

Statistical Advantages of Perturbing Cosine Router in Sparse Mixture of Experts

Huy Nguyen[†] Pedram Akbarian^{†,*} Trang Pham^{◇,*}
Trang Nguyen^{◇,*} Shujian Zhang[†] Nhat Ho[†]

The University of Texas at Austin[†]
VinAI Research[◇]

May 23, 2024

Abstract

The cosine router in sparse Mixture of Experts (MoE) has recently emerged as an attractive alternative to the conventional linear router. Indeed, the cosine router demonstrates favorable performance in image and language tasks and exhibits better ability to mitigate the representation collapse issue, which often leads to parameter redundancy and limited representation potentials. Despite its empirical success, a comprehensive analysis of the cosine router in sparse MoE has been lacking. Considering the least square estimation of the cosine routing sparse MoE, we demonstrate that due to the intrinsic interaction of the model parameters in the cosine router via some partial differential equations, regardless of the structures of the experts, the estimation rates of experts and model parameters can be as slow as $\mathcal{O}(1/\log^\tau(n))$ where $\tau > 0$ is some constant and n is the sample size. Surprisingly, these pessimistic non-polynomial convergence rates can be circumvented by the widely used technique in practice to stabilize the cosine router — simply adding noises to the \mathbb{L}_2 norms in the cosine router, which we refer to as *perturbed cosine router*. Under the strongly identifiable settings of the expert functions, we prove that the estimation rates for both the experts and model parameters under the perturbed cosine routing sparse MoE are significantly improved to polynomial rates. Finally, we conduct extensive simulation studies in both synthetic and real data settings to empirically validate our theoretical results.

1 Introduction

Proposed by Jacobs et. al. [15] and Jordan et. al. [17], a mixture of experts (MoE) has been known as an effective statistical method to incorporate the capabilities of various specialized models called experts. Different from conventional mixture models [22] in which the mixture weights are scalars, the MoE rather utilizes a routing mechanism to determine a set of weights depending on an input token. In particular, the router first computes the similarity scores between each token and experts, and then assign more weights to the more relevant experts determined based on those scores. To further improve the scalability of the MoE, Shazeer et. al. [37] have recently introduced a sparse variant of this model, which routes each input to only a subset of experts. This sparse MoE model allows us to increase the number of learnable parameters with a nearly constant computational overhead. As a consequence, the sparse MoE has been leveraged in several applications, including large language models [16, 33, 19, 43, 7, 4, 32], computer vision [35, 21, 36], speech recognition

* Equal contribution.

[40, 9, 31], multi-task learning [12], and other applications [11, 3].

In the above applications, practitioners often use a linear router which calculates the similarity score by taking the inner product of a token hidden representation and an expert embedding. Nevertheless, Chi et. al. [2] discovered that utilizing the linear router might induce the representation collapse issue. This phenomenon takes place when a fraction of experts govern the decision-making process, leading to the redundancy of other experts. In response, Chi et. al. [2] proposed an alternative known as a cosine router. In particular, this router begins with projecting the token hidden representation into a low-dimensional space, followed by applying \mathbb{L}_2 normalization to both the token representations and expert embeddings. By doing so, the similarity scores become more stable, circumventing the dominance of certain experts. The efficacy of the cosine routing MoE has been experimentally demonstrated in language modeling [2], and domain generalization [20]. On the other hand, a comprehensive theoretical study of the cosine router has remained lacking.

In the literature, there are some attempts to understand the MoE models with different types of gating functions whose outputs are the composition of some functions and the routing scores. First of all, considering the classification problem with cluster structures, Chen et. al. [1] demonstrated that the router operated by a neural network could learn the cluster-center features, which helped divide a complex problem into simpler classification sub-problems that individual experts could handle. Next, Ho et. al. [14] studied the expert estimation under an input-free gating Gaussian MoE model, showing that the rates for estimating experts depend on the algebraic structures among experts. Subsequently, the Gaussian MoE models with softmax gating [29], Gaussian gating [30], Top-K sparse softmax gating [27], and dense-to-sparse gating functions [26] were continuously explored. Those works pointed out that interactions among model parameters via some partial differential equations (PDE) did harm the expert estimation rates, and advocated using Top-1 sparse gating function to eliminate such interactions. Saying that the setting of Gaussian MoE was far from practice, Nguyen et. al. [28] rather took into account a regression framework with the regression function being a deterministic MoE model. They verified the benefits of formulating experts as feed-forward networks with popular activation functions like ReLU and sigmoid from the perspective of the expert estimation problem.

In this paper, our main objective is to investigate the effects of the cosine router on the convergence of expert and parameter estimations since this problem allows us to have useful insights into the design of MoE models as in prior works. For that sake, let us now present the problem setting formally.

Problem setting. We assume that $(X_1, Y_1), (X_2, Y_2), \dots, (X_n, Y_n) \in \mathbb{R}^{d_1} \times \mathbb{R}$ is an i.i.d sample of size n generated according to the following model

$$Y_i = f_{G_*}(X_i) + \varepsilon_i, \quad i = 1, \dots, n, \quad (1)$$

where regression function $f_{G_*}(\cdot)$ takes the following form:

$$f_{G_*}(x) := \sum_{i=1}^{k_*} \text{Softmax} \left(\text{Top}_K \left(\frac{(\beta_{1i}^*)^\top x}{\|\beta_{1i}^*\| \cdot \|x\|}, \beta_{0i}^* \right) \right) \cdot h(x, \eta_i^*). \quad (2)$$

Here, the function $h(x, \eta)$ is known as *the expert function*, which we assumed to be of parametric form. Meanwhile, $(\beta_{0i}^*, \beta_{1i}^*, \eta_i^*)_{i=1}^{k_*}$ are true yet unknown parameters in $\mathbb{R} \times \mathbb{R}^{d_1} \times \mathbb{R}^{d_2}$ and $G_* :=$

$\sum_{i=1}^{k_*} \exp(\beta_{0i}^*) \delta_{(\beta_{1i}^*, \eta_i^*)}$ denotes the associated *mixing measure*, i.e. a weighted sum of Dirac measures δ . Additionally, we define for any vectors $v = (v_1, v_2, \dots, v_{k_*})$ and $u = (u_1, u_2, \dots, u_{k_*})$ in \mathbb{R}^{k_*} that $\text{Softmax}(v_i) := \exp(v_i) / \sum_{j=1}^{k_*} \exp(v_j)$ and

$$\text{Top}_K(v_i, u_i) := \begin{cases} v_i + u_i, & \text{if } v_i \text{ is in the top } K \text{ elements of } v; \\ -\infty, & \text{otherwise.} \end{cases}$$

In the cosine router in equation (2), we omit the step of reducing the dimension of the input token x , and assume that it has already been in a low-dimensional space for simplicity. Furthermore, we assume that X_1, X_2, \dots, X_n are i.i.d. samples from some probability distribution μ . Lastly, $\varepsilon_1, \varepsilon_2, \dots, \varepsilon_n$ are independent Gaussian noise variables such that $\mathbb{E}[\varepsilon_i | X_i] = 0$ and $\text{Var}(\varepsilon_i | X_i) = \sigma^2$ for all $1 \leq i \leq n$. Notably, the Gaussian assumption is just for the simplicity of proof argument.

Least squares estimation (LSE). To estimate the true parameters $(\beta_{0i}^*, \beta_{1i}^*, \eta_i^*)_{i=1}^{k_*}$, we leverage the popular least squares method [38]. Formally, the mixing measure G_* is approximated by

$$\hat{G}_n := \arg \min_G \sum_{i=1}^n (Y_i - f_G(X_i))^2. \quad (3)$$

Under the *exact-specified* setting, i.e., when the true number of expert k_* is known, the minimum in the above equation is subject to the set of all mixing measures with k_* atoms, denoted by $\mathcal{E}_{k_*}(\Theta) := \{G = \sum_{i=1}^{k_*} \exp(\beta_{0i}) \delta_{(\beta_{1i}, \eta_i)} : (\beta_{0i}, \beta_{1i}, \eta_i) \in \Theta\}$. On the other hand, under the *over-specified* setting, i.e., when k_* is unknown and the true model (2) is over-specified by a mixture of k experts where $k > k_*$, the minimum is subject to the set of all mixing measures with at most k atoms, i.e., $\mathcal{G}_k(\Theta) := \{G = \sum_{i=1}^{k'} \exp(\beta_{0i}) \delta_{(\beta_{1i}, \eta_i)} : 1 \leq k' \leq k, (\beta_{0i}, \beta_{1i}, \eta_i) \in \Theta\}$.

Universal assumptions. In the sequel, we implicitly impose four following mild assumptions on the model parameters, which were widely used in previous works [27, 28], unless stating otherwise:

(A.1) *Convergence of LSE:* The parameter space $\Theta \subseteq \mathbb{R} \times \mathbb{R}^{d_1} \times \mathbb{R}^{d_2}$ is compact, while the input space $\mathcal{X} \subseteq \mathbb{R}^{d_1}$ is bounded. This helps ensure the convergence of least squares estimation.

(A.2) *Distinct experts:* The true parameters $\eta_1^*, \dots, \eta_{k_*}^*$ are pair-wise distinct so that the experts $h(\cdot, \eta_1^*), \dots, h(\cdot, \eta_{k_*}^*)$ are different from each other. Furthermore, the expert function $h(\cdot, \eta)$ is Lipschitz continuous w.r.t its parameters and bounded.

(A.3) *Identifiability of the MoE:* In order that the cosine routing MoE is identifiable, i.e., $f_G(x) = f_{G_*}(x)$ for almost every x implies that $G \equiv G_*$, we let $\beta_{0k_*}^* = 0$.

(A.4) *Input-dependent router:* To ensure that the router is input-dependent, we assume that at least one among the parameters $\beta_{11}^*, \dots, \beta_{1k_*}^*$ is non-zero.

Technical challenges. The normalization of parameters in the cosine router leads to a fundamental challenge in theory. In particular, to establish parameter and expert estimation rates based on the convergence rate of regression function estimation, we rely on the decomposition of the regression function discrepancy $f_{\hat{G}_n}(x) - f_{G_*}(x)$ into a combination of linearly independent terms via Taylor expansions to the product of the softmax's numerator and the expert function, i.e. $H(x, \beta_1, \eta) :=$

$\exp(\frac{\beta_1^\top x}{\|\beta_1\| \cdot \|x\|})h(x, \eta)$. However, the normalization of β_1 in the cosine router leads to an intrinsic interaction inside this router via the following PDE:

$$\beta_1^\top \frac{\partial H}{\partial \beta_1}(x, \beta_1, \eta) = 0. \quad (4)$$

This PDE leads to a complex combination of terms in the Taylor expansions of the regression function discrepancy $f_{\widehat{G}_n}(x) - f_{G^*}(x)$, thereby creating several linearly dependent terms. We later demonstrate in Section 2 that these complicated interactions lead to very slow rates of estimating experts and parameters. In particular, these rates could be as slow as $1/\log^\tau(n)$ for some $\tau > 0$, where n denotes the sample size. To the best of our knowledge, such a phenomenon with the cosine router has never been observed in previous works.

Main contributions. In this work, we develop a comprehensive theoretical analysis of regression function estimation as well as parameter and expert estimations under the cosine router MoE model (1). Our contributions are two-fold and can be summarized as follows:

1. Cosine router: Equipped with the cosine router, we demonstrate that under both the exact-specified and the over-specified settings, the rates for estimating ground-truth parameters $\beta_{0i}^*, \beta_{1i}^*$ and η_i^* are slower than any polynomial rates and, therefore, could be as slow as $\mathcal{O}_P(1/\log^\tau(n))$, where $\tau > 0$ is some constant. These slow rates are attributed to the internal interaction among router parameters expressed by the PDE in equation (4). As a result, the estimation rates for experts $h(\cdot, \eta_i^*)$ are also negatively affected, and could be of order $\mathcal{O}_P(1/\log^\tau(n))$.

2. Perturbed cosine router: In response, we propose a novel router called *perturbed cosine router* in which we add noises to the \mathbb{L}_2 norms of the token representations and the expert embeddings. This not only helps stabilize the router but also eliminates the intrinsic interaction in equation (4). Additionally, we also establish identifiability conditions to characterize expert functions that have faster estimation rates than others under the exact-specified and over-specified settings, respectively. Those conditions indicate that the rates for estimating experts, which are formulated as feed-forward networks with widely used activation functions such as ReLU and GeLU, are significantly improved, ranging from $\mathcal{O}_P(\sqrt[4]{\log(n)/n})$ to $\mathcal{O}_P(\sqrt{\log(n)/n})$.

Outline. In Section 2, we establish the convergence rates of parameter and expert estimations under both the exact-specified and over-specified settings of the cosine router MoE. Then, we derive these rates when the cosine router is replaced by the perturbed cosine router in Section 3. Based on these theoretical results, we derive a few practical implications in Section 4. We empirically verify the (theoretical) benefits of the perturbed cosine router over the cosine router under both the synthetic and real data settings in Section 5 before concluding the paper in Section 6. Finally, proofs and additional details of the experiments are deferred to the Appendices.

Notations. We let $[n]$ stand for the set $\{1, 2, \dots, n\}$ for any $n \in \mathbb{N}$. Next, for any set S , we denote $|S|$ as its cardinality. For any vector $v \in \mathbb{R}^d$ and $\alpha := (\alpha_1, \alpha_2, \dots, \alpha_d) \in \mathbb{N}^d$, we let $v^\alpha = v_1^{\alpha_1} v_2^{\alpha_2} \dots v_d^{\alpha_d}$, $|v| := v_1 + v_2 + \dots + v_d$ and $\alpha! := \alpha_1! \alpha_2! \dots \alpha_d!$, while $\|v\|$ stands for its \mathbb{L}_2 -norm value. Lastly, for any 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 $a_n \leq C b_n$ for all $n \in \mathbb{N}$, where $C > 0$ is some universal constant. The notation $a_n = \mathcal{O}_P(b_n)$ indicates that a_n/b_n is stochastically bounded.

2 Cosine Router Mixture of Experts

In this section, we characterize the parameter and expert estimation rates under over-specified setting of the cosine router MoE. We first start with the convergence rate of the regression function estimation $f_{\widehat{G}_n}$ to the true regression function f_{G_*} under the $\mathbb{L}_2(\mu)$ norm in the following theorem:

Theorem 1. *Given the least-square estimator \widehat{G}_n defined in equation (3), the regression estimator $f_{\widehat{G}_n}(\cdot)$ converges to the true regression function $f_{G_*}(\cdot)$ at the following rate:*

$$\|f_{\widehat{G}_n} - f_{G_*}\|_{\mathbb{L}_2(\mu)} = \mathcal{O}_P(\sqrt{\log(n)/n}).$$

The proof of Theorem 1 is in Appendix B.1. The result of Theorem 1 indicates that the regression estimation rate is parametric. Therefore, as long as we can establish the lower bound $\|f_{\widehat{G}_n} - f_{G_*}\|_{\mathbb{L}_2(\mu)} \gtrsim \mathcal{L}(\widehat{G}_n, G_*)$ where \mathcal{L} is some loss function among parameters, we arrive at the parameter estimation rate $\mathcal{L}(\widehat{G}_n, G_*) = \mathcal{O}_P(\sqrt{\log(n)/n})$. This approach is the key component of the convergence rates of parameter and expert estimations under the exact-specified and over-specified settings of the cosine router MoE that we are going to discuss in Section 2.1 and Section 2.2.

2.1 Exact-specified Setting

Recall that under the exact-specified setting, the true number of experts k_* is known. Then, based on the notion of Voronoi cells [24], we will construct a Voronoi loss function among parameters tailored to this setting.

Voronoi loss. Let G be a mixing measure with k' atoms $\omega_i := (\beta_{1i}, \eta_i)$. Then, we distribute these atoms to the Voronoi cells generated by the atoms $\omega_j^* := (\beta_{1j}^*, \eta_j^*)$ of G_* , which are defined as

$$\mathcal{A}_j \equiv \mathcal{A}_j(G) := \{i \in [k'] : \|\omega_i - \omega_j^*\| \leq \|\omega_i - \omega_\ell^*\|, \forall \ell \neq j\}. \quad (5)$$

Since \widehat{G}_n has k_* atoms under this setting, each Voronoi cell $\mathcal{A}_j(\widehat{G}_n)$ has exactly one element when the sample size n is sufficiently large. Then, the Voronoi loss of interest, $\mathcal{L}_{1,r}(G, G_*)$, is given by

$$\max_{\{\ell_1, \dots, \ell_{K'}\} \subset [k_*]} \left\{ \sum_{j=1}^{K'} \left| \sum_{i \in \mathcal{A}_{\ell_j}} \exp(\beta_{0i}) - \exp(\beta_{0\ell_j}^*) \right| + \sum_{j=1}^{K'} \sum_{i \in \mathcal{A}_{\ell_j}} \exp(\beta_{0i}) \left[\|\Delta\beta_{1i\ell_j}\|^r + \|\Delta\eta_{i\ell_j}\|^r \right] \right\},$$

where $r \geq 1$ is some constant, $\Delta\beta_{1i\ell_j} := \beta_{1i} - \beta_{1\ell_j}^*$ and $\Delta\eta_{i\ell_j} := \eta_i - \eta_{\ell_j}^*$. Compared to the Voronoi loss in [28], the maximum operator is included in the above loss to deal with the Top_K function.

Note that, due to the parameter interaction inside the cosine router captured by the PDE (4), the lower bound $\|f_{\widehat{G}_n} - f_{G_*}\|_{\mathbb{L}_2(\mu)} \gtrsim \mathcal{L}_{1,r}(\widehat{G}_n, G_*)$ does not hold true, and thus, we cannot achieve the desired bound $\mathcal{L}_{1,r}(\widehat{G}_n, G_*) = \mathcal{O}_P(\sqrt{\log(n)/n})$ mentioned in Section 2. By contrast, we show in Appendix B.2 an opposed result to the previous lower bound, saying that

$$\lim_{\varepsilon \rightarrow 0} \inf_{G \in \mathcal{E}_{k_*}(\Theta) : \mathcal{L}_{1,r}(G, G_*) \leq \varepsilon} \frac{\|f_G - f_{G_*}\|_{\mathbb{L}_2(\mu)}}{\mathcal{L}_{1,r}(G, G_*)} = 0,$$

for any $r \geq 1$. This result implies the following minimax lower bound of parameter estimation:

Theorem 2. Under the exact-specified setting, the following minimax lower bound of estimating G_*

$$\inf_{\bar{G}_n \in \mathcal{E}_{k_*}(\Theta)} \sup_{G \in \mathcal{E}_{k_*}(\Theta)} \mathbb{E}_{f_G}[\mathcal{L}_{1,r}(\bar{G}_n, G)] \gtrsim n^{-1/2},$$

holds true for any $r \geq 1$, where \mathbb{E}_{f_G} indicates the expectation taken w.r.t the product measure with f_G^n and the infimum is over all estimators taking values in $\mathcal{E}_{k_*}(\Theta)$.

See Appendix B.2 for the proof of Theorem 2. There are two main implications of the above result:

(i) Parameter estimation rates. The above minimax lower bound together with the formulation of $\mathcal{L}_{1,r}$ indicate that the rates for estimating parameters β_{1j}^*, η_j^* are slower than any polynomial rates $\mathcal{O}_P(n^{-1/2r})$ and, thus, could be of order $\mathcal{O}_P(1/\log^\tau(n))$, where $\tau > 0$ is some constant.

(ii) Expert estimation rates. Assume that $\hat{G}_n := \sum_{i=1}^{k_*} \exp(\hat{\beta}_{0i}) \delta_{(\hat{\beta}_{1i}^n, \hat{\eta}_i^n)}$. Since the expert $h(\cdot, \eta)$ is Lipschitz continuous, it follows that

$$\sup_x |h(x, \hat{\eta}_i^n) - h(x, \eta_j^*)| \lesssim \|\hat{\eta}_i^n - \eta_j^*\|. \quad (6)$$

As a consequence, the estimation rates for the experts $h(\cdot, \eta_j^*)$ are no better than those for parameters η_j^* . Therefore, they could also be as slow as $\mathcal{O}_P(1/\log^\tau(n))$ regardless of their structures, including feed-forward networks as suggested in Nguyen et. al. [28].

2.2 Over-specified Setting

Under the over-specified setting, the true number of experts k_* becomes unknown. Then, we seek the LSE \hat{G}_n within the set of all mixing measures with at most k atoms $\mathcal{G}_k(\Theta)$, where $k > k_*$. Thus, there exists some atom (β_{1j}^*, η_j^*) of G_* approximated by at least two atoms $(\hat{\beta}_{1i}^n, \hat{\eta}_i^n)$ of \hat{G}_n . Equivalently, the expert $h(\cdot, \eta_j^*)$ is fitted by at least two experts $h(\cdot, \hat{\eta}_i^n)$. As a result, it is necessary to activate more than K experts in the formulation of the regression estimator $f_{\hat{G}_n}(\cdot)$ to ensure its convergence to the true regression function $f_{G_*}(\cdot)$. To this end, we consider a new regression function

$$\bar{f}_G(x) := \sum_{i=1}^k \text{Softmax} \left(\text{Top}_{\bar{K}} \left(\frac{(\beta_{1i})^\top x}{\|\beta_{1i}\| \cdot \|x\|}, \beta_{0i} \right) \right) \cdot h(x, \eta_i), \quad (7)$$

in which we turn on \bar{K} experts per input, where $\bar{K} > K$. Moreover, we also provide in Proposition 1 the minimum value of \bar{K} such that this new regression function can be used for estimating $f_{G_*}(\cdot)$.

Proposition 1. The following inequality holds true only if $\bar{K} \geq \max_{\{\ell_1, \dots, \ell_K\} \subset [k_*]} \sum_{j=1}^K |\mathcal{A}_{\ell_j}|$:

$$\inf_{G \in \mathcal{G}_k(\Theta)} \|\bar{f}_G - f_{G_*}\|_{\mathbb{L}_2(\mu)} = 0.$$

Similar to the exact-specified setting, the effects of the parameter interaction (4) on the convergence of least squares estimation under the over-specified setting is also illustrated by minimax lower bound for estimating G_* in Theorem 3, whose proof is deferred to Appendix B.3.

Theorem 3. Under the over-specified setting, the following minimax lower bound of estimating G_*

$$\inf_{\bar{G}_n \in \mathcal{G}_k(\Theta)} \sup_{G \in \mathcal{G}_k(\Theta) \setminus \mathcal{G}_{k^*-1}(\Theta)} \mathbb{E}_{f_G}[\mathcal{L}_{1,r}(\bar{G}_n, G)] \gtrsim n^{-1/2},$$

holds true for any $r \geq 1$, where \mathbb{E}_{f_G} indicates the expectation taken w.r.t the product measure with f_G^n and the infimum is over all estimators taking values in \mathcal{G}_k .

It can be seen that the convergence behavior of parameter estimation under the over-specified setting is analogous to that under the exact-specified setting. That is, the rates for estimating parameters β_{1j}^* and η_j^* as well as experts $h(\cdot, \eta_j^*)$ could be as slow as $\mathcal{O}_P(1/\log^7(n))$. This result indicates that the cosine router is not sample efficient for the sparse MoE model, which motivates us to enhance the performance of this router in Section 3.

3 Perturbed Cosine Router Mixture of Experts

In this section, we demonstrate that the pessimistic non-polynomial convergence rates of parameter and expert estimation under the cosine router can be easily circumvented by the widely used technique in practice to stabilize the cosine router: adding noises to the \mathbb{L}_2 norm in the cosine router. We name this new router as *perturbed cosine router*. We now present the formulation of a MoE with the perturbed cosine router under the regression setting.

Problem setup for the perturbed cosine router MoE model. We assume that an i.i.d. sample of size n : $(X_1, Y_1), (X_2, Y_2), \dots, (X_n, Y_n) \in \mathbb{R}^{d_1} \times \mathbb{R}$ is generated according to the model

$$Y_i = g_{G_*}(X_i) + \varepsilon_i, \quad i = 1, \dots, n, \quad (8)$$

where regression function $g_{G_*}(\cdot)$ takes the following form:

$$g_{G_*}(x) := \sum_{i=1}^{k_*} \text{Softmax} \left(\text{Top}_K \left(\frac{(\beta_{1i}^*)^\top x}{(\|\beta_{1i}^*\| + \tau_1) \cdot (\|x\| + \tau_2)}, \beta_{0i}^* \right) \right) \cdot h(x, \eta_i^*). \quad (9)$$

Here, $\tau_1, \tau_2 > 0$ are two noise hyper-parameters. The main difference between the two regression functions f_{G_*} and g_{G_*} is the noise hyper-parameters τ_1, τ_2 that we add to the norms of the expert embeddings β_{1i}^* and the token input x , which leads to the perturbed cosine router. By doing so, the parameter interaction inside the router as in equation (4) does not occur. More specifically, let us denote $\tilde{H}(x, \beta_1, \eta) := \exp\left(\frac{\beta_1^\top x}{(\|\beta_1\| + \tau_1) \cdot (\|x\| + \tau_2)}\right) h(x, \eta)$, then it can be verified that $\beta_1^\top \frac{\partial \tilde{H}}{\partial \beta_1}(x, \beta_1, \eta) \neq 0$.

Least squares estimation. Similar to the cosine router setting, we can estimate the unknown ground-truth parameters $(\beta_{0i}^*, \beta_{1i}^*, \eta_i^*)_{i=1}^{k_*}$ using the least-square estimator, which is given by:

$$\tilde{G}_n := \arg \min_G \sum_{i=1}^n \left(Y_i - g_G(X_i) \right)^2. \quad (10)$$

In the following theory, we provide a convergence rate of regression function estimation under the perturbed cosine router MoE model.

Theorem 4. Given a least squares estimator \tilde{G}_n defined in equation (10), the regression function estimation $g_{\tilde{G}_n}(\cdot)$ admits the following convergence rate:

$$\|g_{\tilde{G}_n} - g_{G_*}\|_{\mathbb{L}_2(\mu)} = \mathcal{O}_P(\sqrt{\log(n)/n}). \quad (11)$$

Proof of Theorem 4 is in Appendix C.1. The result of Theorem 4 proves that the regression function estimation rate $\mathcal{O}_P(\sqrt{\log(n)/n})$ under the perturbed cosine router MoE is of the same order as that with the vanilla cosine router in Theorem 1. Following the similar proof strategy in the cosine router MoE in Section 2 for capturing the convergence rates of parameter and expert estimations under the perturbed cosine router MoE model, it is sufficient to establish the lower bound between the difference of regression functions and the difference of parameters under both the exact-specified and over-specified settings.

In this section, we study the over-specified setting of the perturbed cosine router. The results for the exact-specified setting of the perturbed cosine router is in Appendix A. As explained in Section 2.2, to estimate the regression function $f_{G_*}(\cdot)$ under this setting, we use the following regression function in which we turn on \bar{K} experts rather than K experts per input:

$$\bar{g}_G(x) := \sum_{i=1}^{\bar{k}} \text{Softmax} \left(\text{Top}_{\bar{K}} \left(\frac{(\beta_{1i})^\top x}{(\|\beta_{1i}\| + \tau_1) \cdot (\|x\| + \tau_2)}, \beta_{0i} \right) \right) \cdot h(x, \eta_i), \quad (12)$$

where $\bar{K} \geq \max_{\{\ell_1, \dots, \ell_K\} \subset [k_*]} \sum_{j=1}^K |\mathcal{A}_{\ell_j}|$. We now derive a condition called *strong identifiability* on the expert function $h(\cdot, \eta)$ to identify which experts exhibit faster estimation rates than others under the over-specified setting.

Definition 1 (Strong identifiability). *An expert function $x \mapsto h(x, \eta)$ is called strongly identifiable if it is twice differentiable with respect to its parameter η , and the set of functions in x*

$$\left\{ \frac{\partial^{|\alpha_1| + |\alpha_2|} \tilde{H}}{\partial \beta_1^{\alpha_1} \partial \eta^{\alpha_2}}(x, \beta_{1i}, \eta_i) : \alpha_1 \in \mathbb{N}^{d_1}, \alpha_2 \in \mathbb{N}^{d_2}, 0 \leq |\alpha_1| + |\alpha_2| \leq 2 \right\},$$

is linearly independent for almost every x for any $k \geq 1$ and pair-wise distinct parameters η_1, \dots, η_k , where we denote $\tilde{H}(x, \beta_1, \eta) := \exp(\frac{\beta_1^\top x}{(\|\beta_1\| + \tau_1) \cdot (\|x\| + \tau_2)}) h(x, \eta)$.

Example. For experts formulated as neural networks, i.e. $h(x, (a, b)) = \phi(a^\top x + b)$, if the activation $\phi(\cdot)$ is selected as $\text{ReLU}(\cdot)$ or $\text{tanh}(\cdot)$, then they are strongly identifiable. Conversely, a linear expert $h(x, (a, b)) = a^\top x + b$ does not meet the strong identifiability the condition.

To capture the convergence behavior of expert estimation rate under the over-specified setting in Theorem 5, we will use the Voronoi loss $\mathcal{L}_3(G, G_*)$ defined as follows:

$$\begin{aligned} \max_{\{\ell_1, \dots, \ell_K \subset [k_*]\}} & \left\{ \sum_{j=1}^K \left| \sum_{i \in \mathcal{A}_{\ell_j}} \exp(\beta_{0i}) - \exp(\beta_{0\ell_j}^*) \right| + \sum_{\substack{j \in [K]: \\ |\mathcal{A}_{\ell_j}|=1}} \sum_{i \in \mathcal{A}_{\ell_j}} \exp(\beta_{0i}) \left[\|\Delta \beta_{1i\ell_j}\| + \|\Delta \eta_{i\ell_j}\| \right] \right. \\ & \left. + \sum_{\substack{j \in [K]: \\ |\mathcal{A}_{\ell_j}|>1}} \sum_{i \in \mathcal{A}_{\ell_j}} \exp(\beta_{0i}) \left[\|\Delta \beta_{1i\ell_j}\|^2 + \|\Delta \eta_{i\ell_j}\|^2 \right] \right\}. \quad (13) \end{aligned}$$

Theorem 5. *Suppose that the expert function $h(x, \eta)$ satisfies the condition in Definition 1, then the following L^2 -lower bound holds true for any $G \in \mathcal{G}_k(\Theta)$:*

$$\|\bar{g}_G - g_{G_*}\|_{\mathbb{L}_2(\mu)} \gtrsim \mathcal{L}_3(G, G_*).$$

Furthermore, this bound and the result in Theorem 4 imply that $\mathcal{L}_3(\tilde{G}_n, G_*) = \mathcal{O}_P(\sqrt{\log(n)/n})$.

The proof of Theorem 5 is in Appendix C.3. A few comments regarding this theorem are in order:

(i) Under the over-specified setting, parameters β_{1j}^*, η_j^* which are fitted by one atom, i.e. $|\mathcal{A}_j(\tilde{G}_n)| = 1$, share the same estimation rate of order $\mathcal{O}_P(\sqrt{\log(n)/n})$. Meanwhile, those for parameters fitted by more than one atom, i.e. $|\mathcal{A}_j(\tilde{G}_n)| > 1$, are slightly slower, standing at order $\mathcal{O}_P(\sqrt[4]{\log(n)/n})$.

(ii) Given the above parameter estimation rates and the inequality (6), we observe that the rates for estimating strongly identifiable experts $h(\cdot, \eta_j^*)$ range from $\mathcal{O}_P(\sqrt[4]{\log(n)/n})$ to $\mathcal{O}_P(\sqrt{\log(n)/n})$. Notably, those rates apply for polynomial experts of degree at least two, i.e. $h(x, (a, b)) = (a^\top x + b)^p$ with $p \geq 2$, as they satisfy the strong identifiability condition. By contrast, the estimation rates for those experts when using the vanilla cosine router (see Theorem 3) and the softmax gating (see [Theorem 4.6, [28]]) are significantly slower, and could be of order $\mathcal{O}_P(1/\log^\tau(n))$. This observation demonstrates the sample efficiency of our proposed perturbed cosine router.

4 Practical Implications

We now discuss two important practical implications from the theoretical results of the paper.

1. Misspecified settings. Thus far in the paper, we have only considered well-specified settings, namely, the data are assumed to be sampled from the (perturbed) cosine router MoE. Although it may look restrictive, the results under this setting lay an important foundation for a more realistic misspecified setting where the data are not necessarily generated from those models.

Under that misspecified setting, we assume that the data are generated from a regression framework as in equation (1) but with an arbitrary regression function $q(\cdot)$, which is not a (perturbed) cosine router MoE. Then, we can demonstrate that the LSE \hat{G}_n converges to a mixing measure $\bar{G} \in \arg \min_{G \in \mathcal{G}_k(\Theta)} \|q - f_G\|_{\mathbb{L}_2(\mu)}$, where $f_G(\cdot)$ is a regression function taking the form of the (perturbed) cosine router MoE. Furthermore, the optimal mixing measure will be in the boundary of the parameter space $\mathcal{G}_k(\Theta)$, namely, \bar{G} has k atoms. Thus, as n becomes sufficiently large, \hat{G}_n also has k atoms. The insights from our theories for the well-specified setting indicate that the Voronoi losses can be used to obtain the estimation rates of individual parameters of the LSE \hat{G}_n to those of \bar{G} and, therefore, the expert estimation rates under the (perturbed) cosine router MoE.

(1.1) Cosine router MoE: the worst expert estimation rate could be as slow as $\mathcal{O}_P(1/\log^\tau(n))$ for some $\tau > 0$. It indicates that we still need an exponential number of data (roughly $\exp(1/\epsilon^\tau)$ where ϵ is the desired approximation error) to estimate the experts as well as select important experts.

(1.2) Perturbed cosine router MoE: the slowest expert estimation rate is of order $\mathcal{O}_P(n^{-1/4})$. Thus, we only need a polynomial number of data (roughly ϵ^{-4}) to estimate the experts. This explains

why the perturbed cosine router is a solution to the parameter estimation problem, or more generally, the expert estimation problem of the MoE models.

2. Model design: From the benefits of the perturbed cosine router for the expert estimation of MoE models, our theories suggest that when using the cosine router to avoid the representation collapse, practitioners should add noises to \mathbb{L}_2 norms of the token hidden representations and the expert embeddings to achieve a favorable performance. Additionally, the strong identifiability condition also verifies the advantages of using non-linear expert networks over linear ones.

5 Experiments

In this section, we first conduct numerical experiments on synthetic data (cf. Section 5.1), and then carry out experiments with real data on language modeling (cf. Section 5.2) and domain generalization (cf. Appendix F) tasks. Our main goal is to empirically demonstrate the efficacy of the perturbed cosine router over the vanilla cosine router in MoE models.

5.1 Numerical Experiments

We first perform numerical experiments on synthetic data to empirically verify the theoretical convergence rates of the least squares estimation for both perturbed and vanilla cosine router MoE models. We generate synthetic data based on the model described in equation (1). Specifically, we generate $\{(X_i, Y_i)\}_{i=1}^n \subset \mathbb{R}^d \times \mathbb{R}$ by first sampling $X_i \sim \text{Uniform}([-1, 1]^d)$ for $i = 1, \dots, n$. Then, we generate Y_i according to the following model:


$$Y_i = f_{G_*}(X_i) + \varepsilon_i, \quad i = 1, \dots, n, \quad (14)$$

where the regression function $f_{G_*}(\cdot)$ is defined as:

$$f_{G_*}(x) := \sum_{i=1}^{k_*} \text{Softmax} \left(\text{Top}_2 \left(\frac{(\beta_{1i}^*)^\top x}{(\|\beta_{1i}^*\| + \tau) \cdot (\|x\| + \tau)}, \beta_{0i}^* \right) \right) \cdot \phi \left((a_i^*)^\top x + b_i^* \right). \quad (15)$$

The input data dimension is set at $d = 32$. We employ $k_* = 8$ experts of the form $\phi((a_i^*)^\top x + b_i^*)$, where the activation function ϕ is either the ReLU function or the identity function. Additionally, we activate $K = 2$ experts per input. The details of the values of the parameters as well as the training procedure are in Appendix E.1.

Results. Two experimental scenarios are examined: (1) Exact-specified, and (2) Over-specified. Data for both scenarios are generated according to equation (14). In the exact-specified scenario, the model is fitted with the same number of experts as the data generation model, specifically $k = k_* = 8$. In the over-specified scenario, the model includes one additional expert, totaling $k = k_* + 1 = 9$ experts, and employs the $\text{Top}_{\bar{K}}$ operator with $\bar{K} = 4$. In each scenario, experiments are conducted using both the standard and the perturbed cosine routers, with τ set to zero for the standard router and 0.1 for the perturbed router. For each experiment, we calculate the Voronoi losses for every model and report the mean values for each sample size in Figure 1. Error bars representing two standard deviations are also shown. In Figure 1a, the empirical convergence rates of both the standard and perturbed routers are analyzed under the exact-specified setting. The perturbed router shows a rapid convergence rate of $\mathcal{O}(n^{-0.5})$, while the standard vanilla router has a noticeably slower rate of $\mathcal{O}(n^{-0.11})$. Similarly, in Figure 1b, the convergence rates are assessed for the same routers under the over-specified setting. Here, the perturbed router again shows a faster convergence rate of $\mathcal{O}(n^{-0.47})$, compared to the standard vanilla router’s slower rate of $\mathcal{O}(n^{-0.05})$.



(a) Exact-specified setting with $k = k_* = 8$ experts

(b) Over-specified setting with $k = k_* + 1 = 9$

Figure 1: Logarithmic plots displaying empirical convergence rates. Subfigures 1a and 1b depict the empirical averages of the corresponding Voronoi losses for the exact and over-specified settings, respectively. The blue lines depict the Voronoi loss associated with the perturbed router, whereas the green lines are indicative of the Voronoi loss associated with the standard vanilla router. The red dash-dotted lines are used to illustrate the fitted lines for determining the empirical convergence rate.

5.2 Language Modeling

For language modeling, we apply both perturbed and vanilla cosine router MoEs to character-level and word-level language modeling tasks to compare their performance.

Datasets. We begin by evaluating the model’s pre-training capabilities on character-level language modeling tasks using the Enwik8 and Text8 datasets ([23]). Additionally, we assess its performance on word-level language modeling with the Wikitext-103 dataset [25].

Metrics. To quantify the performance of our perturbed cosine router relative to the original one, we utilize the Bit per character (BPC) metric for character-level language modeling and Perplexity (PPL) for word-level language modeling tasks.

Architecture and training procedure. In order to alleviate the representation collapse issues associated with estimating routing scores in the original space, we first employ the XMoE method [2] to project input representations on lower-dimensional space and parameterize experts with corresponding lower-dimensional embeddings. Subsequently, we calculate the routing scores of inputs and embeddings in this reduced-dimensional space using our proposed perturbed cosine router. Our experiments utilize the Switch Transformer architecture [8], which is fundamentally a sparse variant of the T5 encoder-decoder architecture [34], with MoE layers replacing the MLPs. Detailed information regarding the datasets, metrics, training setup and hyperparameters for this task is provided in Appendix E.2.

Results. The empirical advantage of our proposed router over the vanilla version when applied to language modeling tasks is demonstrated in Table 1. The results indicate that the perturbed cosine router enhances the performance of XMoE on the Enwik8, Text8, and Wikitext-103 datasets across both small and medium configurations. Notably, for the Enwik8 dataset at both small and medium scales, as well as for the Text8 dataset at the medium scale, XMoE with the perturbed cosine router reduces BPC, with reductions ranging from approximately 0.15 to 0.20 BPC. Furthermore, for the Wikitext-103 dataset, our proposed router marginally outperforms the original cosine router.

Table 1: Performance of XMoE with vanilla and perturbed cosine routers on language modeling task.

Model	Enwik8 (BPC ↓)		Text8 (BPC ↓)		Wikitext-103 (PPL ↓)	
	Small	Medium	Small	Medium	Small	Medium
XMoE (Vanilla cosine router)	1.213	1.161	1.310	1.271	90.070	38.018
XMoE (Perturbed cosine router)	1.197	1.147	1.303	1.251	89.910	37.859

6 Conclusion

In this paper, we investigate the impacts of the cosine router on the convergence rates of least squares estimation in MoE models. We figure out that owing to the parameter interaction inside the cosine router expressed by a PDE, the rates for estimating parameters and experts are slower than any polynomial rates and, therefore, could be as slow as $\mathcal{O}_P(1/\log^\tau(n))$. In response to this issue, we propose using the perturbed cosine router where we add noises to the \mathbb{L}_2 norms of the token representations and the expert embeddings in the cosine router in order to eliminate the previous parameter interaction. Equipped with this novel router, we demonstrate that if the expert function satisfies the weak (strong) identifiability condition, then the parameter and expert estimation rates are significantly improved to be of polynomial orders under the exact-specified (over-specified) setting.

Limitations. There are a few limitations in our current analysis. First of all, the assumption that the data are sampled from the (perturbed) cosine router MoE is often violated in real-world settings. However, as discussed in Section 4, our theories can totally be extended to a more realistic misspecified setting where the data are not necessarily generated from those models, which we leave for future development. Second, since the ground-truth parameters are implicitly assumed to be independent of the sample size n , the parameter and expert estimation rates presented in this work are point-wise rather than uniform. To cope with this problem, we can utilize the techniques for characterizing the uniform parameter estimation rates in traditional mixture models (see [13, 5]). Nevertheless, since the adaptation of those techniques to the setting of the (perturbed) cosine router MoE is still challenging due to the complex structures of the (perturbed) cosine router, we believe that further technical tools need to be developed to achieve the desired uniform estimation rates.

In this supplementary material, we first explore the exact-specified setting of the perturbed cosine router MoE model in Appendix A. Next, we provide proofs for theoretical results of Section 2 and Section 3 in Appendix B and Appendix C, respectively. Those proofs are partially supported by auxiliary results presented in Appendix D. Subsequently, in Appendix E, we specify the details for the experiments performed in Section 5. Finally, we conduct further experiments on the applications of MoE models in domain generalization in Appendix F to empirically demonstrate the benefits of using our proposed perturbed cosine router over the vanilla cosine router.

A Exact-specified Setting of the Perturbed Cosine Router MoE

We now consider the exact-specified setting of the perturbed cosine router MoE model (8). To begin with, we introduce a condition called *weak identifiability* on the expert function $h(\cdot, \eta)$ to characterize which experts have faster estimation rates than others under this setting.

Definition 2 (Weak identifiability). *An expert function $x \mapsto h(x, \eta)$ is said to be weakly identifiable if it is differentiable w.r.t its parameter η and the set of functions in x*

$$\left\{ \frac{\partial^{|\alpha_1|+|\alpha_2|} \tilde{H}}{\partial \beta_1^{\alpha_1} \partial \eta^{\alpha_2}}(x, \beta_{1i}, \eta_i) : \alpha_1 \in \mathbb{N}^{d_1}, \alpha_2 \in \mathbb{N}^{d_2}, 0 \leq |\alpha_1| + |\alpha_2| \leq 1 \right\},$$

is linearly independent for almost every x , for any $k \geq 1$ and pair-wise distinct parameters η_1, \dots, η_k , where we denote $\tilde{H}(x, \beta_1, \eta) := \exp\left(\frac{\beta_1^\top x}{(\|\beta_1\| + \tau_1) \cdot (\|x\| + \tau_2)}\right) h(x, \eta)$.

Recall from the ‘‘Technical challenges’’ paragraph in Section 1 that a key step to establish the expert estimation rates is to decompose the difference $f_{\tilde{G}_n}(x) - f_{G_*}(x)$ into a combination of linearly independent terms via Taylor expansions to the function $H(\cdot, \beta_1, \eta)$. Therefore, the purpose of the weak identifiability condition is to avoid all potential parameter interactions as in equation (4), which may lead to undesirable linearly dependent terms.

Example. For simplicity, we consider experts formulated as neural networks, i.e. $h(x, (a, b)) = \phi(a^\top x + b)$. It can be validated that if the function $\phi(\cdot)$ is either a popular activation such as $\text{ReLU}(\cdot)$ and $\tanh(\cdot)$ or a polynomial $\phi(z) = z^p$, for any $p \in \mathbb{N}$, then the expert $h(x, (a, b))$ is weakly identifiable. On the other hand, a constant expert $h(\cdot, \eta) = \text{constant}$ fails to satisfy the weak identifiability the condition.

Next, we will use the Voronoi loss function $\mathcal{L}_2(G, G_*)$ defined below to determine the estimation rates for weakly identifiable experts in Theorem 6, whose proof can be found in Appendix C.2.

$$\max_{\{\ell_1, \dots, \ell_K\} \subset [k_*]} \left\{ \sum_{j=1}^K \left| \sum_{i \in \mathcal{A}_{\ell_j}} \exp(\beta_{0i}) - \exp(\beta_{0\ell_j}^*) \right| + \sum_{j=1}^K \sum_{i \in \mathcal{A}_{\ell_j}} \exp(\beta_{0i}) \left[\|\Delta \beta_{1i\ell_j}\| + \|\Delta \eta_{i\ell_j}\| \right] \right\}. \quad (16)$$

Theorem 6. *Assume that $h(\cdot, \eta)$ is a weakly identifiable expert function, then the following lower bound holds true for any $G \in \mathcal{E}_{k_*}(\Theta)$:*

$$\|g_G - g_{G_*}\|_{\mathbb{L}_2(\mu)} \gtrsim \mathcal{L}_2(G, G_*).$$

Furthermore, this bound and the result in Theorem 4 imply that $\mathcal{L}_2(\tilde{G}_n, G_*) = \mathcal{O}_P(\sqrt{\log(n)/n})$.

The bound $\mathcal{L}_2(\tilde{G}_n, G_*) = \mathcal{O}_P(\sqrt{\log(n)/n})$ all the parameters β_{1j}^*, η_j^* enjoy the same parametric estimation rates, standing at order $\mathcal{O}_P(\sqrt{\log(n)/n})$. Furthermore, by employing the argument in equation (6), we deduce that the rates for estimating experts $h(\cdot, \eta_j^*)$ are also of order $\mathcal{O}_P(\sqrt{\log(n)/n})$. Those rates are substantially faster than their counterparts when using the vanilla cosine router, which could be as slow as $\mathcal{O}_P(1/\log^\tau(n))$ (see Theorem 2). This comparison highlights the benefits of our proposed perturb cosine router over the vanilla cosine router.

B Proof for Results in Section 2

In this appendix, we provide proofs for the theoretical results regarding the cosine router in stated in Section 2, including Theorem 1, Theorem 2, and Theorem 3, in that order.

B.1 Proof of Theorem 1

Prior to presenting the proof, let us introduce some necessary notations. Firstly, we denote by $\mathcal{R}_k(\Theta)$ the set of regression functions w.r.t mixing measures in $\mathcal{G}_k(\Theta)$, that is, $\mathcal{R}_k(\Theta) := \{g_G(x) : G \in \mathcal{G}_k(\Theta)\}$. Additionally, for each $\delta > 0$, the L^2 ball centered around the regression function g_{G_*} and intersected with the set $\mathcal{R}_k(\Theta)$ is defined as

$$\mathcal{R}_k(\Theta, \delta) := \{g \in \mathcal{R}_k(\Theta) : \|g - g_{G_*}\|_{L^2(\mu)} \leq \delta\}.$$

In order to measure the size of the above set, Geer et. al. [38] suggest using the following quantity:

$$\mathcal{J}_B(\delta, \mathcal{R}_k(\Theta, \delta)) := \int_{\delta^2/2^{13}}^{\delta} H_B^{1/2}(t, \mathcal{R}_k(\Theta, t), \|\cdot\|_{L^2(\mu)}) dt \vee \delta, \quad (17)$$

where $H_B(t, \mathcal{R}_k(\Theta, t), \|\cdot\|_{L^2(\mu)})$ stands for the bracketing entropy [38] of $\mathcal{R}_k(\Theta, u)$ under the L^2 -norm, and $t \vee \delta := \max\{t, \delta\}$. By using the similar proof argument of Theorem 7.4 and Theorem 9.2 in [38] with notations being adapted to this work, we obtain the following lemma:

Lemma 1. *Take $\Psi(\delta) \geq \mathcal{J}_B(\delta, \mathcal{R}_k(\Theta, \delta))$ that satisfies $\Psi(\delta)/\delta^2$ is a non-increasing function of δ . Then, for some universal constant c and for some sequence (δ_n) such that $\sqrt{n}\delta_n^2 \geq c\Psi(\delta_n)$, we achieve that*

$$\mathbb{P}\left(\|g_{\tilde{G}_n} - g_{G_*}\|_{L^2(\mu)} > \delta\right) \leq c \exp\left(-\frac{n\delta^2}{c^2}\right),$$

for all $\delta \geq \delta_n$.

General picture. We first show that when the expert functions are Lipschitz continuous, the following bound holds for any $0 < \varepsilon \leq 1/2$:

$$H_B(\varepsilon, \mathcal{R}_k(\Theta), \|\cdot\|_{L^2(\mu)}) \lesssim \log(1/\varepsilon). \quad (18)$$

Given this bound, it follows that

$$\mathcal{J}_B(\delta, \mathcal{R}_k(\Theta, \delta)) = \int_{\delta^2/2^{13}}^{\delta} H_B^{1/2}(t, \mathcal{R}_k(\Theta, t), \|\cdot\|_{L^2(\mu)}) dt \vee \delta \lesssim \int_{\delta^2/2^{13}}^{\delta} \log(1/t) dt \vee \delta. \quad (19)$$

Let $\Psi(\delta) = \delta \cdot [\log(1/\delta)]^{1/2}$, then $\Psi(\delta)/\delta^2$ is a non-increasing function of δ . Furthermore, equation (19) indicates that $\Psi(\delta) \geq \mathcal{J}_B(\delta, \mathcal{R}_k(\Theta, \delta))$. In addition, let $\delta_n = \sqrt{\log(n)/n}$, then we get that $\sqrt{n}\delta_n^2 \geq c\Psi(\delta_n)$ for some universal constant c . Finally, by applying Lemma 1, we achieve the desired conclusion of the theorem. As a consequence, it suffices to demonstrate the bound (30).

Proof for the bound (18). Since the expert functions are Lipschitz continuous, then for any function $f_G \in \mathcal{R}_k(\Theta)$, we have that $f_G(x) \leq M$ for all x where $M > 0$ is some constant.

Let $\tau \leq \varepsilon$ and $\{\pi_1, \dots, \pi_N\}$ be the τ -cover under the $\mathbb{L}_2(\mu)$ norm of the set $\mathcal{R}_k(\Theta)$ where $N := N(\tau, \mathcal{R}_k(\Theta), \|\cdot\|_{\mathbb{L}_2(\mu)})$ is the η -covering number of the metric space $(\mathcal{R}_k(\Theta), \|\cdot\|_{\mathbb{L}_2(\mu)})$. Then, we construct the brackets of the form $[L_i(x), U_i(x)]$ for all $i \in [N]$ as follows:

$$\begin{aligned} L_i(x) &:= \max\{\pi_i(x) - \tau, 0\}, \\ U_i(x) &:= \max\{\pi_i(x) + \tau, M\}. \end{aligned}$$

From the above formulation, it can be checked that $\mathcal{R}_k(\Theta) \subset \cup_{i=1}^N [L_i(x), U_i(x)]$, and $U_i(x) - L_i(x) \leq 2 \min\{2\tau, M\}$. Thus, we get that

$$\|U_i - L_i\|_{\mathbb{L}_2(\mu)}^2 = \int (U_i(x) - L_i(x))^2 d\mu(x) \leq \int 16\tau^2 d\mu(x) = 16\tau^2,$$

which indicates that $\|U_i - L_i\|_{\mathbb{L}_2(\mu)} \leq 4\tau$. By definition of the bracketing entropy, we achieve that

$$H_B(4\tau, \mathcal{R}_k(\Theta), \|\cdot\|_{\mathbb{L}_2(\mu)}) \leq \log N = \log N(\tau, \mathcal{R}_k(\Theta), \|\cdot\|_{\mathbb{L}_2(\mu)}). \quad (20)$$

Therefore, it is necessary to provide an upper bound for the covering number N . Indeed, let us denote $\Delta := \{(\beta_0, \beta_1) \in \mathbb{R} \times \mathbb{R}^{d_1} : (\beta_0, \beta_1, \eta) \in \Theta\}$ and $\Omega := \{\eta \in \mathbb{R}^{d_2} : (\beta_0, \beta_1, \eta) \in \Theta\}$. Since Θ is a compact set, Δ and Ω are also compact. Therefore, we can find τ -covers Δ_τ and Ω_τ for Δ and Ω , respectively. Furthermore, it can be validated that

$$|\Delta_\tau| \leq \mathcal{O}_P(\tau^{-(d_1+1)k}), \quad |\Omega_\tau| \leq \mathcal{O}_P(\tau^{-d_2k}).$$

For each mixing measure $G = \sum_{i=1}^k \exp(\beta_{0i})\delta_{(\beta_{1i}, \eta_i)} \in \mathcal{G}_k(\Theta)$, we consider two other mixing measures G' and \bar{G} defined as

$$G' := \sum_{i=1}^k \exp(\beta_{0i})\delta_{(\beta_{1i}, \bar{\eta}_i)}, \quad \bar{G} := \sum_{i=1}^k \exp(\bar{\beta}_{0i})\delta_{(\bar{\beta}_{1i}, \bar{\eta}_i)}.$$

Here, $\bar{\eta}_i \in \Omega_\tau$ such that $\bar{\eta}_i$ is the closest to η_i in that set, while $(\bar{\beta}_{0i}, \bar{\beta}_{1i}) \in \Delta_\tau$ is the closest to

(β_{0i}, β_{1i}) in that set. Now, we aim to upper bound the term $\|f_G - f_{G'}\|_{\mathbb{L}_2(\mu)}^2$. In particular, we have

$$\begin{aligned}
\|f_G - f_{G'}\|_{\mathbb{L}_2(\mu)}^2 &= \int \left[\sum_{i=1}^k \text{Softmax} \left(\text{Top}_{\overline{K}} \left(\frac{(\beta_{1i})^\top x}{\|\beta_{1i}\| \cdot \|x\|}, \beta_{0i} \right) \right) \cdot [h(x, \eta_i) - h(x, \bar{\eta}_i)] \right]^2 d\mu(x) \\
&\leq \int \left[\sum_{i=1}^k (h(x, \eta_i) - h(x, \bar{\eta}_i)) \right]^2 d\mu(x) \\
&\leq k \int \sum_{i=1}^k [h(x, \eta_i) - h(x, \bar{\eta}_i)]^2 d\mu(x) \\
&\leq k \int \sum_{i=1}^k [L_1 \cdot \|\eta_i - \bar{\eta}_i\|]^2 d\mu(x) \\
&\leq k^2 (L_1 \tau)^2,
\end{aligned}$$

which implies that $\|g_G - g_{G'}\|_{\mathbb{L}_2(\mu)} \lesssim \tau$. Above, the second inequality is obtained by applying the Cauchy-Schwarz inequality, the third inequality holds as the softmax weight is bounded by 1, and the fourth inequality is due to the fact that the expert $h(x, \cdot)$ is a Lipschitz function with some Lipschitz constant L_1 .

Next, we demonstrate that $\|g_{G'} - g_{\overline{G}}\|_1 \lesssim \eta$. To this end, let us consider $q := \binom{k}{K}$ K -element subsets of $[k]$, which are assumed to take the form $\{\ell_1, \ell_2, \dots, \ell_K\}$ for any $\ell \in [q]$. Additionally, we also denote $\{\ell_{K+1}, \dots, \ell_k\} := \{1, \dots, k\} \setminus \{\ell_1, \dots, \ell_K\}$ for any $\ell \in [q]$. Then, we define

$$\begin{aligned}
\mathcal{X}_\ell &:= \{x \in \mathcal{X} : \cos(\beta_{1i}, x) \geq \cos(\beta_{1i'}, x) : i \in \{\ell_1, \dots, \ell_K\}, i' \in \{\ell_{K+1}, \dots, \ell_{k_*}\}\}, \\
\tilde{\mathcal{X}}_\ell &:= \{x \in \mathcal{X} : \cos(\bar{\beta}_{1i}, x) \geq \cos(\bar{\beta}_{1i'}, x) : i \in \{\ell_1, \dots, \ell_K\}, i' \in \{\ell_{K+1}, \dots, \ell_{k_*}\}\}.
\end{aligned}$$

By using the same arguments as in the proof of Lemma 3 in Appendix D, we achieve that either $\mathcal{X}_\ell = \tilde{\mathcal{X}}_\ell$ or \mathcal{X}_ℓ has measure zero for any $\ell \in [q]$. As the Softmax function is differentiable, it is Lipschitz continuous with some Lipschitz constant $L_2 \geq 0$. Since \mathcal{X} is a bounded set, we may assume that $\|x\| \leq B$ for any $x \in \mathcal{X}$. Next, we denote

$$\begin{aligned}
\pi_\ell(x) &:= \left(\frac{(\beta_{1\ell_i})^\top x}{\|\beta_{1\ell_i}\| \cdot \|x\|} + \beta_{0\ell_i} \right)_{i=1}^K, \\
\bar{\pi}_\ell(x) &:= \left(\frac{(\bar{\beta}_{1\ell_i})^\top x}{\|\bar{\beta}_{1\ell_i}\| \cdot \|x\|} + \bar{\beta}_{0\ell_i} \right)_{i=1}^K,
\end{aligned}$$

for any K -element subset $\{\ell_1, \dots, \ell_K\}$ of $\{1, \dots, k_*\}$. Then, we get

$$\begin{aligned}
\|\text{Softmax}(\pi_\ell(x)) - \text{Softmax}(\bar{\pi}_\ell(x))\| &\leq L_2 \cdot \|\pi_\ell(x) - \bar{\pi}_\ell(x)\| \\
&\leq L_2 \cdot \sum_{i=1}^K \left(\|\beta_{1\ell_i} - \bar{\beta}_{1\ell_i}\| \cdot \|x\| + |\beta_{0\ell_i} - \bar{\beta}_{0\ell_i}| \right) \\
&\leq L_2 \cdot \sum_{i=1}^K (\eta B + \eta) \\
&\lesssim \eta.
\end{aligned}$$

Back to the proof for $\|f_{G'} - f_{\bar{G}}\|_1 \lesssim \eta$, it follows from the above results that

$$\begin{aligned}
\|f_{G'} - f_{\bar{G}}\|_1 &= \int_{\mathcal{X}} |f_{G'}(x) - f_{\bar{G}}(x)| \, d\mu(x) \leq \sum_{\ell=1}^q \int_{\mathcal{X}_\ell} |f_{G'}(x) - f_{\bar{G}}(x)| \, d\mu(x) \\
&\leq \sum_{\ell=1}^q \int_{\mathcal{X}_\ell} \sum_{i=1}^K \left| \text{Softmax}(\pi_\ell(x)_i) - \text{Softmax}(\bar{\pi}_\ell(x)_i) \right| \cdot |h(x, \bar{\eta}_{\ell_i})| \, d\mu(x) \\
&\lesssim \eta,
\end{aligned} \tag{21}$$

By the triangle inequality, we have

$$\|f_G - f_{\bar{G}}\|_{\mathbb{L}_2(\mu)} \leq \|f_G - f_{G'}\|_{\mathbb{L}_2(\mu)} + \|f_{G'} - f_{\bar{G}}\|_{\mathbb{L}_2(\mu)} \lesssim \tau.$$

By definition of the covering number, we deduce that

$$\begin{aligned}
N(\tau, \mathcal{R}_k(\Theta), \|\cdot\|_{\mathbb{L}_2(\mu)}) &\leq |\Delta_\tau| \times |\Omega_\tau| \\
&\leq \mathcal{O}_P(n^{-(d_1+1)k}) \times \mathcal{O}(n^{-d_2k}) \\
&\leq \mathcal{O}(n^{-(d_1+1+d_2)k}).
\end{aligned} \tag{22}$$

Putting the results in equations (20) and (22) together, we achieve that

$$H_B(4\tau, \mathcal{R}_k(\Theta), \|\cdot\|_{\mathbb{L}_2(\mu)}) \lesssim \log(1/\tau).$$

By setting $\tau = \varepsilon/4$, we achieve that

$$H_B(\varepsilon, \mathcal{R}_k(\Theta), \|\cdot\|_{\mathbb{L}_2(\mu)}) \lesssim \log(1/\varepsilon),$$

which completes the proof.

B.2 Proof of Theorem 2

Lemma 2. *If the following holds for any $r \geq 1$:*

$$\lim_{\varepsilon \rightarrow 0} \inf_{G \in \mathcal{E}_{k_*}(\Theta) : \mathcal{L}_{1,r}(G, G_*) \leq \varepsilon} \frac{\|f_G - f_{G_*}\|_{\mathbb{L}_2(\mu)}}{\mathcal{L}_{1,r}(G, G_*)} = 0, \tag{23}$$

then we obtain that

$$\inf_{\bar{G}_n \in \mathcal{E}_{k_*}(\Theta)} \sup_{G \in \mathcal{E}_{k_*}(\Theta)} \mathbb{E}_{f_G}[\mathcal{L}_{1,r}(\bar{G}_n, G)] \gtrsim n^{-1/2}. \tag{24}$$

Proof of Lemma 2. Indeed, from the Gaussian assumption on the noise variables ϵ_i , we obtain that $Y_i | X_i \sim \mathcal{N}(f_{G_*}(X_i), \sigma^2)$ for all $i \in [n]$. Next, the assumption in equation (23) indicates for sufficiently small $\varepsilon > 0$ and a fixed constant $C_1 > 0$ which we will choose later, we can find a mixing measure $G'_* \in \mathcal{E}_{k_*}(\Theta)$ such that $\mathcal{L}_{1,r}(G'_*, G_*) = 2\varepsilon$ and $\|f_{G'_*} - f_{G_*}\|_{L^2(\mu)} \leq C_1\varepsilon$. From Le Cam's lemma [41], as the Voronoi loss function $\mathcal{L}_{1,r}$ satisfies the weak triangle inequality, we obtain that

$$\begin{aligned}
&\inf_{\bar{G}_n \in \mathcal{E}_{k_*}(\Theta)} \sup_{G \in \mathcal{E}_{k_*}(\Theta)} \mathbb{E}_{f_G}[\mathcal{L}_{1,r}(\bar{G}_n, G)] \\
&\gtrsim \frac{\mathcal{L}_{1,r}(G'_*, G_*)}{8} \exp(-n \mathbb{E}_{X \sim \mu}[\text{KL}(\mathcal{N}(f_{G'_*}(X), \sigma^2), \mathcal{N}(f_{G_*}(X), \sigma^2))]) \\
&\gtrsim \varepsilon \cdot \exp(-n \|f_{G'_*} - f_{G_*}\|_{L^2(\mu)}^2), \\
&\gtrsim \varepsilon \cdot \exp(-C_1 n \varepsilon^2),
\end{aligned} \tag{25}$$

where the second inequality is due to the fact that

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

By choosing $\varepsilon = n^{-1/2}$, we obtain that $\varepsilon \cdot \exp(-C_1 n \varepsilon^2) = n^{-1/2} \exp(-C_1)$. As a consequence, we achieve the desired minimax lower bound in equation (24). \square

Main proof. It is sufficient to show that the following limit holds true for any $r \geq 1$:

$$\lim_{\varepsilon \rightarrow 0} \inf_{G \in \mathcal{E}_{k_*}(\Theta): \mathcal{L}_{1,r}(G, G_*) \leq \varepsilon} \frac{\|f_G - f_{G_*}\|_{\mathbb{L}_2(\mu)}}{\mathcal{L}_{1,r}(G, G_*)} = 0. \quad (26)$$

To this end, we need to construct a sequence of mixing measures $G_n \in \mathcal{E}_{k_*}(\Theta)$ that satisfies $\mathcal{L}_{1,r}(G_n, G_*) \rightarrow 0$ and

$$\frac{\|f_{G_n} - f_{G_*}\|_{\mathbb{L}_2(\mu)}}{\mathcal{L}_{1,r}(G_n, G_*)} \rightarrow 0,$$

as $n \rightarrow \infty$. Next, let us take into account the sequence $G_n = \sum_{i=1}^{k_*} \exp(\beta_{0i}^n) \delta_{(\beta_{1i}^n, \eta_i^n)}$ in which

- $\exp(\beta_{0i}^n) = \exp(\beta_{0i}^*)$ for any $1 \leq i \leq k_*$;
- $\beta_{11}^n = \left(1 + \frac{1}{n}\right) \beta_{11}^*$ and $\beta_{1i}^n = \beta_{1i}^*$ for any $2 \leq i \leq k_*$;
- $\eta_i^n = \eta_i^*$ for any $1 \leq i \leq k_*$.

Consequently, it can be verified that when $n \rightarrow \infty$, we have

$$\mathcal{L}_{1,r}(G_n, G_*) = \exp(\beta_{01}^*) \left[\|\beta_{11}^n - \beta_{11}^*\|^r \right] = \exp(\beta_{01}^*) \cdot \left(\frac{\sqrt{d}}{n} \right)^r \rightarrow 0,$$

Next, we demonstrate that $\|f_{G_n} - f_{G_*}\|_{\mathbb{L}_2(\mu)} / \mathcal{L}_{1,r}(G_n, G_*) \rightarrow 0$.

Subsequently, to specify the top K selection in the formulations of $f_{G_n}(x)$ and $f_{G_*}(x)$, we divide the input space \mathcal{X} into complement subsets in two ways. In particular, we first consider $q := \binom{k_*}{K}$ different K -element subsets of $[k_*]$, which are assumed to take the form $\{\ell_1, \dots, \ell_K\}$, for $\ell \in [q]$. Additionally, we denote $\{\ell_{K+1}, \dots, \ell_{k_*}\} := [k_*] \setminus \{\ell_1, \dots, \ell_K\}$. Then, we define for each $\ell \in [q]$ two following subsets of \mathcal{X} :

$$\begin{aligned} \mathcal{X}_\ell^n &:= \left\{ x \in \mathcal{X} : \cos(\beta_{1j}^n, x) \geq \cos(\beta_{1j'}^n, x) : \forall j \in \{\ell_1, \dots, \ell_K\}, j' \in \{\ell_{K+1}, \dots, \ell_{k_*}\} \right\}, \\ \mathcal{X}_\ell^* &:= \left\{ x \in \mathcal{X} : \cos(\beta_{1j}^*, x) \geq \cos(\beta_{1j'}^*, x) : \forall j \in \{\ell_1, \dots, \ell_K\}, j' \in \{\ell_{K+1}, \dots, \ell_{k_*}\} \right\}. \end{aligned}$$

Since $\beta_{1j}^n \rightarrow \beta_{1j}^*$ as $n \rightarrow \infty$ for any $j \in [k_*]$, we have for any arbitrarily small $\eta_j > 0$ that $\|\beta_{1j}^n - \beta_{1j}^*\| \leq \eta_j$ for sufficiently large n . By applying Lemma 3, we obtain that $\mathcal{X}_\ell^n = \mathcal{X}_\ell^*$ for any $\ell \in [q]$ for sufficiently large n .

Let us consider an arbitrary subset \mathcal{X}_ℓ^* of the input space \mathcal{X} where ℓ is an arbitrary index in $[q]$.

Case 1: $1 \notin \{\ell_1, \dots, \ell_K\}$. In this case, we can validate that $f_{G_n}(x) - f_{G_*}(x) = 0$ for almost every $x \in \mathcal{X}_\ell^*$.

Case 2: $1 \in \{\ell_1, \dots, \ell_K\}$. WLOG, we assume that $\{\ell_1, \dots, \ell_K\} = \{1, \dots, K\}$. Then, for almost every $x \in \mathcal{X}_\ell^*$, we can represent $f_{G_n}(x)$ and $f_{G_*}(x)$ as

$$\begin{aligned} f_{G_n}(x) &= \sum_{i=1}^K \text{Softmax}\left(\frac{(\beta_{1i}^n)^\top x}{\|\beta_{1i}^n\| \cdot \|x\|}\right) \cdot h(x, \eta_i^n), \\ f_{G_*}(x) &= \sum_{i=1}^K \text{Softmax}\left(\frac{(\beta_{1i}^*)^\top x}{\|\beta_{1i}^*\| \cdot \|x\|}\right) \cdot h(x, \eta_i^*). \end{aligned}$$

Now, we consider the quantity

$$Q_n(x) := \left[\sum_{j=1}^{k_*} \exp\left(\frac{(\beta_{1j}^*)^\top x}{\|\beta_{1j}^*\| \cdot \|x\|} + \beta_{0j}^* \right) \right] \cdot [f_{G_n}(x) - f_{G_*}(x)], \quad (27)$$

which can be decomposed as follows:

$$\begin{aligned} Q_n(x) &= \sum_{j=1}^{k_*} \sum_{i \in \mathcal{A}_j} \exp(\beta_{0i}^n) \left[\exp\left(\frac{(\beta_{1i}^n)^\top x}{\|\beta_{1i}^n\| \cdot \|x\|}\right) h(x, \eta_i^n) - \exp\left(\frac{(\beta_{1j}^*)^\top x}{\|\beta_{1j}^*\| \cdot \|x\|}\right) h(x, \eta_j^*) \right] \\ &\quad - \sum_{j=1}^{k_*} \sum_{i \in \mathcal{A}_j} \exp(\beta_{0i}^n) \left[\exp\left(\frac{(\beta_{1i}^n)^\top x}{\|\beta_{1i}^n\| \cdot \|x\|}\right) f_{G_n}(x) - \exp\left(\frac{(\beta_{1j}^*)^\top x}{\|\beta_{1j}^*\| \cdot \|x\|}\right) f_{G_n}(x) \right] \\ &\quad + \sum_{j=1}^{k_*} \left(\sum_{i \in \mathcal{A}_j} \exp(\beta_{0i}^n) - \exp(\beta_{0j}^*) \right) \exp\left(\frac{(\beta_{1j}^*)^\top x}{\|\beta_{1j}^*\| \cdot \|x\|}\right) [h(x, \eta_j^*) - f_{G_n}(x)] \\ &:= A_n(x) - B_n(x) + C_n(x). \end{aligned}$$

Since $\exp(\beta_{0i}^n) = \exp(\beta_{0i}^*)$ for all $i \in [k_*]$, we deduce that $C_n(x) = 0$. Additionally, from the choices of β_{1i}^n and η_i^n , we can rewrite $A_n(x)$ as

$$A_n(x) = \exp(\beta_{01}^*) \left[\exp\left(\frac{(\beta_{11}^n)^\top x}{\|\beta_{11}^n\| \cdot \|x\|}\right) - \exp\left(\frac{(\beta_{11}^*)^\top x}{\|\beta_{11}^*\| \cdot \|x\|}\right) \right] h(x, \eta_1^*).$$

Let us denote $F(x, \beta_1) := \exp\left(\frac{\beta_1^\top x}{\|\beta_1\| \cdot \|x\|}\right)$. By applying the Taylor expansion of order r , we have

$$\begin{aligned} A_n(x) &= \exp(\beta_{01}^*) h(x, \eta_1^*) \sum_{|\alpha|=1}^r \frac{1}{\alpha!} \cdot (\beta_{11}^n - \beta_{11}^*)^\alpha \cdot \frac{\partial^{|\alpha|} F}{\partial \beta_1^\alpha}(x, \beta_{11}^*) + R(x) \\ &= \exp(\beta_{01}^*) h(x, \eta_1^*) \sum_{|\alpha|=1}^r \frac{1}{\alpha!} \left(1 + \frac{1}{n}\right)^{|\alpha|} (\beta_{11}^*)^\alpha \cdot \frac{\partial^{|\alpha|} F}{\partial \beta_1^\alpha}(x, \beta_{11}^*) + R(x), \end{aligned}$$

where $R(x)$ is a Taylor remainder such that $R(x)/\mathcal{L}_{1,r}(G_n, G_*) \rightarrow 0$ as $n \rightarrow \infty$. It is implied from Lemma 4 (see Appendix D) that

$$\sum_{|\alpha|=t} \frac{1}{\alpha!} (\beta_{11}^*)^\alpha \cdot \frac{\partial^{|\alpha|} F}{\partial \beta_1^\alpha}(x, \beta_{11}^*) = 0,$$

for any $1 \leq t \leq r$, it follows that $A_n(x) = R(x)$. This result indicates that $A_n(x)/\mathcal{L}_{1,r}(G_n, G_*) \rightarrow 0$ as $n \rightarrow \infty$. By arguing similarly, we also obtain that $B_n(x)/\mathcal{L}_{1,r}(G_n, G_*) \rightarrow 0$ as $n \rightarrow \infty$. Combine the previous results together, we achieve that

$$Q_n(x)/\mathcal{L}_{1,r}(G_n, G_*) \rightarrow 0.$$

Since the input space \mathcal{X} and the parameter space Θ are both bounded, the term $\sum_{j=1}^{k_*} \exp\left(\frac{(\beta_{1j}^*)^\top x}{\|\beta_{1j}^*\| \cdot \|x\|} + \beta_{0j}^*\right)$ is also bounded. This result together with the formulation of $Q_n(x)$ in equation (27) suggests that $[f_{G_n}(x) - f_{G_*}(x)]/\mathcal{L}_{1,r}(G_n, G_*) \rightarrow 0$ for almost every $x \in \mathcal{X}_\ell^*$. Note that the subset \mathcal{X}_ℓ^* is chosen arbitrarily, therefore, the previous result holds for almost every $x \in \mathcal{X}$. As a consequence, we get that

$$\frac{\|f_{G_n} - f_{G_*}\|_{\mathbb{L}_2(\mu)}}{\mathcal{L}_{1,r}(G_n, G_*)} \rightarrow 0,$$

as $n \rightarrow \infty$, and hence, achieve the result in equation (26).

B.3 Proof of Theorem 3

Similar to the proof of Theorem 2 in Appendix B.2, we only need to demonstrate that the following limit holds true for any $r \geq 1$:

$$\lim_{\varepsilon \rightarrow 0} \inf_{G \in \mathcal{G}_k(\Theta): \mathcal{L}_{1,r}(G, G_*) \leq \varepsilon} \frac{\|\bar{f}_G - f_{G_*}\|_{\mathbb{L}_2(\mu)}}{\mathcal{L}_{1,r}(G, G_*)} = 0. \quad (28)$$

For that purpose, it suffices to build a sequence of mixing measures $G_n \in \mathcal{G}_k(\Theta)$ that satisfies $\mathcal{L}_{1,r}(G_n, G_*) \rightarrow 0$ and

$$\frac{\|\bar{f}_{G_n} - f_{G_*}\|_{\mathbb{L}_2(\mu)}}{\mathcal{L}_{1,r}(G_n, G_*)} \rightarrow 0,$$

as $n \rightarrow \infty$. Let us consider the sequence $G_n = \sum_{i=1}^{k_*+1} \exp(\beta_{0i}^n) \delta_{(\beta_{1i}^n, \eta_i^n)}$ in which

- $\exp(\beta_{01}^n) = \exp(\beta_{02}^n) = \frac{1}{2} \exp(\beta_{01}^*)$, and $\exp(\beta_{0i}^n) = \exp(\beta_{0(i-1)}^*)$ for any $3 \leq i \leq k_* + 1$;
- $\beta_{11}^n = \left(1 - \frac{1}{n}\right) \beta_{11}^*$, $\beta_{12}^n = \left(1 + \frac{1}{n}\right) \beta_{11}^*$ and $\beta_{1i}^n = \beta_{1(i-1)}^*$ for any $3 \leq i \leq k_* + 1$;
- $\eta_1^n = \eta_2^n = \eta_1^*$, and $\eta_i^n = \eta_{i-1}^*$ for any $3 \leq i \leq k_* + 1$.

Consequently, it can be verified that when $n \rightarrow \infty$, we have

$$\mathcal{L}_{1,r}(G_n, G_*) = \frac{1}{2} \exp(\beta_{01}^*) \left[\|\beta_{11}^n - \beta_{11}^*\|^r + \|\beta_{12}^n - \beta_{11}^*\|^r \right] = \exp(\beta_{01}^*) \cdot \left(\frac{\sqrt{d_1}}{n}\right)^r \rightarrow 0,$$

Next, we demonstrate that $\|\bar{f}_{G_n} - f_{G_*}\|_{\mathbb{L}_2(\mu)}/\mathcal{L}_{1,r}(G_n, G_*) \rightarrow 0$.

Regarding the top- K selection in the conditional density g_{G_*} , we partition the input space \mathcal{X} in a similar fashion to Appendix B.2. More specifically, we consider $q = \binom{k_*}{K}$ subsets $\{\ell_1, \dots, \ell_K\}$ of $\{1, \dots, k_*\}$ for any $\ell \in [q]$, and denote $\{\ell_{K+1}, \dots, \ell_{k_*}\} := [k_*] \setminus \{\ell_1, \dots, \ell_K\}$. Then, we define

$$\mathcal{X}_\ell^* := \left\{ x \in \mathcal{X} : \cos(\beta_{1j}^*, x) \geq \cos(\beta_{1j'}^*, x), \forall j \in \{\ell_1, \dots, \ell_K\}, j' \in \{\ell_{K+1}, \dots, \ell_{k_*}\} \right\},$$

for any $\ell \in [q]$.

On the other hand, we need to introduce a new partition method of the input space for the weight selection in the regression function \bar{f}_{G_n} . In particular, let $\bar{K} \in \mathbb{N}$ such that $\max_{\{\ell_j\}_{j=1}^K \subset [k_*]} \sum_{j=1}^K |\mathcal{A}_{\ell_j}| \leq \bar{K} \leq k$ and $\bar{q} := \binom{k}{\bar{K}}$. Then, for any $\bar{\ell} \in [\bar{q}]$, we denote $(\bar{\ell}_1, \dots, \bar{\ell}_k)$ as a subset of $[k]$ and $\{\bar{\ell}_{\bar{K}+1}, \dots, \bar{\ell}_k\} := [k] \setminus \{\bar{\ell}_1, \dots, \bar{\ell}_{\bar{K}}\}$. Additionally, we define

$$\mathcal{X}_{\bar{\ell}}^n := \left\{ x \in \mathcal{X} : \cos(\beta_{1i}^n, x) \geq \cos(\beta_{1i'}^n, x), \forall i \in \{\bar{\ell}_1, \dots, \bar{\ell}_{\bar{K}}\}, i' \in \{\bar{\ell}_{\bar{K}+1}, \dots, \bar{\ell}_k\} \right\}.$$

Recall that we have $\beta_{1i}^n \rightarrow \beta_{1j}^*$ as $n \rightarrow \infty$ for any $j \in [k_*]$ and $i \in \mathcal{A}_j$. Thus, for any arbitrarily small $\zeta_j > 0$, we have that $\|\beta_{1i}^n - \beta_{1j}^*\| \leq \zeta_j$ for sufficiently large n . Then, by employing arguments as in Lemma 3 that $\mathcal{X}_\ell^* = \mathcal{X}_{\bar{\ell}}^n$ for sufficiently large n , where $\bar{\ell} \in [\bar{q}]$ such that $\{\bar{\ell}_1, \dots, \bar{\ell}_{\bar{K}}\} = \mathcal{A}_1 \cup \dots \cup \mathcal{A}_K$.

Let us consider an arbitrary subset \mathcal{X}_ℓ^* of the input space \mathcal{X} where ℓ is an arbitrary index in $[q]$.

Case 1: $1 \notin \{\ell_1, \dots, \ell_K\}$. In this case, we can validate that $f_{G_n}(x) - f_{G_*}(x) = 0$ for almost every $x \in \mathcal{X}_\ell^*$.

Case 2: $1 \in \{\ell_1, \dots, \ell_K\}$. WLOG, we assume that $\{\ell_1, \dots, \ell_K\} = \{1, \dots, K\}$. Then, for almost every $x \in \mathcal{X}_\ell^*$, we can represent $f_{G_n}(x)$ and $f_{G_*}(x)$ as follows:

$$\begin{aligned} f_{G_*}(x) &= \sum_{j=1}^K \text{Softmax}\left(\frac{(\beta_{1j}^*)^\top x}{\|\beta_{1j}^*\| \cdot \|x\|}\right) \cdot h(x, \eta_j^*), \\ \bar{f}_{G_n}(x) &= \sum_{j=1}^K \sum_{i \in \mathcal{A}_j} \text{Softmax}\left(\frac{(\beta_{1i}^n)^\top x}{\|\beta_{1i}^n\| \cdot \|x\|}\right) \cdot h(x, \eta_i^n). \end{aligned}$$

Now, we consider the quantity

$$Q_n(x) := \left[\sum_{j=1}^{k_*} \exp\left(\frac{(\beta_{1j}^*)^\top x}{\|\beta_{1j}^*\| \cdot \|x\|} + \beta_{0j}^* \right) \right] \cdot [f_{G_n}(x) - f_{G_*}(x)], \quad (29)$$

which can be decomposed as follows:

$$\begin{aligned}
Q_n(x) &= \sum_{j=1}^{k_*} \sum_{i \in \mathcal{A}_j} \exp(\beta_{0i}^n) \left[\exp\left(\frac{(\beta_{1i}^n)^\top x}{\|\beta_{1i}^n\| \cdot \|x\|}\right) h(x, \eta_i^n) - \exp\left(\frac{(\beta_{1j}^*)^\top x}{\|\beta_{1j}^*\| \cdot \|x\|}\right) h(x, \eta_j^*) \right] \\
&\quad - \sum_{j=1}^{k_*} \sum_{i \in \mathcal{A}_j} \exp(\beta_{0i}^n) \left[\exp\left(\frac{(\beta_{1i}^n)^\top x}{\|\beta_{1i}^n\| \cdot \|x\|}\right) f_{G_n}(x) - \exp\left(\frac{(\beta_{1j}^*)^\top x}{\|\beta_{1j}^*\| \cdot \|x\|}\right) f_{G_n}(x) \right] \\
&\quad + \sum_{j=1}^{k_*} \left(\sum_{i \in \mathcal{A}_j} \exp(\beta_{0i}^n) - \exp(\beta_{0j}^*) \right) \exp\left(\frac{(\beta_{1j}^*)^\top x}{\|\beta_{1j}^*\| \cdot \|x\|}\right) \left[h(x, \eta_j^*) - f_{G_n}(x) \right] \\
&:= A_n(x) - B_n(x) + C_n(x).
\end{aligned}$$

From the choices of β_{1i}^n and η_i^n , we can rewrite $A_n(x)$ as

$$A_n(x) = \frac{1}{2} \exp(\beta_{01}^*) h(x, \eta_1^*) \sum_{i=1}^2 \left[\exp\left(\frac{(\beta_{1i}^n)^\top x}{\|\beta_{1i}^n\| \cdot \|x\|}\right) - \exp\left(\frac{(\beta_{11}^*)^\top x}{\|\beta_{11}^*\| \cdot \|x\|}\right) \right].$$

Let us denote $F(x, \beta_1) := \exp\left(\frac{\beta_1^\top x}{\|\beta_1\| \cdot \|x\|}\right)$. By applying the Taylor expansion of order r , we have

$$\begin{aligned}
A_n(x) &= \frac{1}{2} \exp(\beta_{01}^*) h(x, \eta_1^*) \sum_{i=1}^2 \sum_{|\alpha|=1}^r \frac{1}{\alpha!} \cdot (\beta_{1i}^n - \beta_{11}^*)^\alpha \cdot \frac{\partial^{|\alpha|} F}{\partial \beta_1^\alpha}(x, \beta_{1i}^*) + R(x) \\
&= \frac{1}{2} \exp(\beta_{01}^*) h(x, \eta_1^*) \sum_{i=1}^2 \sum_{|\alpha|=1}^r \frac{1}{\alpha!} \left(1 + \frac{(-1)^i}{n}\right)^{|\alpha|} (\beta_{11}^*)^\alpha \cdot \frac{\partial^{|\alpha|} F}{\partial \beta_1^\alpha}(x, \beta_{11}^*) + R(x),
\end{aligned}$$

where $R(x)$ is a Taylor remainder such that $R(x)/\mathcal{L}_{1,r}(G_n, G_*) \rightarrow 0$ as $n \rightarrow \infty$. It follows from Lemma 4 (see Appendix D) that

$$\begin{aligned}
&\sum_{|\alpha|=t} \frac{1}{\alpha!} \left(1 + \frac{(-1)^i}{n}\right)^{|\alpha|} (\beta_{11}^*)^\alpha \cdot \frac{\partial^{|\alpha|} F}{\partial \beta_1^\alpha}(x, \beta_{11}^*) \\
&= \left(1 + \frac{(-1)^i}{n}\right)^t \sum_{|\alpha|=t} \frac{1}{\alpha!} (\beta_{11}^*)^\alpha \cdot \frac{\partial^{|\alpha|} F}{\partial \beta_1^\alpha}(x, \beta_{11}^*) = 0.
\end{aligned}$$

for any $1 \leq t \leq r$. Thus, we get that $A_n(x) = R(x)$, which implies that $A_n(x)/\mathcal{L}_{1,r}(G_n, G_*) \rightarrow 0$ as $n \rightarrow \infty$. By arguing similarly, we also obtain that $B_n(x)/\mathcal{L}_{1,r}(G_n, G_*) \rightarrow 0$ as $n \rightarrow \infty$. Furthermore, we have

$$\begin{aligned}
C_n(x) &= \left(\sum_{i=1}^2 \exp(\beta_{0i}^n) - \exp(\beta_{01}^*) \right) \exp\left(\frac{(\beta_{11}^*)^\top x}{\|\beta_{11}^*\| \cdot \|x\|}\right) \left[h(x, \eta_1^*) - \bar{f}_{G_n}(x) \right] \\
&\quad + \sum_{j=2}^{k_*} \left(\exp(\beta_{0(j+1)}^n) - \exp(\beta_{0j}^*) \right) \exp\left(\frac{(\beta_{1j}^*)^\top x}{\|\beta_{1j}^*\| \cdot \|x\|}\right) \left[h(x, \eta_j^*) - \bar{f}_{G_n}(x) \right] \\
&= 0.
\end{aligned}$$

Putting the previous results together, we achieve that

$$Q_n(x)/\mathcal{L}_{1,r}(G_n, G_*) \rightarrow 0.$$

Since the input space \mathcal{X} and the parameter space Θ are both bounded, the term $\sum_{j=1}^{k_*} \exp\left(\frac{(\beta_{1j}^*)^\top x}{\|\beta_{1j}^*\| \cdot \|x\|} + \beta_{0j}^*\right)$ is also bounded. This result together with the formulation of $Q_n(x)$ in equation (29) suggests that $[f_{G_n}(x) - f_{G_*}(x)]/\mathcal{L}_{1,r}(G_n, G_*) \rightarrow 0$ for almost every $x \in \mathcal{X}_\ell^*$. Note that the subset \mathcal{X}_ℓ^* is chosen arbitrarily, therefore, the previous result holds for almost every $x \in \mathcal{X}$. As a consequence, we get that

$$\frac{\|f_{G_n} - f_{G_*}\|_{\mathbb{L}_2(\mu)}}{\mathcal{L}_{1,r}(G_n, G_*)} \rightarrow 0,$$

as $n \rightarrow \infty$, and hence, achieve the result in equation (28).

C Proof for Results in Section 3

In this appendix, we provide proofs for the theoretical results regarding the perturbed cosine router, namely Theorem 4, Theorem 6, and Theorem 5, in that order.

C.1 Proof of Theorem 4

Similar to the proof of Theorem 1 in Appendix B.1, it is sufficient to demonstrate that when the expert functions are Lipschitz continuous, the following bound holds for any $0 < \varepsilon \leq 1/2$:

$$H_B(\varepsilon, \mathcal{R}_k(\Theta), \|\cdot\|_{\mathbb{L}_2(\mu)}) \lesssim \log(1/\varepsilon). \quad (30)$$

Indeed, for any function $g_G \in \mathcal{R}_k(\Theta)$, since the expert functions are bounded, we obtain that $g_G(x) \leq M$ for all x where M is bounded constant of the expert functions. Let $\tau \leq \varepsilon$ and $\{\pi_1, \dots, \pi_N\}$ be the τ -cover under the L^2 norm of the set $\mathcal{R}_k(\Theta)$ where $N := N(\tau, \mathcal{R}_k(\Theta), \|\cdot\|_{\mathbb{L}_2(\mu)})$ is the η -covering number of the metric space $(\mathcal{R}_k(\Theta), \|\cdot\|_{\mathbb{L}_2(\mu)})$. Then, we construct the brackets of the form $[L_i(x), U_i(x)]$ for all $i \in [N]$ as follows:

$$\begin{aligned} L_i(x) &:= \max\{\pi_i(x) - \tau, 0\}, \\ U_i(x) &:= \max\{\pi_i(x) + \tau, M\}. \end{aligned}$$

From the above construction, we can validate that $\mathcal{R}_k(\Theta) \subset \cup_{i=1}^N [L_i(x), U_i(x)]$ and $U_i(x) - L_i(x) \leq 2 \min\{2\tau, M\}$. Therefore, it follows that

$$\|U_i - L_i\|_{\mathbb{L}_2(\mu)}^2 = \int (U_i(x) - L_i(x))^2 d\mu(x) \leq \int 16\tau^2 d\mu(x) = 16\tau^2,$$

which implies that $\|U_i - L_i\|_{\mathbb{L}_2(\mu)} \leq 4\tau$. By definition of the bracketing entropy, we deduce that

$$H_B(4\tau, \mathcal{R}_k(\Theta), \|\cdot\|_{\mathbb{L}_2(\mu)}) \leq \log N = \log N(\tau, \mathcal{R}_k(\Theta), \|\cdot\|_{\mathbb{L}_2(\mu)}). \quad (31)$$

Therefore, we need to provide an upper bound for the covering number N . In particular, we denote $\Delta := \{(\beta_0, \beta_1) \in \mathbb{R} \times \mathbb{R}^{d_1} : (\beta_0, \beta_1, \eta) \in \Theta\}$ and $\Omega := \{\eta \in \mathbb{R}^{d_2} : (\beta_0, \beta_1, \eta) \in \Theta\}$. Since Θ is a

compact set, Δ and Ω are also compact. Therefore, we can find τ -covers Δ_τ and Ω_τ for Δ and Ω , respectively. We can check that

$$|\Delta_\tau| \leq \mathcal{O}_P(\tau^{-(d_1+1)k}), \quad |\Omega_\tau| \leq \mathcal{O}_P(\tau^{-d_2k}).$$

For each mixing measure $G = \sum_{i=1}^k \exp(\beta_{0i}) \delta_{(\beta_{1i}, \eta_i)} \in \mathcal{G}_k(\Theta)$, we consider other two mixing measures:

$$G' := \sum_{i=1}^k \exp(\beta_{0i}) \delta_{(\beta_{1i}, \bar{\eta}_i)}, \quad \bar{G} := \sum_{i=1}^k \exp(\bar{\beta}_{0i}) \delta_{(\bar{\beta}_{1i}, \bar{\eta}_i)}.$$

Here, $\bar{\eta}_i \in \Omega_\tau$ such that $\bar{\eta}_i$ is the closest to η_i in that set, while $(\bar{\beta}_{0i}, \bar{\beta}_{1i}) \in \Delta_\tau$ is the closest to (β_{0i}, β_{1i}) in that set. From the above formulations, we get that

$$\begin{aligned} & \|g_G - g_{G'}\|_{\mathbb{L}_2(\mu)}^2 \\ &= \int \left[\sum_{i=1}^k \text{Softmax} \left(\text{Top}_K \left(\frac{(\beta_{1i})^\top x}{(\|\beta_{1i}\| + \tau_1) \cdot (\|x\| + \tau_2)}, \beta_{0i} \right) \right) \cdot [h(x, \eta_i) - h(x, \bar{\eta}_i)] \right]^2 d\mu(x) \\ &\leq \int \left[\sum_{i=1}^k (h(x, \eta_i) - h(x, \bar{\eta}_i)) \right]^2 d\mu(x) \\ &\leq k \int \sum_{i=1}^k [h(x, \eta_i) - h(x, \bar{\eta}_i)]^2 d\mu(x) \\ &\leq k \int \sum_{i=1}^k [L_1 \cdot \|\eta_i - \bar{\eta}_i\|]^2 d\mu(x) \\ &\leq k^2 (L_1 \tau)^2, \end{aligned}$$

which indicates that $\|g_G - g_{G'}\|_{\mathbb{L}_2(\mu)} \lesssim \tau$. Here, the second inequality is according to the Cauchy-Schwarz inequality, the third inequality occurs as the softmax weight is bounded by 1, and the fourth inequality follows from the fact that the expert $h(x, \cdot)$ is a Lipschitz function with Lipschitz constant L_1 .

Next, we will also demonstrate that $\|g_{G'} - g_{\bar{G}}\|_1 \lesssim \eta$. For that purpose, let us consider $q := \binom{k}{K}$ K -element subsets of $\{1, \dots, k\}$, which are assumed to take the form $\{\ell_1, \ell_2, \dots, \ell_K\}$ for any $\ell \in [q]$. Additionally, we also denote $\{\ell_{K+1}, \dots, \ell_k\} := \{1, \dots, k\} \setminus \{\ell_1, \dots, \ell_K\}$ for any $\ell \in [q]$. Then, we define

$$\begin{aligned} \mathcal{X}_\ell &:= \left\{ x \in \mathcal{X} : \frac{(\beta_{1\ell_i})^\top x}{(\|\beta_{1\ell_i}\| + \tau_1) \cdot (\|x\| + \tau_2)} \geq \frac{(\beta_{1\ell_{i'}})^\top x}{(\|\beta_{1\ell_{i'}}\| + \tau_1) \cdot (\|x\| + \tau_2)}, \right. \\ &\quad \left. i \in \{\ell_1, \dots, \ell_K\}, i' \in \{\ell_{K+1}, \dots, \ell_{k_*}\} \right\}, \\ \tilde{\mathcal{X}}_\ell &:= \left\{ x \in \mathcal{X} : \frac{(\bar{\beta}_{1\ell_i})^\top x}{(\|\bar{\beta}_{1\ell_i}\| + \tau_1) \cdot (\|x\| + \tau_2)} \geq \frac{(\bar{\beta}_{1\ell_{i'}})^\top x}{(\|\bar{\beta}_{1\ell_{i'}}\| + \tau_1) \cdot (\|x\| + \tau_2)}, \right. \\ &\quad \left. i \in \{\ell_1, \dots, \ell_K\}, i' \in \{\ell_{K+1}, \dots, \ell_{k_*}\} \right\}. \end{aligned}$$

By using the same arguments as in the proof of Lemma 3 in Appendix D, we achieve that either $\mathcal{X}_\ell = \tilde{\mathcal{X}}_\ell$ or \mathcal{X}_ℓ has measure zero for any $\ell \in [q]$. As the Softmax function is differentiable, it is a Lipschitz function with some Lipschitz constant $L \geq 0$. Since \mathcal{X} is a bounded set, we may assume that $\|x\| \leq B$ for any $x \in \mathcal{X}$. Next, we denote

$$\begin{aligned}\pi_\ell(x) &:= \left(\frac{(\beta_{1\ell_i})^\top x}{(\|\beta_{1\ell_i}\| + \tau_1) \cdot (\|x\| + \tau_2)} + \beta_{0\ell_i} \right)_{i=1}^K, \\ \bar{\pi}_\ell(x) &:= \left(\frac{(\bar{\beta}_{1\ell_i})^\top x}{(\|\bar{\beta}_{1\ell_i}\| + \tau_1) \cdot (\|x\| + \tau_2)} + \bar{\beta}_{0\ell_i} \right)_{i=1}^K,\end{aligned}$$

for any K -element subset $\{\ell_1, \dots, \ell_K\}$ of $\{1, \dots, k_*\}$. Then, we get

$$\begin{aligned}\|\text{Softmax}(\pi_\ell(x)) - \text{Softmax}(\bar{\pi}_\ell(x))\| &\leq L \cdot \|\pi_\ell(x) - \bar{\pi}_\ell(x)\| \\ &\leq L \cdot \sum_{i=1}^K \left(\|\beta_{1\ell_i} - \bar{\beta}_{1\ell_i}\| \cdot \|x\| + |\beta_{0\ell_i} - \bar{\beta}_{0\ell_i}| \right) \\ &\leq L \cdot \sum_{i=1}^K (\eta B + \eta) \\ &\lesssim \eta.\end{aligned}$$

Back to the proof for $\|g_{G'} - g_{\bar{G}}\|_1 \lesssim \eta$, it follows from the above results that

$$\begin{aligned}\|g_{G'} - g_{\bar{G}}\|_1 &= \int_{\mathcal{X}} |g_{G'}(x) - g_{\bar{G}}(x)| \, d\mu(x) \leq \sum_{\ell=1}^q \int_{\mathcal{X}_\ell} |g_{G'}(x) - g_{\bar{G}}(x)| \, d\mu(x) \\ &\leq \sum_{\ell=1}^q \int_{\mathcal{X}_\ell} \sum_{i=1}^K \left| \text{Softmax}(\pi_\ell(x)_i) - \text{Softmax}(\bar{\pi}_\ell(x)_i) \right| \cdot |h(x, \bar{\eta}_{\ell_i})| \, d\mu(x) \\ &\lesssim \eta,\end{aligned}\tag{32}$$

According to the triangle inequality, we have

$$\|g_G - g_{\bar{G}}\|_{\mathbb{L}_2(\mu)} \leq \|g_G - g_{G'}\|_{\mathbb{L}_2(\mu)} + \|g_{G'} - g_{\bar{G}}\|_{\mathbb{L}_2(\mu)} \lesssim \tau.$$

By definition of the covering number, we deduce that

$$N(\tau, \mathcal{R}_k(\Theta), \|\cdot\|_{\mathbb{L}_2(\mu)}) \leq |\Delta_\tau| \times |\Omega_\tau| \leq \mathcal{O}_P(n^{-(d_1+1)k}) \times \mathcal{O}(n^{-d_2k}) \leq \mathcal{O}(n^{-(d_1+1+d_2)k}).\tag{33}$$

Combine equations (31) and (33), we achieve that

$$H_B(4\tau, \mathcal{R}_k(\Theta), \|\cdot\|_{\mathbb{L}_2(\mu)}) \lesssim \log(1/\tau).$$

Let $\tau = \varepsilon/4$, then we obtain that

$$H_B(\varepsilon, \mathcal{R}_k(\Theta), \|\cdot\|_{\mathbb{L}_2(\mu)}) \lesssim \log(1/\varepsilon).$$

Hence, the proof is completed.

C.2 Proof of Theorem 6

In this proof, we aim to establish the following inequality:

$$\inf_{G \in \mathcal{E}_{k^*}(\Theta)} \|g_G - g_{G^*}\|_{\mathbb{L}_2(\mu)} / \mathcal{L}_2(G, G^*) > 0. \quad (34)$$

For that purpose, we divide the proof of the above inequality into local and global parts in the sequel.

Local part: In this part, we demonstrate that

$$\lim_{\varepsilon \rightarrow 0} \inf_{G \in \mathcal{G}_k(\Theta) : \mathcal{L}_2(G, G^*) \leq \varepsilon} \|g_G - g_{G^*}\|_{\mathbb{L}_2(\mu)} / \mathcal{L}_2(G, G^*) > 0. \quad (35)$$

Assume by contrary that the above inequality does not hold true, then there exists a sequence of mixing measures $G_n = \sum_{i=1}^{k^*} \exp(\beta_{0i}^n) \delta_{(\beta_{1i}^n, \eta_i^n)}$ in $\mathcal{G}_k(\Theta)$ such that $\mathcal{L}_{2n} := \mathcal{L}_2(G_n, G^*) \rightarrow 0$ and

$$\|g_{G_n} - g_{G^*}\|_{\mathbb{L}_2(\mu)} / \mathcal{L}_{2n} \rightarrow 0, \quad (36)$$

as $n \rightarrow \infty$. Let us denote by $\mathcal{A}_j^n := \mathcal{A}_j(G_n)$ a Voronoi cell of G_n generated by the j -th components of G_n . Since our arguments are asymptotic, we may assume that those Voronoi cells do not depend on the sample size, i.e. $\mathcal{A}_j = \mathcal{A}_j^n$. Moreover, recall that under the exact-specified setting, each Voronoi cell has only one element. Therefore, we may assume WLOG that $\mathcal{A}_j = \{j\}$, and

$$\mathcal{L}_{2n} := \sum_{i=1}^K \left| \exp(\beta_{0i}^n) - \exp(\beta_{0i}^*) \right| + \sum_{i=1}^K \exp(\beta_{0i}^n) \left[\|\Delta\beta_{1i}^n\| + \|\Delta\eta_i^n\| \right],$$

where we denote $\Delta\beta_{1i}^n := \beta_{1i}^n - \beta_{1i}^*$ and $\Delta\eta_i^n := \eta_i^n - \eta_i^*$.

Let $\ell \in [q]$ such that $\{\ell_1, \dots, \ell_K\} = \{1, \dots, K\}$. Note that Lemma 3 indicates that $\mathcal{X}_\ell^n = \mathcal{X}_\ell^*$ for sufficiently large n . Then, for almost every $x \in \mathcal{X}_\ell^*$, we can rewrite $g_{G_n}(x)$ and $g_{G^*}(x)$ as

$$\begin{aligned} g_{G_n}(x) &= \sum_{i=1}^K \text{Softmax} \left(\frac{(\beta_{1i}^n)^\top x}{(\|\beta_{1i}^n\| + \tau_1) \cdot (\|x\| + \tau_2)} \right) \cdot h(x, \eta_i^n), \\ g_{G^*}(x) &= \sum_{i=1}^K \text{Softmax} \left(\frac{(\beta_{1i}^*)^\top x}{(\|\beta_{1i}^*\| + \tau_1) \cdot (\|x\| + \tau_2)} \right) \cdot h(x, \eta_i^*). \end{aligned}$$

Since $\mathcal{L}_{2n} \rightarrow 0$, we get that $(\beta_{1i}^n, \eta_i^n) \rightarrow (\beta_{1i}^*, \eta_i^*)$ and $\exp(\beta_{0i}^n) \rightarrow \exp(\beta_{0i}^*)$ as $n \rightarrow \infty$ for any $i \in [K]$. Now, we divide the proof of local part into three steps as follows:

Step 1: Taylor expansion. In this step, we decompose the term

$$Q_n(x) := \left[\sum_{j=1}^K \exp \left(\frac{(\beta_{1j}^*)^\top x}{(\|\beta_{1j}^*\| + \tau_1) \cdot (\|x\| + \tau_2)} + \beta_{0j}^* \right) \right] \cdot [g_{G_n}(x) - g_{G^*}(x)]$$

into a combination of linearly independent elements using Taylor expansion. In particular, let us denote $F(x, \beta_1) := \exp\left(\frac{\beta_1^\top x}{(\|\beta_1\| + \tau_1) \cdot (\|x\| + \tau_2)}\right)$, then we have

$$\begin{aligned}
Q_n(x) &= \sum_{i=1}^K \exp(\beta_{0i}^n) \left[F(x, \beta_{1i}^n) h(x, \eta_i^n) - F(x, \beta_{1i}^*) h(x, \eta_i^*) \right] \\
&\quad - \sum_{i=1}^K \exp(\beta_{0i}^n) \left[F(x, \beta_{1i}^n) - F(x, \beta_{1i}^*) \right] g_{G_n}(x) \\
&\quad + \sum_{i=1}^K \left(\exp(\beta_{0i}^n) - \exp(\beta_{0i}^*) \right) \left[F(x, \beta_{1i}^*) h(x, \eta_i^*) - F(x, \beta_{1i}^*) g_{G_n}(x) \right] \\
&:= A_n(x) - B_n(x) + C_n(x).
\end{aligned} \tag{37}$$

By means of the first-order Taylor expansion, we have

$$A_n(x) = \sum_{i=1}^K \sum_{|\alpha|=1} \frac{\exp(\beta_{0i}^n)}{\alpha!} (\Delta\beta_{1i}^n)^{\alpha_1} (\Delta\eta_i^n)^{\alpha_2} \cdot \frac{\partial^{|\alpha_1|} F}{\partial \beta_1^{\alpha_1}}(x, \beta_{1i}^*) \frac{\partial^{|\alpha_2|} h}{\partial \eta^{\alpha_2}}(x, \eta_i^*) + R_1(x), \tag{38}$$

where $R_1(x)$ is a Taylor remainder such that $R_1(x)/\mathcal{L}_{2n} \rightarrow 0$ as $n \rightarrow \infty$. Similarly, we also get that

$$B_n(x) = \sum_{i=1}^K \sum_{|\gamma|=1} \frac{\exp(\beta_{0i}^n)}{\gamma!} (\Delta\beta_{1i}^n)^\gamma \cdot \frac{\partial^{|\gamma|} F}{\partial \beta_1^\gamma}(x, \beta_{1i}^*) g_{G_n}(x) + R_2(x),$$

where $R_2(x)$ is a Taylor remainder such that $R_2(x)/\mathcal{L}_{2n} \rightarrow 0$ as $n \rightarrow \infty$. As a result, we deduce that

$$\begin{aligned}
Q_n(x) &= \sum_{i=1}^K \sum_{|\alpha|=0}^1 T_{i, \alpha_1, \alpha_2}^n \cdot \frac{\partial^{|\alpha_1|} F}{\partial \beta_1^{\alpha_1}}(x, \beta_{1i}^*) \frac{\partial^{|\alpha_2|} h}{\partial \eta^{\alpha_2}}(x, \eta_i^*) + R_1(x) \\
&\quad - \sum_{i=1}^K \sum_{|\gamma|=0}^1 S_{i, \gamma}^n \cdot \frac{\partial^{|\gamma|} F}{\partial \beta_1^\gamma}(x, \beta_{1i}^*) g_{G_n}(x) - R_2(x),
\end{aligned} \tag{39}$$

where we define

$$\begin{aligned}
T_{i, \alpha_1, \alpha_2}^n &:= \frac{\exp(\beta_{0i}^n)}{\alpha!} (\Delta\beta_{1i}^n)^{\alpha_1} (\Delta\eta_i^n)^{\alpha_2}, \\
S_{i, \gamma}^n &:= \frac{\exp(\beta_{0i}^n)}{\gamma!} (\Delta\beta_{1ij}^n)^\gamma,
\end{aligned}$$

for any $(\alpha_1, \alpha_2) \neq (\mathbf{0}_d, 0)$ and $\gamma \neq \mathbf{0}_d$. Otherwise, $T_{i, \mathbf{0}_d, 0}^n = S_{i, \mathbf{0}_d}^n := \exp(\beta_{0i}^n) - \exp(\beta_{0i}^*)$.

Step 2: Non-vanishing coefficients. In this step, we show that not all the ratios $T_{i, \alpha_1, \alpha_2}^n/\mathcal{L}_{2n}$, and $S_{i, \gamma}^n/\mathcal{L}_{2n}$ converge to zero. Indeed, assume by contrary that all of them converge to zero, i.e.

$$\frac{T_{i, \alpha_1, \alpha_2}^n}{\mathcal{L}_{2n}} \rightarrow 0, \quad \frac{S_{i, \gamma}^n}{\mathcal{L}_{2n}} \rightarrow 0$$

as $n \rightarrow \infty$. Then, it follows that

- $\frac{1}{\mathcal{L}_{2n}} \cdot \sum_{i=1}^K \left| \exp(\beta_{0i}^n) - \exp(\beta_{0i}^*) \right| = \frac{1}{\mathcal{L}_{2n}} \cdot \sum_{i=1}^K |S_{i, \mathbf{0}_d}^n| \rightarrow 0;$
- $\frac{1}{\mathcal{L}_{2n}} \cdot \sum_{i=1}^K \exp(\beta_{0i}^n) \|\Delta \beta_{1i}^n\|_1 = \frac{1}{\mathcal{L}_{2n}} \sum_{i=1}^K \sum_{u=1}^{d_1} |T_{i, e_{d_1, u}, \mathbf{0}_d}^n| \rightarrow 0;$
- $\frac{1}{\mathcal{L}_{2n}} \cdot \sum_{i=1}^K \exp(\beta_{0i}^n) \|\Delta \eta_i^n\|_1 = \frac{1}{\mathcal{L}_{2n}} \sum_{i=1}^K \sum_{v=1}^{d_2} |T_{i, \mathbf{0}_d, e_{d_2, v}}^n| \rightarrow 0.$

Due to the topological equivalence of the norm-1 and norm-2, we deduce that

$$\frac{1}{\mathcal{L}_{2n}} \cdot \sum_{i=1}^K \exp(\beta_{0i}^n) \|\Delta \beta_{1i}^n\| \rightarrow 0, \quad \frac{1}{\mathcal{L}_{2n}} \cdot \sum_{i=1}^K \exp(\beta_{0i}^n) \|\Delta \eta_i^n\| \rightarrow 0.$$

As a result, we obtain that

$$1 = \frac{\mathcal{L}_{2n}}{\mathcal{L}_{2n}} = \frac{1}{\mathcal{L}_{2n}} \left\{ \sum_{i=1}^K \left| \exp(\beta_{0i}^n) - \exp(\beta_{0i}^*) \right| + \sum_{i=1}^K \exp(\beta_{0i}^n) \left[\|\Delta \beta_{1i}^n\| + \|\Delta \eta_i^n\| \right] \right\} \rightarrow 0,$$

which is a contradiction. Thus, at least one among the ratios $T_{i, \alpha_1, \alpha_2}^n / \mathcal{L}_{2n}$, and $S_{i, \gamma}^n / \mathcal{L}_{2n}$ must not go to zero as $n \rightarrow \infty$.

Step 3: Application of Fatou's lemma. In this step, we demonstrate a result opposed to that in Step 2, i.e. the ratios $T_{i, \alpha_1, \alpha_2}^n / \mathcal{L}_{2n}$, and $S_{i, \gamma}^n / \mathcal{L}_{2n}$ all converge to zero.

In particular, let us denote by m_n the maximum of the absolute values of $T_{i, \alpha_1, \alpha_2}^n / \mathcal{L}_{2n}$, and $S_{i, \gamma}^n / \mathcal{L}_{2n}$. Since at least one among those ratios must not approach zero as $n \rightarrow \infty$, we get that $1/m_n \not\rightarrow \infty$ as $n \rightarrow \infty$.

Recall from the hypothesis in equation (36) that $\|g_{G_n} - g_{G_*}\|_{\mathbb{L}_2(\mu)} / \mathcal{L}_{2n} \rightarrow 0$ as $n \rightarrow \infty$, which indicates that $\|g_{G_n} - g_{G_*}\|_{\mathbb{L}_1(\mu)} / \mathcal{L}_{2n} \rightarrow 0$ due to the equivalence between $\mathbb{L}_1(\mu)$ -norm and $L^2(\mu)$ -norm. By means of the Fatou's lemma, we have

$$0 = \lim_{n \rightarrow \infty} \frac{\|g_{G_n} - g_{G_*}\|_{\mathbb{L}_1(\mu)}}{m_n \mathcal{L}_{2n}} \geq \int \liminf_{n \rightarrow \infty} \frac{|g_{G_n}(x) - g_{G_*}(x)|}{m_n \mathcal{L}_{2n}} d\mu(x) \geq 0.$$

This result implies that $[g_{G_n}(x) - g_{G_*}(x)] / [m_n \mathcal{L}_{2n}] \rightarrow 0$ for almost every x .

Let us denote

$$\begin{aligned} T_{i, \alpha_1, \alpha_2}^n / m_n \mathcal{L}_{2n} &\rightarrow t_{i, \alpha_1, \alpha_2}, \\ S_{i, \gamma}^n / m_n \mathcal{L}_{2n} &\rightarrow s_{i, \gamma} \end{aligned}$$

with a note that at least one among the limits $t_{i, \alpha_1, \alpha_2}$, $s_{i, \gamma}$ is non-zero. Then, from the decomposition in equation (39), we deduce that

$$\sum_{i=1}^K \sum_{|\alpha|=0}^1 t_{i, \alpha_1, \alpha_2} \cdot \frac{\partial^{|\alpha_1|} F}{\partial \beta_1^{\alpha_1}}(x, \beta_{1i}^*) \frac{\partial^{|\alpha_2|} h}{\partial \eta^{\alpha_2}}(x, \eta_i^*) - \sum_{i=1}^K \sum_{|\gamma|=0}^1 s_{i, \gamma} \cdot \frac{\partial^{|\gamma|} F}{\partial \beta_1^\gamma}(x, \beta_{1i}^*) g_{G_n}(x) = 0,$$

for almost every x . Note that the expert function $h(\cdot, \eta)$ satisfies the condition in Definition 2, then the above equation implies that $t_{i, \alpha_1, \alpha_2} = s_{i, \gamma} = 0$, for any $i \in [K]$, $\alpha_1 \in \mathbb{N}^{d_1}$, $\alpha_2 \in \mathbb{N}^{d_2}$ and $\gamma \in \mathbb{N}^{d_1}$ such that $0 \leq |\alpha_1| + |\alpha_2|, |\gamma| \leq 2$. This contradicts the fact that at least one among the limits $t_{i, \alpha_1, \alpha_2}, s_{i, \gamma}$ is different from zero.

Hence, we obtain the local inequality in equation (35). Thus, we can find an $\varepsilon' > 0$ such that

$$\inf_{G \in \mathcal{E}_{k_*}(\Theta): \mathcal{L}_2(G, G_*) \leq \varepsilon'} \|g_G - g_{G_*}\|_{\mathbb{L}_2(\mu)} / \mathcal{L}_2(G, G_*) > 0.$$

Global part: Given the above result, it suffices to demonstrate that

$$\inf_{G \in \mathcal{E}_{k_*}(\Theta): \mathcal{L}_2(G, G_*) > \varepsilon'} \|g_G - g_{G_*}\|_{\mathbb{L}_2(\mu)} / \mathcal{L}_2(G, G_*) > 0. \quad (40)$$

Assume by contrary that the inequality (40) does not hold true, then we can find a sequence of mixing measures $G'_n \in \mathcal{E}_{k_*}(\Theta)$ such that $\mathcal{L}_2(G'_n, G_*) > \varepsilon'$ and

$$\lim_{n \rightarrow \infty} \frac{\|g_{G'_n} - g_{G_*}\|_{\mathbb{L}_2(\mu)}}{\mathcal{L}_2(G'_n, G_*)} = 0,$$

which indicates that $\|g_{G'_n} - g_{G_*}\|_{\mathbb{L}_2(\mu)} \rightarrow 0$ as $n \rightarrow \infty$. Recall that Θ is a compact set, therefore, we can replace the sequence G'_n by one of its subsequences that converges to a mixing measure $G' \in \mathcal{E}_{k_*}(\Omega)$. Since $\mathcal{L}_2(G'_n, G_*) > \varepsilon'$, we deduce that $\mathcal{L}_2(G', G_*) > \varepsilon'$.

Next, by invoking the Fatou's lemma, we have that

$$0 = \lim_{n \rightarrow \infty} \|g_{G'_n} - g_{G_*}\|_{\mathbb{L}_2(\mu)}^2 \geq \int \liminf_{n \rightarrow \infty} |g_{G'_n}(x) - g_{G_*}(x)|^2 d\mu(x).$$

Thus, we get that $g_{G'}(x) = g_{G_*}(x)$ for almost every x . From Proposition 2, we deduce that $G' \equiv G_*$. Consequently, it follows that $\mathcal{L}_2(G', G_*) = 0$, contradicting the fact that $\mathcal{L}_2(G', G_*) > \varepsilon' > 0$.

Hence, the proof is completed.

C.3 Proof of Theorem 5

In this proof, it is sufficient to demonstrate the following inequality:

$$\inf_{G \in \mathcal{E}_{k_*}(\Theta)} \|\bar{g}_G - g_{G_*}\|_{\mathbb{L}_2(\mu)} / \mathcal{L}_3(G, G_*) > 0. \quad (41)$$

This can be done by deriving its local part and the global part as in Appendix C.2. Since the global part can be argued in a similar fashion, our main goal is to prove the local part:

$$\lim_{\varepsilon \rightarrow 0} \inf_{G \in \mathcal{G}_k(\Theta): \mathcal{L}_3(G, G_*) \leq \varepsilon} \|\bar{g}_G - g_{G_*}\|_{\mathbb{L}_2(\mu)} / \mathcal{L}_3(G, G_*) > 0. \quad (42)$$

Assume by contrary that the above inequality does not hold true, then there exists a sequence of mixing measures $G_n = \sum_{i=1}^{k_*} \exp(\beta_{0i}^n) \delta_{(\beta_{1i}^n, \eta_i^n)}$ in $\mathcal{G}_k(\Theta)$ such that $\mathcal{L}_{3n} := \mathcal{L}_3(G_n, G_*) \rightarrow 0$ and

$$\|\bar{g}_{G_n} - g_{G_*}\|_{\mathbb{L}_2(\mu)} / \mathcal{L}_{3n} \rightarrow 0, \quad (43)$$

as $n \rightarrow \infty$. Let us denote by $\mathcal{A}_j^n := \mathcal{A}_j(G_n)$ a Voronoi cell of G_n generated by the j -th components of G_* . Since our arguments are asymptotic, we may assume that those Voronoi cells do not depend on the sample size, i.e., $\mathcal{A}_j = \mathcal{A}_j^n$. Therefore, we may assume WLOG that

$$\begin{aligned} \mathcal{L}_{3n} := & \sum_{i=1}^K \left| \sum_{i \in \mathcal{A}_j} \exp(\beta_{0i}^n) - \exp(\beta_{0j}^*) \right| + \sum_{j \in [K]: |\mathcal{A}_j| > 1} \sum_{i \in \mathcal{A}_j} \exp(\beta_{0i}^n) \left[\|\Delta\beta_{1ij}^n\|^2 + \|\Delta\eta_{ij}^n\|^2 \right] \\ & + \sum_{j \in [K]: |\mathcal{A}_j| = 1} \sum_{i \in \mathcal{A}_j} \exp(\beta_{0i}^n) \left[\|\Delta\beta_{1ij}^n\| + \|\Delta\eta_{ij}^n\| \right], \end{aligned}$$

where we denote $\Delta\beta_{1ij}^n := \beta_{1i}^n - \beta_{1j}^*$ and $\Delta\eta_{ij}^n := \eta_i^n - \eta_j^*$.

Let us consider $\ell \in [q]$ such that $\{\ell_1, \dots, \ell_K\} = \{1, \dots, K\}$. If $\{\bar{\ell}_1, \dots, \bar{\ell}_K\} \neq \mathcal{A}_1 \cup \dots \cup \mathcal{A}_K$ for any $\bar{\ell} \in [\bar{q}]$, then $\|\bar{g}_{G_n} - g_{G_*}\|_{\mathbb{L}_2(\mu)} / \mathcal{L}_3(G_n, G_*) \not\rightarrow 0$ as $n \rightarrow \infty$. This contradicts the hypothesis that this term must approach zero. Therefore, we only need to consider the scenario when there exists $\bar{\ell} \in [\bar{q}]$ such that $\{\bar{\ell}_1, \dots, \bar{\ell}_K\} = \mathcal{A}_1 \cup \dots \cup \mathcal{A}_K$. Note that we have $\beta_{1i}^n \rightarrow \beta_{1j}^*$ as $n \rightarrow \infty$ for any $j \in [K]$ and $i \in \mathcal{A}_j$. Thus, for any arbitrarily small $\eta_j > 0$, we have that $\|\beta_{1i}^n - \beta_{1j}^*\| \leq \eta_j$ for sufficiently large n . Then, by employing arguments in Lemma 3, we have that $\mathcal{X}_\ell^* = \mathcal{X}_{\bar{\ell}}^n$ for sufficiently large n . Thus, for almost every $x \in \mathcal{X}_\ell^*$, we can rewrite $g_{G_n}(x)$ and $g_{G_*}(x)$ as

$$\begin{aligned} \bar{g}_{G_n}(x) &= \sum_{j=1}^K \sum_{i \in \mathcal{A}_j} \text{Softmax} \left(\frac{(\beta_{1i}^n)^\top x}{(\|\beta_{1i}^n\| + \tau_1) \cdot (\|x\| + \tau_2)} \right) \cdot h(x, \eta_i^n), \\ g_{G_*}(x) &= \sum_{j=1}^K \text{Softmax} \left(\frac{(\beta_{1j}^*)^\top x}{(\|\beta_{1j}^*\| + \tau_1) \cdot (\|x\| + \tau_2)} \right) \cdot h(x, \eta_j^*). \end{aligned}$$

Now, we divide the proof of local part into three steps as follows:

Step 1: Taylor expansion. In this step, we decompose the term

$$Q_n(x) := \left[\sum_{j=1}^K \exp \left(\frac{(\beta_{1j}^*)^\top x}{(\|\beta_{1j}^*\| + \tau_1) \cdot (\|x\| + \tau_2)} + \beta_{0j}^* \right) \right] \cdot [\bar{g}_{G_n}(x) - g_{G_*}(x)]$$

into a combination of linearly independent elements using Taylor expansion. In particular, let us denote $F(x, \beta_1) := \exp \left(\frac{\beta_1^\top x}{(\|\beta_1\| + \tau_1) \cdot (\|x\| + \tau_2)} \right)$, then we have

$$\begin{aligned} Q_n(x) &= \sum_{j=1}^K \sum_{i \in \mathcal{A}_j} \exp(\beta_{0i}^n) \left[F(x, \beta_{1i}^n) h(x, \eta_i^n) - F(x, \beta_{1j}^*) h(x, \eta_j^*) \right] \\ &\quad - \sum_{j=1}^K \sum_{i \in \mathcal{A}_j} \exp(\beta_{0i}^n) \left[F(x, \beta_{1i}^n) - F(x, \beta_{1j}^*) \right] g_{G_n}(x) \\ &\quad + \sum_{j=1}^K \left(\sum_{i \in \mathcal{A}_j} \exp(\beta_{0i}^n) - \exp(\beta_{0j}^*) \right) \left[F(x, \beta_{1j}^*) h(x, \eta_j^*) - F(x, \beta_{1j}^*) g_{G_n}(x) \right] \\ &:= A_n(x) - B_n(x) + C_n(x). \end{aligned} \tag{44}$$

Next, we continue to separate the term $A_n(x)$ into two parts as

$$\begin{aligned} A_n(x) &:= \sum_{j \in [K]: |\mathcal{A}_j|=1} \sum_{i \in \mathcal{A}_j} \exp(\beta_{0i}^n) \left[F(x, \beta_{1i}^n) h(x, \eta_i^n) - F(x, \beta_{1j}^*) h(x, \eta_j^*) \right] \\ &\quad + \sum_{j \in [K]: |\mathcal{A}_j|>1} \sum_{i \in \mathcal{A}_j} \exp(\beta_{0i}^n) \left[F(x, \beta_{1i}^n) h(x, \eta_i^n) - F(x, \beta_{1j}^*) h(x, \eta_j^*) \right] \\ &:= A_{n,1}(x) + A_{n,2}(x) \end{aligned}$$

Similar to equation (38), by applying the first-order and the second-order Taylor expansions to $A_{n,1}(x)$ and $A_{n,2}(x)$, respectively, we have

$$\begin{aligned} A_{n,1}(x) &= \sum_{j \in [K]: |\mathcal{A}_j|=1} \sum_{|\alpha|=1} \frac{\exp(\beta_{0i}^n)}{\alpha!} (\Delta \beta_{1ij}^n)^{\alpha_1} (\Delta \eta_{ij}^n)^{\alpha_2} \cdot \frac{\partial^{|\alpha_1|} F}{\partial \beta_1^{\alpha_1}}(x, \beta_{1j}^*) \frac{\partial^{|\alpha_2|} h}{\partial \eta^{\alpha_2}}(x, \eta_j^*) + R_1(x), \\ A_{n,2}(x) &= \sum_{j \in [K]: |\mathcal{A}_j|>1} \sum_{|\alpha|=1}^2 \frac{\exp(\beta_{0i}^n)}{\alpha!} (\Delta \beta_{1ij}^n)^{\alpha_1} (\Delta \eta_{ij}^n)^{\alpha_2} \cdot \frac{\partial^{|\alpha_1|} F}{\partial \beta_1^{\alpha_1}}(x, \beta_{1j}^*) \frac{\partial^{|\alpha_2|} h}{\partial \eta^{\alpha_2}}(x, \eta_j^*) + R_2(x), \end{aligned}$$

where $R_i(x)$ is a Taylor remainder such that $R_i(x)/\mathcal{L}_{3n} \rightarrow 0$ as $n \rightarrow \infty$, for $i \in \{1, 2\}$. Analogously, we also get that $B_n(x) = B_{n,1}(x) + B_{n,2}(x)$ where

$$\begin{aligned} B_{n,1}(x) &= \sum_{j \in [K]: |\mathcal{A}_j|=1} \sum_{|\gamma|=1} \frac{\exp(\beta_{0i}^n)}{\gamma!} (\Delta \beta_{1ij}^n)^\gamma \cdot \frac{\partial^{|\gamma|} F}{\partial \beta_1^\gamma}(x, \beta_{1j}^*) g_{G_n}(x) + R_3(x), \\ B_{n,2}(x) &= \sum_{j \in [K]: |\mathcal{A}_j|>1} \sum_{|\gamma|=1}^2 \frac{\exp(\beta_{0i}^n)}{\gamma!} (\Delta \beta_{1ij}^n)^\gamma \cdot \frac{\partial^{|\gamma|} F}{\partial \beta_1^\gamma}(x, \beta_{1j}^*) g_{G_n}(x) + R_4(x), \end{aligned}$$

in which $R_i(x)$ is a Taylor remainder such that $R_i(x)/\mathcal{L}_{3n} \rightarrow 0$ as $n \rightarrow \infty$, for $i \in \{3, 4\}$.

As a result, we deduce that

$$\begin{aligned} Q_n(x) &= \sum_{j=1}^K \sum_{|\alpha|=0}^2 T_{j,\alpha_1,\alpha_2}^n \cdot \frac{\partial^{|\alpha_1|} F}{\partial \beta_1^{\alpha_1}}(x, \beta_{1j}^*) \frac{\partial^{|\alpha_2|} h}{\partial \eta^{\alpha_2}}(x, \eta_j^*) + R_1(x) + R_2(x) \\ &\quad - \sum_{j=1}^K \sum_{|\gamma|=0}^2 S_{j,\gamma}^n \cdot \frac{\partial^{|\gamma|} F}{\partial \beta_1^\gamma}(x, \beta_{1j}^*) g_{G_n}(x) - R_3(x) - R_4(x), \end{aligned} \tag{45}$$

where we define

$$\begin{aligned} T_{j,\alpha_1,\alpha_2}^n &:= \sum_{i \in \mathcal{A}_j} \frac{\exp(\beta_{0i}^n)}{\alpha!} (\Delta \beta_{1ij}^n)^{\alpha_1} (\Delta \eta_{ij}^n)^{\alpha_2}, \\ S_{j,\gamma}^n &:= \sum_{i \in \mathcal{A}_j} \frac{\exp(\beta_{0i}^n)}{\gamma!} (\Delta \beta_{1ij}^n)^\gamma, \end{aligned}$$

for any $(\alpha_1, \alpha_2) \neq (\mathbf{0}_d, 0)$ and $\gamma \neq \mathbf{0}_d$. Otherwise, $T_{j,\mathbf{0}_d,0}^n = S_{j,\mathbf{0}_d}^n := \sum_{i \in \mathcal{A}_j} \exp(\beta_{0i}^n) - \exp(\beta_{0j}^*)$.

Step 2: Non-vanishing coefficients. In this step, we show that not all the ratios $T_{j,\alpha_1,\alpha_2}^n/\mathcal{L}_{3n}$, and $S_{j,\gamma}^n/\mathcal{L}_{3n}$ converge to zero. Indeed, assume by contrary that all of them converge to zero, i.e.

$$\frac{T_{j,\alpha_1,\alpha_2}^n}{\mathcal{L}_{3n}} \rightarrow 0, \quad \frac{S_{j,\gamma}^n}{\mathcal{L}_{3n}} \rightarrow 0$$

as $n \rightarrow \infty$. Then, it follows that

- $\frac{1}{\mathcal{L}_{3n}} \cdot \sum_{j=1}^K \left| \sum_{i \in \mathcal{A}_j} \exp(\beta_{0i}^n) - \exp(\beta_{0j}^*) \right| = \frac{1}{\mathcal{L}_{3n}} \cdot \sum_{j=1}^K |S_{j,\mathbf{0}_d}^n| \rightarrow 0;$
- $\frac{1}{\mathcal{L}_{3n}} \cdot \sum_{j=1}^K \sum_{i \in \mathcal{A}_j} \exp(\beta_{0i}^n) \|\Delta \beta_{1ij}^n\|_1 = \frac{1}{\mathcal{L}_{3n}} \sum_{j=1}^K \sum_{u=1}^{d_1} |T_{j,e_{d_1,u},\mathbf{0}_d}^n| \rightarrow 0;$
- $\frac{1}{\mathcal{L}_{3n}} \cdot \sum_{j=1}^K \sum_{i \in \mathcal{A}_j} \exp(\beta_{0i}^n) \|\Delta \eta_{ij}^n\|_1 = \frac{1}{\mathcal{L}_{3n}} \sum_{j=1}^K \sum_{v=1}^{d_2} |T_{j,\mathbf{0}_d,e_{d_2,v}}^n| \rightarrow 0;$
- $\frac{1}{\mathcal{L}_{3n}} \cdot \sum_{j=1}^K \sum_{i \in \mathcal{A}_j} \exp(\beta_{0i}^n) \|\Delta \beta_{1ij}^n\|^2 = \frac{1}{\mathcal{L}_{3n}} \sum_{j=1}^K \sum_{u=1}^{d_1} |T_{j,2e_{d_1,u},\mathbf{0}_d}^n| \rightarrow 0;$
- $\frac{1}{\mathcal{L}_{3n}} \cdot \sum_{j=1}^K \sum_{i \in \mathcal{A}_j} \exp(\beta_{0i}^n) \|\Delta \eta_{ij}^n\|^2 = \frac{1}{\mathcal{L}_{3n}} \sum_{j=1}^K \sum_{v=1}^{d_2} |T_{j,\mathbf{0}_d,2e_{d_2,v}}^n| \rightarrow 0.$

Due to the topological equivalence of the norm-1 and norm-2, we deduce that

$$\frac{1}{\mathcal{L}_{3n}} \cdot \sum_{i=1}^K \exp(\beta_{0i}^n) \|\Delta \beta_{1i}^n\| \rightarrow 0, \quad \frac{1}{\mathcal{L}_{3n}} \cdot \sum_{i=1}^K \exp(\beta_{0i}^n) \|\Delta \eta_i^n\| \rightarrow 0.$$

Thus, by taking the summation of the above limits, we obtain that $1 = \mathcal{L}_{3n}/\mathcal{L}_{3n} \rightarrow 0$ as $n \rightarrow \infty$, which is a contradiction. Consequently, at least one among the ratios $T_{j,\alpha_1,\alpha_2}^n/\mathcal{L}_{3n}$, and $S_{j,\gamma}^n/\mathcal{L}_{3n}$ must not go to zero as $n \rightarrow \infty$.

Step 3: Application of Fatou's lemma. In this step, we demonstrate a result opposed to that in Step 2, i.e. the ratios $T_{j,\alpha_1,\alpha_2}^n/\mathcal{L}_{3n}$, and $S_{j,\gamma}^n/\mathcal{L}_{3n}$ all converge to zero.

In particular, let us denote by m_n the maximum of the absolute values of $T_{j,\alpha_1,\alpha_2}^n/\mathcal{L}_{3n}$, and $S_{j,\gamma}^n/\mathcal{L}_{3n}$. Since at least one among those ratios must not approach zero as $n \rightarrow \infty$, we get that $1/m_n \not\rightarrow \infty$ as $n \rightarrow \infty$.

Recall from the hypothesis in equation (43) that $\|\bar{g}_{G_n} - g_{G_*}\|_{\mathbb{L}_2(\mu)}/\mathcal{L}_{3n} \rightarrow 0$ as $n \rightarrow \infty$, which indicates that $\|\bar{g}_{G_n} - g_{G_*}\|_{\mathbb{L}_1(\mu)}/\mathcal{L}_{3n} \rightarrow 0$ due to the equivalence between $\mathbb{L}_1(\mu)$ -norm and $L^2(\mu)$ -norm. By means of the Fatou's lemma, we have

$$0 = \lim_{n \rightarrow \infty} \frac{\|\bar{g}_{G_n} - g_{G_*}\|_{\mathbb{L}_1(\mu)}}{m_n \mathcal{L}_{3n}} \geq \int \liminf_{n \rightarrow \infty} \frac{|\bar{g}_{G_n}(x) - g_{G_*}(x)|}{m_n \mathcal{L}_{3n}} d\mu(x) \geq 0.$$

This result implies that $[\bar{g}_{G_n}(x) - g_{G_*}(x)]/[m_n \mathcal{L}_{3n}] \rightarrow 0$ for almost every x .

Let us denote

$$\begin{aligned} T_{j,\alpha_1,\alpha_2}^n/m_n \mathcal{L}_{3n} &\rightarrow t_{j,\alpha_1,\alpha_2}, \\ S_{j,\gamma}^n/m_n \mathcal{L}_{3n} &\rightarrow s_{j,\gamma} \end{aligned}$$

with a note that at least one among the limits t_{j,α_1,α_2} , $s_{j,\gamma}$ is non-zero. Then, from the decomposition in equation (45), we deduce that

$$\sum_{j=1}^K \sum_{|\alpha|=0}^2 t_{j,\alpha_1,\alpha_2} \cdot \frac{\partial^{|\alpha_1|} F}{\partial \beta_1^{\alpha_1}}(x, \beta_{1j}^*) \frac{\partial^{|\alpha_2|} h}{\partial \eta^{\alpha_2}}(x, \eta_j^*) - \sum_{j=1}^K \sum_{|\gamma|=0}^2 s_{j,\gamma} \cdot \frac{\partial^{|\gamma|} F}{\partial \beta_1^\gamma}(x, \beta_{1j}^*) \bar{g}_{G_n}(x) = 0,$$

for almost every x . Note that the expert function $h(\cdot, \eta)$ satisfies the condition in Definition 1, then the above equation implies that $t_{j,\alpha_1,\alpha_2} = s_{j,\gamma} = 0$, for any $j \in [K]$, $\alpha_1 \in \mathbb{N}^{d_1}$, $\alpha_2 \in \mathbb{N}^{d_2}$ and $\gamma \in \mathbb{N}^{d_1}$ such that $0 \leq |\alpha_1| + |\alpha_2|, |\gamma| \leq 2$. This contradicts the fact that at least one among the limits t_{j,α_1,α_2} , $s_{j,\gamma}$ is different from zero.

Hence, we obtain the local inequality in equation (42), and completes the proof.

D Auxiliary Results

In this appendix, we present three additional results to facilitate the proofs in Appendix B and Appendix C.

Proposition 2 (Identifiability). *If $g_G(x) = g_{G^*}(x)$ holds true for almost every x , then it follows that $G \equiv G'$.*

Proof of Proposition 2. For almost every x , since $g_G(x) = g_{G^*}(x)$, then we have

$$\begin{aligned} \sum_{i=1}^k \text{Softmax}\left(\text{Top}_K\left(\frac{(\beta_{1i})^\top x}{(\|\beta_{1i}\| + \tau_1) \cdot (\|x\| + \tau_2)}, \beta_{0i}\right)\right) \cdot h(x, \zeta_i) \\ = \sum_{i=1}^{k_*} \text{Softmax}\left(\text{Top}_K\left(\frac{(\beta_{1i}^*)^\top x}{(\|\beta_{1i}^*\| + \tau_1) \cdot (\|x\| + \tau_2)}, \beta_{0i}^*\right)\right) \cdot h(x, \eta_i^*). \end{aligned} \quad (46)$$

Due to the identifiability of the expert function $h(\cdot, \eta)$, the set $\{h(x, \eta'_i) : i \in [k']\}$, where $\eta'_1, \dots, \eta'_{k'}$ are pair-wise different vectors for some $k' \in \mathbb{N}$, is linearly independent for almost every x .

Additionally, note that if $k \neq k_*$, then there exists some index $i \in [k]$ such that $\zeta_i \neq \eta_j^*$ for any $j \in [k_*]$. This result implies that $\text{Softmax}\left(\text{Top}_K\left(\frac{(\beta_{1i})^\top x}{(\|\beta_{1i}\| + \tau_1) \cdot (\|x\| + \tau_2)} + \beta_{0i}\right)\right) = 0$ for almost every x , which is a contradiction. Thus, we must have $k = k_*$. As a result, it follows that

$$\left\{ \text{Softmax}\left((\beta_{1i})^\top X + \beta_{0i}\right) : i \in [k] \right\} = \left\{ \text{Softmax}\left((\beta_{1i}^*)^\top X + \beta_{0i}^*\right) : i \in [k_*] \right\},$$

for almost every X . WLOG, we may assume that

$$\begin{aligned} \text{Softmax}\left(\text{Top}_K\left(\frac{(\beta_{1i})^\top x}{(\|\beta_{1i}\| + \tau_1) \cdot (\|x\| + \tau_2)}, \beta_{0i}\right)\right) \\ = \text{Softmax}\left(\text{Top}_K\left(\frac{(\beta_{1i}^*)^\top x}{(\|\beta_{1i}^*\| + \tau_1) \cdot (\|x\| + \tau_2)}, \beta_{0i}^*\right)\right), \end{aligned} \quad (47)$$

for almost every x , for any $i \in [k_*]$. As the Softmax function is invariant to translations, then the equation (47) indicates that

$$\frac{(\beta_{1i})^\top x}{(\|\beta_{1i}\| + \tau_1) \cdot (\|x\| + \tau_2)} = \frac{(\beta_{1i}^*)^\top x}{(\|\beta_{1i}^*\| + \tau_1) \cdot (\|x\| + \tau_2)}$$

$$\beta_{0i} = \beta_{0i}^* + v_0,$$

for some $v_0 \in \mathbb{R}$. The first equation implies that $\beta_{1i} = \beta_{1i}^*$, while the second equation together with the assumption $\beta_{0k} = \beta_{0k}^* = 0$ lead to $\beta_{0i} = \beta_{0i}^*$ for any $i \in [k_*]$.

Let us consider $x \in \mathcal{X}_\ell$, where $\ell \in [q]$ such that $\{\ell_1, \dots, \ell_K\} = \{1, \dots, K\}$. Then, the equation (46) can be rewritten as

$$\sum_{i=1}^{k_*} \exp(\beta_{0i}) \exp\left(\frac{(\beta_{1i}^*)^\top x}{(\|\beta_{1i}^*\| + \tau_1) \cdot (\|x\| + \tau_2)}\right) h(x, \zeta_i)$$

$$= \sum_{i=1}^{k_*} \exp(\beta_{0i}^*) \exp\left(\frac{(\beta_{1i}^*)^\top x}{(\|\beta_{1i}^*\| + \tau_1) \cdot (\|x\| + \tau_2)}\right) h(x, \eta_i^*), \quad (48)$$

for almost every $x \in \mathcal{X}_\ell$. Next, we denote P_1, P_2, \dots, P_m as a partition of the index set $[k_*]$, where $m \leq k$, such that $\exp(\beta_{0i}) = \exp(\beta_{0i'})$ for any $i, i' \in P_j$ and $j \in [m]$. On the other hand, when i and i' do not belong to the same set P_j , we let $\exp(\beta_{0i}) \neq \exp(\beta_{0i'})$. Thus, we can represent equation (48) as

$$\sum_{j=1}^m \sum_{i \in P_j} \exp(\beta_{0i}) \exp\left(\frac{(\beta_{1i}^*)^\top x}{(\|\beta_{1i}^*\| + \tau_1) \cdot (\|x\| + \tau_2)}\right) h(x, \zeta_i)$$

$$= \sum_{j=1}^m \sum_{i \in P_j} \exp(\beta_{0i}^*) \exp\left(\frac{(\beta_{1i}^*)^\top x}{(\|\beta_{1i}^*\| + \tau_1) \cdot (\|x\| + \tau_2)}\right) h(x, \eta_i^*),$$

for almost every $x \in \mathcal{X}_\ell$. Recall that we have $\beta_{1i} = \beta_{1i}^*$ and $\beta_{0i} = \beta_{0i}^*$, for any $i \in [k_*]$, then the above result leads to

$$\{\zeta_i : i \in P_j\} \equiv \{\eta_i^* : i \in P_j\},$$

for any $j \in [m]$. As a consequence, we obtain that

$$G = \sum_{j=1}^m \sum_{i \in P_j} \exp(\beta_{0i}) \delta_{(\beta_{1i}, \zeta_i)} = \sum_{j=1}^m \sum_{i \in P_j} \exp(\beta_{0i}^*) \delta_{(\beta_{1i}^*, \eta_i^*)} = G_*.$$

Hence, we reach the conclusion of this proposition. \square

Lemma 3. For any $i \in [k_*]$, let $\beta_{1i}, \beta_{1i}^* \in \mathbb{R}^{d_1}$ such that $\|\beta_{1i} - \beta_{1i}^*\| \leq \zeta_i$ for some sufficiently small $\zeta_i > 0$. Let us denote for any $\ell \in [q]$ that

$$\mathcal{X}_\ell^* := \left\{ x \in \mathcal{X} : \cos(\beta_{1j}^*, x) \geq \cos(\beta_{1j'}, x) : \forall j \in \{\ell_1, \dots, \ell_K\}, j' \in \{\ell_{K+1}, \dots, \ell_{k_*}\} \right\},$$

$$\mathcal{X}_\ell := \left\{ x \in \mathcal{X} : \cos(\beta_{1i}, x) \geq \cos(\beta_{1i'}, x), \forall i \in \{\ell_1, \dots, \ell_K\}, i' \in \{\ell_{K+1}, \dots, \ell_{k_*}\} \right\}.$$

Then, unless the set \mathcal{X}_ℓ^* has measure zero, we obtain that $\mathcal{X}_\ell^* = \mathcal{X}_\ell$.

Proof of Lemma 3. Let $\zeta_i = M_i\varepsilon$, where $\varepsilon > 0$ is a given constant and M_i will be chosen later. For an arbitrary $\ell \in [q]$, since \mathcal{X} and Θ are bounded sets, we can find some constant $c_\ell^* \geq 0$ such that

$$\min_{x,i,i'} \left[\cos(\beta_{1i}^*, x) - \cos(\beta_{1i'}^*, x) \right] = c_\ell^* \varepsilon, \quad (49)$$

where the minimum is subject to $x \in \mathcal{X}_\ell^*$, $i \in \{\ell_1, \dots, \ell_K\}$ and $i' \in \{\ell_{K+1}, \dots, \ell_{k_*}\}$. We will point out that $c_\ell^* > 0$. Assume by contrary that $c_\ell^* = 0$. For $x \in \mathcal{X}_\ell^*$, we may assume for any $1 \leq i < j \leq k_*$ that

$$\cos(\beta_{1\ell_i}^*, x) \geq \cos(\beta_{1\ell_j}^*, x).$$

Since $c_\ell^* = 0$, it follows from equation (49) that $\cos(\beta_{1\ell_K}^*, x) - \cos(\beta_{1\ell_{K+1}}^*, x) = 0$, or equivalently

$$\left(\frac{\beta_{1\ell_K}^*}{\|\beta_{1\ell_K}^*\|} - \frac{\beta_{1\ell_{K+1}}^*}{\|\beta_{1\ell_{K+1}}^*\|} \right)^\top x = 0$$

Thus, \mathcal{X}_ℓ^* is a subset of

$$\mathcal{Z} := \left\{ x \in \mathcal{X} : \left(\frac{\beta_{1\ell_K}^*}{\|\beta_{1\ell_K}^*\|} - \frac{\beta_{1\ell_{K+1}}^*}{\|\beta_{1\ell_{K+1}}^*\|} \right)^\top x = 0 \right\}.$$

Since $\beta_{1\ell_K}^* - \beta_{1\ell_{K+1}}^* \neq \mathbf{0}_d$ and the input distribution μ is continuous, it follows that the set \mathcal{Z} has measure zero. Since $\mathcal{X}_\ell^* \subseteq \mathcal{Z}$, we can conclude that \mathcal{X}_ℓ^* also has measure zero, which contradicts the hypothesis of Lemma 3. Therefore, we must have $c_\ell^* > 0$.

Now, we show that $\mathcal{X}_\ell^* \subseteq \mathcal{X}_\ell$. For any vector u , let us denote $\bar{u} = \frac{u}{\|u\|}$. Let $x \in \mathcal{X}_\ell^*$, then we have for any $i \in \{\ell_1, \dots, \ell_K\}$ and $i' \in \{\ell_{K+1}, \dots, \ell_{k_*}\}$ that

$$\begin{aligned} \bar{\beta}_{1i}^\top \bar{x} &= (\bar{\beta}_{1i} - \bar{\beta}_{1i}^*)^\top \bar{x} + (\bar{\beta}_{1i}^*)^\top \bar{x} \\ &\geq -M_i \varepsilon B + (\bar{\beta}_{1i'}^*)^\top \bar{x} + c_\ell^* \varepsilon \\ &= -M_i \varepsilon B + c_\ell^* \varepsilon + (\bar{\beta}_{1i'}^* - \bar{\beta}_{1i'}^\top)^\top \bar{x} + \bar{\beta}_{1i'}^\top \bar{x} \\ &\geq -2M_i \varepsilon B + c_\ell^* \varepsilon + \bar{\beta}_{1i'}^\top \bar{x}. \end{aligned}$$

By setting $M_i \leq \frac{c_\ell^*}{2B}$, we get that $\bar{\beta}_{1i}^\top \bar{x} \geq \bar{\beta}_{1i'}^\top \bar{x}$. Thus, it follows that $x \in \mathcal{X}_\ell$, which means $\mathcal{X}_\ell^* \subseteq \mathcal{X}_\ell$. Similarly, assume that there exists some constant $c_\ell \geq 0$ that satisfies

$$\min_{x,i,i'} \left[\cos(\beta_{1i}^*, x) - \cos(\beta_{1i'}^*, x) \right] = c_\ell^* \varepsilon.$$

Here, the above minimum is subject to $x \in \mathcal{X}_\ell$, $i \in \{\ell_1, \dots, \ell_K\}$ and $i' \in \{\ell_{K+1}, \dots, \ell_{k_*}\}$. If $M_i \leq \frac{c_\ell}{2B}$, then we also receive that $\mathcal{X}_\ell \subseteq \mathcal{X}_\ell^*$.

Hence, if we set $M_i = \frac{1}{2B} \min\{c_\ell^*, c_\ell\}$, we reach the conclusion that $\mathcal{X}_\ell^* = \mathcal{X}_\ell$. \square

Lemma 4. Let $F(x, \beta_1) := \exp\left(\frac{\beta_1^\top x}{\|\beta_1\| \cdot \|x\|}\right)$. For any vector $\beta_1 \in \mathbb{R}^{d_1}$ and $t \in \mathbb{N}$, we have

$$\sum_{u_1, \dots, u_t=1}^{d_1} \beta_1^{(u_1)} \dots \beta_1^{(u_t)} \cdot \frac{\partial^t F}{\partial \beta_1^{(u_1)} \dots \partial \beta_1^{(u_t)}}(x, \beta_1) = 0. \quad (50)$$

Proof of Lemma 4. We will prove the above result by using the induction method. In particular, we first show that it holds for $t = 1$. By taking the first derivative of F w.r.t β_1 , we have

$$\frac{\partial F}{\partial \beta_1}(x, \beta_1) = \frac{x \cdot \|\beta_1\| \cdot \|x\| - \frac{\beta_1}{\|\beta_1\|} \cdot \|x\| \cdot \beta_1^\top x}{\|\beta_1\|^2 \|x\|^2} \cdot F(x, \beta_1).$$

Then, it follows that

$$\beta_1^\top \frac{\partial F}{\partial \beta_1}(x, \beta_1) = \frac{\beta_1^\top x \cdot \|\beta_1\| \cdot \|x\| - \frac{\beta_1^\top \beta_1}{\|\beta_1\|} \cdot \|x\| \cdot \beta_1^\top x}{\|\beta_1\|^2 \|x\|^2} = 0,$$

or equivalently,

$$\sum_{u_1=1}^{d_1} \beta_1^{(u_1)} \cdot \frac{\partial F}{\partial \beta_1^{(u_1)}}(x, \beta_1) = 0.$$

Subsequently, assume that the equation (50) holds for $t - 1$, i.e.

$$\sum_{u_1, \dots, u_{t-1}=1}^{d_1} \beta_1^{(u_1)} \dots \beta_1^{(u_{t-1})} \cdot \frac{\partial^{t-1} F}{\partial \beta_1^{(u_1)} \dots \partial \beta_1^{(u_{t-1})}}(x, \beta_1) = 0,$$

we will demonstrate that it also holds for t . Note that the above left hand side can be decomposed as

$$\begin{aligned} & \sum_{u_1, \dots, u_{t-1}=1}^{d_1} \beta_1^{(u_1)} \dots \beta_1^{(u_{t-1})} \cdot \frac{\partial^{t-1} F}{\partial \beta_1^{(u_1)} \dots \partial \beta_1^{(u_{t-1})}}(x, \beta_1) = \sum_{u_1, \dots, u_{t-1} \neq u_t} \beta_1^{(u_1)} \dots \beta_1^{(u_{t-1})} \cdot \frac{\partial^{t-1} F}{\partial \beta_1^{(u_1)} \dots \partial \beta_1^{(u_{t-1})}} \\ & + \binom{t-1}{1} \sum_{u_2, \dots, u_{t-1} \neq u_t} \beta_1^{(u_2)} \dots \beta_1^{(u_{t-1})} \beta_1^{(u_t)} \cdot \frac{\partial^{t-1} F}{\partial \beta_1^{(u_2)} \dots \partial \beta_1^{(u_{t-1})} \partial \beta_1^{(u_t)}}(x, \beta_1) \\ & + \dots \\ & + \binom{t-1}{t-2} \sum_{u_{t-1} \neq u_t} \beta_1^{(u_{t-1})} (\beta_1^{(u_t)})^{t-2} \cdot \frac{\partial^{t-1} F}{\partial \beta_1^{(u_{t-1})} \partial (\beta_1^{(u_t)})^{t-2}}(x, \beta_1) + (\beta_1^{(u_t)})^{t-1} \cdot \frac{\partial^{t-1} F}{\partial (\beta_1^{(u_t)})^{t-1}}(x, \beta_1), \end{aligned}$$

where u_t is some index in $[d]$. By taking the derivatives of both sides w.r.t β_1 , we get

$$\begin{aligned}
0 &= \sum_{u_1, \dots, u_{t-1} \neq u_t} \beta_1^{(u_1)} \dots \beta_1^{(u_{t-1})} \cdot \frac{\partial^t F}{\partial \beta_1^{(u_1)} \dots \partial \beta_1^{(u_t)}} \\
&+ \binom{t-1}{1} \sum_{u_2, \dots, u_{t-1} \neq u_t} \beta_1^{(u_2)} \dots \beta_1^{(u_{t-1})} \cdot \frac{\partial^{t-1} F}{\partial \beta_1^{(u_2)} \dots \partial \beta_1^{(u_{t-1})} \partial \beta_1^{(u_t)}}(x, \beta_1) \\
&+ \binom{t-1}{1} \sum_{u_2, \dots, u_{t-1} \neq u_t} \beta_1^{(u_2)} \dots \beta_1^{(u_{t-1})} \beta_1^{(u_t)} \cdot \frac{\partial^t F}{\partial \beta_1^{(u_2)} \dots \partial \beta_1^{(u_{t-1})} \partial (\beta_1^{(u_t)})^2}(x, \beta_1) \\
&+ \dots \\
&+ \binom{t-1}{t-2} (t-2) \sum_{u_{t-1} \neq u_t} \beta_1^{(u_{t-1})} (\beta_1^{(u_t)})^{t-3} \cdot \frac{\partial^{t-1} F}{\partial \beta_1^{(u_{t-1})} \partial (\beta_1^{(u_t)})^{t-2}}(x, \beta_1) \\
&+ \binom{t-1}{t-2} \sum_{u_{t-1} \neq u_t} \beta_1^{(u_{t-1})} (\beta_1^{(u_t)})^{t-2} \cdot \frac{\partial^t F}{\partial \beta_1^{(u_{t-1})} \partial (\beta_1^{(u_t)})^{t-1}}(x, \beta_1) \\
&+ (t-1) (\beta_1^{(u_t)})^{t-2} \cdot \frac{\partial^{t-1} F}{\partial (\beta_1^{(u_t)})^{t-1}}(x, \beta_1) + (\beta_1^{(u_t)})^{t-1} \cdot \frac{\partial^t F}{\partial (\beta_1^{(u_t)})^t}(x, \beta_1) \\
&= \sum_{u_1, \dots, u_{t-1}=1}^{d_1} \beta_1^{(u_1)} \dots \beta_1^{(u_{t-1})} \cdot \frac{\partial^t F}{\partial \beta_1^{(u_1)} \dots \partial \beta_1^{(u_t)}}(x, \beta_1) \\
&\quad + (t-1) \sum_{u_2, \dots, u_{t-1}=1}^{d_1} \beta_1^{(u_2)} \dots \beta_1^{(u_{t-1})} \cdot \frac{\partial^{t-1} F}{\partial \beta_1^{(u_2)} \dots \partial \beta_1^{(u_t)}}(x, \beta_1).
\end{aligned}$$

Therefore, it follows that

$$\begin{aligned}
\sum_{u_t=1}^{d_1} \beta_1^{(u_t)} \cdot 0 &= \sum_{u_1, \dots, u_t=1}^{d_1} \beta_1^{(u_1)} \dots \beta_1^{(u_{t-1})} \beta_1^{(u_t)} \cdot \frac{\partial^t F}{\partial \beta_1^{(u_1)} \dots \partial \beta_1^{(u_t)}}(x, \beta_1) \\
&\quad + (t-1) \sum_{u_2, \dots, u_t=1}^{d_1} \beta_1^{(u_2)} \dots \beta_1^{(u_{t-1})} \beta_1^{(u_t)} \cdot \frac{\partial^{t-1} F}{\partial \beta_1^{(u_2)} \dots \partial \beta_1^{(u_t)}}(x, \beta_1).
\end{aligned}$$

It is worth noting that

$$\begin{aligned}
&\sum_{u_2, \dots, u_t=1}^{d_1} \beta_1^{(u_2)} \dots \beta_1^{(u_{t-1})} \beta_1^{(u_t)} \cdot \frac{\partial^{t-1} F}{\partial \beta_1^{(u_2)} \dots \partial \beta_1^{(u_t)}}(x, \beta_1) \\
&= \sum_{u_1, \dots, u_{t-1}=1}^{d_1} \beta_1^{(u_1)} \dots \beta_1^{(u_{t-1})} \cdot \frac{\partial^{t-1} F}{\partial \beta_1^{(u_1)} \dots \partial \beta_1^{(u_{t-1})}}(x, \beta_1) = 0.
\end{aligned}$$

Consequently, we deduce that

$$\sum_{u_1, \dots, u_t=1}^{d_1} \beta_1^{(u_1)} \dots \beta_1^{(u_{t-1})} \beta_1^{(u_t)} \cdot \frac{\partial^t F}{\partial \beta_1^{(u_1)} \dots \partial \beta_1^{(u_t)}}(x, \beta_1) = 0.$$

Hence, we reach the conclusion in equation (50). \square

E Experimental Details

In this appendix, we provide the details for the numerical experiments on synthetic data, and the experiments with real data on language modeling conducted in Section 5.

E.1 Experimental Details for Synthetic Data

Model details. We now provide the details for the model parameters in model (15). The variance of Gaussian noise is specified as $\sigma^2 = 0.01$. For simplicity, the perturbations for both $\|x\|$ and $\|\beta_{1i}^*\|$ are considered identical, denoted by $\tau_1 = \tau_2 = \tau$. The true parameters for the router, $(\beta_{1i}^*, \beta_{0i}^*) \in \mathbb{R}^d \times \mathbb{R}$, are drawn independently from an isotropic Gaussian distribution with zero mean and variance $\sigma_r^2 = 0.01/d$ for $1 \leq i \leq 6$, and otherwise are set to zero. Similarly, the true parameters of the experts, $(a_i^*, b_i^*) \in \mathbb{R}^d \times \mathbb{R}$, are drawn independently of an isotropic Gaussian distribution with zero mean and variance $\sigma_e^2 = 1/d$ for all experts. These parameters remain unchanged for all experiments.

Training procedure. For each sample size n , spanning from 10^3 to 10^5 , we perform 20 experiments. In every experiment, the parameters initialization for the router’s and experts’ parameters are adjusted to be near the true parameters, minimizing potential instabilities from the optimization process. Subsequently, we execute SGD across 10 epochs, employing a learning rate of $\eta = 0.1$ to fit a model to the synthetic data. All the numerical experiments are conducted on a MacBook Air equipped with an M1 chip CPU.

E.2 Experimental Details for Language Modeling Task

Datasets. We use the Enwik8 and Text8 datasets ([23]) for our character-level language modeling task. The Enwik8 dataset comprises 100 million bytes of unprocessed Wikipedia text, while the Text8 dataset contains 100 million processed Wikipedia characters. We further evaluate the word-level language modeling task on the Wikitext-103 dataset [25], which is the largest available word-level language modeling benchmark with long-term dependency. It contains 103M training tokens from 28K articles, with an average length of 3.6K tokens per article, which allows us to test the ability of long-term dependency modeling.

Metrics. In the main paper, we employ the Bit per character (BPC) metric to assess the performance of character-level language modeling tasks. This metric is used to measure the average number of bits needed to encode each character in the dataset. It is calculated as follows:

$$\text{BPC}(X) = -\frac{1}{T} \sum_{t=1}^T \log_2 \hat{P}_t(x_t)$$

where T is the length of the input string X , \hat{P}_t is the approximate distribution and x_t is the character in the input string at location t .

For the word-level language modeling task on the Wikitext-103 dataset, we utilize Perplexity (PPL) as our evaluation metric. It represents the exponentiated average negative log-likelihood of a sequence and demonstrate how well the model predicts the next word in a sequence. More specifically, if we have a tokenized sequence $X = (x_0, x_1, \dots, x_t)$, the perplexity of X is:

$$\text{PPL}(X) = \exp \left\{ -\frac{1}{t} \sum_{i=1}^t \log p_\theta(x_i | x_{<i}) \right\}$$

Table 2: Average out-of-distribution test accuracies.

	PACS	VLCS	OfficeHome	TerraIncognita	DomainNet
Vanilla	87.22	78.99	73.27	45.55	48.45
Pertubed	89.36	80.01	74.09	49.87	48.51

Table 3: Per-domain performance of PACS, VLCS, OfficeHome, TerraIncognita.

		A	C	P	S
PACS	Vanilla	89.24	86.11	97.60	75.92
	Pertubed	89.87	86.97	97.90	82.68
		C	L	S	V
VLCS	Vanilla	98.59	67.42	70.88	79.07
	Pertubed	98.59	67.80	74.70	78.95
		A	C	P	R
OfficeHome	Vanilla	73.40	57.27	78.69	83.70
	Pertubed	74.64	57.85	79.59	84.27
		L100	L30	L43	L46
TerraIncognita	Vanilla	50.00	37.49	53.02	41.67
	Pertubed	57.59	43.30	56.93	41.67

where $p_{\theta}(x_i|x_{<i})$ is the log-likelihood of the i^{th} token conditioned on the preceding tokens $x_{<i}$ according to our model.

Training setup and hyperparameters. We consider two model configurations: the *small* and *medium* setups. The *small* setup has a total of 15 million parameters with 6 SMoE layers [6], each layer has the capability to learn spatial structure in the input domain and routing experts at a fine-grained level to utilize it. Similarly, the *medium* setup consists of 36 parameters with 8 SMoE layers [6]. During training, we use Adam optimizer [18] with default parameters. We set the number of training steps to 60000 and 80000 for small and medium configurations, respectively. The results are averaged over three runs for fair comparisons.

All language modeling experiments are conducted on NVIDIA A100 GPUs. Training the small configuration of the Text8 and Enwik8 datasets takes 11 hours, whereas Wikitext-103 requires 5 hours. For medium configurations, training Text8 and Enwik8 takes 17 hours, while Wikitext-103 training takes 8 hours.

F Domain Generalization

In this appendix, we carry our several experiments on the applications of MoE models in the field of domain generalization. Our objective is to empirically demonstrate the efficacy of our proposed

Table 4: Per-domain performance of DomainNet.

		clipart	infograph	painting	quickdraw	real	sketch
DomainNet	Vanilla	68.05	24.48	55.75	17.39	69.41	55.59
	Pertubed	68.31	24.52	55.03	17.90	69.46	55.83

perturbed cosine router over the vanilla cosine router in this field.

Domain generalization [42] aims to generalize a model’s performance to unseen test domains with distributions different from those encountered during training. Specifically, in domain generalization, a model is expected to leverage multiple training datasets gathered from various domains and exhibit robustness to domain shifts during testing. Such ability of out-of-distribution generalization largely hinges on the model’s capability to incorporate invariances across multiple domains [20]. Given that distribution shifts in data correspond to distribution shifts in (visual) attributes [39], capturing these diverse attributes and aligning them with invariant correlations is crucial. Mixture of Experts emerges as a powerful tool for efficiently capturing these visual attributes, and it has been proven effective in enhancing performance in domain generalization [20]. Therefore, we further justify the effectiveness of our perturbed cosine router in domain generalization.

Datasets. We followed the experimental setting of [20] and evaluated our method using 5 benchmark datasets in DomainBed [2]: PACS, VLCS, OfficeHome, TerraIncognita, and DomainNet. Each dataset is comprised of images for classification tasks from different domains.

Architecture. Following [10], we conduct experiments on ViT-S/16, which has an input patch size of 16×16 , comprising 6 heads in multi-head attention layers, and a total of 12 transformer blocks. We adopt a *last-two* two-layer configuration, where each MoE block comprises 6 experts. The router selects the top 2 out of 6 experts for each image patch.

Training procedure and result. We follow the training-domain validation procedure outlined in [2, 10], where each training domain is split into training and validation subsets. The final overall validation set consists of the validation subsets from all training domains. Subsequently, we select the model with the highest performance on the overall validation set. To ensure fair comparisons, the results are averaged over three runs. All DG experiments are run on NVIDIA A100 GPUs with 15,000 iterations. The training time on PACS, VLCS, OfficeHome, and TerraIncognita is 2 hours, while the training time for DomainNet is 7 hours.

Table 2 summarizes the results of our experiments. For each dataset, we report the average results across test domains. The results demonstrate that our perturbed cosine router consistently outperforms the vanilla cosine router across all datasets, thereby convincingly justifying the effectiveness of adding noise to cosine routers. Detailed performance metrics for each domain are reported in Tables 3 and 4.

References

- [1] Z. Chen, Y. Deng, Y. Wu, Q. Gu, and Y. Li. Towards understanding the mixture-of-experts layer in deep learning. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh,

- editors, *Advances in Neural Information Processing Systems*, volume 35, pages 23049–23062. Curran Associates, Inc., 2022. (Cited on page 2.)
- [2] Z. Chi, L. Dong, S. Huang, D. Dai, S. Ma, B. Patra, S. Singhal, P. Bajaj, X. Song, X.-L. Mao, H. Huang, and F. Wei. On the representation collapse of sparse mixture of experts. In A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, editors, *Advances in Neural Information Processing Systems*, 2022. (Cited on pages 2, 11, and 40.)
- [3] Y. Chow, A. Tulepbergenov, O. Nachum, D. Gupta, M. Ryu, M. Ghavamzadeh, and C. Boutilier. A Mixture-of-Expert Approach to RL-based Dialogue Management. In *The Eleventh International Conference on Learning Representations*, 2023. (Cited on page 2.)
- [4] D. Dai, C. Deng, C. Zhao, R. X. Xu, H. Gao, D. Chen, J. Li, W. Zeng, X. Yu, Y. Wu, Z. Xie, Y. K. Li, P. Huang, F. Luo, C. Ruan, Z. Sui, and W. Liang. Deepseekmoe: Towards ultimate expert specialization in mixture-of-experts language models. *arXiv preprint arXiv:2401.04088*, 2024. (Cited on page 1.)
- [5] D. Do, H. Nguyen, K. Nguyen, and N. Ho. Minimax optimal rate for parameter estimation in multivariate deviated models. In *Advances in Neural Information Processing Systems*, volume 36, pages 30096–30133. Curran Associates, Inc., 2023. (Cited on page 12.)
- [6] N. Dryden and T. Hoefler. Spatial mixture-of-experts. In A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, editors, *Advances in Neural Information Processing Systems*, 2022. (Cited on page 39.)
- [7] N. Du, Y. Huang, A. M. Dai, S. Tong, D. Lepikhin, Y. Xu, M. Krikun, Y. Zhou, A. Yu, O. Firat, B. Zoph, L. Fedus, M. Bosma, Z. Zhou, T. Wang, E. Wang, K. Webster, M. Pellat, K. Robinson, K. Meier-Hellstern, T. Duke, L. Dixon, K. Zhang, Q. Le, Y. Wu, Z. Chen, and C. Cui. Glam: Efficient scaling of language models with mixture-of-experts. In *ICML*, 2022. (Cited on page 1.)
- [8] W. Fedus, B. Zoph, and N. M. Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *J. Mach. Learn. Res.*, 23:120:1–120:39, 2021. (Cited on page 11.)
- [9] A. Gulati, J. Qin, C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang. Conformer: Convolution-augmented Transformer for Speech Recognition. In *Proc. Interspeech 2020*, pages 5036–5040, 2020. (Cited on page 2.)
- [10] I. Gulrajani and D. Lopez-Paz. In search of lost domain generalization. In *International Conference on Learning Representations*, 2021. (Cited on page 40.)
- [11] X. Han, H. Nguyen, C. Harris, N. Ho, and S. Saria. Fusemoe: Mixture-of-experts transformers for fleximodal fusion. *arXiv preprint arXiv:2402.03226*, 2024. (Cited on page 2.)
- [12] H. Hazimeh, Z. Zhao, A. Chowdhery, M. Sathiamoorthy, Y. Chen, R. Mazumder, L. Hong, and E. Chi. DSelect-k: Differentiable Selection in the Mixture of Experts with Applications to Multi-Task Learning. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P. S. Liang, and J. W. Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 29335–29347. Curran Associates, Inc., 2021. (Cited on page 2.)

- [13] P. Heinrich and J. Kahn. Strong identifiability and optimal minimax rates for finite mixture estimation. *The Annals of Statistics*, 46(6):2844–2870, 2018. (Cited on page 12.)
- [14] N. Ho, C.-Y. Yang, and M. I. Jordan. Convergence rates for Gaussian mixtures of experts. *Journal of Machine Learning Research*, 23(323):1–81, 2022. (Cited on page 2.)
- [15] R. A. Jacobs, M. I. Jordan, S. J. Nowlan, and G. E. Hinton. Adaptive mixtures of local experts. *Neural Computation*, 3, 1991. (Cited on page 1.)
- [16] A. Q. Jiang, A. Sablayrolles, A. Roux, A. Mensch, B. Savary, C. Bamford, D. S. Chaplot, D. de las Casas, E. B. Hanna, F. Bressand, G. Lengyel, G. Bour, G. Lample, L. R. Lavaud, L. Saulnier, M.-A. Lachaux, P. Stock, S. Subramanian, S. Yang, S. Antoniak, T. L. Scao, T. Gervet, T. Lavril, T. Wang, T. Lacroix, and W. E. Sayed. Mixtral of experts. *arxiv preprint arxiv 2401.04088*, 2024. (Cited on page 1.)
- [17] M. I. Jordan and R. A. Jacobs. Hierarchical mixtures of experts and the EM algorithm. *Neural Computation*, 6:181–214, 1994. (Cited on page 1.)
- [18] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization, 2017. (Cited on page 39.)
- [19] D. Lepikhin, H. Lee, Y. Xu, D. Chen, O. Firat, Y. Huang, M. Krikun, N. Shazeer, and Z. Chen. GShard: Scaling Giant Models with Conditional Computation and Automatic Sharding. In *International Conference on Learning Representations*, 2021. (Cited on page 1.)
- [20] B. Li, Y. Shen, J. Yang, Y. Wang, J. Ren, T. Che, J. Zhang, and Z. Liu. Sparse mixture-of-experts are domain generalizable learners. In *The Eleventh International Conference on Learning Representations*, 2023. (Cited on pages 2 and 40.)
- [21] H. Liang, Z. Fan, R. Sarkar, Z. Jiang, T. Chen, K. Zou, Y. Cheng, C. Hao, and Z. Wang. M³ViT: Mixture-of-Experts Vision Transformer for Efficient Multi-task Learning with Model-Accelerator Co-design. In *NeurIPS*, 2022. (Cited on page 1.)
- [22] B. Lindsay. *Mixture models: Theory, geometry and applications*. In NSF-CBMS Regional Conference Series in Probability and Statistics. IMS, Hayward, CA., 1995. (Cited on page 1.)
- [23] M. Mahoney. Large text compression benchmark. <http://www.matmahoney.net/dc/text.html>, 2011. (Cited on pages 11 and 38.)
- [24] T. Manole and N. Ho. Refined convergence rates for maximum likelihood estimation under finite mixture models. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 14979–15006. PMLR, 17–23 Jul 2022. (Cited on page 5.)
- [25] S. Merity, C. Xiong, J. Bradbury, and R. Socher. Pointer sentinel mixture models, 2016. (Cited on pages 11 and 38.)
- [26] H. Nguyen, P. Akbarian, and N. Ho. Is temperature sample efficient for softmax Gaussian mixture of experts? In *Proceedings of the ICML*, 2024. (Cited on page 2.)

- [27] H. Nguyen, P. Akbarian, F. Yan, and N. Ho. Statistical perspective of top-k sparse softmax gating mixture of experts. In *International Conference on Learning Representations*, 2024. (Cited on pages 2 and 3.)
- [28] H. Nguyen, N. Ho, and A. Rinaldo. On least squares estimation in softmax gating mixture of experts. In *Proceedings of the ICML*, 2024. (Cited on pages 2, 3, 5, 6, and 9.)
- [29] H. Nguyen, T. Nguyen, and N. Ho. Demystifying softmax gating function in Gaussian mixture of experts. In *Advances in Neural Information Processing Systems*, 2023. (Cited on page 2.)
- [30] H. Nguyen, T. Nguyen, K. Nguyen, and N. Ho. Towards convergence rates for parameter estimation in Gaussian-gated mixture of experts. In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, 2024. (Cited on page 2.)
- [31] F. Peng, R. A. Jacobs, and M. A. Tanner. Bayesian Inference in Mixtures-of-Experts and Hierarchical Mixtures-of-Experts Models With an Application to Speech Recognition. *Journal of the American Statistical Association*, 91(435):953–960, 1996. (Cited on page 2.)
- [32] Q. Pham, G. Do, H. Nguyen, T. Nguyen, C. Liu, M. Sartipi, B. T. Nguyen, S. Ramasamy, X. Li, S. Hoi, and N. Ho. Competesmoe – effective training of sparse mixture of experts via competition. *arXiv preprint arXiv:2402.02526*, 2024. (Cited on page 1.)
- [33] J. Puigcerver, C. Riquelme, B. Mustafa, and N. Houlsby. From sparse to soft mixtures of experts. In *The Twelfth International Conference on Learning Representations*, 2024. (Cited on page 1.)
- [34] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(1), jan 2020. (Cited on page 11.)
- [35] C. Riquelme, J. Puigcerver, B. Mustafa, M. Neumann, R. Jenatton, A. S. Pinto, D. Keysers, and N. Houlsby. Scaling vision with sparse mixture of experts. In *Advances in Neural Information Processing Systems*, volume 34, pages 8583–8595. Curran Associates, Inc., 2021. (Cited on page 1.)
- [36] C. Ruiz, J. Puigcerver, B. Mustafa, M. Neumann, R. Jenatton, A. Pinto, D. Keysers, and N. Houlsby. Scaling vision with sparse mixture of experts. In *NeurIPS*, 2021. (Cited on page 1.)
- [37] N. Shazeer, A. Mirhoseini, K. Maziarz, A. Davis, Q. Le, G. Hinton, and J. Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *International Conference on Learning Representations (ICLR)*, 2017. (Cited on page 1.)
- [38] S. van de Geer. *Empirical processes in M-estimation*. Cambridge University Press, 2000. (Cited on pages 3 and 14.)
- [39] O. Wiles, S. Goyal, F. Stimberg, S.-A. Rebuffi, I. Ktena, K. D. Dvijotham, and A. T. Cemgil. A fine-grained analysis on distribution shift. In *International Conference on Learning Representations*, 2022. (Cited on page 40.)
- [40] Z. You, S. Feng, D. Su, and D. Yu. Speechmoe: Scaling to large acoustic models with dynamic routing mixture of experts. In *Interspeech*, 2021. (Cited on page 2.)

- [41] B. Yu. Assouad, Fano, and Le Cam. *Festschrift for Lucien Le Cam*, pages 423–435, 1997. (Cited on page 17.)
- [42] K. Zhou, Z. Liu, Y. Qiao, T. Xiang, and C. C. Loy. Domain generalization: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):4396–4415, 2023. (Cited on page 40.)
- [43] Y. Zhou, N. Du, Y. Huang, D. Peng, C. Lan, D. Huang, S. Shakeri, D. So, A. Dai, Y. Lu, Z. Chen, Q. Le, C. Cui, J. Laudon, and J. Dean. Brainformers: Trading simplicity for efficiency. In *International Conference on Machine Learning*, pages 42531–42542. PMLR, 2023. (Cited on page 1.)