbf{X} \right) N_j^* \mathbf{X} \tilde{M}_{n,j,u+v,0} m_{Q,j} m_{V,j} \\
&+ \sum_{u,v} \left(\mathbf{X}^\top \frac{\partial^2 M_j^*}{\partial \mathbf{B}^{(u)} \partial \mathbf{B}^{(v)}} \mathbf{X} \right) N_j^* \mathbf{X} \tilde{M}_{n,j,0,u+v} m_{Q,j} m_{V,j} \\
&+ \sum_{u,v} \left(\mathbf{X}^\top \frac{\partial M_j^*}{\partial \mathbf{B}^{(u)}} \mathbf{X} \mathbf{X}^\top \frac{\partial M_j^*}{\partial \mathbf{B}^{(v)}} \mathbf{X} \right) N_j^* \mathbf{X} \tilde{M}_{n,j,0,u+v} m_{Q,j} m_{V,j} \\
&+ \sum_{u,v} \left(\mathbf{X}^\top \frac{\partial^2 M_j^*}{\partial \mathbf{A}^{(u)} \partial \mathbf{B}^{(v)}} \mathbf{X} \right) N_j^* \mathbf{X} \tilde{M}_{n,j,u,v} m_{Q,j} m_{V,j} + \sum_{u,v} \left(\mathbf{X}^\top \frac{\partial M_j^*}{\partial \mathbf{A}^{(u)}} \mathbf{X} \mathbf{X}^\top \frac{\partial M_j^*}{\partial \mathbf{B}^{(v)}} \mathbf{X} \right) N_j^* \mathbf{X} \tilde{M}_{n,j,u,v} m_{Q,j} m_{V,j} \\
&+ \sum_{u,v} \frac{\partial^2 N_j^*}{\partial \mathbf{A}^{(u)} \partial \mathbf{A}^{(v)}} \mathbf{X} \tilde{M}_{n,j,u+v,0} m_{V,j} + \sum_{u,v} \frac{\partial^2 N_j^*}{\partial \mathbf{B}^{(u)} \partial \mathbf{B}^{(v)}} \mathbf{X} \tilde{M}_{n,j,0,u+v} m_{V,j} + \sum_{u,v} \frac{\partial^2 N_j^*}{\partial \mathbf{A}^{(u)} \partial \mathbf{B}^{(v)}} \mathbf{X} \tilde{M}_{n,j,u,v} m_{V,j} \\
&+ \sum_{u,v} \left(\mathbf{X}^\top \frac{\partial M_j^*}{\partial \mathbf{A}^{(u)}} \mathbf{X} \right) \frac{\partial N_j^*}{\partial \mathbf{A}^{(v)}} \mathbf{X} \tilde{M}_{n,j,u,0,v,0} m_{Q,j} m_{V,j} + \sum_{u,v} \left(\mathbf{X}^\top \frac{\partial M_j^*}{\partial \mathbf{B}^{(u)}} \mathbf{X} \right) \frac{\partial N_j^*}{\partial \mathbf{B}^{(v)}} \mathbf{X} \tilde{M}_{n,j,0,u,0,v} m_{Q,j} m_{V,j} \\
&+ \sum_{u,v} \left(\mathbf{X}^\top \frac{\partial M_j^*}{\partial \mathbf{A}^{(u)}} \mathbf{X} \right) \frac{\partial N_j^*}{\partial \mathbf{B}^{(v)}} \mathbf{X} \tilde{M}_{n,j,u,0,0,v} m_{Q,j} m_{V,j} + \sum_{u,v} \left(\mathbf{X}^\top \frac{\partial M_j^*}{\partial \mathbf{B}^{(u)}} \mathbf{X} \right) \frac{\partial N_j^*}{\partial \mathbf{A}^{(v)}} \mathbf{X} \tilde{M}_{n,j,0,u,v,0} m_{Q,j} m_{V,j} \left. \right] \\
&+ \tilde{R}_{n,2}(\mathbf{X}),
\end{aligned}$$

here we use the notation $\frac{\partial M_j^*}{\partial \mathbf{A}^u}$ to denote for the value of $\frac{\partial M}{\partial \mathbf{A}^u}$ at $(\mathbf{A}_j^*, \mathbf{B}_j^*)$.

Decompose $\tilde{B}_n(\mathbf{X})$. Moving to $\tilde{B}_n(\mathbf{X})$, we can decompose this function as follows:

$$\begin{aligned} \tilde{B}_n(\mathbf{X}) &= \sum_{j:|\mathcal{V}_j|=1} \sum_{i \in \mathcal{V}_j} \exp(c_{n,i}) [\tilde{U}_{n,i} - \tilde{U}_j^*] f_{\tilde{G}_n}(\mathbf{X}) + \sum_{j:|\mathcal{V}_j|>1} \sum_{i \in \mathcal{V}_j} \exp(c_{n,i}) [\tilde{U}_{n,i} - \tilde{U}_j^*] f_{\tilde{G}_n}(\mathbf{X}) \\ &:= \tilde{B}_{n,1}(\mathbf{X}) + \tilde{B}_{n,2}(\mathbf{X}). \end{aligned}$$

We perform the Taylor expansions up to the first order for $\tilde{B}_{n,1}(\mathbf{X})$ and the second order for $\tilde{B}_{n,2}(\mathbf{X})$ leads to

$$\begin{aligned} \tilde{B}_{n,1}(\mathbf{X}) &= \sum_{j:|\mathcal{V}_j|=1} \sum_{|\alpha|=1} \tilde{M}_{n,j,\alpha_1,\alpha_2} \frac{\partial^{|\alpha|} \tilde{U}_j^*}{\partial \mathbf{A}^{\alpha_1} \partial \mathbf{B}^{\alpha_2}} f_{\tilde{G}_n}(\mathbf{X}) + \tilde{R}_{n,3}(\mathbf{X}), \\ \tilde{B}_{n,2}(\mathbf{X}) &= \sum_{j:|\mathcal{V}_j|=1} \sum_{1 \leq |\alpha| \leq 2} \tilde{M}_{n,j,\alpha_1,\alpha_2} \frac{\partial^{|\alpha|} \tilde{U}_j^*}{\partial \mathbf{A}^{\alpha_1} \partial \mathbf{B}^{\alpha_2}} f_{\tilde{G}_n}(\mathbf{X}) + \tilde{R}_{n,4}(\mathbf{X}), \end{aligned}$$

where the Taylor remainders $\tilde{R}_{n,3}(\mathbb{X}), \tilde{R}_{n,4}(\mathbb{X})$ satisfy that $\tilde{R}_{n,3}(\mathbb{X})/\mathcal{D}_2(\tilde{G}_n, G_*) \rightarrow 0$ and $\tilde{R}_{n,4}(\mathbb{X})/\mathcal{D}_2(\tilde{G}_n, G_*) \rightarrow 0$. Direct calculation leads to

$$\begin{aligned} \tilde{B}_{n,1}(\mathbf{X}) &= \sum_{j:|\mathcal{V}_j|=1} \exp(m_{Q,j} \mathbf{X}^\top \mathbf{Q}_j^* \mathbf{X}) \sum_{u=(1,1)}^{(d,d)} \left(\mathbf{X}^\top \frac{\partial \mathbf{Q}_j^*}{\partial \mathbf{A}^u} \mathbf{X} \right) f_{\tilde{G}_n}(\mathbf{X}) \tilde{M}_{n,j,u,0} m_{Q,j} \\ &\quad + \sum_u \left(\mathbf{X}^\top \frac{\partial \mathbf{Q}_j^*}{\partial \mathbf{B}^u} \mathbf{X} \right) f_{\tilde{G}_n}(\mathbf{X}) \tilde{M}_{n,j,0,u} m_{Q,j} + \tilde{R}_{n,3}(\mathbf{X}), \\ \tilde{B}_{n,2}(\mathbb{X}) &= \sum_{j:|\mathcal{V}_j|>1} \exp(m_{Q,j} \mathbf{X}^\top \mathbf{M}_j^* \mathbf{X}) \sum_{u=(1,1)}^{(d,d)} \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{A}^u} \mathbf{X} \right) f_{\tilde{G}_n}(\mathbf{X}) \tilde{M}_{n,j,u,0} m_{Q,j} \\ &\quad + \sum_u \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{B}^u} \mathbf{X} \right) f_{\tilde{G}_n}(\mathbf{X}) \tilde{M}_{n,j,0,u} m_{Q,j} \\ &\quad + \sum_{u,v} \left(\mathbf{X}^\top \frac{\partial^2 \mathbf{M}_j^*}{\partial \mathbf{A}^{(u)} \partial \mathbf{A}^{(v)}} \mathbf{X} \right) f_{\tilde{G}_n}(\mathbf{X}) \tilde{M}_{n,j,u+v,0} m_{Q,j} \\ &\quad + \sum_{u,v} \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{A}^{(u)}} \mathbf{X} \mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{A}^{(v)}} \mathbf{X} \right) f_{\tilde{G}_n}(\mathbf{X}) \tilde{M}_{n,j,u+v,0} m_{Q,j} \\ &\quad + \sum_{u,v} \left(\mathbf{X}^\top \frac{\partial^2 \mathbf{M}_j^*}{\partial \mathbf{B}^{(u)} \partial \mathbf{B}^{(v)}} \mathbf{X} \right) f_{\tilde{G}_n}(\mathbf{X}) \tilde{M}_{n,j,0,u+v} m_{Q,j} \\ &\quad + \sum_{u,v} \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{B}^{(u)}} \mathbf{X} \mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{B}^{(v)}} \mathbf{X} \right) f_{\tilde{G}_n}(\mathbf{X}) \tilde{M}_{n,j,0,u+v} m_{Q,j} \\ &\quad + \sum_{u,v} \left(\mathbf{X}^\top \frac{\partial^2 \mathbf{M}_j^*}{\partial \mathbf{A}^{(u)} \partial \mathbf{B}^{(v)}} \mathbf{X} \right) f_{\tilde{G}_n}(\mathbf{X}) \tilde{M}_{n,j,u,v} m_{Q,j} \\ &\quad + \sum_{u,v} \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{A}^{(u)}} \mathbf{X} \mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{B}^{(v)}} \mathbf{X} \right) f_{\tilde{G}_n}(\mathbf{X}) \tilde{M}_{n,j,u,v} m_{Q,j}. \end{aligned}$$

Putting all the above results together, we can represent the function $T_n(\mathbf{X})$ as follows:

$$\begin{aligned} Q_n(\mathbb{X}) &= \tilde{A}_{n,1} + \tilde{A}_{n,2} - \tilde{B}_{n,1} - \tilde{B}_{n,2} \\ &\quad + \sum_{j=1}^L \tilde{N}_{n,j} \exp(\mathbf{X}^\top m_{Q,j} \mathbf{M}_j^* \mathbf{X}) \left[(m_{V,j} \mathbf{N}_j^*) \mathbb{X} - f_{\tilde{G}_n}(\mathbb{X}) \right], \end{aligned} \quad (14)$$

where $\tilde{N}_{n,j} := \sum_{i \in \mathcal{V}_j} \exp(c_{n,i}) - \exp(c_j^*)$ for any $j \in [L]$.

Step 2 - Non-vanishing coefficients. As indicated in equation (14), the ratio $Q_n(\mathbb{X})/\mathcal{D}_2(\tilde{G}_n, G_*)$ can be expressed as a linear combination of the following independent functions:

$$\begin{aligned}
& \tilde{U}_j^*(\mathbf{X}) \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{A}^{(u)}} \mathbf{X} \right) \mathbf{N}_j^* \mathbf{X}, \quad \tilde{U}_j^*(\mathbf{X}) \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{B}^{(u)}} \mathbf{X} \right) \mathbf{N}_j^* \mathbf{X}, \\
& \tilde{U}_j^*(\mathbf{X}) \frac{\partial \mathbf{N}_j^*}{\partial \mathbf{A}^{(u)}} \mathbf{X}, \quad \tilde{U}_j^*(\mathbf{X}) \frac{\partial \mathbf{N}_j^*}{\partial \mathbf{B}^{(u)}} \mathbf{X}, \\
& \tilde{U}_j^*(\mathbf{X}) \left(\mathbf{X}^\top \frac{\partial^2 \mathbf{M}_j^*}{\partial \mathbf{A}^{(u)} \partial \mathbf{A}^{(v)}} \mathbf{X} \right) \mathbf{N}_j^* \mathbf{X}, \quad \tilde{U}_j^*(\mathbf{X}) \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{A}^{(u)}} \mathbf{X} \mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{A}^{(v)}} \mathbf{X} \right) \mathbf{N}_j^* \mathbf{X}, \\
& \tilde{U}_j^*(\mathbf{X}) \left(\mathbf{X}^\top \frac{\partial^2 \mathbf{M}_j^*}{\partial \mathbf{B}^{(u)} \partial \mathbf{B}^{(v)}} \mathbf{X} \right) \mathbf{N}_j^* \mathbf{X}, \quad \tilde{U}_j^*(\mathbf{X}) \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{B}^{(u)}} \mathbf{X} \mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{B}^{(v)}} \mathbf{X} \right) \mathbf{N}_j^* \mathbf{X}, \\
& \tilde{U}_j^*(\mathbf{X}) \left(\mathbf{X}^\top \frac{\partial^2 \mathbf{M}_j^*}{\partial \mathbf{A}^{(u)} \partial \mathbf{B}^{(v)}} \mathbf{X} \right) \mathbf{N}_j^* \mathbf{X}, \quad \tilde{U}_j^*(\mathbf{X}) \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{A}^{(u)}} \mathbf{X} \mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{B}^{(v)}} \mathbf{X} \right) \mathbf{N}_j^* \mathbf{X}, \\
& \tilde{U}_j^*(\mathbf{X}) \frac{\partial^2 \mathbf{N}_j^*}{\partial \mathbf{A}^{(u)} \partial \mathbf{A}^{(v)}} \mathbf{X}, \quad \tilde{U}_j^*(\mathbf{X}) \frac{\partial^2 \mathbf{N}_j^*}{\partial \mathbf{B}^{(u)} \partial \mathbf{B}^{(v)}}, \quad \tilde{U}_j^*(\mathbf{X}) \frac{\partial^2 \mathbf{N}_j^*}{\partial \mathbf{A}^{(u)} \partial \mathbf{B}^{(v)}}, \\
& \tilde{U}_j^*(\mathbf{X}) \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{A}^{(u)}} \mathbf{X} \right) \frac{\partial \mathbf{N}_j^*}{\partial \mathbf{A}^{(v)}} \mathbf{X}, \quad \tilde{U}_j^*(\mathbf{X}) \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{B}^{(u)}} \mathbf{X} \right) \frac{\partial \mathbf{N}_j^*}{\partial \mathbf{B}^{(v)}} \mathbf{X}, \\
& \tilde{U}_j^*(\mathbf{X}) \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{A}^{(u)}} \mathbf{X} \right) \frac{\partial \mathbf{N}_j^*}{\partial \mathbf{B}^{(v)}} \mathbf{X}, \quad \tilde{U}_j^*(\mathbf{X}) \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{B}^{(u)}} \mathbf{X} \right) \frac{\partial \mathbf{N}_j^*}{\partial \mathbf{A}^{(v)}} \mathbf{X}, \\
& \tilde{U}_j^*(\mathbf{X}) \mathbf{N}_j^* \mathbf{X}, \quad \tilde{U}_j^*(\mathbf{X}) f_{\tilde{G}_n}(\mathbf{X}), \\
& \tilde{U}_j^*(\mathbf{X}) \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{A}^u} \mathbf{X} \right) f_{\tilde{G}_n}(\mathbf{X}), \quad \tilde{U}_j^*(\mathbf{X}) \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{B}^u} \mathbf{X} \right) f_{\tilde{G}_n}(\mathbf{X}), \\
& \tilde{U}_j^*(\mathbf{X}) \left(\mathbf{X}^\top \frac{\partial^2 \mathbf{M}_j^*}{\partial \mathbf{A}^{(u)} \partial \mathbf{A}^{(v)}} \mathbf{X} \right) f_{\tilde{G}_n}(\mathbf{X}), \quad \tilde{U}_j^*(\mathbf{X}) \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{A}^{(u)}} \mathbf{X} \mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{A}^{(v)}} \mathbf{X} \right) f_{\tilde{G}_n}(\mathbf{X}), \\
& \tilde{U}_j^*(\mathbf{X}) \left(\mathbf{X}^\top \frac{\partial^2 \mathbf{M}_j^*}{\partial \mathbf{B}^{(u)} \partial \mathbf{B}^{(v)}} \mathbf{X} \right) f_{\tilde{G}_n}(\mathbf{X}), \quad \tilde{U}_j^*(\mathbf{X}) \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{B}^{(u)}} \mathbf{X} \mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{B}^{(v)}} \mathbf{X} \right) f_{\tilde{G}_n}(\mathbf{X}), \\
& \tilde{U}_j^*(\mathbf{X}) \left(\mathbf{X}^\top \frac{\partial^2 \mathbf{M}_j^*}{\partial \mathbf{A}^{(u)} \partial \mathbf{B}^{(v)}} \mathbf{X} \right) f_{\tilde{G}_n}(\mathbf{X}), \quad \tilde{U}_j^*(\mathbf{X}) \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{A}^{(u)}} \mathbf{X} \mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{B}^{(v)}} \mathbf{X} \right) f_{\tilde{G}_n}(\mathbf{X}),
\end{aligned}$$

for any indices $1 \leq j \leq L$ and $u = (u_1, u_2), v = (v_1, v_2)$ with $1 \leq u_1, v_1, u_2, v_2 \leq d$.

We will show that at least one of the coefficients of these independent functions does not go to 0 as $n \rightarrow \infty$. Assume by contrary that all these coefficients of these linear independent functions go to 0 when $n \rightarrow \infty$. From equation (14), we have that $\tilde{M}_{n,j,\alpha_1,\alpha_2}/\mathcal{D}_2(\tilde{G}_n, G_*)$, $\tilde{M}_{n,j,\alpha_1,\beta_1,\alpha_2,\beta_2}/\mathcal{D}_2(\tilde{G}_n, G_*)$, and $\tilde{N}_{n,j}/\mathcal{D}_2(\tilde{G}_n, G_*)$ go to 0 for all the coefficients $\alpha_1, \alpha_2, \beta_1, \beta_2 \in \mathbb{N}^{d \times d}$ satisfying that $1 \leq |\alpha_1| + |\beta_1| + |\alpha_2| + |\beta_2| \leq 2$.

Since $\tilde{N}_{n,j}/\mathcal{D}_2(\tilde{G}_n, G_*) \rightarrow 0$, we find that for any $j \in [L]$

$$\frac{|\sum_{i \in \mathcal{V}_j} \exp(c_{n,i}) - \exp(c_j^*)|}{\mathcal{D}_2(\tilde{G}_n, \tilde{G}_*)} = \frac{|\tilde{N}_{n,j}|}{\mathcal{D}_2(\tilde{G}_n, \tilde{G}_*)} \rightarrow 0.$$

Taking the summation of these limits over $j \in [L]$ yields

$$\frac{\sum_{j=1}^L |\sum_{i \in \mathcal{V}_j} \exp(c_{n,i}) - \exp(c_j^*)|}{\mathcal{D}_2(\tilde{G}_n, \tilde{G}_*)} \rightarrow 0. \tag{15}$$

Now, for any index $j \in [L]$ such that $|\mathcal{V}_j| = 1$, the limits $\tilde{M}_{n,j,e_u,0_d}/\mathcal{D}_2(\tilde{G}_n, \tilde{G}_*) \rightarrow 0$ lead to $\frac{\sum_{i \in \mathcal{V}_j} \exp(c_{n,i}) \|\Delta \mathbf{A}_{n,i,j}\|_1}{\mathcal{D}_2(\tilde{G}_n, \tilde{G}_*)} \rightarrow 0$ as $n \rightarrow \infty$. Because the ℓ_1 -norm and ℓ_2 -norm are equivalent, this

1620 result implies that

$$1621 \frac{\sum_{j:|\mathcal{V}_j|=1} \sum_{i \in \mathcal{V}_j} \exp(c_{n,i}) \|\Delta \mathbf{A}_{n,ij}\|}{\mathcal{D}_2(\tilde{G}_n, \tilde{G}_*)} \rightarrow 0. \quad 1622$$

1623 Similarly, since $\frac{\tilde{M}_{n,j,0_d,e_u}/\mathcal{D}_2(\tilde{G}_n, \tilde{G}_*)}{\sum_{j:|\mathcal{V}_j|=1} \sum_{i \in \mathcal{V}_j} \exp(c_{n,i}) \|\Delta \mathbf{B}_{n,ij}\|} \rightarrow 0$, we also have that

$$1624 \frac{\sum_{j:|\mathcal{V}_j|=1} \sum_{i \in \mathcal{V}_j} \exp(c_{n,i}) (\|\Delta \mathbf{A}_{n,ij}\| + \|\Delta \mathbf{B}_{n,ij}\|)}{\mathcal{D}_2(\tilde{G}_n, \tilde{G}_*)} \rightarrow 0. \quad 1625 \quad (16)$$

1626 Moving to indices $j \in [L]$ such that their corresponding Voronoi cells have more than one element, i.e., $|\mathcal{V}_j| > 1$. The limits $\frac{\tilde{M}_{n,j,2e_u,0_d}}{\mathcal{D}_2(\tilde{G}_n, \tilde{G}_*)} \rightarrow 0$ and $\frac{\tilde{M}_{n,j,0_d,2e_u}}{\mathcal{D}_2(\tilde{G}_n, \tilde{G}_*)} \rightarrow 0$ induces that

$$1627 \frac{\sum_{j:|\mathcal{V}_j|>1} \sum_{i \in \mathcal{V}_j} \exp(c_{n,i}) (\|\Delta \mathbf{A}_{n,ij}\|^2 + \|\Delta \mathbf{B}_{n,ij}\|^2)}{\mathcal{D}_2(\tilde{G}_n, G_*)} \rightarrow 0. \quad 1628 \quad (17)$$

1629 By putting the results in equations (15), (16), and (17) together, we arrive at $1 = \frac{\mathcal{D}_2(\tilde{G}_n, \tilde{G}_*)}{\mathcal{D}_2(\tilde{G}_n, G_*)} \rightarrow 0$ as $n \rightarrow \infty$, which is a contradiction. As a consequence, at least one of the coefficients of the linear independent functions in $T_n(\mathbf{X})/\mathcal{D}_2(\tilde{G}_n, G_*)$ does not go to 0 as $n \rightarrow \infty$.

1630 **Step 3 - Application of the Fatou's lemma.** We define \tilde{m}_n to be the maximum of the absolute values of the coefficients of the linear independent functions in $T_n(\mathbf{X})/\mathcal{D}_2(\tilde{G}_n, G_*)$. As at least one of these coefficients does not go to 0, it indicates that $1/\tilde{m}_n \not\rightarrow \infty$ as $n \rightarrow \infty$.

1631 Since $\|f_{\tilde{G}_n} - f_{\tilde{G}_*}\|_{L^2(\mu)}/\mathcal{D}_2(\tilde{G}_n, \tilde{G}_*) \rightarrow 0$ as $n \rightarrow \infty$, we obtain $\|f_{\tilde{G}_n} - f_{\tilde{G}_*}\|_{L^2(\mu)}/(\tilde{m}_n \mathcal{D}_2(\tilde{G}_n, G_*)) \rightarrow 0$. An application of the Fatou's lemma leads to:

$$1632 0 = \lim_{n \rightarrow \infty} \frac{\|f_{\tilde{G}_n} - f_{\tilde{G}_*}\|_{L^2(\mu)}}{\tilde{m}_n \mathcal{D}_2(\tilde{G}_n, G_*)} \geq \int \liminf_{n \rightarrow \infty} \frac{|f_{\tilde{G}_n}(\mathbf{X}) - f_{\tilde{G}_*}(\mathbf{X})|}{\tilde{m}_n \mathcal{D}_2(\tilde{G}_n, G_*)} d\mu(\mathbf{X}) \geq 0. \quad 1633$$

1634 This inequality implies that $\liminf_{n \rightarrow \infty} \frac{|f_{\tilde{G}_n}(\mathbf{X}) - f_{\tilde{G}_*}(\mathbf{X})|}{\tilde{m}_n \mathcal{D}_2(\tilde{G}_n, G_*)} = 0$ for almost surely \mathbf{X} . As $n \rightarrow \infty$, we denote

$$1635 \frac{\tilde{M}_{n,j,\alpha_1,\alpha_2}}{\tilde{m}_n \mathcal{D}_2(\tilde{G}_n, G_*)} \rightarrow \tilde{\tau}_{j,\alpha_1,\alpha_2}, \quad \frac{\tilde{M}_{n,j,\alpha_1,\alpha_2,\beta_1,\beta_2}}{\tilde{m}_n \mathcal{D}_2(\tilde{G}_n, G_*)} \rightarrow \tilde{\xi}_{j,\alpha_1,\alpha_2,\beta_1,\beta_2}, \quad \frac{\tilde{N}_{n,j}}{\tilde{m}_n \mathcal{D}_2(\tilde{G}_n, G_*)} \rightarrow \tilde{\lambda}_{0,j}, \quad 1636$$

1637 for any indices $j \in [L]$ and any coefficients $\alpha_1, \alpha_2, \beta_1, \beta_2$ such that $1 \leq |\alpha_1| + |\alpha_2| + |\beta_1| + |\beta_2| \leq 2$. Here, at least one element of the set $\{\tilde{\tau}_{j,\alpha_1,\alpha_2}, \tilde{\xi}_{j,\alpha_1,\alpha_2,\beta_1,\beta_2}, \tilde{\lambda}_{0,j}\}$ is different from 0. Given the

1638 above notations, the limit $\liminf_{n \rightarrow \infty} \frac{|f_{\tilde{G}_n}(\mathbf{X}) - f_{\tilde{G}_*}(\mathbf{X})|}{\tilde{m}_n \mathcal{D}_2(\tilde{G}_n, G_*)} = 0$ implies that,

$$1639 0 = \liminf_{n \rightarrow \infty} \frac{T_n(\mathbf{X})}{\tilde{m}_n \mathcal{D}_2(\tilde{G}_n, G_*)} \quad 1640$$

$$1641 = \liminf_{n \rightarrow \infty} \frac{\tilde{A}_{n,1} + \tilde{A}_{n,2} - \tilde{B}_{n,1} - \tilde{B}_{n,2} + \sum_{j=1}^L \tilde{N}_{n,j} \exp(\mathbf{X}^\top m_{Q,j} \mathbf{M}_j^* \mathbf{X}) \left[(m_{V,j} \mathbf{N}_j^*) \mathbb{X} - f_{\tilde{G}_n}(\mathbb{X}) \right]}{\tilde{m}_n \mathcal{D}_2(\tilde{G}_n, G_*)}, \quad 1642$$

for almost surely \mathbb{X} . For example, we look at the limit of $\tilde{A}_{n,1}/\tilde{m}_n \mathcal{D}_2(\tilde{G}_n, G_*)$,

$$\begin{aligned} \liminf_{n \rightarrow \infty} \frac{\tilde{A}_{n,1}}{\tilde{m}_n \mathcal{D}_2(\tilde{G}_n, G_*)} &= \sum_{j:|\mathcal{V}_j|=1} \left(\exp(m_{Q,j} \mathbf{X}^\top \mathbf{M}_j^* \mathbf{X}) \left[\sum_{u=(1,1)}^{(d,d)} \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{A}^u} \mathbf{X} \right) \mathbf{N}_j^* \mathbf{X} \times \tilde{\tau}_{j,u,0} m_{Q,j} m_{V,j} \right. \right. \\ &+ \sum_{u=(1,1)}^{(d,d)} \left(\mathbf{X}^\top \frac{\partial \mathbf{M}_j^*}{\partial \mathbf{B}^u} \mathbf{X} \right) \mathbf{N}_j^* \mathbf{X} \times \tilde{\tau}_{j,0,u} m_{Q,j} m_{V,j} + \sum_{u=(1,1)}^{(d,d)} \frac{\partial \mathbf{N}_j^*}{\partial \mathbf{A}^u} \mathbf{X} \times \tilde{\tau}_{j,u,0} m_{V,j} \\ &\left. \left. + \frac{\partial \mathbf{N}_j^*}{\partial \mathbf{B}^u} \mathbf{X} \times \tilde{\tau}_{j,0,u} m_{V,j} \right] \right). \end{aligned}$$

From the equation $0 = \liminf_{n \rightarrow \infty} \frac{T_n(\mathbf{X})}{\tilde{m}_n \mathcal{D}_2(\tilde{G}_n, G_*)}$ and the linear independence of the functions imply that all the coefficients $\{\tilde{\tau}_{j,\alpha_1,\alpha_2}, \tilde{\xi}_{j,\alpha_1,\alpha_2,\beta_1,\beta_2}, \tilde{\lambda}_{0,j}\}$ are 0. It is a contradiction. As a consequence, we obtain that

$$\lim_{\varepsilon \rightarrow 0} \inf_{\tilde{G} \in \tilde{\mathcal{G}}_{L'}(\Theta): \mathcal{D}_2(\tilde{G}, \tilde{G}_*) \leq \varepsilon} \|f_{\tilde{G}} - f_{\tilde{G}_*}\|_{L^2(\mu)} / \mathcal{D}_2(\tilde{G}_n, G_*) > 0.$$

D.2.2 GLOBAL PART

The result of the local part implies that there exists a positive constant ε' such that

$$\inf_{\tilde{G} \in \tilde{\mathcal{G}}_{L'}(\Theta): \mathcal{D}_2(\tilde{G}, \tilde{G}_*) \leq \varepsilon'} \|f_{\tilde{G}} - f_{\tilde{G}_*}\|_{L^2(\mu)} / \mathcal{D}_2(\tilde{G}, \tilde{G}_*) > 0.$$

Therefore the remaining part is to prove

$$\inf_{\tilde{G} \in \tilde{\mathcal{G}}_{L'}(\Theta): \mathcal{D}_2(\tilde{G}, \tilde{G}_*) > \varepsilon'} \|f_{\tilde{G}} - f_{\tilde{G}_*}\|_{L^2(\mu)} / \mathcal{D}_2(\tilde{G}, \tilde{G}_*) > 0.$$

We assume by contradiction that the above claim does not hold. It means there exists a sequence of measures $\tilde{G}'_n := \sum_{j=1}^L \exp(c_{n,j}) \delta_{(\mathbf{B}_{n,j}, \mathbf{A}_{n,j})}$ in $\tilde{\mathcal{G}}_{L'}(\Theta)$ such that

$$\begin{cases} \mathcal{D}_2(\tilde{G}', \tilde{G}_*) > \varepsilon' \\ \|f_{\tilde{G}'_n} - f_{\tilde{G}_*}\|_{L^2(\mu)} / \mathcal{D}_2(\tilde{G}', \tilde{G}_*) \rightarrow 0 \end{cases}$$

as $n \rightarrow \infty$, which implies that $\|f_{\tilde{G}'_n} - f_{\tilde{G}_*}\|_{L^2(\mu)} \rightarrow 0$ as $n \rightarrow \infty$.

Given that Θ is a compact set, there exists a mixing measure \tilde{G}' in $\tilde{\mathcal{G}}_{L'}(\Theta)$ such that one of the \tilde{G}'_n 's subsequences converges to \tilde{G}' . Since $\mathcal{D}_2(\tilde{G}'_n, \tilde{G}_*) > \varepsilon'$, we obtain that $\mathcal{D}_2(\tilde{G}', \tilde{G}_*) > \varepsilon'$. An application of the Fatou's lemma leads to

$$0 = \lim_{n \rightarrow \infty} \|f_{\tilde{G}'_n} - f_{\tilde{G}_*}\|_{L^2(\mu)} \geq \int \liminf_{n \rightarrow \infty} \|f_{\tilde{G}'_n}(\mathbf{X}) - f_{\tilde{G}_*}(\mathbf{X})\|^2 d\mu(\mathbf{X}).$$

The above inequality indicates that $f_{\tilde{G}'} = f_{\tilde{G}_*}$ for almost surely \mathbb{X} . From the identifiability property that we will prove shortly below, we deduce that $\tilde{G}' \equiv \tilde{G}_*$. It follows that $\mathcal{D}_2(\tilde{G}', \tilde{G}_*) = 0$, which is opposed to the fact that $\mathcal{D}_2(\tilde{G}', \tilde{G}_*) > \varepsilon' > 0$. Hence, the proof is completed. **Proof for the identifiability property.** We will prove that if $f_{\tilde{G}}(\mathbf{X}) = f_{\tilde{G}_*}(\mathbf{X})$ for almost surely \mathbf{X} , then $\tilde{G} \equiv \tilde{G}_*$. To ease the presentation, for any mixing measure $\tilde{G} = \sum_{j=1}^{\tilde{L}} \exp(c_j) \delta_{(\mathbf{B}_j, \mathbf{A}_j)} \in \tilde{\mathcal{G}}_{L'}(\Theta)$, we denote

$$\text{softmax}_G(u) = \frac{\exp(u)}{\sum_{k=1}^L \exp\left(\mathbf{X}^\top m_{Q,k} \frac{\mathbf{C}_Q + \sigma_2(\mathbf{B}_k) \sigma_1(\mathbf{A}_k)}{\|\mathbf{C}_Q + \sigma_2(\mathbf{B}_k) \sigma_1(\mathbf{A}_k)\| + \tau_Q} \mathbf{X} + c_k\right)},$$

where $u \in \left\{ \mathbf{X}^\top m_{Q,j} \frac{\mathbf{C}_Q + \sigma_2(\mathbf{B}_j) \sigma_1(\mathbf{A}_j)}{\|\mathbf{C}_Q + \sigma_2(\mathbf{B}_j) \sigma_1(\mathbf{A}_j)\| + \tau_Q} \mathbf{X} + c_j : j \in [\tilde{L}] \right\}$.

The equation $f_{\tilde{G}}(\mathbb{X}) = f_{\tilde{G}_*}(\mathbb{X})$ indicates that

$$\begin{aligned} & \sum_{j=1}^{\tilde{L}} \operatorname{softmax} \left(\mathbf{X}^\top m_{Q,j} \frac{\mathbf{C}_Q + \sigma_2(\mathbf{B}_j)\sigma_1(\mathbf{A}_j)}{\|\mathbf{C}_Q + \sigma_2(\mathbf{B}_j)\sigma_1(\mathbf{A}_j)\| + \tau_Q} \mathbf{X} + c_j \right) \left(m_{V,j} \frac{\mathbf{C}_V + \sigma_2(\mathbf{B}_j)\sigma_1(\mathbf{A}_j)}{\|\mathbf{C}_V + \sigma_2(\mathbf{B}_j)\sigma_1(\mathbf{A}_j)\| + \tau_V} \right) \mathbf{X} \\ &= \sum_{j=1}^L \operatorname{softmax} \left(\mathbf{X}^\top m_{Q,j} \frac{\mathbf{C}_Q + \sigma_2(\mathbf{B}_j^*)\sigma_1(\mathbf{A}_j^*)}{\|\mathbf{C}_Q + \sigma_2(\mathbf{B}_j^*)\sigma_1(\mathbf{A}_j^*)\| + \tau_Q} \mathbf{X} + c_j^* \right) \left(m_{V,j} \frac{\mathbf{C}_V + \sigma_2(\mathbf{B}_j^*)\sigma_1(\mathbf{A}_j^*)}{\|\mathbf{C}_V + \sigma_2(\mathbf{B}_j^*)\sigma_1(\mathbf{A}_j^*)\| + \tau_V} \right) \mathbf{X} \end{aligned} \quad (18)$$

That equation implies that $\tilde{L} = L$. As a consequence, we find that

$$\begin{aligned} & \left\{ \operatorname{softmax} \left(\mathbf{X}^\top m_{Q,j} \frac{\mathbf{C}_Q + \sigma_2(\mathbf{B}_j)\sigma_1(\mathbf{A}_j)}{\|\mathbf{C}_Q + \sigma_2(\mathbf{B}_j)\sigma_1(\mathbf{A}_j)\| + \tau_Q} \mathbf{X} + c_j \right) : j \in [\tilde{L}] \right\} \\ &= \left\{ \operatorname{softmax} \left(\mathbf{X}^\top m_{Q,j} \frac{\mathbf{C}_Q + \sigma_2(\mathbf{B}_j^*)\sigma_1(\mathbf{A}_j^*)}{\|\mathbf{C}_Q + \sigma_2(\mathbf{B}_j^*)\sigma_1(\mathbf{A}_j^*)\| + \tau_Q} \mathbf{X} + c_j^* \right) : j \in [L] \right\} \end{aligned}$$

for almost surely \mathbb{X} . By relabelling the indices, we can assume without loss of generality that for any $j \in [L]$

$$\operatorname{softmax} \left(\mathbf{X}^\top m_{Q,j} \frac{\mathbf{C}_Q + \sigma_2(\mathbf{B}_j^*)\sigma_1(\mathbf{A}_j^*)}{\|\mathbf{C}_Q + \sigma_2(\mathbf{B}_j^*)\sigma_1(\mathbf{A}_j^*)\| + \tau_Q} \mathbf{X} + c_j^* \right) = \operatorname{softmax} \left(\mathbf{X}^\top m_{Q,j} \frac{\mathbf{C}_Q + \sigma_2(\mathbf{B}_j)\sigma_1(\mathbf{A}_j)}{\|\mathbf{C}_Q + \sigma_2(\mathbf{B}_j)\sigma_1(\mathbf{A}_j)\| + \tau_Q} \mathbf{X} + c_j \right),$$

for almost surely \mathbb{X} . Given the invariance to translation of the softmax function, the equation (18) leads to

$$\begin{aligned} & \sum_{j=1}^{\tilde{L}} \exp(c_j) \exp \left(\mathbf{X}^\top m_{Q,j} \frac{\mathbf{C}_Q + \sigma_2(\mathbf{B}_j^*)\sigma_1(\mathbf{A}_j^*)}{\|\mathbf{C}_Q + \sigma_2(\mathbf{B}_j^*)\sigma_1(\mathbf{A}_j^*)\| + \tau_Q} \mathbf{X} \right) \left(m_{V,j} \frac{\mathbf{C}_V + \sigma_2(\mathbf{B}_j)\sigma_1(\mathbf{A}_j)}{\|\mathbf{C}_V + \sigma_2(\mathbf{B}_j)\sigma_1(\mathbf{A}_j)\| + \tau_V} \right) \mathbf{X} \\ &= \sum_{j=1}^L \exp(c_j^*) \exp \left(\mathbf{X}^\top m_{Q,j} \frac{\mathbf{C}_Q + \sigma_2(\mathbf{B}_j^*)\sigma_1(\mathbf{A}_j^*)}{\|\mathbf{C}_Q + \sigma_2(\mathbf{B}_j^*)\sigma_1(\mathbf{A}_j^*)\| + \tau_Q} \mathbf{X} \right) \left(m_{V,j}^* \frac{\mathbf{C}_V + \sigma_2(\mathbf{B}_j^*)\sigma_1(\mathbf{A}_j^*)}{\|\mathbf{C}_V + \sigma_2(\mathbf{B}_j^*)\sigma_1(\mathbf{A}_j^*)\| + \tau_V} \right) \mathbf{X}, \end{aligned} \quad (19)$$

for almost surely \mathbb{X} .

Now, the index set $[L]$ can be partitioned into \tilde{m} subsets $\tilde{K}_1, \tilde{K}_2, \dots, \tilde{K}_{\tilde{m}}$ where $\tilde{m} \leq L$, such that $\exp(c_j) = \exp(c_{j'})$ for any indices $j, j' \in \tilde{K}_i$ and $i \in [\tilde{m}]$. Thus, equation (19) can be rewritten as follows:

$$\begin{aligned} & \sum_{i=1}^{\tilde{m}} \sum_{j \in \tilde{K}_i} \exp(c_j^*) \exp \left(\mathbf{X}^\top m_{Q,j} \frac{\mathbf{C}_Q + \sigma_2(\mathbf{B}_j^*)\sigma_1(\mathbf{A}_j^*)}{\|\mathbf{C}_Q + \sigma_2(\mathbf{B}_j^*)\sigma_1(\mathbf{A}_j^*)\| + \tau_Q} \mathbf{X} \right) \left(m_{V,j} \frac{\mathbf{C}_V + \sigma_2(\mathbf{B}_j)\sigma_1(\mathbf{A}_j)}{\|\mathbf{C}_V + \sigma_2(\mathbf{B}_j)\sigma_1(\mathbf{A}_j)\| + \tau_V} \right) \mathbf{X} \\ &= \sum_{i=1}^{\tilde{m}} \sum_{j \in \tilde{K}_i} \exp(c_j^*) \exp \left(\mathbf{X}^\top m_{Q,j} \frac{\mathbf{C}_Q + \sigma_2(\mathbf{B}_j^*)\sigma_1(\mathbf{A}_j^*)}{\|\mathbf{C}_Q + \sigma_2(\mathbf{B}_j^*)\sigma_1(\mathbf{A}_j^*)\| + \tau_Q} \mathbf{X} \right) \left(m_{V,j} \frac{\mathbf{C}_V + \sigma_2(\mathbf{B}_j^*)\sigma_1(\mathbf{A}_j^*)}{\|\mathbf{C}_V + \sigma_2(\mathbf{B}_j^*)\sigma_1(\mathbf{A}_j^*)\| + \tau_V} \right) \mathbf{X}, \end{aligned}$$

for almost surely \mathbb{X} . The above equation implies that

$$\left\{ \frac{\mathbf{C}_V + \sigma_2(\mathbf{B}_j)\sigma_1(\mathbf{A}_j)}{\|\mathbf{C}_V + \sigma_2(\mathbf{B}_j)\sigma_1(\mathbf{A}_j)\| + \tau_V} : j \in \tilde{K}_i \right\} = \left\{ \frac{\mathbf{C}_V + \sigma_2(\mathbf{B}_j^*)\sigma_1(\mathbf{A}_j^*)}{\|\mathbf{C}_V + \sigma_2(\mathbf{B}_j^*)\sigma_1(\mathbf{A}_j^*)\| + \tau_V} : j \in \tilde{K}_i \right\},$$

for any $i \in [\tilde{m}]$ and for almost surely \mathbb{X} . Since the activation functions σ_1 and σ_2 are algebraically independent, the above result indicates that

$$\sum_{i=1}^{\tilde{m}} \sum_{j \in \tilde{K}_i} \exp(c_j) \delta_{(\mathbf{B}_j, \mathbf{A}_j)} = \sum_{i=1}^{\tilde{m}} \sum_{j \in \tilde{K}_i} \exp(c_j^*) \delta_{(\mathbf{B}_j^*, \mathbf{A}_j^*)}.$$

As a consequence, $\tilde{G} \equiv \tilde{G}_*$ and the proof is completed.

D.3 PROOF OF PROPOSITION 1

We recall that $(\mathbf{X}_1, \mathbf{Y}_1), (\mathbf{X}_2, \mathbf{Y}_2), \dots, (\mathbf{X}_n, \mathbf{Y}_n) \in \mathbb{R}^d \times \mathbb{R}^d$ are i.i.d. samples from the following regression model:

$$\mathbf{Y}_i = f_{G_*}(\mathbf{X}_i) + \varepsilon_i, \quad i = 1, 2, \dots, n,$$

where the Gaussian noises $\varepsilon_1, \dots, \varepsilon_n$ are i.i.d. and satisfy that $\mathbb{E}[\varepsilon_i | \mathbf{X}_i] = 0$ and $\text{Var}(\varepsilon_i | \mathbf{X}_i) = \sigma^2 I_{\bar{d}}$ for all $i \in [n]$ and $f_{G_*}(\cdot)$ admits the following form:

$$f_{G_*}(\mathbf{X}) := \sum_{j=1}^L \frac{\exp\left(\mathbf{X}^\top m_{Q,j} \frac{C_Q + \sigma_2(\mathbf{W}_{2,j}^* \mathbf{B}_j^*) \sigma_1(\mathbf{W}_{1,j}^* \mathbf{A}_j^*)}{\|C_Q + \sigma_2(\mathbf{W}_{2,j}^* \mathbf{B}_j^*) \sigma_1(\mathbf{W}_{1,j}^* \mathbf{A}_j^*)\| + \tau_Q} C_K \mathbf{X} + c_j^*\right)}{S(\mathbf{X})} \\ \times \left(m_{V,j} \frac{C_V + \sigma_2(\mathbf{W}_{2,j}^* \mathbf{B}_j^*) \sigma_1(\mathbf{W}_{1,j}^* \mathbf{A}_j^*)}{\|C_V + \sigma_2(\mathbf{W}_{2,j}^* \mathbf{B}_j^*) \sigma_1(\mathbf{W}_{1,j}^* \mathbf{A}_j^*)\| + \tau_V} \right) \mathbf{X}, \quad (20)$$

where

$$S(\mathbf{X}) = \sum_{k=1}^L \exp\left(\mathbf{X}^\top m_{Q,k} \frac{C_Q + \sigma_2(\mathbf{W}_{2,j}^* \mathbf{B}_j^*) \sigma_1(\mathbf{W}_{1,j}^* \mathbf{A}_j^*)}{\|C_Q + \sigma_2(\mathbf{W}_{2,j}^* \mathbf{B}_j^*) \sigma_1(\mathbf{W}_{1,j}^* \mathbf{A}_j^*)\| + \tau_Q} C_K \mathbf{X} + c_k^*\right). \quad (21)$$

The least-square estimator G_n takes the following form:

$$G_n := \arg \min_{G \in \mathcal{G}_{L'}(\Theta)} \sum_{i=1}^n \|\mathbf{Y}_i - f_G(\mathbf{X}_i)\|^2,$$

From the Gaussianity assumption of $\varepsilon_i | \mathbf{X}_i$ for all $i \in [n]$, we have $\mathbf{Y}_i | \mathbf{X}_i \sim \mathcal{N}(f_{G_*}(\mathbf{X}_i), \sigma^2 I_d)$ for all $i \in [n]$. Therefore, the least square estimator G_n is indeed equivalent to a maximum likelihood estimator with respect to the data $Y_1 | \mathbf{X}_1, \dots, Y_n | \mathbf{X}_n$, which has the following form:

$$G_n \in \arg \max_{G \in \mathcal{G}_{L'}(\Theta)} \frac{1}{n} \sum_{i=1}^n \log(p(\mathbf{Y}_i | f_G(\mathbf{X}_i), \sigma^2 I_d))$$

where $p(\mathbf{Y}_i | f_G(\mathbf{X}_i), \sigma^2 I_d)$ stands for multivariate Gaussian distribution with mean $f_G(\mathbf{X}_i)$ and covariance matrix $\sigma^2 I_d$.

Next, we will denote some useful notations to help us quantify the distance between the two density functions, $p(\mathbf{Y} | f_{G_n}(\mathbb{X}), \sigma^2 I_d)$ and $p(\mathbf{Y} | f_{G_*}(\mathbb{X}), \sigma^2 I_d)$. Recall the set of all mixing measure $\mathcal{G}_L(\Theta)$ and denote the set of regression function $\mathcal{F}_L := \{f_G(x) : G \in \mathcal{G}_L(\Theta)\}$. Let \mathcal{P}_L denotes the set of conditional density of all mixing measures in $\mathcal{G}_L(\Theta)$, and we further denote,

$$\tilde{\mathcal{P}}_L(\Theta) := \{p_{(G+G_*)/2}(\mathbf{Y} | \mathbb{X}) : G \in \mathcal{G}_L(\Theta)\}, \\ \tilde{\mathcal{P}}_L^{1/2}(\Theta) := \{p_{(G+G_*)/2}^{1/2}(\mathbf{Y} | \mathbb{X}) : G \in \mathcal{G}_L(\Theta)\}.$$

Also, for each $\delta > 0$, we define the Hellinger ball in $\tilde{\mathcal{P}}_L^{1/2}(\Theta)$ centered around the conditional density p_{G_*} ,

$$\tilde{\mathcal{P}}_L^{1/2}(\Theta, \delta) := \{p^{1/2} \in \tilde{\mathcal{P}}_L^{1/2}(\Theta) : h(p, p_{G_*}) \leq \delta\}.$$

To measure the size of the above sets, we leverage the following quantity from van de Geer (2000):

$$\mathcal{J}(\delta, \tilde{\mathcal{P}}_L^{1/2}(\Theta, \delta)) := \int_{\delta^2/2^{13}}^{\delta} H_B^{1/2}(t, \tilde{\mathcal{P}}_L^{1/2}(\Theta, t), \|\cdot\|_{\mathcal{L}^2(\mu)}) dt \vee \delta, \quad (22)$$

where $H_B(t, \tilde{\mathcal{P}}_L^{1/2}(\Theta, t), \|\cdot\|_{\mathcal{L}^2(\mu)})$ denotes the bracketing entropy of $\tilde{\mathcal{P}}_L^{1/2}(\Theta, t)$ under \mathcal{L}^2 -norm, while $t \vee \delta = \max(t, \delta)$. By adapting the proof argument of Theorem 7.4 and Theorem 9.2 in van de Geer (2000), we have the following lemma.

Lemma 2. Consider the function $\Psi(\delta) \geq \mathcal{J}(\delta, \mathcal{P}_{HL}^{1/2}(\Theta, \delta))$ such that $\Psi(\delta)/\delta^2$ is a non-increasing function of δ . Then, for some universal constant c and sequence (δ_n) such that $\sqrt{n}\delta_n^2 \geq c\Psi(\delta_n)$, we have that

$$\mathbb{P}\left(\mathbb{E}_{\mathbb{X}}[h(p_{\hat{G}_n}(\cdot|\mathbb{X}), p_{G_*}(\cdot|\mathbb{X}))] > \delta\right) \leq c \exp\left(-\frac{n\delta^2}{\nu^2}\right)$$

for all $\delta \geq \delta_n$.

The goal is to upper bound the bracketing entropy for any $0 < \epsilon \leq 1/2$, i.e.,

$$H_B(\epsilon, \tilde{\mathcal{P}}_{HL}^{1/2}(\Theta, \epsilon), \|\cdot\|_{\mathcal{L}^2(\mu)}) \lesssim \log(1/\epsilon). \quad (23)$$

We will prove this bound later. Given this bound, it follows that,

$$\mathcal{J}_B(\delta, \tilde{\mathcal{P}}_{HL}^{1/2}(\Theta, \delta)) \lesssim \int_{\delta^2/2^{13}}^{\delta} \log(1/t) dt \vee \delta. \quad (24)$$

Consider $\Psi(\delta) = \delta \cdot [\log(1/\delta)]^{1/2}$, the it is obvious that $\Psi(\delta)/\delta^2$ is non-increasing function of δ . In addition, equation 24 implies that $\Psi(\delta) \geq \mathcal{J}_B(\delta, \tilde{\mathcal{P}}_{HL}^{1/2}(\Theta, \delta))$. By choosing $\delta_n = \sqrt{\log(n)/n}$, we have $\sqrt{n}\delta_n^2 \geq c\Psi(\delta_n)$ for some universal constant c . An application of Lemma 2 leads us to the conclusion of Proposition 1:

$$h(p(\mathbf{Y}|g_{\hat{G}_n}(\mathbb{X}), \sigma^2 I_d), p(\mathbf{Y}|g_{G_*}(\mathbb{X}), \sigma^2 I_d)) = \mathcal{O}(\sqrt{\log(n)/n}), \quad (25)$$

where h denotes the Hellinger distance. Hellinger distance between two multivariate normal distributions has the following closed-form,

$$h(p(\mathbf{Y}|g_{\hat{G}_n}(\mathbb{X}), \sigma^2 I_d), p(\mathbf{Y}|g_{G_*}(\mathbb{X}), \sigma^2 I_d)) = 1 - \exp\left\{-\frac{1}{8\sigma^2} \|g_{\hat{G}_n}(\mathbb{X}) - g_{G_*}(\mathbb{X})\|^2\right\}.$$

Therefore, for sufficient large n , there exists some universal constant C such that

$$\|g_{\hat{G}_n}(\mathbb{X}) - g_{G_*}(\mathbb{X})\|^2 \leq 8\sigma^2 \log\left(\frac{1}{1 - C \log(n)/n}\right) \leq 16\sigma^2 C \log(n)/n.$$

As a consequence, we have

$$\|g_{\hat{G}_n}(\mathbb{X}) - g_{G_*}(\mathbb{X})\| = \mathcal{O}(\sqrt{\log(n)/n}),$$

or $\|g_{\hat{G}_n} - g_{G_*}\|_{L^2(\mu)} = \mathcal{O}_P(\sqrt{\log(n)/n})$. This concludes the proof of this proposition.

Now we go back to prove the upper bound at equation 23. First we derive an upper bound for the multivariate Gaussian density $p_G(\cdot|\mathbb{X})$. Since the variance effect σ^2 is fixed, we have

$$p_G(\mathbf{Y}|\mathbb{X}) = \frac{1}{(2\pi\sigma^2)^{d/2}} \exp\left(-\frac{\|\mathbf{Y} - f_G(\mathbb{X})\|^2}{2\sigma^2}\right) \leq \frac{1}{(2\pi\sigma^2)^{d/2}}.$$

Because the input space \mathcal{X} and parameter space Θ are both bounded, there exists a constant M such that $\|f_G(\mathbb{X})\| \leq M$ for $G \in \mathcal{G}_L$ and $\mathbb{X} \in \mathcal{X}$. Thus, for any $\|\mathbf{Y}\| \geq 2M$, we have $\frac{\|\mathbf{Y} - g_G(\mathbb{X})\|^2}{2\sigma^2} \geq \frac{\|\mathbf{Y}\|^2}{8\sigma^2}$, which leads to

$$p(\mathbf{Y}|g_G(\mathbb{X}), \sigma^2 I_d) = \frac{1}{(2\pi\sigma^2)^{d/2}} \exp\left(-\frac{\|\mathbf{Y} - g_G(\mathbb{X})\|^2}{2\sigma^2}\right) \leq \frac{1}{(2\pi\sigma^2)^{d/2}} \exp\left(-\frac{\|\mathbf{Y}\|^2}{8\sigma^2}\right).$$

We define the integrable function

$$K(\mathbf{Y}|\mathbb{X}) = \begin{cases} (2\pi\sigma^2)^{-d/2} & \text{for } \|\mathbf{Y}\| \leq 2M, \\ (2\pi\sigma^2)^{-d/2} \exp\left(-\frac{\|\mathbf{Y}\|^2}{8\sigma^2}\right) & \text{for } \|\mathbf{Y}\| > 2M, \end{cases}$$

and thus $p(\mathbf{Y}|g_G(\mathbb{X}), \sigma^2 I_d) \leq K(\mathbf{Y}|\mathbb{X})$ for all \mathbf{Y} and $\mathbb{X} \in \mathcal{X}$.

Let $\tau < \epsilon$ and $\{e_1, \dots, e_n\}$ be the τ -cover of $\mathcal{P}_L(\Theta)$ under ℓ_1 -norm such that the covering number $N := N(\tau, \mathcal{P}_L(\Theta), \|\cdot\|_1)$. Next, we construct the brackets of the form $[L_i(\mathbf{Y}|\mathbb{X}), U_i(\mathbf{Y}|\mathbb{X})]$, for $1 \leq i \leq N$ with

$$\begin{aligned} L_i(\mathbf{Y}|\mathbb{X}) &:= \max\{e_i(\mathbf{Y}|\mathbb{X}) - \tau, 0\}, \\ U_i(\mathbf{Y}|\mathbb{X}) &:= \max\{e_i(\mathbf{Y}|\mathbb{X}) + \tau, K(\mathbf{Y}|\mathbb{X})\}. \end{aligned}$$

It is straightforward to check that $\mathcal{P}_L(\Theta) \subset \bigcup_{i=1}^N [L_i(\mathbf{Y}|\mathbb{X}), U_i(\mathbf{Y}|\mathbb{X})]$ and $U_i(\mathbf{Y}|\mathbb{X}) - L_i(\mathbf{Y}|\mathbb{X}) \leq \min\{\eta, K(\mathbf{Y}|\mathbb{X})\}$. From this, we can achieve the following upper bound

$$\begin{aligned} \|U_i - L_i\|_1 &= \int_{\|\mathbf{Y}\| \leq 2M} |U_i(\mathbf{Y}|\mathbb{X}) - L_i(\mathbf{Y}|\mathbb{X})| d(\mathbb{X}, \mathbf{Y}) + \int_{\|\mathbf{Y}\| > 2M} |U_i(\mathbf{Y}|\mathbb{X}) - L_i(\mathbf{Y}|\mathbb{X})| d(\mathbb{X}, \mathbf{Y}) \\ &\leq K\tau + \exp\left(-\frac{K^2}{2\sigma^2}\right) \leq K'\tau, \end{aligned}$$

where $K := \max\{2M, \sqrt{8\sigma^2}\} \log(1/\tau)$ and K' be a positive constant. From the definition of bracket entropy, given that $H_B(K'\tau, \mathcal{P}_L(\Theta), \|\cdot\|_1)$ is the logarithm of the smallest number of bracket of size $K'\tau$ necessary to cover $\mathcal{P}_L(\Theta)$, we have

$$H_B(K'\tau, \mathcal{P}_L(\Theta), \|\cdot\|_1) \leq \log(N) = \log N(\tau, \mathcal{P}_L(\Theta), \|\cdot\|_1). \quad (26)$$

Now we want to bound the covering number N . Specifically, we denote $\Delta = \{\mathbf{B}_j, \mathbf{A}_j, c_j : (\mathbf{B}_j, \mathbf{A}_j, c_j) \in \Theta\}$. Because the parameter space Θ is compact, Δ is also a compact set. Thus, we can find τ -covers Δ_τ such that $|\Delta_\tau| \leq \mathcal{O}(\tau^{-2rdL-L})$.

For each mixing measure $G = \sum_{j=1}^L \exp(c_j) \delta_{(\mathbf{B}_j, \mathbf{A}_j)} \in \mathcal{G}_L$, we consider an other mixing measure as follows,

$$G' = \sum_{j=1}^L \exp(c'_j) \delta_{(\mathbf{B}'_j, \mathbf{A}'_j)}, \quad (27)$$

where $(c'_j, \mathbf{B}'_j, \mathbf{A}'_j) \in \Delta_\tau$ is the closest to $(c_j, \mathbf{B}_j, \mathbf{A}_j)$. We define some following intermediate functions to alleviate the bound:

$$\begin{aligned} f_1(\mathbf{X}) &:= \sum_{j=1}^L \text{softmax} \left(\mathbf{X}^\top m_{Q,j} \frac{C_Q + \sigma_2(\mathbf{B}_j) \sigma_1(\mathbf{A}_j)}{\|C_Q + \sigma_2(\mathbf{B}_j) \sigma_1(\mathbf{A}_j)\| + \tau_Q} \mathbf{X} + c_j \right) \left(m_{V,j} \frac{C_V + \sigma_2(\mathbf{B}'_j) \sigma_1(\mathbf{A}'_j)}{\|C_V + \sigma_2(\mathbf{B}'_j) \sigma_1(\mathbf{A}'_j)\| + \tau_V} \right) \mathbf{X}, \\ f_2(\mathbf{X}) &:= \sum_{j=1}^L \text{softmax} \left(\mathbf{X}^\top m_{Q,j} \frac{C_Q + \sigma_2(\mathbf{B}'_j) \sigma_1(\mathbf{A}'_j)}{\|C_Q + \sigma_2(\mathbf{B}'_j) \sigma_1(\mathbf{A}'_j)\| + \tau_Q} \mathbf{X} + c'_j \right) \left(m_{V,j} \frac{C_V + \sigma_2(\mathbf{B}'_j) \sigma_1(\mathbf{A}'_j)}{\|C_V + \sigma_2(\mathbf{B}'_j) \sigma_1(\mathbf{A}'_j)\| + \tau_V} \right) \mathbf{X}, \end{aligned}$$

Now, we have,

$$\begin{aligned} \|f - f_1\|_\infty &\leq \sum_{j=1}^L \sup_{\mathbf{X} \in \mathcal{X}} \text{softmax} \left(\mathbf{X}^\top m_{Q,j} \frac{C_Q + \sigma_2(\mathbf{B}_j) \sigma_1(\mathbf{A}_j)}{\|C_Q + \sigma_2(\mathbf{B}_j) \sigma_1(\mathbf{A}_j)\| + \tau_Q} \mathbf{X} + c_j \right) \cdot m_{V,j} \|\mathbf{V}_j \mathbf{X} - \mathbf{V}'_j \mathbf{X}\| \\ &\leq \sum_{j=1}^L \sup_{\mathbf{X} \in \mathcal{X}} m_{V,j} \|\mathbf{V}_j \mathbf{X} - \mathbf{V}'_j \mathbf{X}\| \\ &\leq \sum_{j=1}^L \sup_{\mathbf{X} \in \mathcal{X}} m_{V,j} \left\| \left(\frac{C_V + \sigma_2(\mathbf{B}_j) \sigma_1(\mathbf{A}_j)}{\|C_V + \sigma_2(\mathbf{B}_j) \sigma_1(\mathbf{A}_j)\| + \tau_V} - \frac{C_V + \sigma_2(\mathbf{B}'_j) \sigma_1(\mathbf{A}'_j)}{\|C_V + \sigma_2(\mathbf{B}'_j) \sigma_1(\mathbf{A}'_j)\| + \tau_V} \right) \mathbf{X} \right\| \\ &\lesssim \sum_{j=1}^L \sup_{\mathbf{X} \in \mathcal{X}} \|(\mathbf{B}_j, \mathbf{A}_j) - (\mathbf{B}'_j, \mathbf{A}'_j)\| \cdot \|\mathbb{X}\| \\ &\lesssim \sum_{j=1}^L \tau \cdot B \lesssim \tau, \end{aligned} \quad (28)$$

where the second last inequality holds due to the fact that the input space is bounded: $\|\mathbf{X}\| \leq B$ for all $\mathbb{X} \in \mathcal{X}$.

Similarly, we also have,

$$\begin{aligned}
\|f_1 - f_2\|_\infty &\leq \sum_{j=1}^L \sup_{\mathbb{X} \in \mathcal{X}} \left| \text{softmax} \left(\mathbf{X}^\top m_{Q,j} \frac{C_Q + \sigma_2(\mathbf{B}_j)\sigma_1(\mathbf{A}_j)}{\|C_Q + \sigma_2(\mathbf{B}_j)\sigma_1(\mathbf{A}_j)\| + \tau_Q} \mathbf{X} + c_j \right) \right. \\
&\quad \left. - \text{softmax} \left(\mathbf{X}^\top m_{Q,j} \frac{C_Q + \sigma_2(\mathbf{B}'_j)\sigma_1(\mathbf{A}'_j)}{\|C_Q + \sigma_2(\mathbf{B}'_j)\sigma_1(\mathbf{A}'_j)\| + \tau_Q} \mathbf{X} + c'_j \right) \right| \cdot m_{V,j} \|\mathbf{V}'_j \mathbf{X}\| \\
&\lesssim \sum_{j=1}^L \sup_{\mathbb{X} \in \mathcal{X}} \left| \text{softmax} \left(\mathbf{X}^\top m_{Q,j} \frac{C_Q + \sigma_2(\mathbf{B}_j)\sigma_1(\mathbf{A}_j)}{\|C_Q + \sigma_2(\mathbf{B}_j)\sigma_1(\mathbf{A}_j)\| + \tau_Q} \mathbf{X} + c_j \right) \right. \\
&\quad \left. - \text{softmax} \left(\mathbf{X}^\top m_{Q,j} \frac{C_Q + \sigma_2(\mathbf{B}'_j)\sigma_1(\mathbf{A}'_j)}{\|C_Q + \sigma_2(\mathbf{B}'_j)\sigma_1(\mathbf{A}'_j)\| + \tau_Q} \mathbf{X} + c'_j \right) \right| \\
&\lesssim \sum_{j=1}^L \|(\mathbf{B}_j, \mathbf{A}_j, c_j) - (\mathbf{B}'_j, \mathbf{A}'_j, c'_j)\| \\
&\lesssim \tau.
\end{aligned} \tag{29}$$

Therefore, by triangle inequality, we have $\|f - f_2\|_\infty \lesssim \tau$. Then, by noting that the Gaussian density function $f(x) = (2\pi\sigma^2)^{-d/2} \exp(-\|x\|^2/2\sigma^2)$ is a global Lipschitz function, we have,

$$\|p(\mathbf{Y}|f_G(\mathbb{X}), \sigma^2 I_d) - p(\mathbf{Y}|f_{G'}(\mathbb{X}), \sigma^2 I_d)\|_1 \lesssim \|g_G(\mathbb{X}) - g_{G'}(\mathbb{X})\|_\infty \lesssim \tau. \tag{30}$$

From the definition of covering number, we get

$$\log N(\tau, \mathcal{P}_{HL}(\Theta), \|\cdot\|_1) \leq |\Delta_\tau| \leq \mathcal{O}(\tau^{-2rdL-L}) \tag{31}$$

From equations 26 and 31, we have

$$H_B(\tau, \mathcal{P}_L(\Theta), \|\cdot\|_1) \lesssim \log(1/\xi).$$

By choosing $\xi = \epsilon/2$, we achieve that

$$H_B(\epsilon, \mathcal{P}_{HL}(\Theta), h) \lesssim \log(1/\epsilon).$$

This concludes our proof.

E USE OF LARGE LANGUAGE MODELS

In this paper, large language models were utilized exclusively for editorial assistance, such as grammar correction and spelling improvements, and were not applied to content creation, data analysis, or experimental design.