

FIRST STEPS TOWARD UNDERSTANDING THE EXTRAPOLATION OF NONLINEAR MODELS TO UNSEEN DOMAINS

Kefan Dong
Stanford University
kefandong@stanford.edu

Tengyu Ma
Stanford University
tengyuma@stanford.edu

ABSTRACT

Real-world machine learning applications often involve deploying neural networks to domains that are not seen in the training time. Hence, we need to understand the extrapolation of *nonlinear* models—under what conditions on the distributions and function class, models can be guaranteed to extrapolate to new test distributions. The question is very challenging because even two-layer neural networks cannot be guaranteed to extrapolate outside the support of the training distribution without further assumptions on the domain shift. This paper makes some initial steps towards analyzing the extrapolation of nonlinear models for structured domain shift. We primarily consider settings where the *marginal* distribution of each coordinate of the data (or subset of coordinates) do not shift significantly across the training and test distributions, but the joint distribution may have a much bigger shift. We prove that the family of nonlinear models of the form $f(x) = \sum f_i(x_i)$, where f_i is an *arbitrary* function on the subset of features x_i , can extrapolate to unseen distributions, if the covariance of the features is well-conditioned. To the best of our knowledge, this is the first result that goes beyond linear models and the bounded density ratio assumption, even though the assumptions on the distribution shift and function class are stylized.

1 INTRODUCTION

In real-world applications, machine learning models are often deployed on domains that are not seen in the training time. For example, we may train machine learning models for medical diagnosis on data from hospitals in Europe and then deploy them to hospitals in Asia.

Thus, we need to understand the extrapolation of models to new test distributions — how the model trained on one distribution behaves on another unseen distribution. This extrapolation of neural networks is central to various robustness questions such as domain generalization (Gulrajani & Lopez-Paz (2020); Ganin et al. (2016); Peters et al. (2016) and references therein) and adversarial robustness (Goodfellow et al., 2014; Kurakin et al., 2018), and also plays a critical role in nonlinear bandits and reinforcement learning where the distribution is constantly changing during training (Dong et al., 2021; Agarwal et al., 2019; Lattimore & Szepesvári, 2020; Sutton & Barto, 2018).

This paper focuses on the following mathematical abstraction of this extrapolation question:

Under what conditions on the source distribution P , target distribution Q , and function class \mathcal{F} do we have that any functions $f, g \in \mathcal{F}$ that agree on P are also guaranteed to agree on Q ?

Here we can measure the agreement of two functions on P by the ℓ_2 distance between f and g under distribution P , that is, $\|f - g\|_P \triangleq \mathbb{E}_{x \sim P}[(f(x) - g(x))^2]^{1/2}$. The function f can be thought of as the learned model, g as the ground-truth function, and thus $\|f - g\|_P$ as the error on the source distribution P .

This question is well-understood for linear function class \mathcal{F} . Essentially, if the covariance of Q can be bounded from above by the covariance of P (in any direction), then the error on Q is guaranteed

to be bounded by the error on P . We refer the reader to Lei et al. (2021); Mousavi Kalan et al. (2020) and references therein for more recent advances along this line.

By contrast, theoretical results for extrapolation of *nonlinear* models is rather limited. Classical results have long settled the case where P and Q have bounded density ratios (Ben-David & Uner, 2014; Sugiyama et al., 2007). Bounded density ratio implies that the support of Q must be a subset of the support of P , and thus arguably these results do not capture the extrapolation behavior of models *outside* the training domain.

Without the bounded density ratio assumption, there was limited prior *positive* result for characterizing the extrapolation power of neural networks. Ben-David et al. (2010) show that the model can extrapolate when the $\mathcal{H}\Delta\mathcal{H}$ -distance between training and test distribution is small. However, it remains unclear for what distributions and function class, the $\mathcal{H}\Delta\mathcal{H}$ -distance can be bounded.¹ In general, the question is challenging partly because of the existence of such a strong impossibility result. As soon as the support of Q is not contained in the support of P (and they satisfy some non-degeneracy condition), it turns out that even two-layer neural networks cannot extrapolate—there are two-layer neural networks f and g that agree on P perfectly but behave very differently on Q (See Proposition 5 for a formal statement.)

The impossibility result suggests that any positive results on the extrapolation of nonlinear models require more fine-grained structures on the relationship between P and Q (which are common in practice Koh et al. (2021); Sagawa et al. (2022)) as well as the function class \mathcal{F} . The structure in the domain shift between P and Q may also need to be compatible with the assumption on the function class \mathcal{F} . This paper makes some first steps towards proving certain family of nonlinear models can extrapolate to a new test domain with structured shift.

We consider a setting where the joint distribution of the data can does not have much overlap across P and Q (and thus bounded density ratio assumption does not hold), whereas the marginal distributions for each coordinate of the data does overlap. Such a scenario may practically happen when the features (coordinates of the data) exhibit different correlations on the source and target distribution. For example, consider the task of predicting the probability of a lightning storm from basic meteorological information such as precipitation, temperature, etc. We learn models from some cities on the west coast of United States and deploy them to the east coast. In this case, the joint test distribution of the features may not necessarily have much overlap with the training distribution—correlation between precipitation and temperature could be vastly different across regions, e.g., the rainy season coincides with the winter’s low temperature on the west coast, but not so much on the east coast. However, the individual feature’s marginal distribution is much more likely to overlap between the source and target—the possible ranges of temperature on east and west coasts are similar.

Concretely, we assume that the features $x \in \mathbb{R}^{s_1+s_2}$ have Gaussian distributions and can be divided into two subsets $x_1 \in \mathbb{R}^{s_1}$ and $x_2 \in \mathbb{R}^{s_2}$ such that each set of feature x_i ($i \in \{1, 2\}$) has the same marginal distributions on P and Q . Moreover, we assume that x_1 and x_2 are not exactly correlated on P —the covariance of features x on distribution P has a strictly positive minimum eigenvalue.

As argued before, restricted assumptions on the function class \mathcal{F} are still necessary (for almost any P and Q without the bounded density ratio property). Here, we assume that \mathcal{F} consists of all functions of the form $f(x) = f_1(x_1) + f_2(x_2)$ for *arbitrary* functions $f_1 : \mathcal{R}^{s_1} \rightarrow \mathbb{R}$ and $f_2 : \mathcal{R}^{s_2} \rightarrow \mathbb{R}$. The function class \mathcal{F} does not contain all two-layer neural networks (so that the impossibility result does not apply), but still consists of a rich set of functions where each subset of features independently contribute to the prediction with arbitrary nonlinear transformations. We show that under these assumptions, if any two models approximately agree on P , they must also approximately agree on Q —formally speaking, $\forall f, g \in \mathcal{F}, \|f - g\|_Q \lesssim \|f - g\|_P$ (Theorem 4).

We also prove a variant of the result above where we divide features vector $x \in \mathbb{R}^d$ into d coordinate, denoted by $x = (x_1, \dots, x_d)$ where $x_i \in \mathbb{R}$. The function class consists of all combinations of *nonlinear* transformations of x_i ’s, that is, $\mathcal{F} = \{\sum_{i=1}^d f_i(x_i)\}$. Assuming coordinates of x are pairwise Gaussian and a non-degenerate covariance matrix, the nonlinear model $f \in \mathcal{F}$ can extrapolate to any distribution Q that has the same marginals as P (Theorem 3).

¹In fact, the $\mathcal{H}\Delta\mathcal{H}$ -distance likely cannot be bounded when the function class contains two-layer neural networks, and the supports of the training and test distributions do not overlap—when there exists a function that can distinguish the source and target domain, the $\mathcal{H}\Delta\mathcal{H}$ divergence will be large.

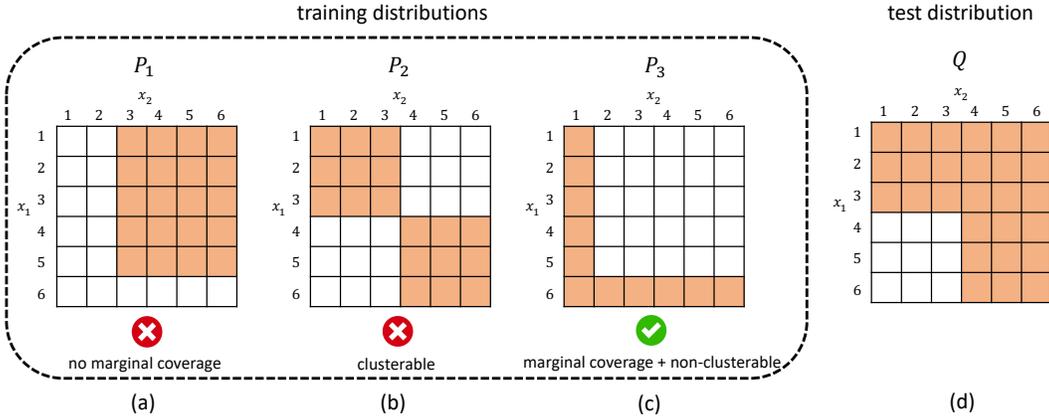


Figure 1: Visualization of three different training distributions P_1, P_2, P_3 and a test distribution Q , where the orange color blocks marks the support. **(a)** and **(b)**: distributions P_1, P_2 that do not satisfy our conditions and cannot extrapolate. **(c)**: a distribution P_3 that satisfies our conditions even though the support of P_3 is sparse. **(d)**: the test distribution.

These results can be viewed as first steps for analyzing extrapolation beyond linear models. Compared the works of Lei et al. (2021) on linear models, our assumptions on the covariance of P are qualitatively similar. We additionally require P and Q have overlapping marginal distributions because it is even necessary for the extrapolation of one-dimensional functions on a single coordinate. Our results work for a more expressive family of nonlinear functions, that is, $\mathcal{F} = \{\sum_{i=1}^d f_i(x_i)\}$, than the linear functions.

We also present a result on the case where x_i 's are discrete variables, which demonstrates the key intuition and also may be of its own interest. Suppose we have two discrete random variable x_1 and x_2 . In this case, the joint distribution of P and Q can be both presented by a matrix (as visualized in Figure 1), and the marginal distributions are the column and row sums of this joint probability matrix. We prove that extrapolation occurs when (1) the support of the *marginals* of Q is covered by P , and (2) the density matrix of P is non-clusterable — we cannot find a subset of rows and columns such the support of P in these rows and columns lies within their intersections.

In Figure 1, we visualize a few interesting cases. First, distributions P_1, P_2 visualized in Figures 1a and 1b, respectively, do not satisfy our conditions. Fundamentally, there are two models that agree on the support of distribution P_1 (or P_2), but still differ much on the non-support. In contrast, our result predict that models trained on distribution P_3 in Figure 1c can extrapolate to the distribution Q in Figure 1d, despite that the support of P_3 is sparse and their support have very little overlap.

We also note that the failure of P_1 and P_2 demonstrate the non-triviality of our results. The overlapping marginal assumption by itself does not guarantees extrapolation, and condition (2), or analogously the minimal eigenvalue condition for the Gaussian cases, is critical for extrapolation.

Our proof technique for the theorems above is generally viewing $\|h\|_P^2$ as $K_P(h, h)$ for some kernel $K_P : \mathcal{F} \times \mathcal{F} \rightarrow \mathbb{R}$. Here h is a shorthand for the error function $f - g$. Note that the kernel takes in two functions in \mathcal{F} as inputs and captures relationship between the functions. Hence, the extrapolation of \mathcal{F} (i.e., proving $\|h\|_Q \lesssim \|h\|_P$ for all $h \in \mathcal{F}$) reduces to the relationship of the kernels (i.e., whether $K_Q(h, h) \lesssim K_P(h, h)$ for all $h \in \mathcal{F}$), which is then governed by properties of the eigenspaces of kernel K_P, K_Q . Thanks to the special structure of our model class $\mathcal{F} = \sum f_i(x_i)$, we can analytically relate the eigenspace of the kernels K_P, K_Q to more explicitly and interpretable property of the data distribution of P and Q .

2 PROBLEM SETUP

We use P and Q to denote the source and target distribution over the space of features $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_k$, respectively. We measure the extrapolation of a model class $\mathcal{F} \subseteq \mathbb{R}^{\mathcal{X}}$ from P to Q

by the \mathcal{F} -restricted error ratio, or \mathcal{F} -RER as a shorthand:²

$$\tau(P, Q, \mathcal{F}) \triangleq \sup_{f, g \in \mathcal{F}} \frac{\mathbb{E}_Q[(f(x) - g(x))^2]}{\mathbb{E}_P[(f(x) - g(x))^2]}. \quad (1)$$

When $\tau(P, Q, \mathcal{F})$ is small, if two models $f, g \in \mathcal{F}$ approximately agree on P (meaning that $\mathbb{E}_P[(f(x) - g(x))^2]$ is small), they must approximately agree on Q because $\mathbb{E}_Q[(f(x) - g(x))^2] \leq \tau(P, Q, \mathcal{F})\mathbb{E}_P[(f(x) - g(x))^2]$.

The \mathcal{F} -restricted error ratio monotonically increases as we enlarge the model class \mathcal{F} , and eventually, $\tau(P, Q, \mathcal{F})$ becomes the density ratio between Q and P if the model \mathcal{F} contains all functions with bounded output. To go beyond the bounded density ratio assumption, in this paper we focus on the structured model class $\mathcal{F} = \{\sum_{i=1}^k f_i(x_i) : \mathbb{E}_P[f_i(x_i)^2] < \infty, \forall i \in [k]\}$ where $f_i : \mathcal{X}_i \rightarrow \mathbb{R}$ is an *arbitrary* function. Since \mathcal{F} is closed under addition, we can simplify Eq. (1) to $\tau(P, Q, \mathcal{F}) = \sup_{f \in \mathcal{F}} \frac{\mathbb{E}_Q[f(x)^2]}{\mathbb{E}_P[f(x)^2]}$. For simplicity, we omit the dependency on P, Q, \mathcal{F} when the context is clear.

If the model class \mathcal{F} includes the ground-truth labeling function, $\tau(P, Q, \mathcal{F})$ upperbounds the ratio between the error on distribution Q and the error on distribution P (formally stated in Proposition 1), which provides the robustness guarantee of the trained model. This is because when g corresponds to the ground-truth label, $\mathbb{E}_P[(f(x) - g(x))^2]$ becomes the ℓ_2 error of model f .

Proposition 1. *Let τ be the \mathcal{F} -RER defined in Eq. (1). For any distribution P, Q and model class \mathcal{F} , if there exists a model f^* in \mathcal{F} that can represent the true labeling $y : \mathcal{X} \rightarrow \mathbb{R}$ on both P and Q :*

$$\mathbb{E}_{\frac{1}{2}(P+Q)}[(y(x) - f^*(x))^2] \leq \epsilon_{\mathcal{F}}, \quad (2)$$

then we have

$$\forall f \in \mathcal{F}, \quad \mathbb{E}_Q[(y(x) - f(x))^2] \leq (8\tau + 4)\epsilon_{\mathcal{F}} + 4\tau\mathbb{E}_P[(y(x) - f(x))^2]. \quad (3)$$

Proof of this proposition is deferred to Appendix A.1

Relationship to the $\mathcal{H}\Delta\mathcal{H}$ -distance. Compared with the $\mathcal{H}\Delta\mathcal{H}$ -distance (Ben-David et al., 2010):

$$d_{\mathcal{H}\Delta\mathcal{H}}(P, Q) = 2 \sup_{f, g \in \mathcal{F}} |\Pr_{x \sim P}[f(x) \neq g(x)] - \Pr_{x \sim Q}[f(x) \neq g(x)]|, \quad (4)$$

our differences are: (1) we consider ℓ_2 loss instead of classification loss, and (2) τ focuses on the ratio of losses whereas $d_{\mathcal{H}\Delta\mathcal{H}}$ focuses on the absolute difference. As we will see later, these differences bring the mathematical simplicity to prove concrete conditions for model extrapolation.

Additional notations. Let $\mathbb{I}[E]$ be the indicator function that equals 1 if the condition E is true, and 0 otherwise. For an integer n , let $[n]$ be the set $\{1, 2, \dots, n\}$. For a vector $x \in \mathbb{R}^d$, we use $[x]_i$ to denote its i -th coordinate. Similarly, $[M]_{i,j}$ denotes the (i, j) -th element of a matrix M . We use $M^{\odot n}$ to represent the element-wise n -th power of the matrix M (i.e., $[M^{\odot n}]_{i,j} = ([M]_{i,j})^n$). Let $I_d \in \mathbb{R}^{d \times d}$ be the identity matrix, $\mathbf{1}_d \in \mathbb{R}^d$ the all-1 vector and $e_{i,d}$ the i -th base vector. We omit the subscript d when the context is clear. For a square matrix $P \in \mathbb{R}^{d \times d}$, we use $\text{diag}(P) \in \mathbb{R}^{d \times d}$ to denote the matrix generated by masking out all non-diagonal terms of P . For list $\sigma_1, \dots, \sigma_d$, let $\text{diag}(\{\sigma_1, \dots, \sigma_d\}) \in \mathbb{R}^{d \times d}$ be the diagonal matrix whose diagonal terms are $\sigma_1, \dots, \sigma_d$.

For a symmetric matrix $M \in \mathbb{R}^{d \times d}$, let $\lambda_1(M) \leq \lambda_2(M) \leq \dots \leq \lambda_d(M)$ be its eigenvalues in ascending order, and $\lambda_{\max}(M), \lambda_{\min}(M)$ the maximum and minimum eigenvalue, respectively. Similarly, we use $\sigma_1(M), \dots, \sigma_{\min(d_1, d_2)}(M)$ to denote the singular values of $M \in \mathbb{R}^{d_1 \times d_2}$.

3 MAIN RESULTS

In this section, we present our main results. Section 3.1 discusses the case where the features have discrete values. In Section 3.2 and 3.3, we extend our analysis to two other settings with real-valued features.

²For simplicity, we set $0/0 = 0$.

3.1 FEATURES WITH DISCRETE VALUES

For better exposition, we discuss the case that $x = (x_1, x_2)$ here and defer the discussion of the general case to Appendix A.5. We assume that x_i takes the value in $\{1, 2, \dots, r_i\}$ for $i \in [2]$. Hence, the density of distribution P can be written in a matrix with dimension $r_1 \times r_2$.

We measure the (approximate) clusterability by eigenvalues of the Laplacian matrix of a bipartite graph associated with the density matrix $P \in \mathbb{R}^{r_1 \times r_2}$, which is known to capture the clustering structure of the graph (Chung, 1996; Alon, 1986). Let G_P be a weighted bipartite graph whose adjacency matrix equals the density matrix $P \in \mathbb{R}^{r_1 \times r_2}$ —concretely, $U = \{u_1, \dots, u_{r_1}\}$ and $V = \{v_1, \dots, v_{r_2}\}$ are the sets of vertices, and the weight between u_i, v_j is $P(x_1 = i, x_2 = j)$. To define the signless Laplacian of G_P , let $d_1 \in \mathbb{R}^{r_1}$ and $d_2 \in \mathbb{R}^{r_2}$ be the row and column sums of the weight matrix P (in other words, degree of the vertices), and $D_1 = \text{diag}(d_1) \in \mathbb{R}^{r_1 \times r_1}$, $D_2 = \text{diag}(d_2) \in \mathbb{R}^{r_2 \times r_2}$ the diagonal matrices induced by d_1, d_2 , respectively. The signless Laplacian K_P and normalized signless Laplacian \bar{K}_P are:

$$K_P = \begin{pmatrix} D_1 & P \\ P^\top & D_2 \end{pmatrix}, \quad \bar{K}_P = \text{diag}(K_P)^{-1/2} K_P \text{diag}(K_P)^{-1/2}. \quad (5)$$

Compared with the standard Laplacian matrix, the non-diagonal terms in the signless Laplacian K_P are positive and equal to the absolute value of corresponding terms in the standard Laplacian matrix. In the following theorem, we bound the \mathcal{F} -RER from above by the eigenvalues of \bar{K}_P and the density ratio of the marginal distributions of Q, P .

Theorem 2. *For any distributions P, Q over discrete random variables x_1, x_2 , and the model class $\mathcal{F} = \{f_1(x_1) + f_2(x_2)\}$ where $f_i : \mathcal{X}_i \rightarrow \mathbb{R}$ is an arbitrary function, the \mathcal{F} -RER can be bounded above:*

$$\tau(P, Q, \mathcal{F}) \leq 2\lambda_2(\bar{K}_P)^{-1} \max_{i \in [2], t \in [r_i]} \frac{Q(x_i = t)}{P(x_i = t)}. \quad (6)$$

Compared with prior works that assumes a bounded density ratio on the entire distribution (e.g., Ben-David & Urner (2014); Sugiyama et al. (2007)), we only require a bounded density ratio of the marginal distributions. In other words, the model class $f(x) = f_1(x_1) + f_2(x_2)$ can extrapolate to distributions with a larger support (see Figure 1c). In contrast, for an unstructured model (i.e., $f(x)$ is an arbitrary function of the entire input x), the model can behave arbitrarily on data points outside the support of P .

Qualitatively, Theorem 2 proves sufficient conditions under which the structured model class can extrapolate to unseen distributions (as visualized in Figure 1)—In particular, Theorem 2 implies that for non-trivial extrapolation, that is, $\tau(P, Q, \mathcal{F}) < \infty$, we need (a) $\max_{i \in [2], t \in [r_i]} \frac{Q(x_i = t)}{P(x_i = t)} < \infty$ (i.e., the support of the *marginals* of Q is covered by P), and (b) $\lambda_2(\bar{K}_P) > 0$. To interpret condition (b), note that Cheeger’s inequality implies that $\lambda_2(\bar{K}_P) > 0$ if and only if the bipartite graph G_P is connected (Chung, 1996; Alon, 1986)³, that is, there does not exist non-empty strict subsets of vertices $U' \subset U, V' \subset V$, such that $P(x_1 \in U', x_2 \notin V') = 0$ and $P(x_1 \notin U', x_2 \in V') = 0$. Equivalent, we cannot shuffle the rows and columns of P to form a block diagonal matrix where each block is a strict sub-matrix of P . In other words, the density matrix P is non-clusterable as discussed in Section 1.

Proof sketch of Theorem 2. In the following we present a proof sketch of Theorem 2. We start with a high-level proof strategy and then instantiate the proof on the setting of Theorem 2.

Suppose we can find a set of (not necessarily orthogonal) basis $\{b_1, \dots, b_r\}$ where $b_i : \mathcal{X} \rightarrow \mathbb{R}$, such that any model $f \in \mathcal{F}$ can be represented as a linear combination of basis, that is, $f = \sum_{i=1}^r v_i b_i$. Since the model family \mathcal{F} is closed under subtraction, we have

$$\tau(P, Q, \mathcal{F}) = \sup_{f \in \mathcal{F}} \frac{\|f\|_Q^2}{\|f\|_P^2} = \sup_{v \in \mathbb{R}^r} \frac{\mathbb{E}_Q[(\sum_{i=1}^r v_i b_i(x))^2]}{\mathbb{E}_P[(\sum_{i=1}^r v_i b_i(x))^2]} = \sup_{v \in \mathbb{R}^r} \frac{\sum_{i,j=1}^r [v]_i [v]_j \mathbb{E}_Q[b_i(x) b_j(x)]}{\sum_{i,j=1}^r [v]_i [v]_j \mathbb{E}_P[b_i(x) b_j(x)]}. \quad (7)$$

³Cheeger’s inequality measures the clustering structure of a graph by the eigenvalues of its standard Laplacian. However, the signless Laplacian and standard Laplacian have the same eigenvalues for bipartite graphs Cvetković et al. (2007); Grone et al. (1990).

If we define the kernel matrices $[K_P]_{i,j} = \mathbb{E}_P[b_i(x)b_j(x)]$ and $[K_Q]_{i,j} = \mathbb{E}_Q[b_i(x)b_j(x)]$ (we use the same notation for the kernel matrix and signless Laplacian because later we will show that the kernels K_P, K_Q equal to the signless Laplacian of the bipartite graphs G_P, G_Q with specific choice of the basis b_i), it follows that

$$\sup_{v \in \mathbb{R}^r} \frac{\sum_{i,j=1}^r v_i v_j \mathbb{E}_Q[b_i(x)b_j(x)]}{\sum_{i,j=1}^r v_i v_j \mathbb{E}_P[b_i(x)b_j(x)]} = \sup_{v \in \mathbb{R}^r} \frac{v^\top K_Q v}{v^\top K_P v}. \quad (8)$$

Hence, upper bounding $\tau(P, Q, \mathcal{F})$ reduces to bounding the eigenvalues of kernel matrices K_P, K_Q .

Since the model has the structure $f(x) = f_1(x_1) + f_2(x_2)$, we can construct the basis $\{b_t\}_{t=1}^r$ explicitly. For any $i \in [2], t \in [r_i]$, with little abuse of notation, let $b_{i,t}(x) = \mathbb{I}[x_i = t]$. We can verify that the set $\{b_{i,t}\}_{i \in [2], t \in [r_i]}$ is indeed a complete set of basis. As a result, the kernel matrices K_P can be computed directly using its definition:

$$\mathbb{E}_P[b_{i,t}(x)b_{j,s}(x)] = \begin{cases} P(x_i = t, x_j = s), & \text{when } i \neq j, \\ P(x_i = t)\mathbb{I}[s = t], & \text{when } i = j, \end{cases} \quad (9)$$

which is exactly the Laplacian matrix defined in Eq. (5).

To prove Eq. (6), we need to upperbound the eigenvalues of K_Q . Since the eigenvalues of the normalized signless Laplacian \bar{K}_Q is universally upper bounded by 2 for every distribution Q , we first write K_P, K_Q in terms of \bar{K}_P, \bar{K}_Q . Formally, let $D_P = \text{diag}(K_P)$ and $D_Q = \text{diag}(K_Q)$ and we have

$$\sup_{v \in \mathbb{R}^r} \frac{v^\top K_Q v}{v^\top K_P v} = \sup_{v \in \mathbb{R}^r} \frac{v^\top D_Q^{1/2} \bar{K}_Q D_Q^{1/2} v}{v^\top D_P^{1/2} \bar{K}_P D_P^{1/2} v} \leq \frac{\lambda_{\max}(\bar{K}_Q)}{\lambda_{\min}(\bar{K}_P)} \sup_{v \in \mathbb{R}^r} \frac{\|D_Q^{1/2} v\|_2^2}{\|D_P^{1/2} v\|_2^2}. \quad (10)$$

However, this naive bound is vacuous because for any P we have $\lambda_{\min}(\bar{K}_P) = 0$. In fact, K_P and K_Q share the eigenvalue 0 and the corresponding eigenvector $u \in \mathbb{R}^{r_1+r_2}$ with $[u]_t = (-1)^{\mathbb{I}[t > r_1]}$. Therefore we can restrict to the subspace orthogonal to the direction u , and then $\lambda_{\min}(\bar{K}_P)$ becomes $\lambda_2(\bar{K}_P)$ in Eq. (10). Finally, by basic algebra we also have $\lambda_{\max}(\bar{K}_Q) \leq 2$ and $\sup_{v \in \mathbb{R}^r} \frac{\|D_Q^{1/2} v\|_2^2}{\|D_P^{1/2} v\|_2^2} \leq \max_{i \in [2], t \in [r_i]} \frac{Q(x_i=t)}{P(x_i=t)}$, which complete the proof sketch. The full proof of Theorem 2 is deferred to Appendix A.3.

3.2 FEATURES WITH REAL VALUES

In this section we extend our analysis to the case where x_1, x_2, \dots, x_d are real-valued random variables. Recall that our model has the structure $f(x) = \sum_{i=1}^d f_i(x_i)$ where f_i is an arbitrary one-dimensional function.

When $d = 2$, we can view this setting as a direct extension of the setting in Section 3.1 where each x_i 's has infinite number of possible values (instead of finite number), and thus the Laplacian ‘matrix’ becomes infinite-dimensional. Nonetheless, we can still bound the \mathcal{F} -RER from above, as stated in the following theorem.

Theorem 3. *For any distributions P, Q over variables $x = (x_1, \dots, x_d)$ with matching marginals, assume that (x_i, x_j) has the distribution of a two-dimensional Gaussian random variable for every $i, j \in [d]$. Let $\tilde{x} = (\tilde{x}_1, \dots, \tilde{x}_d)$ be the normalized input where $\tilde{x}_i \triangleq (x_i - \mathbb{E}_P[x_i])\text{Var}(x_i)^{-1/2}$ has zero mean and unit variance for every $i \in [d]$, and $\tilde{\Sigma}_P \triangleq \mathbb{E}_P[\tilde{x}\tilde{x}^\top]$ the covariance matrix of \tilde{x} . Then*

$$\tau(P, Q, \mathcal{F}) \leq \frac{d}{\lambda_{\min}(\tilde{\Sigma}_P)}. \quad (11)$$

For better exposition, we first focus on the case where every x_i has zero mean and unit variance, hence $\tilde{x} = x$ and $\tilde{\Sigma}_P = \Sigma_P \triangleq \mathbb{E}_P[xx^\top]$. Compared with linear models, Theorem 3 proves that the structured nonlinear model class $f(x) = \sum_{i=1}^d f_i(x_i)$ can extrapolate with similar conditions—for linear models $\mathcal{F}_{\text{linear}} \triangleq \{v^\top x : v \in \mathbb{R}^d\}$ we have

$$\tau(P, Q, \mathcal{F}_{\text{linear}}) = \sup_{v \in \mathbb{R}^d} \frac{\|v^\top x\|_Q^2}{\|v^\top x\|_P^2} = \sup_{v \in \mathbb{R}^d} \frac{v^\top \mathbb{E}_Q[xx^\top] v}{v^\top \mathbb{E}_P[xx^\top] v} \lesssim \lambda_{\min}(\Sigma_P)^{-1} = \lambda_{\min}(\tilde{\Sigma}_P)^{-1},$$

which equals to the RHS of Eq. (11) up to factors that does not depend on the covariance Σ_P, Σ_Q .

We emphasize that we only assume the marginals on every *pair* of features x_i, x_j is Gaussian, which does not imply the Gaussianity of the joint distribution of x . In fact, there exists a non-Gaussian distribution that satisfies our assumption.

Proof sketch of Theorem 3. On a high level, we treat the features x_i 's as discrete random variables, and follow the same proof strategy as in Theorem 2. For better exposition, we first assume that x_i has zero mean and unit variance for every $i \in [d]$, hence $\tilde{\Sigma}_P = \Sigma_P \triangleq \mathbb{E}_P[xx^\top]$.

First we consider a simplified case when $d = 2$. Because x_1, x_2 are continuous, the normalized signless Laplacian \bar{K}_P is infinite dimensional, and has the form $\bar{K}_P = \begin{pmatrix} I & A \\ A^\top & I \end{pmatrix}$, where A is an infinite dimensional ‘‘matrix’’ indexed by real numbers $x_1, x_2 \in \mathbb{R}$ with values $[A]_{x_1, x_2} = P(x_1, x_2)/\sqrt{P(x_1)P(x_2)}$, and I is the identity ‘‘matrix’’. Recall the proof of Theorem 2 gives

$$\tau(P, Q, \mathcal{F}) \leq 2\lambda_2(\bar{K}_P)^{-1} \max_{i \in [2], t} \frac{Q(x_i = t)}{P(x_i = t)}. \quad (12)$$

By the assumption that P, Q have matching marginals, we get $\max_{i \in [2], t} \frac{Q(x_i = t)}{P(x_i = t)} = 1$. As result, we only need to lowerbound the second smallest eigenvalue of \bar{K}_P . To this end, we first write A in its singular value decomposition form $A = U\Lambda V^\top$, where $UU^\top = I, VV^\top = I$ and $\Lambda = \text{diag}\{\sigma_n\}_{n \geq 0}$ with $\sigma_0 \geq \sigma_1 \geq \dots$. Then we get

$$\bar{K}_P = \begin{pmatrix} I & A \\ A^\top & I \end{pmatrix} = \begin{pmatrix} U & 0 \\ 0 & V \end{pmatrix} \begin{pmatrix} I & \Lambda \\ \Lambda^\top & I \end{pmatrix} \begin{pmatrix} U^\top & 0 \\ 0 & V^\top \end{pmatrix}. \quad (13)$$

Since the matrix $\hat{K}_P \triangleq \begin{pmatrix} I & \Lambda \\ \Lambda^\top & I \end{pmatrix}$ consists of four diagonal sub-matrices, we can shuffle the rows

and columns of \hat{K}_P to form a block-diagonal matrix with blocks $\left\{ \begin{pmatrix} 1 & \sigma_n \\ \sigma_n & 1 \end{pmatrix} \right\}_{n=0,1,2,\dots}$. As a

result, the eigenvalues of \hat{K}_P are $1 \pm \sigma_0, 1 \pm \sigma_1, \dots$. Because $1 = \sigma_0 \geq \sigma_1 \geq \dots \geq 0$, the smallest and second smallest eigenvalues of \hat{K}_P are $1 - \sigma_0$ and $1 - \sigma_1$, respectively, meaning that $\lambda_2(\bar{K}_P) = \lambda_2(\hat{K}_P) = 1 - \sigma_1$. By the assumption that (x_1, x_2) follows from Gaussian distribution, the ‘‘matrix’’ A is a Gaussian kernel, whose eigenvalues and eigenfunctions can be computed analytically—Theorem 11 proves that $\sigma_1 = |\mathbb{E}_P[x_1 x_2]|$ if x_1, x_2 have zero mean and unit variance. Consequently, $\lambda_2(\bar{K}_P) = 1 - \sigma_1 = 1 - |\mathbb{E}_P[x_1 x_2]| = \lambda_{\min}(\Sigma_P)$.

Now we briefly discuss the case when $d = 3$, and the most general cases (i.e., $d > 3$) are proved similarly. When $d = 3$, the normalized kernel will have the following form

$$\bar{K}_P = \begin{pmatrix} I & A & B \\ A^\top & I & C \\ B^\top & C^\top & I \end{pmatrix}. \quad (14)$$

By the assumption that x_1, x_2 are zero mean and unit variance with joint Gaussian distribution, matrices A is symmetric because $[A]_{x_1, x_2} = P(x_1, x_2)/\sqrt{P(x_1)P(x_2)} = P(x_2, x_1)/\sqrt{P(x_1)P(x_2)} = [A]_{x_2, x_1}$. Similarly, matrices B, C are symmetric. In addition, Theorem 11 shows that the eigenfunctions of the Gaussian kernel is *independent* of the value $\mathbb{E}_P[x_i x_j]$. Hence, the matrices A, B, C shares the same eigenspace and can be diagonalized simultaneously:

$$\bar{K}_P = \begin{pmatrix} I & A & B \\ A^\top & I & C \\ B^\top & C^\top & I \end{pmatrix} = \begin{pmatrix} U & 0 & 0 \\ 0 & U & 0 \\ 0 & 0 & U \end{pmatrix} \begin{pmatrix} I & \Lambda_A & \Lambda_B \\ \Lambda_A^\top & I & \Lambda_C \\ \Lambda_B^\top & \Lambda_C^\top & I \end{pmatrix} \begin{pmatrix} U^\top & 0 & 0 \\ 0 & U^\top & 0 \\ 0 & 0 & U^\top \end{pmatrix}. \quad (15)$$

By reshuffling the columns and rows, the eigenvalues of \bar{K}_P are the union of the eigenvalues of following matrices

$$\{\hat{K}_P^{(n)}\}_{n=0,1,2,\dots} \triangleq \left\{ \begin{pmatrix} 1 & \sigma_n(A) & \sigma_n(B) \\ \sigma_n(A) & 1 & \sigma_n(C) \\ \sigma_n(B) & \sigma_n(C) & 1 \end{pmatrix} \right\}_{n=0,1,2,\dots}. \quad (16)$$

Theorem 11 implies that $\sigma_n(A) = ([\Sigma_P]_{1,2})^n$, $\sigma_n(B) = ([\Sigma_P]_{1,3})^n$ and $\sigma_n(C) = ([\Sigma_P]_{2,3})^n$. Consequently we get $\hat{K}_P^{(n)} = \Sigma_P^{\odot n}$. Then, this theorem follows directly by noticing $\lambda_{\min}(\Sigma_P^{\odot n}) \geq \lambda_{\min}(\Sigma_P)^n$ for $n \geq 1$ (Lemma 13).

Finally, the general case where x_i has arbitrary mean and variance can be reduced to the case where x_i has zero mean and unit variance (Lemma 8). The full proof of Theorem 11 is deferred to Appendix A.6.

3.3 TWO FEATURES WITH MULTI-DIMENSIONAL GAUSSIAN DISTRIBUTION

Now we extend Theorem 3 to the case where $x_1 \in \mathbb{R}^{d_1}$, $x_2 \in \mathbb{R}^{d_2}$ are two subset of features with dimensions $d_1, d_2 > 1$, respectively, and the input $x = (x_1, x_2)$ has Gaussian distribution. Recall that the model class is $\mathcal{F} = \{f_1(x_1) + f_2(x_2) : \mathbb{E}_P[f_i(x_i)^2] < \infty, \forall i \in [2]\}$. In this case, we can still upper bound the \mathcal{F} -RER by the eigenvalues of the covariance matrix Σ_P , which is stated in the following theorem.

Theorem 4. *For any distributions P, Q over variables $x = (x_1, x_2)$ where $x_1 \in \mathbb{R}^{d_1}$, $x_2 \in \mathbb{R}^{d_2}$, let $\Sigma_P = \mathbb{E}_P[xx^\top]$. If $x = (x_1, x_2)$ has Gaussian distribution on both P and Q with zero mean and matching marginals, and $\mathbb{E}_P[x_1x_1^\top] = I$, $\mathbb{E}_P[x_2x_2^\top] = I$, then*

$$\tau(P, Q, \mathcal{F}) \leq \frac{2}{\lambda_{\min}(\Sigma_P)}. \quad (17)$$

We defer the proof of Theorem 4 to Appendix A.7.

Compared with Theorem 3 where the features x_1, \dots, x_d are one-dimensional, our condition for the covariance is almost the same: $\lambda_{\min}(\Sigma_P) > 0$. However, the model class considered Theorem 4 is more powerful because it captures nonlinear interactions between features within the same subset. As a compromise, the assumption on the marginals of P and Q is stronger because Theorem 4 requires matching marginals on each subset of the features, whereas Theorem 3 only requires matching marginals on each individual feature.

Remarks. Our current techniques can only handle the case when the input is divided into $k = 2$ subsets. This is because for $k \geq 3$ we must diagonalize multiple multi-dimensional Gaussian kernels simultaneously using the same set of eigenfunctions, as required in the proof of Theorem 3. However, these multi-dimensional Gaussian kernels do not share the same eigenfunctions because the rotation matrix U, V depends on the covariance $\mathbb{E}_P[x_i x_j^\top]$. Hence, the proof strategy for Theorem 3 fails for the case $k \geq 3$.

4 LOWER BOUNDS

In this section, we prove a lower bound as a motivation to consider structured distributions shifts. The following proposition shows that in the worst case, models learned on P cannot extrapolate to Q when the support of distribution Q is not contained in the support of P .

Proposition 5. *Let the model class \mathcal{F} be the family of two-layer neural networks with ReLU activation: $\mathcal{F} = \{\sum_i a_i \text{ReLU}(w_i^\top x + b_i) : w_i \in \mathbb{R}^d, a_i, b_i \in \mathbb{R}\}$. Suppose for simplicity that all the inputs have unit norm (i.e., $\|x\|_2 = 1$). If Q has non-zero probability mass on the set of points well-separated from the support of P in the sense that*

$$\exists \epsilon > 0, \quad Q(\{x : \|x\|_2 = 1, \text{dist}(x, \text{supp}(P)) \geq \epsilon\}) > 0, \quad (18)$$

we can construct a model $f \in \mathcal{F}$ such that $\|f\|_P = 0$ but $\|f\|_Q$ can be arbitrarily large.

A complete proof of this proposition is deferred to Appendix A.8. On a high level, we prove this proposition by construct a two-layer neural network g_t that represents a bump function around any given input $t \in S^{d-1}$. As a result, when t is a point in $\text{supp}(Q) \setminus \text{supp}(P)$, the model $g_t(x)$ will have zero ℓ_2 norm on P but have a positive ℓ_2 norm on Q . This construction is inspired by Dong et al. (2021, Theorem 5.1).

5 RELATED WORKS

The most related work is Ben-David et al. (2010), where they use the $\mathcal{H}\Delta\mathcal{H}$ -distance to measure the maximum discrepancy of two models $f, g \in \mathcal{F}$ on any distributions P, Q . However, it remains an open question to determine when $\mathcal{H}\Delta\mathcal{H}$ -distance is small for *concrete* nonlinear model classes and distributions. On the technical side, the quantity τ is an analog of the $\mathcal{H}\Delta\mathcal{H}$ -distance for regression problems, and we provide concrete examples where τ is upper bounded even if the distributions P, Q have significantly different support.

Another closely related settings are domain adaptation (Ganin & Lempitsky, 2015; Ghifary et al., 2016; Ganin et al., 2016) and domain generalization (Gulrajani & Lopez-Paz, 2020; Peters et al., 2016), where the algorithms actively improve the extrapolation of learned model either by using unlabeled data from the test domain (Sun & Saenko, 2016; Li et al., 2020a;b; Zhang et al., 2019), or learn an invariant model across different domains (Arjovsky et al., 2019; Peters et al., 2016; Gulrajani & Lopez-Paz, 2020). There are also algorithms that learn features whose distributions on the source and target domain have a small discrepancy measured by the maximum mean discrepancy (Donahue et al., 2014; Long et al., 2015), or the Wasserstein distances (Shen et al., 2018; Courty et al., 2017). In comparison, this paper studies whether a model trained on one distribution (without any implicit bias and unlabeled data from test domain) extrapolates to new distributions. There are also prior works that theoretically analyze algorithms that use additional (unlabeled) data from the test distribution, such as self-training (Wei et al., 2020; Chen et al., 2020), contrastive learning (Shen et al., 2022; HaoChen et al., 2022), label propagation (Cai et al., 2021), etc.

6 CONCLUSIONS

In this paper, we propose to study domain shifts between P and Q with the structure that each feature’s marginal distribution has good overlap between source and target domain but the joint distribution of the features may have a much bigger shift. As a first step toward understanding the extrapolation of nonlinear models, we prove sufficient conditions for the model $f(x) = \sum_{i=1}^k f_i(x_i)$ to extrapolate where f_i is an arbitrary function of a single feature.

Even though the assumptions on the shift and function class is stylized, to the best of our knowledge, this is the first analysis of how nonlinear models extrapolate when source and target distribution *do not* have shared support in concrete settings. There still remain many interesting open questions, which we leave as future works:

- (a) Our current proof can only deal with a restricted nonlinear model family of the special form $f(x) = \sum_{i=1}^k f_i(x_i)$. Can we extend to a more general model class?
- (b) In this paper, we focus on regression tasks with ℓ_2 loss for mathematical simplicity, whereas majority of the prior works focus on the classification problems. Do similar results also hold for classification problem?

ACKNOWLEDGMENT

The authors would like to thank Yuanhao Wang, Yuhao Zhou, Hong Liu, Ananya Kumar, Jason D. Lee, and Kendrick Shen for helpful discussions. The authors would also like to thank the support from NSF CIF 2212263.

REFERENCES

- Alekh Agarwal, Nan Jiang, and Sham M Kakade. Reinforcement learning: Theory and algorithms. *CS Dept., UW Seattle, Seattle, WA, USA, Tech. Rep*, 2019.
- Noga Alon. Eigenvalues and expanders. *Combinatorica*, 6(2):83–96, 1986.
- Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.
- Sara Beery, Arushi Agarwal, Elijah Cole, and Vighnesh Birodkar. The iwildcam 2021 competition dataset. *arXiv preprint arXiv:2105.03494*, 2021.

- Shai Ben-David and Ruth Urner. Domain adaptation—can quantity compensate for quality? *Annals of Mathematics and Artificial Intelligence*, 70(3):185–202, 2014.
- Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79(1-2):151–175, 2010.
- Tianle Cai, Ruiqi Gao, Jason Lee, and Qi Lei. A theory of label propagation for subpopulation shift. In *International Conference on Machine Learning*, pp. 1170–1182. PMLR, 2021.
- Enrico Celeghini, Manuel Gadella, and Mariano A del Olmo. Hermite functions and fourier series. *Symmetry*, 13(5):853, 2021.
- Yining Chen, Colin Wei, Ananya Kumar, and Tengyu Ma. Self-training avoids using spurious features under domain shift. *arXiv preprint arXiv:2006.10032*, 2020.
- Fan RK Chung. Laplacians of graphs and cheeger’s inequalities. *Combinatorics, Paul Erdos is Eighty*, 2(157-172):13–2, 1996.
- Nicolas Courty, Rémi Flamary, Amaury Habrard, and Alain Rakotomamonjy. Joint distribution optimal transportation for domain adaptation. *Advances in Neural Information Processing Systems*, 30, 2017.
- Dragoš Cvetković, Peter Rowlinson, and Slobodan K Simić. Signless laplacians of finite graphs. *Linear Algebra and its applications*, 423(1):155–171, 2007.
- Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. In *International Conference on Machine Learning (ICML)*, volume 32, pp. 647–655, 2014.
- Kefan Dong, Jiaqi Yang, and Tengyu Ma. Provable model-based nonlinear bandit and reinforcement learning: Shelve optimism, embrace virtual curvature. *arXiv preprint arXiv:2102.04168*, 2021.
- Gregory E Fasshauer. Positive definite kernels: past, present and future. *Dolomites Research Notes on Approximation*, 4:21–63, 2011.
- Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pp. 1180–1189. PMLR, 2015.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016.
- Muhammad Ghifary, W Bastiaan Kleijn, Mengjie Zhang, David Balduzzi, and Wen Li. Deep reconstruction-classification networks for unsupervised domain adaptation. In *European conference on computer vision*, pp. 597–613. Springer, 2016.
- Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- Robert Grone, Russell Merris, and VS. Sunder. The laplacian spectrum of a graph. *SIAM Journal on matrix analysis and applications*, 11(2):218–238, 1990.
- Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. *arXiv preprint arXiv:2007.01434*, 2020.
- Jeff Z HaoChen, Colin Wei, Ananya Kumar, and Tengyu Ma. Beyond separability: Analyzing the linear transferability of contrastive representations to related subpopulations. *arXiv preprint arXiv:2204.02683*, 2022.
- Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, Tony Lee, Etienne David, Ian Stavness, Wei Guo, Berton A. Earnshaw, Imran S. Haque, Sara Beery, Jure Leskovec, Anshul Kundaje, Emma Pierson, Sergey Levine, Chelsea Finn, and Percy Liang. WILDS: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning (ICML)*, 2021.

- Ananya Kumar, Tengyu Ma, and Percy Liang. Understanding self-training for gradual domain adaptation. In *International Conference on Machine Learning (ICML)*, 2020.
- Alexey Kurakin, Ian J Goodfellow, and Samy Bengio. Adversarial examples in the physical world. In *Artificial intelligence safety and security*, pp. 99–112. Chapman and Hall/CRC, 2018.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Qi Lei, Wei Hu, and Jason Lee. Near-optimal linear regression under distribution shift. In *International Conference on Machine Learning*, pp. 6164–6174. PMLR, 2021.
- Jingjing Li, Erpeng Chen, Zhengming Ding, Lei Zhu, Ke Lu, and Heng Tao Shen. Maximum density divergence for domain adaptation. *IEEE transactions on pattern analysis and machine intelligence*, 43(11):3918–3930, 2020a.
- Mengxue Li, Yi-Ming Zhai, You-Wei Luo, Peng-Fei Ge, and Chuan-Xian Ren. Enhanced transport distance for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13936–13944, 2020b.
- Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *International conference on machine learning*, pp. 97–105, 2015.
- Mohammadreza Mousavi Kalan, Zalan Fabian, Salman Avestimehr, and Mahdi Soltanolkotabi. Minimax lower bounds for transfer learning with linear and one-hidden layer neural networks. *Advances in Neural Information Processing Systems*, 33:1959–1969, 2020.
- Jonas Peters, Peter Bühlmann, and Nicolai Meinshausen. Causal inference by using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 78:947–1012, 2016.
- Alexander D Poularikas. *Handbook of formulas and tables for signal processing*. CRC press, 2018.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning (ICML)*, volume 139, pp. 8748–8763, 2021.
- Shiori Sagawa, Pang Wei Koh, Tony Lee, Irena Gao, Kendrick Shen Sang Michael Xie, Ananya Kumar, Weihua Hu, Michihiro Yasunaga, Sara Beery Henrik Marklund, Etienne David, Ian Stavness, Wei Guo, Jure Leskovec, Tatsunori Hashimoto Kate Saenko, Sergey Levine, Chelsea Finn, and Percy Liang. Extending the wilds benchmark for unsupervised adaptation. In *International Conference on Learning Representations*, 2022.
- Jian Shen, Yanru Qu, Weinan Zhang, and Yong Yu. Wasserstein distance guided representation learning for domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- Kendrick Shen, Robbie M Jones, Ananya Kumar, Sang Michael Xie, Jeff Z HaoChen, Tengyu Ma, and Percy Liang. Connect, not collapse: Explaining contrastive learning for unsupervised domain adaptation. In *International Conference on Machine Learning*, pp. 19847–19878. PMLR, 2022.
- Masashi Sugiyama, Matthias Krauledat, and Klaus-Robert Müller. Covariate shift adaptation by importance weighted cross validation. *Journal of Machine Learning Research*, 8(May):985–1005, 2007.
- Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *European conference on computer vision*, pp. 443–450. Springer, 2016.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- Colin Wei, Kendrick Shen, Yining Chen, and Tengyu Ma. Theoretical analysis of self-training with deep networks on unlabeled data, 2020. URL <https://openreview.net/forum?id=rC8sJ4i6kaH>.

Yuchen Zhang, Tianle Liu, Mingsheng Long, and Michael I Jordan. Bridging theory and algorithm for domain adaptation. *arXiv preprint arXiv:1904.05801*, pp. 7404–7413, 2019.

Huaiyu Zhu, Christopher KI Williams, Richard Rohwer, and Michal Morciniec. Gaussian regression and optimal finite dimensional linear models. 1997.

A MISSING PROOFS

In the following, we present the missing proofs.

A.1 PROOF OF PROPOSITION 1

In this section, we prove Proposition 1.

Proof of Proposition 1. By the definition of τ , for any $f \in \mathcal{F}$ we get

$$\frac{\mathbb{E}_Q[(f(x) - f^*(x))^2]}{\mathbb{E}_P[(f(x) - f^*(x))^2]} \leq \tau. \quad (19)$$

Or equivalently,

$$\mathbb{E}_Q[(f(x) - f^*(x))^2] \leq \tau \mathbb{E}_P[(f(x) - f^*(x))^2]. \quad (20)$$

As a result,

$$\mathbb{E}_Q[(y - f(x))^2] \leq 2\mathbb{E}_Q[(y - f^*(x))^2] + 2\mathbb{E}_Q[(f(x) - f^*(x))^2] \quad (21)$$

$$\leq 4\epsilon_{\mathcal{F}} + 2\tau \mathbb{E}_P[(f(x) - f^*(x))^2] \leq 4\epsilon_{\mathcal{F}} + 4\tau (\mathbb{E}_P[(y - f(x))^2] + \mathbb{E}_P[(y - f^*(x))^2]) \quad (22)$$

$$\leq (8\tau + 4)\epsilon_{\mathcal{F}} + 4\tau \mathbb{E}_P[(y - f(x))^2]. \quad (23)$$

□

A.2 PROOF OF PROPOSITION 6

In this section, we state and prove Proposition 6.

Proposition 6. *Let $P(x_1, x_2)$ be a distribution over discrete random variables $x_1 \in \mathcal{X}_1, x_2 \in \mathcal{X}_2$, and recall that $\mathcal{F} = \{f_1(x_1) + f_2(x_2)\}$. If there exists subsets $U_1 \subseteq \mathcal{X}_1, U_2 \subseteq \mathcal{X}_2$ such that (1) $0 < |U_1| < |\mathcal{X}_1|$ and $0 < |U_2| < |\mathcal{X}_2|$, and (2) $P(x_1, x_2) = 0$ when $\mathbb{I}[x_1 \in U_1] \neq \mathbb{I}[x_2 \in U_2]$ (e.g., the distribution visualized in Figure 1b), then we can construct models $f, g \in \mathcal{F}$ and $x^\dagger \in \mathcal{X}_1 \times \mathcal{X}_2$ where $\|f - g\|_P = 0$ and $f(x^\dagger) \neq g(x^\dagger)$.*

Proof. Since the model class \mathcal{F} is closed under addition, we only need to construct $h \in \mathcal{F}$ such that $\|h\|_P = 0$ but $h(x^\dagger) \neq 0$.

Let $h_1(x_1) = (-1)^{\mathbb{I}[x_1 \in U_1]}$ and $h_2(x_2) = -(-1)^{\mathbb{I}[x_2 \in U_2]}$. Consider the model $h(x) = h_1(x_1) + h_2(x_2)$. For any x_1, x_2 where $\mathbb{I}[x_1 \in U_1] = \mathbb{I}[x_2 \in U_2]$, we get $h(x) = 0$. Combining with condition (2) we get $\|h\|_P = 0$.

On the other hand, let $x^\dagger = (x_1^\dagger, x_2^\dagger)$ with $x_1^\dagger \in U_1$ and $x_2^\dagger \in \mathcal{X}_2 \setminus U_2$. Such x^\dagger exists because of condition (1). Then $h(x^\dagger) = -2$. □

A.3 PROOF OF THEOREM 2

In this section, we prove Theorem 2.

Proof of Theorem 2. Let

$$b_t(x_1, x_2) = \begin{cases} \mathbb{I}[x_1 = t], & \text{when } t \leq r_1, \\ \mathbb{I}[x_2 = t - r_1], & \text{when } r_1 < t \leq r_1 + r_2. \end{cases} \quad (24)$$

Then for any $f, f' \in \mathcal{F}$, we can always find $v \in \mathbb{R}^{r_1+r_2}$ such that $f(x_1, x_2) - f'(x_1, x_2) = \sum_{i=1}^{r_1+r_2} v_i b_i(x_1, x_2)$ for all x_1, x_2 . Indeed, for any $f, f' \in \mathcal{F}$ the architecture of our model implies $(f - f')(x_1, x_2) = (f_1 - f'_1)(x_1) + (f_2 - f'_2)(x_2)$. Therefore, we can simply set $v_t = (f_1 - f'_1)(t)$ for $1 \leq t \leq r_1$ and $v_t = (f_2 - f'_2)(t - r_1)$ for $r_1 < t \leq r_2$.

Consequently,

$$\tau = \sup_{f, f' \in \mathcal{F}} \frac{\mathbb{E}_Q[(f(x) - f'(x))^2]}{\mathbb{E}_P[(f(x) - f'(x))^2]} = \sup_{v \in \mathbb{R}^{r_1+r_2}} \frac{\mathbb{E}_Q[(\sum_{i=1}^{r_1+r_2} b_i(x_1, x_2) v_i)^2]}{\mathbb{E}_P[(\sum_{i=1}^{r_1+r_2} b_i(x_1, x_2) v_i)^2]} \quad (25)$$

$$= \sup_{v \in \mathbb{R}^{r_1+r_2}} \frac{\sum_{i,j=1}^{r_1+r_2} v_i v_j \mathbb{E}_Q[b_i(x_1, x_2) b_j(x_1, x_2)]}{\sum_{i,j=1}^{r_1+r_2} v_i v_j \mathbb{E}_P[b_i(x_1, x_2) b_j(x_1, x_2)]}. \quad (26)$$

The definition of $b_t(x_1, x_2)$ implies that for any distribution P ,

$$\mathbb{E}_P[b_i(x_1, x_2) b_j(x_1, x_2)] = \begin{cases} \mathbb{I}[i = j] P(x_1 = i), & \text{when } 1 \leq i, j \leq r_1, \\ \mathbb{I}[i = j] P(x_2 = j), & \text{when } r_1 < i, j \leq r_1 + r_2, \\ P(x_1 = i, x_2 = j), & \text{when } 1 \leq i \leq r_1 < j \leq r_1 + r_2, \\ P(x_1 = j, x_2 = i), & \text{when } 1 \leq j \leq r_1 < i \leq r_1 + r_2. \end{cases} \quad (27)$$

Then we have $[K_P]_{i,j} = \mathbb{E}_P[b_i(x_1, x_2) b_j(x_1, x_2)]$. Consequently,

$$\tau = \sup_{v \in \mathbb{R}^{r_1+r_2}} \frac{\sum_{i,j=1}^{r_1+r_2} v_i v_j \mathbb{E}_Q[b_j(x_1, x_2) b_j(x_1, x_2)]}{\sum_{i,j=1}^{r_1+r_2} v_i v_j \mathbb{E}_P[b_j(x_1, x_2) b_j(x_1, x_2)]} = \sup_{v \in \mathbb{R}^{r_1+r_2}} \frac{v^\top K_Q v}{v^\top K_P v}, \quad (28)$$

Let $u \in \mathbb{R}^{r_1+r_2}$ such that

$$[u]_i = \begin{cases} 1, & \text{when } i \leq r_1, \\ -1, & \text{when } r_1 < i \leq r_2. \end{cases}$$

We claim that u is an eigenvector to both K_Q and K_P with eigenvalue 0. To see this, for any distribution P and $i \in [r_1]$ we have

$$[K_P u]_i = P(x_1 = i)[u]_i + \sum_{j=1}^{r_2} P(x_1 = i, x_2 = j)[u]_{r_1+j} \quad (29)$$

$$= P(x_1 = i) - \sum_{j=1}^{r_2} P(x_1 = i, x_2 = j) = 0. \quad (30)$$

Similarly for $i \in [r_2]$ we have $[K_P u]_{r_1+i} = 0$. Combining these two cases together we prove $K_P u = 0$.

Then, by algebraic manipulation,

$$\sup_{v \in \mathbb{R}^{r_1+r_2}} \frac{v^\top K_Q v}{v^\top K_P v} = \sup_{v \in \mathbb{R}^{r_1+r_2}, v \perp u} \frac{v^\top K_Q v}{v^\top K_P v} = \sup_{v \in \mathbb{R}^{r_1+r_2}, v \perp u} \frac{v^\top D_Q^{1/2} \bar{K}_Q D_Q^{1/2} v}{v^\top D_P^{1/2} \bar{K}_P D_P^{1/2} v} \quad (31)$$

$$\leq \frac{\lambda_{\max}(\bar{K}_Q)}{\lambda_2(\bar{K}_P)} \sup_v \frac{\|D_Q^{1/2} v\|_2^2}{\|D_P^{1/2} v\|_2^2} \leq \frac{\lambda_{\max}(\bar{K}_Q)}{\lambda_2(\bar{K}_P)} \max_{i \in [2], t \in [r_i]} \frac{Q(x_i = t)}{P(x_i = t)}. \quad (32)$$

As a result, we only need to prove $\lambda_{\max}(\bar{K}_Q) \leq 2$. To this end, note that

$$\lambda_{\max}(\bar{K}_Q) = \sup_{w \in \mathbb{R}^{r_1+r_2}} \frac{w^\top \bar{K}_Q w}{w^\top w} = \sup_{v \in \mathbb{R}^{r_1+r_2}} \frac{v^\top K_Q v}{v^\top D_Q v} \quad (33)$$

$$= \sup_{v \in \mathbb{R}^{r_1+r_2}} \frac{\sum_{i \in [r_1], j \in [r_2]} P(x_1 = i, x_2 = j) ([v]_i + [v]_{r_1+j})^2}{\sum_{i \in [r_1]} P(x_1 = i) [v]_i^2 + \sum_{j \in [r_2]} P(x_2 = j) [v]_{r_1+j}^2} \quad (34)$$

$$\leq \sup_{v \in \mathbb{R}^{r_1+r_2}} \frac{2 \sum_{i \in [r_1], j \in [r_2]} P(x_1 = i, x_2 = j) ([v]_i^2 + [v]_{r_1+j}^2)}{\sum_{i \in [r_1]} P(x_1 = i) [v]_i^2 + \sum_{j \in [r_2]} P(x_2 = j) [v]_{r_1+j}^2} \quad (35)$$

$$\leq 2. \quad (36)$$

□

A.4 PROOF OF COROLLARY ??

In this section, we prove Corollary ??.

Proof of Corollary ??. Recall that Theorem 2 proves

$$\tau \leq 2\lambda_2(\bar{K}_P)^{-1} \max_{i \in [2], t \in [r_i]} \frac{Q(x_i = t)}{P(x_i = t)}. \quad (37)$$

As a result, we only need to prove conditions (a) and (b) implies the RHS of Eq. (37) is finite.

First of all, by definition condition (a) implies $\max_{i \in [2], t \in [r_i]} \frac{Q(x_i = t)}{P(x_i = t)} < \infty$. Secondly, we prove that $\lambda_2(\bar{K}_P) = 0$ implies the negation of condition (b), which is equivalent to the statement “condition (b) implies $\lambda_2(\bar{K}_P) > 0$ ”. By Cheeger’s inequality (Eq. (??)), $\lambda_2(\bar{K}_P) = 0$ implies that $\phi(G_P) = 0$. By the definition of $\phi(G_P)$, there exists subsets $U_1 \subseteq [r_1], U_2 \subseteq [r_2]$ such that

- (a) $P(x_1 = u_1, x_2 = u_2) = 0$ when $u_1 \in U_1, u_2 \notin U_2$, and
- (b) $P(x_1 = u_1, x_2 = u_2) = 0$ when $u_1 \notin U_1, u_2 \in U_2$.

Then, we can reorganize $P(x_1, x_2)$ by listing the rows corresponding to U_1 first, and the columns corresponding to U_2 first. As a result, $P(x_1, x_2)$ will be a block diagonal matrix where the first block has dimension $|U_1| \times |U_2|$ and the second block $(r_1 - |U_1|) \times (r_2 - |U_2|)$. Consequently, we prove that $\lambda_2(\bar{K}_P) = 0$ implies the negation of condition (b).

Hence, with condition (b) we get $\lambda_2(\bar{K}_P) > 0$. It follows that $\lambda_2(\bar{K}_P)^{-1} < \infty$.

Combining everything together we have $\tau < \infty$. \square

A.5 EXTENDING THEOREM 2 TO MULTIPLE DIMENSIONS ($k > 2$).

In this section, we extend Theorem 2 to the case where $x = (x_1, x_2, \dots, x_k)$ with $k > 2$. Following the same approach, we first define the (signless) Laplacian matrix K_P .

Without loss of generality, we assume that x_i takes the value in $\{1, 2, \dots, r_i\}$ for all $i \in [k]$. Define $r = \sum_{i=1}^k r_i$. For a distribution P , define the matrix $P_{ij} \in \mathbb{R}^{r_i \times r_j}$ with entry $[P_{ij}]_{a,b}$ equals the probability mass $P(x_i = a, x_j = b)$ when $i \neq j$, and diagonal matrix $P_{ii} \in \mathbb{R}^{r_i \times r_i}$ with entry $[P_{ii}]_{a,a}$ equals the probability mass $P(x_i = a)$. In addition, let $K_P, \bar{K}_P \in \mathbb{R}^{r \times r}$ be the matrices defined as follows:

$$K_P = \begin{pmatrix} P_{11} & P_{12} & \cdots & P_{1k} \\ P_{21} & P_{22} & \cdots & P_{2k} \\ \vdots & \vdots & \ddots & \vdots \\ P_{k1} & P_{k2} & \cdots & P_{kk} \end{pmatrix}, \quad \bar{K}_P = \text{diag}(K_P)^{-1/2} K_P \text{diag}(K_P)^{-1/2}. \quad (38)$$

Then the extrapolation power of the nonlinear model class \mathcal{F} is summarized in the following theorem.

Theorem 7. For any two dimensional distribution P, Q over discrete random variables x_1, x_2, \dots, x_k , we have

$$\tau = \sup_{v \in \mathbb{R}^r} \frac{v^\top K_Q v}{v^\top K_P v} \leq k\lambda_k(\bar{K}_P)^{-1} \max_{i \in [k], t \in [r_i]} \frac{Q(x_i = t)}{P(x_i = t)}. \quad (39)$$

Proof of Theorem 7. As discusses in the proof sketch, we first construct a set of basis for the model class $\mathcal{F} = \{\sum_{i=1}^k f_i(x_i)\}$. For any $t \in [r]$, let $\iota(t) \in [k]$ be the index such that $\sum_{j=1}^{\iota(t)-1} r_j < t \leq \sum_{j=1}^{\iota(t)} r_j$ and define $\tau(t) = t - \sum_{j=1}^{\iota(t)-1} r_j$. Then we define

$$b_t(x) = \mathbb{I}[x_{\iota(t)} = \tau(t)]. \quad (40)$$

Then for any $f \in \mathcal{F}$, we can always find $v \in \mathbb{R}^r$ such that $f(x) = \sum_{i=1}^r [v]_i b_i(x)$ for all $x \in \mathcal{X}$. Indeed, for any $f \in \mathcal{F}$ the architecture of our model implies $f(x) = \sum_{i=1}^k f_i(x_i)$. Therefore, we can simply set $v_t = f_{\iota(t)}(\tau(t))$ for $t \in [r]$.

Since \mathcal{F} is closed under addition, we get

$$\tau = \sup_{f, f' \in \mathcal{F}} \frac{\mathbb{E}_Q[(f(x) - f'(x))^2]}{\mathbb{E}_P[(f(x) - f'(x))^2]} = \sup_{f \in \mathcal{F}} \frac{\mathbb{E}_Q[(f(x))^2]}{\mathbb{E}_P[(f(x))^2]} \quad (41)$$

$$= \sup_{v \in \mathbb{R}^r} \frac{\mathbb{E}_Q[(\sum_{i=1}^r b_i(x)[v]_i)^2]}{\mathbb{E}_P[(\sum_{i=1}^r b_i(x)[v]_i)^2]} \quad (42)$$

$$= \sup_{v \in \mathbb{R}^r} \frac{\sum_{i,j=1}^r [v]_i [v]_j \mathbb{E}_Q[b_i(x)b_j(x)]}{\sum_{i,j=1}^r [v]_i [v]_j \mathbb{E}_P[b_i(x)b_j(x)]}. \quad (43)$$

The definition of $b_t(x)$ implies that for any distribution P ,

$$\mathbb{E}_P[b_i(x)b_j(x)] = P(x_{\iota(i)} = \tau(i), x_{\iota(j)} = \tau(j)). \quad (44)$$

Then the kernel matrix is given by $[K_P]_{i,j} = \mathbb{E}_P[b_i(x)b_j(x)]$, which is exactly the definition in Eq. (38). Consequently,

$$\tau = \sup_{v \in \mathbb{R}^r} \frac{\sum_{i,j=1}^r [v]_i [v]_j \mathbb{E}_Q[b_i(x)b_j(x)]}{\sum_{i,j=1}^r [v]_i [v]_j \mathbb{E}_P[b_i(x)b_j(x)]} = \sup_{v \in \mathbb{R}^r} \frac{v^\top K_Q v}{v^\top K_P v}, \quad (45)$$

which proves the first part of Eq. (39).

Now we prove the second part the theorem. To this end, we first characterize the shared eigenvectors of K_P and K_Q .

For any $1 \leq t \leq r-1$, define the vector $u_t \in \mathbb{R}^r$ where

$$[u_t]_i = \begin{cases} 1, & \text{when } \iota(i) = t, \\ -1, & \text{when } \iota(i) = t+1 \\ 0, & \text{o.w.} \end{cases}$$

We claim that u_t is a eigenvector to both K_Q and K_P with eigenvalue 0. To see this, for any distribution P and $i \in [r]$ we have

$$[K_P u_t]_i = \sum_{j:\iota(j)=t} P(x_{\iota(i)} = \tau(i), x_t = \tau(j)) - \sum_{j:\iota(j)=t+1} P(x_{\iota(i)} = \tau(i), x_{t+1} = \tau(j)) \quad (46)$$

$$= P(x_{\iota(i)} = \tau(i)) - P(x_{\iota(i)} = \tau(i)) = 0. \quad (47)$$

Let $U = \text{span}(u_1, \dots, u_{r-1})$ be the linear subspace of such eigenvectors. It follows that

$$u^\top K_P u = u^\top K_Q u = 0, \quad \forall u \in U. \quad (48)$$

Because U is the linear subspace of eigenvectors of K_P (and K_Q) corresponding to eigenvalues 0, which is the minimum eigenvalue of K_P (and K_Q), we get

$$\begin{aligned} \sup_{v \in \mathbb{R}^r} \frac{v^\top K_Q v}{v^\top K_P v} &= \sup_{v \in \mathbb{R}^r, v \perp U} \frac{v^\top K_Q v}{v^\top K_P v} = \sup_{v \in \mathbb{R}^r, v \perp U} \frac{v^\top D_Q^{1/2} \bar{K}_Q D_Q^{1/2} v}{v^\top D_P^{1/2} \bar{K}_P D_P^{1/2} v} \quad (49) \\ &\leq \frac{\lambda_{\max}(\bar{K}_Q)}{\lambda_k(\bar{K}_P)} \sup_v \frac{\|D_Q^{1/2} v\|_2^2}{\|D_P^{1/2} v\|_2^2} \leq \frac{\lambda_{\max}(\bar{K}_Q)}{\lambda_k(\bar{K}_P)} \max_{i \in [k], t \in [r_i]} \frac{Q(x_i = t)}{P(x_i = t)}. \quad (50) \end{aligned}$$

As a result, we only need to prove $\lambda_{\max}(\bar{K}_Q) \leq k$. To this end, for any $i \in [k]$, let $s_i = \sum_{j=1}^{i-1} r_j$. Note that

$$\lambda_{\max}(\bar{K}_Q) = \sup_{u \in \mathbb{R}^r} \frac{u^\top \bar{K}_Q u}{u^\top u} = \sup_{v \in \mathbb{R}^r} \frac{v^\top K_Q v}{v^\top D_Q v} \quad (51)$$

$$= \sup_{v \in \mathbb{R}^r} \frac{\sum_{j_1 \in [r_1], j_2 \in [r_2], \dots, j_k \in [r_k]} P(x_1 = j_1, x_2 = j_2, \dots, x_k = j_k) \left(\sum_{i=1}^k [v]_{s_i+j_i} \mathbb{I}[x_i = j_i] \right)^2}{\sum_{i \in [k], j \in [r_k]} P(x_i = j) [v]_{s_i+j}^2} \quad (52)$$

$$\leq \sup_{v \in \mathbb{R}^r} \frac{k \sum_{j_1 \in [r_1], j_2 \in [r_2], \dots, j_k \in [r_k]} P(x_1 = j_1, x_2 = j_2, \dots, x_k = j_k) \left(\sum_{i=1}^k \mathbb{I}[x_i = j_i] [v]_{s_i+j_i}^2 \right)}{\sum_{i \in [k], j \in [r_k]} P(x_i = j) [v]_{s_i+j}^2} \quad (53)$$

$$= \sup_{v \in \mathbb{R}^r} \frac{k \sum_{i \in [k], j \in [r_k]} P(x_i = j) [v]_{s_i+j}^2}{\sum_{i \in [k], j \in [r_k]} P(x_i = j) [v]_{s_i+j}^2} \quad (54)$$

$$\leq k. \quad (55)$$

□

A.6 PROOF OF THEOREM 3

In this section, we prove Theorem 3.

Proof of Theorem 3. Now consider any fixed model pairs $f, f' \in \mathcal{F}$. Let $g_i(x_i) = f_i(x_i) - f'_i(x_i)$ and $g = \sum_{i=1}^d g_i(x_i)$. Recall that we assume the marginals satisfies $x_i \sim \mathcal{N}(0, 1)$ for every $i \in [d]$. As a result, Lemma 9 implies that there exists coefficients $\{\alpha_i^{(n)}\}_{i \in [d], n \in \mathbb{Z}_+}$ with

$$g_i(x_i)\sqrt{P(x_i)} = \sum_{n \geq 0} \alpha_i^{(n)} \psi_n(x_i), \quad \forall i \in [d], \quad (56)$$

where $\{\psi_n(\cdot)\}_{n \in \mathbb{Z}_+}$ are a set of orthonormal basis of $L^2(\mathbb{R})$ defined by

$$\psi_n(x) \triangleq H_n \left(\frac{1}{\sqrt{2}} x \right) \exp \left(-\frac{1}{4} x^2 \right) (2\pi)^{-1/4} (2^n n!)^{-1/2}.$$

As a result,

$$\mathbb{E}_P \left[\left(\sum_{i=1}^d g_i(x_i) \right)^2 \right] = \int \sum_{i,j=1}^d g_i(x_i) g_j(x_j) P(x_1, \dots, x_d) dx_1 \cdots x_d \quad (57)$$

$$= \sum_{i=1}^d \int g_i(x_i)^2 P(x_i) dx_i + 2 \sum_{1 \leq i < j \leq d} \int g_i(x_i) g_j(x_j) P(x_i, x_j) dx_i x_j \quad (58)$$

$$= \sum_{i=1}^d \int \left(\sum_{n \geq 0} \alpha_i^{(n)} \psi_n(x_i) \right)^2 dx_i \quad (59)$$

$$+ 2 \sum_{1 \leq i < j \leq d} \int (g_i(x_i) \sqrt{P(x_i)}) (g_j(x_j) \sqrt{P(x_j)}) \frac{P(x_i, x_j)}{\sqrt{P(x_i)P(x_j)}} dx_i x_j \quad (60)$$

$$= \sum_{i=1}^d \sum_{n \geq 0} (\alpha_i^{(n)})^2 + 2 \sum_{1 \leq i < j \leq d} \int \left(\sum_{n \geq 0} \alpha_i^{(n)} \psi_n(x_i) \right) \left(\sum_{n \geq 0} \alpha_j^{(n)} \psi_n(x_j) \right) \frac{P(x_i, x_j)}{\sqrt{P(x_i)P(x_j)}} dx_i x_j. \quad (61)$$

By Theorem 11, we have

$$\frac{P(x_i, x_j)}{\sqrt{P(x_i)P(x_j)}} = \sum_{n \geq 0} ([\Sigma_P]_{i,j})^n \psi_n(x_i) \psi_n(x_j). \quad (62)$$

Continuing the second term of Eq. (61) we get

$$\sum_{1 \leq i < j \leq d} \int \left(\sum_{n \geq 0} \alpha_i^{(n)} \psi_n(x_i) \right) \left(\sum_{n \geq 0} \alpha_j^{(n)} \psi_n(x_i) \right) \frac{P(x_i, x_j)}{\sqrt{P(x_i)P(x_j)}} dx_i x_j \quad (63)$$

$$= \sum_{1 \leq i < j \leq d} \int \left(\sum_{n \geq 0} \alpha_i^{(n)} \psi_n(x_i) \right) \left(\sum_{n \geq 0} \alpha_j^{(n)} \psi_n(x_i) \right) \left(\sum_{n \geq 0} ([\Sigma_P]_{i,j})^n \psi_n(x_i) \psi_n(x_j) \right) dx_i x_j \quad (64)$$

$$= \sum_{1 \leq i < j \leq d} \sum_{n \geq 0} \alpha_i^{(n)} \alpha_j^{(n)} ([\Sigma_P]_{i,j})^n. \quad (65)$$

As a result,

$$\mathbb{E}_P \left[\left(\sum_{i=1}^d g_i(x_i) \right)^2 \right] = \sum_{i=1}^d \sum_{n \geq 0} (\alpha_i^{(n)})^2 + 2 \sum_{1 \leq i < j \leq d} \sum_{n \geq 0} \alpha_i^{(n)} \alpha_j^{(n)} ([\Sigma_P]_{i,j})^n. \quad (66)$$

Define $\kappa = \sup_{n \geq 0} \sup_{v \in \mathbb{R}^d} \frac{v^\top (\Sigma_Q)^{\odot n} v}{v^\top (\Sigma_P)^{\odot n} v}$. In the following, we prove

$$\mathbb{E}_Q \left[\left(\sum_{i=1}^d g_i(x_i) \right)^2 \right] \leq \kappa \mathbb{E}_P \left[\left(\sum_{i=1}^d g_i(x_i) \right)^2 \right]. \quad (67)$$

For any fixed $n \geq 0$, let $\Sigma_P^{(n)} \in \mathbb{R}^{d \times d}$ be the matrix where $[\Sigma_P^{(n)}]_{i,j} = ([\Sigma_P]_{i,j})^n$ (define $\Sigma_Q^{(n)}$ similarly), and $\vec{\alpha}^{(n)} \in \mathbb{R}^d$ the vector where $[\vec{\alpha}^{(n)}]_i = \alpha_i^{(n)}$. Then

$$\mathbb{E}_P \left[\left(\sum_{i=1}^d g_i(x_i) \right)^2 \right] = \sum_{n \geq 0} (\vec{\alpha}^{(n)})^\top \Sigma_P^{(n)} \vec{\alpha}^{(n)}, \quad \mathbb{E}_Q \left[\left(\sum_{i=1}^d g_i(x_i) \right)^2 \right] = \sum_{n \geq 0} (\vec{\alpha}^{(n)})^\top \Sigma_Q^{(n)} \vec{\alpha}^{(n)}.$$

By the definition of κ we get $\Sigma_Q^{(n)} \preceq \kappa \Sigma_P^{(n)}$ for all $n \geq 0$. Consequently,

$$\sum_{n \geq 0} (\vec{\alpha}^{(n)})^\top \Sigma_Q^{(n)} \vec{\alpha}^{(n)} \leq \kappa \left(\sum_{n \geq 0} (\vec{\alpha}^{(n)})^\top \Sigma_P^{(n)} \vec{\alpha}^{(n)} \right), \quad (68)$$

which implies Eq. (67).

Now we prove $\kappa \leq \frac{d}{\lambda_{\min}(\Sigma_P)}$. Since Σ_Q is a covariance matrix with $\text{diag}(\Sigma_Q) = I$, we have $\|\Sigma_Q^{\odot n}\|_2 \leq d \|\Sigma_Q^{\odot n}\|_\infty \leq d$. In addition, Lemma 13 proves that $\lambda_{\min}(\Sigma_P^{\odot n}) \geq \lambda_{\min}(\Sigma_P)$ for every $n \geq 1$. As a result,

$$\inf_{n \geq 1} \inf_{v \in \mathbb{R}^d} (v^\top (\Sigma_P)^{\odot n} v)^{-1} \geq \inf_{n \geq 1} \lambda_{\min}(\Sigma_P^{\odot n}) \geq \lambda_{\min}(\Sigma_P).$$

Since $n = 0 \implies (\Sigma_P)^{\odot n} = (\Sigma_Q)^{\odot n}$, we get

$$\kappa = \sup_{n \geq 0} \sup_{v \in \mathbb{R}^d} \frac{v^\top (\Sigma_Q)^{\odot n} v}{v^\top (\Sigma_P)^{\odot n} v} = \max \left\{ 1, \sup_{n \geq 1} \sup_{v \in \mathbb{R}^d} \frac{v^\top (\Sigma_Q)^{\odot n} v}{v^\top (\Sigma_P)^{\odot n} v} \right\} \leq \max \left\{ 1, \frac{d}{\lambda_{\min}(\Sigma_P)} \right\}. \quad (69)$$

By the fact that $\lambda_{\min}(\Sigma_P) \leq d \|\Sigma_P\|_\infty \leq d$, we prove the desired result. \square

In the following, we state and proof Lemma 8.

Lemma 8. *For any distributions P, Q with matching marginals, define P' to be the density of the random variable $t = (t_1, \dots, t_d)$ where $t_i \triangleq \text{var}_P(x_i)^{-1/2}(x_i - \mathbb{E}_P[x_i])$, $x \sim P$ (and define Q' similarly). Then we have*

$$\tau(P, Q, \mathcal{F}) = \tau(P', Q', \mathcal{F}). \quad (70)$$

Proof. Let $\{h_i : \mathbb{R} \rightarrow \mathbb{R}\}_{i=1,2,\dots,d}$ be any set of invertible functions. Since $f = \sum_i f_i(x_i) \in \mathcal{F}$ is equivalent to $f \circ h \triangleq \sum_i f_i(h_i(x_i)) \in \mathcal{F}$, we get

$$\sup_{f \in \mathcal{F}} \frac{\mathbb{E}_Q[\sum_i f_i(x_i)^2]}{\mathbb{E}_P[\sum_i f_i(x_i)^2]} = \sup_{f \in \mathcal{F}} \frac{\mathbb{E}_Q[\sum_i f_i(h_i(x_i))^2]}{\mathbb{E}_P[\sum_i f_i(h_i(x_i))^2]}. \quad (71)$$

Let P' be the density of the random variable $(h_1(x_1), \dots, h_d(x_d))$ when $x \sim P$ (and define Q' similarly). Then we have

$$\frac{\mathbb{E}_{x \sim Q}[\sum_i f_i(h_i(x_i))^2]}{\mathbb{E}_{x \sim P}[\sum_i f_i(h_i(x_i))^2]} = \frac{\mathbb{E}_{t \sim Q'}[\sum_i f_i(t_i)^2]}{\mathbb{E}_{t \sim P'}[\sum_i f_i(t_i)^2]}. \quad (72)$$

Consequently,

$$\tau(P, Q, \mathcal{F}) = \sup_{f \in \mathcal{F}} \frac{\mathbb{E}_Q[\sum_i f_i(x_i)^2]}{\mathbb{E}_P[\sum_i f_i(x_i)^2]} = \sup_{f \in \mathcal{F}} \frac{\mathbb{E}_Q[\sum_i f_i(h_i(x_i))^2]}{\mathbb{E}_P[\sum_i f_i(h_i(x_i))^2]} \quad (73)$$

$$= \sup_{f \in \mathcal{F}} \frac{\mathbb{E}_{t \sim Q'}[\sum_i f_i(t_i)^2]}{\mathbb{E}_{t \sim P'}[\sum_i f_i(t_i)^2]} = \tau(P', Q', \mathcal{F}). \quad (74)$$

Finally, this lemma is proved by taking $h_i(x) = \text{var}_P(x_i)^{-1/2}(x - \mathbb{E}_P[x_i])$. \square

A.7 PROOF OF THEOREM 4

In this section, we prove Theorem 4. We first present a proof sketch of Theorem 4.

Proof sketch of Theorem 4. We start by considering a simpler case when $d_1 = d_2$ and $\Sigma_{12} \triangleq \mathbb{E}_P[x_1 x_2^\top] \in \mathbb{R}^{d_1 \times d_2}$ has only diagonal terms. In other words, the two dimensional random variable $([x_1]_i, [x_2]_i)$ are independent with every other coordinates in the input. In this case, we can decompose the multi-dimensional Gaussian kernel into products of one-dimensional Gaussian kernels in the following way

$$[A]_{x_1, x_2} \triangleq P(x_1, x_2) / \sqrt{P(x_1)P(x_2)} = \prod_{i=1}^{d_1} P([x_1]_i, [x_2]_i) / \sqrt{P([x_1]_i)P([x_2]_i)}. \quad (75)$$

Consequently, the singular values of A will be the products of singular values of these one-dimensional Gaussian kernels: $\{\prod_{i=1}^{d_1} |[\Sigma_{12}]_{i,i}|^{k_i}\}_{k_1, k_2, \dots, k_{d_1} \geq 0}$. Therefore, the second largest singular value of A will be $\max_{i \in [d_1]} |[\Sigma_{12}]_{i,i}|$. Following the same reasoning as Theorem 3 we get

$$\tau(P, Q, \mathcal{F}) \leq \frac{2}{1 - \max_{i \in [d_1]} |[\Sigma_{12}]_{i,i}|} = \frac{2}{1 - \sigma_{\max}(\Sigma_{12})}$$

where $\sigma_{\max}(\Sigma_{12})$ is the largest singular value of Σ_{12} . By noticing that $1 - \sigma_{\max}(\Sigma_{12}) = \lambda_{\min}(\Sigma_P)$ (Lemma 15) we get Eq. (17).

Now we turn to the general cases where Σ_{12} is not diagonal. Recall that our model is $f(x) = f_1(x_1) + f_2(x_2)$ for some arbitrary functions f_1, f_2 . Hence, we can rotate the inputs x_1, x_2 without affecting the model \mathcal{F} -RER. Formally speaking, for any orthogonal matrices $U \in \mathbb{R}^{d_1 \times d_1}$ and $V \in \mathbb{R}^{d_2 \times d_2}$,

$$\sup_{f \in \mathcal{F}} \frac{\mathbb{E}_Q[(f_1(x_1) + f_2(x_2))^2]}{\mathbb{E}_P[(f_1(x_1) + f_2(x_2))^2]} = \sup_{f \in \mathcal{F}} \frac{\mathbb{E}_Q[(f_1(Ux_1) + f_2(Vx_2))^2]}{\mathbb{E}_P[(f_1(Ux_1) + f_2(Vx_2))^2]}. \quad (76)$$

If U, V are the orthonormal matrices in the singular value decomposition $\Sigma_{12} = U^\top \Lambda_{12} V$ where Λ_{12} is a diagonal matrix, we get

$$\mathbb{E}_P[(Ux_1)(Vx_2)^\top] = \Lambda_{12}, \quad \mathbb{E}_P[(Ux_1)^\top(Ux_1)] = I, \quad \mathbb{E}_P[(Vx_2)^\top(Vx_2)] = I. \quad (77)$$

As a result, reusing the result from the previous case on inputs Ux_1, Vx_2 proves the desired result. \square

In the following, we present the full proof.

Proof of Theorem 4. Without loss of generality, we assume $d_2 \leq d_1$. Since \mathcal{F} is closed under addition, we have

$$\tau = \max_{g \in \mathcal{F}} \frac{\mathbb{E}_Q[(g_1(x_1) + g_2(x_2))^2]}{\mathbb{E}_P[(g_1(x_1) + g_2(x_2))^2]}.$$

In the following, we first prove

$$\tau \leq \frac{2}{1 - \sigma_{\max}(\Sigma_{12})} \quad (78)$$

by considering two cases separately.

Case 1: when $g = (g_1, g_2) \in \mathcal{F}$ satisfies $\mathbb{E}_P[g_1(x_1)] = \mathbb{E}_P[g_2(x_2)] = 0$. Since P and Q have matching marginals on x_1 and x_2 respectively, we get

$$\mathbb{E}_Q[g_1(x_1)^2 + g_2(x_2)^2] = \mathbb{E}_Q[g_1(x_1)^2] + \mathbb{E}_Q[g_2(x_2)^2] \quad (79)$$

$$= \mathbb{E}_P[g_1(x_1)^2] + \mathbb{E}_P[g_2(x_2)^2] = \mathbb{E}_P[(g_1(x_1) + g_2(x_2))^2]. \quad (80)$$

As a result, by the definition of τ we have

$$\begin{aligned} & \frac{\mathbb{E}_Q[(g_1(x_1) + g_2(x_2))^2]}{\mathbb{E}_P[(g_1(x_1) + g_2(x_2))^2]} = \frac{\mathbb{E}_Q[(g_1(x_1) + g_2(x_2))^2]}{\mathbb{E}_Q[g_1(x_1)^2 + g_2(x_2)^2]} \frac{\mathbb{E}_P[g_1(x_1)^2 + g_2(x_2)^2]}{\mathbb{E}_P[(g_1(x_1) + g_2(x_2))^2]} \\ & \leq \frac{\mathbb{E}_Q[2(g_1(x_1)^2 + g_2(x_2)^2)]}{\mathbb{E}_Q[g_1(x_1)^2 + g_2(x_2)^2]} \frac{\mathbb{E}_P[g_1(x_1)^2 + g_2(x_2)^2]}{\mathbb{E}_P[(g_1(x_1) + g_2(x_2))^2]} = 2 \frac{\mathbb{E}_P[g_1(x_1)^2 + g_2(x_2)^2]}{\mathbb{E}_P[(g_1(x_1) + g_2(x_2))^2]}. \end{aligned}$$

Hence, we only need to prove

$$(1 - \sigma_{\max}(\Sigma_{12}))\mathbb{E}_P[g_1(x_1)^2 + g_2(x_2)^2] \leq \mathbb{E}_P[(g_1(x_1) + g_2(x_2))^2]. \quad (81)$$

Let $\Sigma_{12} = U\Lambda V^\top$ be the singular value decomposition of Σ_{12} , where $\Lambda \in \mathbb{R}^{d_1 \times d_2}$ is a diagonal matrix with entries $\sigma_1, \dots, \sigma_{d_2}$.

Following Theorem 12, we define $t_1 = U^\top x_1 \in \mathbb{R}^{d_1}$, $t_2 = V^\top x_2 \in \mathbb{R}^{d_2}$. Consequently, $\mathbb{E}_P[t_1 t_2^\top] = \Lambda$ and

$$(t_1, t_2) \sim \mathcal{N}\left(0, \begin{pmatrix} I_{d_1} & \Lambda \\ \Lambda^\top & I_{d_2} \end{pmatrix}\right). \quad (82)$$

Let $P'(t_1, t_2)$ be the density of the distribution of (t_1, t_2) , and $P'(t_1), P'(t_2)$ the density of corresponding marginals. As a result,

$$\mathbb{E}_P[(g_1(x_1) + g_2(x_2))^2] = \mathbb{E}_{P'}[(g_1(t_1) + g_2(t_2))^2] \quad (83)$$

$$= \sum_{i=1}^2 \int g_i(t_i)^2 P'(t_i) dt_i + 2 \int g_1(t_1) g_2(t_2) P'(t_1, t_2) dt_1 dt_2. \quad (84)$$

Since $g_i(t_i) \sqrt{P'(t_i)}$ is square integrable, there exists weights $\{\alpha_i^{(n_1, \dots, n_{d_i})}\}_{n_1, \dots, n_{d_i} \geq 0}$ such that

$$\forall i \in [2], \quad g_i(t_i) \sqrt{P'(t_i)} = \sum_{n_1, \dots, n_{d_i} \geq 0} \alpha_i^{(n_1, \dots, n_{d_i})} \prod_{j=1}^{d_i} \psi_{n_j}([t_i]_j), \quad (85)$$

where $\{\prod_{j=1}^{d_i} \psi_{n_j}([t_i]_j)\}_{n_1, \dots, n_{d_i} \geq 0}$ forms an orthonormal basis of $L^2(\mathbb{R}^{d_i})$. Consequently,

$$\mathbb{E}_P[g_1(x_1)^2 + g_2(x_2)^2] = \mathbb{E}_{P'}[g_1(t_1)^2 + g_2(t_2)^2] \quad (86)$$

$$= \sum_{n_1, \dots, n_{d_1}} \left(\alpha_1^{(n_1, \dots, n_{d_1})}\right)^2 + \sum_{n_1, \dots, n_{d_2}} \left(\alpha_2^{(n_1, \dots, n_{d_2})}\right)^2. \quad (87)$$

Now we turn to the RHS of Eq. (81). Continuing Eq. (84) and apply Theorem 12 we get,

$$\sum_{i=1}^2 \int g_i(t_i)^2 P'(t_i) dt_i = \sum_{i=1}^2 \sum_{n_1, \dots, n_{d_i}} \left(\alpha_i^{(n_1, \dots, n_{d_i})}\right)^2$$

and

$$\begin{aligned} & \int g_1(t_1) g_2(t_2) P'(t_1, t_2) dt_1 dt_2 \\ &= \int \prod_{i=1}^2 \left(\sum_{n_1, \dots, n_{d_1} \geq 0} \alpha_i^{(n_1, \dots, n_{d_i})} \prod_{j=1}^{d_i} \psi_{n_j}([t_i]_j) \right) \prod_{i=1}^{d_2} \left(\sum_{n=0}^{\infty} \sigma_i^n \psi_n([t_1]_i) \psi_n([t_2]_i) \right) \prod_{i=d_2+1}^{d_1} \psi_0([t_1]_i) dt_1 dt_2 \\ &= \sum_{n_1, \dots, n_{d_2} \geq 0} \alpha_1^{(n_1, \dots, n_{d_2}, 0, \dots, 0)} \alpha_2^{(n_1, \dots, n_{d_2})} \prod_{i=1}^{d_2} \sigma_i^{n_i}. \end{aligned}$$

As a result,

$$\mathbb{E}_{P'}[(g_1(t_1) + g_2(t_2))^2] \quad (88)$$

$$= \sum_{i=1}^2 \sum_{n_1, \dots, n_{d_i}} \left(\alpha_i^{(n_1, \dots, n_{d_i})}\right)^2 + 2 \sum_{n_1, \dots, n_{d_2} \geq 0} \alpha_1^{(n_1, \dots, n_{d_2}, 0, \dots, 0)} \alpha_2^{(n_1, \dots, n_{d_2})} \prod_{i=1}^{d_2} \sigma_i^{n_i} \quad (89)$$

$$= \sum_{n_1, \dots, n_{d_2}} \begin{pmatrix} \alpha_1^{(n_1, \dots, n_{d_2}, 0, \dots, 0)} \\ \alpha_2^{(n_1, \dots, n_{d_2})} \end{pmatrix}^\top \begin{pmatrix} 1 & \prod_{i=1}^{d_2} \sigma_i^{n_i} \\ \prod_{i=1}^{d_2} \sigma_i^{n_i} & 1 \end{pmatrix} \begin{pmatrix} \alpha_1^{(n_1, \dots, n_{d_2}, 0, \dots, 0)} \\ \alpha_2^{(n_1, \dots, n_{d_2})} \end{pmatrix}$$

$$+ \sum_{n_1, \dots, n_{d_1}} \mathbb{I}[n_{d_2+1} + \dots + n_{d_1} \neq 0] \left(\alpha_i^{(n_1, \dots, n_{d_i})} \right)^2. \quad (90)$$

Now for any $n_1, \dots, n_{d_2} \geq 0$, consider the matrix

$$\begin{pmatrix} 1 & \prod_{i=1}^{d_2} \sigma_i^{n_i} \\ \prod_{i=1}^{d_2} \sigma_i^{n_i} & 1 \end{pmatrix}.$$

When $n_1 + \dots + n_{d_2} \neq 0$, we have $\prod_{i=1}^{d_2} \sigma_i^{n_i} \leq \max_{i \in [d_2]} \sigma_i$ because Lemma 14 implies $\sigma_i \leq 1$ for every $i \in [d_2]$. Therefore in this case,

$$\begin{pmatrix} \alpha_1^{(n_1, \dots, n_{d_2}, 0, \dots, 0)} \\ \alpha_2^{(n_1, \dots, n_{d_2})} \end{pmatrix}^\top \begin{pmatrix} 1 & \prod_{i=1}^{d_2} \sigma_i^{n_i} \\ \prod_{i=1}^{d_2} \sigma_i^{n_i} & 1 \end{pmatrix} \begin{pmatrix} \alpha_1^{(n_1, \dots, n_{d_2}, 0, \dots, 0)} \\ \alpha_2^{(n_1, \dots, n_{d_2})} \end{pmatrix} \quad (91)$$

$$\geq (1 - \sigma_{\max}(\Sigma_{12})) \left(\left(\alpha_1^{(n_1, \dots, n_{d_2}, 0, \dots, 0)} \right)^2 + \left(\alpha_2^{(n_1, \dots, n_{d_2})} \right)^2 \right). \quad (92)$$

When $n_1 + \dots + n_{d_2} = 0$, by Eq. (85) we get

$$\alpha_1^{(0, \dots, 0)} = \int g_1(t_1) \sqrt{P'(t_1)} \prod_{j=1}^{d_1} \psi_0([t_1]_j) dt_1 \quad (93)$$

$$= \int g_1(t_1) \sqrt{P'(t_1)} \prod_{j=1}^{d_1} \sqrt{P'([t_1]_j)} dt_1 \quad (94)$$

$$= \int g_1(t_1) P'(t_1) dt_1 = \mathbb{E}_{P'}[g_1(t_1)] = \mathbb{E}_P[g_1(x_1)] = 0. \quad (95)$$

Similarly, $\alpha_2^{(0, \dots, 0)} = 0$. As a result, when $n_1 + \dots + n_{d_2} = 0$ we have

$$\begin{pmatrix} \alpha_1^{(n_1, \dots, n_{d_2}, 0, \dots, 0)} \\ \alpha_2^{(n_1, \dots, n_{d_2})} \end{pmatrix}^\top \begin{pmatrix} 1 & \prod_{i=1}^{d_2} \sigma_i^{n_i} \\ \prod_{i=1}^{d_2} \sigma_i^{n_i} & 1 \end{pmatrix} \begin{pmatrix} \alpha_1^{(n_1, \dots, n_{d_2}, 0, \dots, 0)} \\ \alpha_2^{(n_1, \dots, n_{d_2})} \end{pmatrix} \quad (96)$$

$$= \left(\left(\alpha_1^{(n_1, \dots, n_{d_2}, 0, \dots, 0)} \right)^2 + \left(\alpha_2^{(n_1, \dots, n_{d_2})} \right)^2 \right) = 0. \quad (97)$$

Combining these two cases together, we get

$$\sum_{n_1, \dots, n_{d_2}} \begin{pmatrix} \alpha_1^{(n_1, \dots, n_{d_2}, 0, \dots, 0)} \\ \alpha_2^{(n_1, \dots, n_{d_2})} \end{pmatrix}^\top \begin{pmatrix} 1 & \prod_{i=1}^{d_2} \sigma_i^{n_i} \\ \prod_{i=1}^{d_2} \sigma_i^{n_i} & 1 \end{pmatrix} \begin{pmatrix} \alpha_1^{(n_1, \dots, n_{d_2}, 0, \dots, 0)} \\ \alpha_2^{(n_1, \dots, n_{d_2})} \end{pmatrix} \quad (98)$$

$$\geq (1 - \sigma_{\max}(\Sigma_{12})) \left(\sum_{n_1, \dots, n_{d_2}} \left(\alpha_1^{(n_1, \dots, n_{d_2}, 0, \dots, 0)} \right)^2 + \left(\alpha_2^{(n_1, \dots, n_{d_2})} \right)^2 \right). \quad (99)$$

Plug in to Eq. (90) we get

$$\mathbb{E}_P[(g_1(x_1) + g_2(x_2))^2] = \mathbb{E}_{P'}[(g_1(t_1) + g_2(t_2))^2] \quad (100)$$

$$\geq (1 - \sigma_{\max}(\Sigma_{12})) \left(\sum_{n_1, \dots, n_{d_1}} \left(\alpha_1^{(n_1, \dots, n_{d_1})} \right)^2 + \sum_{n_1, \dots, n_{d_2}} \left(\alpha_2^{(n_1, \dots, n_{d_2})} \right)^2 \right) \quad (101)$$

$$= (1 - \sigma_{\max}(\Sigma_{12})) \mathbb{E}_P[g_1(x_1)^2 + g_2(x_2)^2], \quad (102)$$

which finished the proof for Eq. (81).

Case 2: when $g = (g_1, g_2) \in \mathcal{F}$ satisfies $\mathbb{E}_P[g_1(x_1) + g_2(x_2)] = 0$. Let $\mu = \mathbb{E}_P[g_1(x_1)]$. Define $g' = (g'_1, g'_2)$ where $g'_1(x_1) = g_1(x_1) - \mu$, $g'_2(x_2) = g_2(x_2) + \mu$. By the definition of \mathcal{F} we get $g' \in \mathcal{F}$, and g' satisfies $\mathbb{E}_P[g_1(x_1)] = \mathbb{E}_P[g_2(x_2)] = 0$. Because $g_1(x_1) + g_2(x_2) = g'_1(x_1) + g'_2(x_2)$ for all $x_1 \in \mathbb{R}^{d_1}$, $x_2 \in \mathbb{R}^{+d_2}$, plugging in the result of case 1 we get

$$\frac{\mathbb{E}_Q[(g_1(x_1) + g_2(x_2))^2]}{\mathbb{E}_P[(g_1(x_1) + g_2(x_2))^2]} = \frac{\mathbb{E}_Q[(g'_1(x_1) + g'_2(x_2))^2]}{\mathbb{E}_P[(g'_1(x_1) + g'_2(x_2))^2]} \leq \frac{1}{1 - \sigma_{\max}(\Sigma_{12})}. \quad (103)$$

Case 3: when $g = (g_1, g_2) \in \mathcal{F}$. Now we consider the most general case. Let $\nu = \mathbb{E}_P[g_1(x_1) + g_2(x_2)]$. Since P and Q have matching marginals on x_1 and x_2 respectively, we get

$$\mathbb{E}_Q[g_1(x_1) + g_2(x_2)] = \mathbb{E}_Q[g_1(x_1)] + \mathbb{E}_Q[g_2(x_2)] \quad (104)$$

$$= \mathbb{E}_P[g_1(x_1)] + \mathbb{E}_P[g_2(x_2)] = \mathbb{E}_P[g_1(x_1) + g_2(x_2)] = \nu. \quad (105)$$

Define $g' = (g'_1, g'_2)$ where $g'_1(x_1) = g_1(x_1) - \nu$, $g'_2(x_2) = g_2(x_2)$. It follows that

$$\frac{\mathbb{E}_Q[(g_1(x_1) + g_2(x_2))^2]}{\mathbb{E}_P[(g_1(x_1) + g_2(x_2))^2]} = \frac{\mathbb{E}_Q[(g'_1(x_1) + g'_2(x_2))^2] + \nu^2}{\mathbb{E}_P[(g'_1(x_1) + g'_2(x_2))^2] + \nu^2} \quad (106)$$

$$\leq \max \left\{ \frac{\mathbb{E}_Q[(g'_1(x_1) + g'_2(x_2))^2]}{\mathbb{E}_P[(g'_1(x_1) + g'_2(x_2))^2]}, 1 \right\} \leq \frac{1}{1 - \sigma_{\max}(\Sigma_{12})}, \quad (107)$$

where the last inequality comes from applying case 2's result to g' .

Combining these three cases together we prove Eq. (78). Then Eq. (17) follows directly from Eq. (78) and Lemma 15. \square

A.8 PROOF OF PROPOSITION 5

In this section, we prove Proposition 5.

Proof of Proposition 5. First, for any $t \in S^{d-1}$ and $\epsilon > 0$ we construct a two-layer neural network $h_{t,\epsilon}(x)$ such that

- (a) for all $x \in S^{d-1}$, $h_{t,\epsilon}(x) = 0$ if $\|x - t\|_2 \geq \epsilon$,
- (b) for all $x \in S^{d-1}$, $h_{t,\epsilon}(x) \geq 1$ if $\|x - t\|_2 \leq \epsilon/2$, and
- (c) for all $x \in S^{d-1}$, $h_{t,\epsilon}(x) \geq 0$.

Recall that a two layer neural network is parameterized as $\sum_i a_i \text{ReLU}(w_i^\top x + b_i)$. Then $h_{t,\epsilon}(x)$ can be constructed using one neuron by setting $w_1 = t$, $b_1 = -1 + \epsilon^2/2$ and $a_1 = 8/(3\epsilon^2)$.

We can verify the construction as follows. When $\|x - t\|_2 \geq \epsilon$, we get

$$\|x - t\|_2^2 \geq \epsilon^2 \implies 2 - 2x^\top t \geq \epsilon^2 \implies x^\top t \leq 1 - \frac{\epsilon^2}{2}. \quad (108)$$

As a result,

$$a_1 \text{ReLU}(w_1^\top x + b_1) = \frac{8}{3\epsilon^2} \text{ReLU}(t^\top x - 1 + \epsilon^2/2) = 0.$$

When $\|x - t\|_2 \leq \epsilon/2$, we get

$$\|x - t\|_2^2 \leq \epsilon^2/4 \implies 2 - 2x^\top t \leq \epsilon^2/4 \implies x^\top t \geq 1 - \frac{\epsilon^2}{8}. \quad (109)$$

As a result,

$$a_1 \text{ReLU}(w_1^\top x + b_1) = \frac{8}{3\epsilon^2} \text{ReLU}(t^\top x - 1 + \epsilon^2/2) \geq 1.$$

Now we construct a two layer neural network f such that $\|f\|_P = 0$ but $\|f\|_Q = 0$. Let $X = \{x : x \in S^{d-1}, \text{dist}(x, \text{supp}(P)) \geq \epsilon\}$. By the condition of this proposition we have $Q(X) > 0$.

Let C be the minimum $\epsilon/2$ -covering of the set X , and thus $|C| < \infty$. Consider the 2-layer neural network $f(x) = \sum_{c \in C} h_{c,\epsilon}$. It follows from the definition of X that $\|x - c\|_2 \geq \epsilon$ for every $x \in \text{supp}(P)$. Consequently, $f(x) = \sum_{c \in C} h_{c,\epsilon} = 0$ for every $x \in \text{supp}(P)$, which implies that $\|f\|_P = 0$.

Now for every $x \in \mathcal{X}$, there exists $t \in C$ such that $\|x - t\|_2 \leq \epsilon/2$. As a result, $h_{t,\epsilon}(x) \geq 1$, which implies that

$$f(x) = \sum_{c \in C} h_{c,\epsilon} \geq h_{t,\epsilon}(x) \geq 1. \quad (110)$$

Hence,

$$\mathbb{E}_Q[f(x)^2] \geq \mathbb{E}_Q[\mathbb{I}[x \in X]] = Q(X) > 0. \quad (111)$$

Since $cf(x) \in \mathcal{F}$ for every $c > 0$, we prove the desired result. \square

B HERMITE POLYNOMIAL AND GAUSSIAN KERNEL

The Hermite polynomial $H_n(x)$ is an degree n polynomial defined as follows,

$$H_n(x) = (-1)^n \exp(x^2) \frac{d^n \exp(-t^2)}{dt^n}, \quad \forall n = 1, 2, \dots \quad (112)$$

with the following orthogonality property (Poularikas, 2018):

$$\int \exp(-x^2/2) H_n(x/\sqrt{2}) H_m(x/\sqrt{2}) dx = \mathbb{I}[n = m] 2^n n! \sqrt{2\pi}. \quad (113)$$

As a result, we can construct a set of orthonormal basis of square integrable functions $L^2(\mathbb{R})$ using the Hermite polynomial.

Lemma 9 (see e.g., Celeghini et al. (2021)). *The set of functions $\{\psi_n(x)\}_{n \geq 0}$ form an orthonormal basis for $L^2(\mathbb{R})$, where*

$$\psi_n(x) \triangleq H_n\left(\frac{1}{\sqrt{2}}x\right) \exp\left(-\frac{1}{4}x^2\right) (2\pi)^{-1/4} (2^n n!)^{-1/2}.$$

Theorem 10 (Chapter 6.2 of Fasshauer (2011), also see Zhu et al. (1997)). *For any $\epsilon > 0, \alpha > 0$, the eigenfunction expansion for the Gaussian $x, z \in \mathbb{R}$ is*

$$\exp(-\epsilon^2(x-z)^2) = \sum_{n=0}^{\infty} \lambda_n \phi_n(x) \phi_n(z), \quad (114)$$

where

$$\lambda_n = \frac{\alpha \epsilon^{2n}}{\left(\frac{\alpha^2}{2} \left(1 + \sqrt{1 + (2\epsilon/\alpha)^2}\right) + \epsilon^2\right)^{n+1/2}}, \quad n = 0, 1, 2, \dots,$$

$$\phi_n(x) = \frac{(1 + (2\epsilon/\alpha)^2)^{1/8}}{\sqrt{2^n n!}} \exp\left(-\left(\sqrt{1 + (2\epsilon/\alpha)^2} - 1\right) \frac{\alpha^2 x^2}{2}\right) H_n\left(\left(1 + (2\epsilon/\alpha)^2\right)^{1/4} \alpha x\right).$$

And the eigenfunctions ϕ_n forms an orthonormal basis under weighted $L^2(\mathbb{R})$ space:

$$\int \phi_m(x) \phi_n(x) \frac{\alpha}{\sqrt{\pi}} \exp(-\alpha^2 x^2) dx = \mathbb{I}[m = n]. \quad (115)$$

Theorem 11. *For any $\rho \in [-1, 1]$, let $P(x_1, x_2)$ be the density of $(x_1, x_2) \sim \mathcal{N}\left(0, \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}\right)$ and $P(x_1), P(x_2)$ the density of corresponding marginals. Then the Gaussian kernel*

$$K_\rho(x_1, x_2) = \frac{P(x_1, x_2)}{\sqrt{P(x_1)P(x_2)}} \quad (116)$$

$$= \frac{1}{\sqrt{2\pi(1-\rho^2)}} \exp\left(-\frac{1}{2(1-\rho^2)}(x_1^2 + x_2^2 - 2\rho x_1 x_2) + \frac{x_1^2}{4} + \frac{x_2^2}{4}\right) \quad (117)$$

has the following eigen-decomposition

$$K_\rho(x_1, x_2) = \sum_{k=0}^{\infty} \rho^k \psi_k(x_1) \psi_k(x_2), \quad (118)$$

where $\psi_k(\cdot)$ is defined in Lemma 9.

Proof. We prove this theorem by considering the following two cases separately.

Case 1: $\rho \geq 0$. By algebraic manipulation we get

$$\sqrt{2\pi}K_\rho(x_1, x_2) = \frac{1}{(1-\rho^2)^{1/2}} \exp\left(-\frac{1}{2(1-\rho^2)}\left(\rho(x_1-x_2)^2 + \frac{(\rho-1)^2}{2}x_1^2 + \frac{(\rho-1)^2}{2}x_2^2\right)\right).$$

Let $\epsilon^2 = \frac{\rho}{2(1-\rho^2)}$ and $\alpha^2 = \frac{(\rho-1)^2}{2(1-\rho^2)}$, we can equivalent write

$$\sqrt{2\pi}K_\rho(x_1, x_2) = \frac{1}{(1-\rho^2)^{1/2}} \exp(-\epsilon^2(x_1-x_2)^2) \exp\left(-\frac{\alpha^2}{2}x_1^2\right) \exp\left(-\frac{\alpha^2}{2}x_2^2\right). \quad (119)$$

By Theorem 10 we get

$$\exp(-\epsilon^2(x_1-x_2)^2) = \sum_n \lambda_n \phi_n(x_1) \phi_n(x_2) \quad (120)$$

with $\int \phi_m(x) \phi_n(x) \frac{\alpha}{\sqrt{\pi}} \exp(-\alpha^2 x^2) dx = \mathbb{I}[m=n]$. Then we can define the function

$$\psi_n(x) = \left(\frac{\alpha}{\sqrt{\pi}}\right)^{1/2} \phi_n(x) \exp\left(-\frac{\alpha^2}{2}x^2\right) \quad (121)$$

such that $\int \psi_n(x) \psi_m(x) dx = \mathbb{I}[m=n]$. Combining Eqs. (119), (120), and (121) we get

$$K_\rho(x_1, x_2) = \sum_n \left(\lambda_n \frac{1}{\sqrt{2\alpha}(1-\rho^2)^{1/2}}\right) \psi_n(x_1) \psi_n(x_2). \quad (122)$$

Now we only need to prove that $\psi_n(\cdot)$ defined in Eq. (121) has the same form as those in Lemma 9, and $\lambda_n \frac{1}{\sqrt{2\alpha}(1-\rho^2)^{1/2}} = \rho^n$.

Recall that

$$\phi_n(x) = \frac{(1+(2\epsilon/\alpha)^2)^{1/8}}{\sqrt{2^n n!}} \exp\left(-\left(\sqrt{1+(2\epsilon/\alpha)^2}-1\right)\frac{\alpha^2 x^2}{2}\right) H_n\left(\left(1+(2\epsilon/\alpha)^2\right)^{1/4} \alpha x\right).$$

Plugin $\epsilon^2 = \frac{\rho}{2(1-\rho^2)}$ and $\alpha^2 = \frac{(\rho-1)^2}{2(1-\rho^2)}$ we get

$$(1+(2\epsilon/\alpha)^2)^{1/8} \left(\frac{\alpha}{\sqrt{\pi}}\right)^{1/2} = \frac{1}{\pi^{1/4}} (\alpha^4 + 4\epsilon^2 \alpha^2)^{1/8} \quad (123)$$

$$= \frac{1}{\pi^{1/4}} \left(\frac{(\rho-1)^4 + 4\rho(\rho-1)^2}{4(1-\rho^2)^2}\right)^{1/8} = \frac{1}{(2\pi)^{1/4}} \quad (124)$$

$$\exp\left(-\frac{\alpha^2}{2}x^2\right) \exp\left(-\left(\sqrt{1+(2\epsilon/\alpha)^2}-1\right)\frac{\alpha^2 x^2}{2}\right) \quad (125)$$

$$= \exp\left(-\sqrt{\alpha^4 + 2\epsilon^2 \alpha^2} \frac{x^2}{2}\right) = \exp\left(-\sqrt{\frac{(\rho-1)^4 + 4\rho(\rho-1)^2}{4(1-\rho^2)^2}} \frac{x^2}{2}\right) \quad (126)$$

$$= \exp\left(-\frac{x^2}{4}\right), \quad (127)$$

$$(1+(2\epsilon/\alpha)^2)^{1/4} \alpha = \left(1 + \frac{4\rho}{(\rho-1)^2}\right)^{1/4} \left(\frac{(\rho-1)^2}{2(1-\rho^2)}\right)^{1/2} \quad (128)$$

$$= \left(\frac{(1+\rho)^2 (\rho-1)^4}{(\rho-1)^2 (1-\rho^2)^2}\right)^{1/4} \frac{1}{\sqrt{2}} = \frac{1}{\sqrt{2}}. \quad (129)$$

As a result,

$$\psi_n(x) = \left(\frac{\alpha}{\sqrt{\pi}}\right)^{1/2} \phi_n(x) \exp\left(-\frac{\alpha^2}{2}x^2\right) \quad (130)$$

$$= H_n \left(\frac{1}{\sqrt{2}} x \right) \exp \left(-\frac{1}{4} x^2 \right) (2\pi)^{-1/4} (2^n n!)^{-1/2}. \quad (131)$$

Now we turn to the eigenvalues. Recall that

$$\lambda_n = \frac{\alpha \epsilon^{2n}}{\left(\frac{\alpha^2}{2} \left(1 + \sqrt{1 + (2\epsilon/\alpha)^2} \right) + \epsilon^2 \right)^{n+1/2}}.$$

Plugin $\epsilon^2 = \frac{\rho}{2(1-\rho^2)}$ and $\alpha^2 = \frac{(\rho-1)^2}{2(1-\rho^2)}$ we get

$$\frac{\alpha^2}{2} \left(1 + \sqrt{1 + (2\epsilon/\alpha)^2} \right) + \epsilon^2 = \frac{1}{2} \left(\alpha^2 + \sqrt{\alpha^4 + 4\epsilon^2 \alpha^2} \right) + \epsilon^2 \quad (132)$$

$$= \frac{1}{2} \left(\frac{(\rho-1)^2}{2(1-\rho^2)} + \frac{1}{2} \right) + \frac{\rho}{2(1-\rho^2)} = \frac{1}{2(1-\rho^2)}. \quad (133)$$

Consequently,

$$\lambda_n \frac{1}{\sqrt{2}\alpha(1-\rho^2)^{1/2}} = \frac{\left(\frac{\rho}{2(1-\rho^2)} \right)^n}{\left(\frac{1}{2(1-\rho^2)} \right)^{n+1/2}} \frac{1}{\sqrt{2}(1-\rho^2)^{1/2}} = \rho^n. \quad (134)$$

Case 2: $\rho < 0$. Recall that

$$K_\rho(x_1, x_2) = \frac{1}{\sqrt{2\pi(1-\rho^2)}} \exp \left(-\frac{1}{2(1-\rho^2)} (x_1^2 + x_2^2 - 2\rho x_1 x_2) + \frac{x_1^2}{4} + \frac{x_2^2}{4} \right).$$

In this case we reuse the results from Case 1 to get

$$K_{-\rho}(x_1, x_2) = \sum_{k=0}^{\infty} (-\rho)^k \psi_k(x_1) \psi_k(x_2). \quad (135)$$

By definition, we get $K_\rho(x_1, x_2) = K_{-\rho}(x_1, -x_2)$. As a result,

$$K_\rho(x_1, x_2) = \sum_{k=0}^{\infty} (-\rho)^k \psi_k(x_1) \psi_k(-x_2) \quad (136)$$

$$= \sum_{k=0}^{\infty} \rho^k \psi_k(x_1) \psi_k(x_2), \quad (137)$$

where the last equation follows from the fact that $\psi_k(x_2) = (-1)^k \psi_k(-x_2)$. \square

Theorem 12. For any $\Sigma \in \mathbb{R}^{d_1 \times d_2}$, let $P(x_1, x_2)$ be the density of $(x_1, x_2) \sim \mathcal{N} \left(0, \begin{pmatrix} I_{d_1} & \Sigma \\ \Sigma^\top & I_{d_2} \end{pmatrix} \right)$ and $P(x_1), P(x_2)$ the density of corresponding marginals, where $x_1 \in \mathbb{R}^{d_1}, x_2 \in \mathbb{R}^{d_2}$. Without loss of generality, assume $d_1 \geq d_2$. Let $\Sigma = U\Lambda V^\top$ be the singular value decomposition of Σ , where $U \in \mathbb{R}^{d_1 \times d_2}, \Lambda \in \mathbb{R}^{d_1 \times d_2}, V \in \mathbb{R}^{d_2 \times d_2}$. Let $\sigma_1, \dots, \sigma_{d_2}$ be the singular values of Σ (i.e., $[\Lambda]_{i,i} = \sigma_i$). Then the Gaussian kernel

$$K_\Sigma(x_1, x_2) = \frac{P(x_1, x_2)}{\sqrt{P(x_1)P(x_2)}} \quad (138)$$

has the following eigen-decomposition

$$K_\Sigma(x_1, x_2) = \prod_{i=1}^{d_2} \left(\sum_{n=0}^{\infty} \sigma_i^n \psi_n([t_1]_i) \psi_n([t_2]_i) \right) \prod_{i=d_2+1}^{d_1} \psi_0([t_1]_i). \quad (139)$$

where $t_1 = U^\top x_1, t_2 = V^\top x_2$, and $\psi_k(\cdot)$ is defined in Lemma 9.

Proof. We prove this theorem by decomposing the multi-dimensional kernel into products of several one-dimensional kernels described in Theorem 11.

Recall that $t_1 = U^\top x_1 \in \mathbb{R}^{d_1}, t_2 = V^\top x_2 \in \mathbb{R}^{d_2}$, where $\Sigma = U\Lambda V^\top$ is the singular value decomposition of Σ . Let $t \in \mathbb{R}^{d_1+d_2}$ be the concatenation of t_1, t_2 . Then we have

$$\mathbb{E}_P[tt^\top] = \begin{pmatrix} I_{d_1} & \Lambda \\ \Lambda^\top & I_{d_2} \end{pmatrix}. \quad (140)$$

Let $P'(t_1, t_2)$ be the density of $(t_1, t_2) \sim \mathcal{N}\left(0, \begin{pmatrix} I_{d_1} & \Lambda \\ \Lambda^\top & I_{d_2} \end{pmatrix}\right)$. By the fact that Λ is a diagonal matrix, we know that variables $\{[t_1]_i, [t_2]_i\}$ is independent from $\{[t_1]_j\}_{j \neq i} \cup \{[t_2]_j\}_{j \neq i}$ for every $1 \leq i \leq d_2$, and variable $[t_1]_i$ is independent from $\{[t_1]_j\}_{j \neq i} \cup \{[t_2]_j\}_{1 \leq j \leq d_2}$ for every $d_2 + 1 \leq i \leq d_1$. In addition, $([t_1]_i, [t_2]_i) \sim \mathcal{N}\left(0, \begin{pmatrix} 1 & \sigma_i \\ \sigma_i & 1 \end{pmatrix}\right)$. As a result,

$$K_\Sigma(x_1, x_2) = \frac{P(x_1, x_2)}{\sqrt{P(x_1)P(x_2)}} = \frac{P'(t_1, t_2)}{\sqrt{P'(t_1)P'(t_2)}} \quad (141)$$

$$= \prod_{i=1}^{d_2} \left(\frac{P'([t_1]_i, [t_2]_i)}{\sqrt{P'([t_1]_i)P'([t_2]_i)}} \right) \prod_{i=d_2+1}^{d_1} \sqrt{P'([t_1]_i)} \quad (142)$$

$$= \prod_{i=1}^{d_2} \left(\sum_{n=0}^{\infty} \sigma_i^n \psi_n([t_1]_i) \psi_n([t_2]_i) \right) \prod_{i=d_2+1}^{d_1} \sqrt{P'([t_1]_i)} \quad (143)$$

$$= \prod_{i=1}^{d_2} \left(\sum_{n=0}^{\infty} \sigma_i^n \psi_n([t_1]_i) \psi_n([t_2]_i) \right) \prod_{i=d_2+1}^{d_1} \psi_0([t_1]_i), \quad (144)$$

where Eq. (143) follows from Theorem 12, and Eq. (144) follows from the definition of ψ_0 . \square

C HELPER LEMMAS

Lemma 13. Let $\Sigma \in \mathbb{R}^{d \times d}$ be a positive semi-definite matrix with $\text{diag}(\Sigma) = I$. Then we have

$$\forall k \geq 1, \quad \lambda_{\min}(\Sigma^{\odot k}) \geq \lambda_{\min}(\Sigma), \quad (145)$$

where $\Sigma^{\odot k}$ denotes the element-wise k -th power of the matrix Σ .

Proof. Let $\lambda = \lambda_{\min}(\Sigma)$ and we have $\Sigma - \lambda I \succeq 0$. As a result, $(\Sigma - \lambda I)^{\odot k} = \Sigma^{\odot k} + (1 - \lambda)^k I - I \succeq 0$. Note that $0 \leq \lambda \leq 1$ because $\text{Tr}(\Sigma - I) = 0$ implies $\lambda_{\min}(\Sigma - I) \leq 0$. It follows that for any $k \geq 1$,

$$\Sigma^{\odot k} \succeq I - (1 - \lambda)^k I \succeq \lambda I, \quad (146)$$

where the last inequality follows from the fact that $1 - \lambda \geq (1 - \lambda)^k$ when $k \geq 1$. Consequently,

$$\lambda_{\min}(\Sigma^{\odot k}) \geq \lambda_{\min}(\Sigma). \quad (147)$$

\square

Lemma 14. Let $\Sigma \in \mathbb{R}^{d_1 \times d_2}$ be a matrix such that

$$\begin{pmatrix} I_{d_1} & \Sigma \\ \Sigma^\top & I_{d_2} \end{pmatrix} \succeq 0. \quad (148)$$

Then we have $\lambda_{\max}(\Sigma^\top \Sigma) \leq 1$.

Proof. We prove by contradiction. Suppose otherwise $\lambda_{\max}(\Sigma^\top \Sigma) > 1$. Let $v \in \mathbb{R}^{d_2}$ be the eigenvector of $\Sigma^\top \Sigma$ corresponds to its maximum eigenvalue. Then

$$\begin{pmatrix} -\Sigma v \\ v \end{pmatrix}^\top \begin{pmatrix} I_{d_1} & \Sigma \\ \Sigma^\top & I_{d_2} \end{pmatrix} \begin{pmatrix} -\Sigma v \\ v \end{pmatrix} = \begin{pmatrix} -\Sigma v \\ v \end{pmatrix}^\top \begin{pmatrix} 0 \\ -\Sigma^\top \Sigma v + v \end{pmatrix} = -v^\top \Sigma^\top \Sigma v + v^\top v. \quad (149)$$

By the assumption that $\lambda_{\max}(\Sigma^\top \Sigma) > 1$, we get

$$-v^\top \Sigma^\top \Sigma v + v^\top v = \|v\|_2^2 (1 - \lambda_{\max}(\Sigma^\top \Sigma)) < 0, \quad (150)$$

which contradicts to Eq. (148). Therefore, we must have $\lambda_{\max}(\Sigma^\top \Sigma) \leq 1$. \square

Lemma 15. *Let M be a positive semi-definite matrix with the following form*

$$M = \begin{pmatrix} I_{d_1} & \Sigma \\ \Sigma^\top & I_{d_2} \end{pmatrix} \quad (151)$$

for some $d_1, d_2 > 0$. Then we have

$$\lambda_{\min}(M) = 1 - \sigma_{\max}(\Sigma), \quad (152)$$

where σ_{\max} is the largest singular value of Σ .

Proof. Without loss of generality, we assume $d_1 \geq d_2$. Let $\Sigma = U\Lambda V^\top$ be the singular value decomposition of Σ , with $U \in \mathbb{R}^{d_1 \times d_1}$, $V \in \mathbb{R}^{d_2 \times d_2}$, $\Lambda \in \mathbb{R}^{d_1 \times d_2}$. Then we have

$$M = \begin{pmatrix} U & 0 \\ 0 & V \end{pmatrix} \begin{pmatrix} I_{d_1} & \Lambda \\ \Lambda^\top & I_{d_2} \end{pmatrix} \begin{pmatrix} U^\top & 0 \\ 0 & V^\top \end{pmatrix}. \quad (153)$$

Let $\bar{M} = \begin{pmatrix} I_{d_1} & \Lambda \\ \Lambda^\top & I_{d_2} \end{pmatrix}$. Because U, V are orthonormal matrices, we get $\lambda_{\min}(M) = \lambda_{\min}(\bar{M})$. Note that $[\Lambda]_{i,j} = 0$ whenever $i \neq j$ and $[\Lambda]_{i,i} = \sigma_i(\Sigma)$. Consequently, the eigenvalues of \bar{M} is

$$1 \pm \sigma_1(\Sigma), 1 \pm \sigma_2(\Sigma), \dots, 1 \pm \sigma_{d_2}(\Sigma) \quad (154)$$

with multiplicity 1, and 1 with multiplicity $d_1 - d_2$. As a result,

$$\lambda_{\min}(M) = \lambda_{\min}(\bar{M})1 - \sigma_{\max}(\Sigma). \quad \square$$

D EXPERIMENT DETAILS

In this section, we present the experiments that support our theory. In section D.1, we substantiate the setting of Theorem 4 using synthetic datasets. Section D.2 shows empirical evidences that supports the Gaussianity assumption required by Theorem 3.

D.1 EXPERIMENTS WITH SYNTHETIC DATASETS

We focus on the setting of Theorem 5 where $x = (x_1, x_2)$ and $x_1 \sim \mathbb{R}^{d_1}, x_2 \sim \mathbb{R}^{d_2}$. We compare the extrapolation of the structured model $f_1(x_1) + f_2(x_2)$ (where f_1, f_2 are two-layer neural networks) and unstructured model $f(x)$ (where f is a two-layer neural network with input $[x_1, x_2] \in \mathbb{R}^{d_1+d_2}$). The ground-truth label is generated by a random structured model $f_1^*(x_1) + f_2^*(x_2)$. Recall that our theory shows that the structured model provably extrapolates to the target distribution. Indeed, empirically we found that the structured model has a significantly smaller *OOD loss* than the unstructured model (even though the *ID losses* are comparable), as shown in Table 1. Therefore, the structured model family indeed has a better extrapolation than the unstructured family.

Model Class	ID loss	OOD loss
Structured model $f_1(x_1) + f_2(x_2)$	0.00100	0.00118
Unstructured model $f(x)$	0.00105	0.01235

Table 1: The ID and OOD losses of different model classes. The structured model has a significantly smaller *OOD loss* than the unstructured model even though the *ID losses* are comparable.

D.1.1 IMPLEMENTATION DETAILS.

The synthetic dataset. The features to the dataset is $x = (x_1, x_2)$ where $x_1 \sim \mathbb{R}^{d_1}, x_2 \sim \mathbb{R}^{d_2}$. We set $d_1 = d_2 = 256$. The ground-truth labeling function is $f^*(x) = f_1^*(x_1) + f_2^*(x_2)$, where f_1^*, f_2^* are two layer neural networks with ReLU activation, hidden dimension 128, and random initialization. For simplicity, labels are noiseless. During the training, we draw fresh samples on every iteration.

For both the ID and OOD distribution, the feature vector x has Gaussian distribution with covariance Σ_P, Σ_Q . To match the setting of Theorem 4, the covariance Σ_P, Σ_Q are generated in the following way:

$$\Sigma_P = \begin{pmatrix} I & \gamma O_P \\ \gamma O_P^\top & I \end{pmatrix}, \Sigma_Q = \begin{pmatrix} I & \gamma O_Q \\ \gamma O_Q^\top & I \end{pmatrix}, \quad (155)$$

where O_P, O_Q are two random orthonormal matrices, and $\gamma = 0.9$. Consequently, $\Sigma_P \succeq O(1)\Sigma_Q$.

The model classes. The structured model is of the form $f_1(x_1) + f_2(x_2)$, where f_1, f_2 are two-layer neural networks with hidden dimension 512 and ReLU activation. The parameters of f_1, f_2 are initialized by $\mathcal{N}(0, \sigma^2)$ with $\sigma = 10^{-3}$. We train the model with batch size 1024, 500 batches per epoch, and 200 epochs in total. We use SGD with learning rate $3 * 10^{-3}$ and momentum 0.9.

The unstructured model is a two layer neural network on the concatenation of x_1 and x_2 with hidden dimension 1024 and ReLU activation. Therefore, the unstructured model is a superset of the structured model (with hidden dimension 512). We train the unstructured model for 235 epochs to get a similar ID loss.⁴ We keep other hyper-parameters exactly the same as the structured model case.

D.1.2 ABLATION STUDY.

The choice of hidden dimensions. We first vary the hidden dimensions of the unstructured model to test whether the differences of structured and unstructured model come from the number of parameters. Table 2 summarizes the result — with a smaller model, the ground-truth labeling function cannot be expressed by the model class (since the ID loss doesn’t converge to 0); with a larger model, the OOD loss becomes larger.

Model Class	ID loss	OOD loss
Unstructured model $f(x)$, hidden dim = 256	0.26957	0.8759
Unstructured model $f(x)$, hidden dim = 512	0.01688	0.07774
Unstructured model $f(x)$, hidden dim = 1024	0.00105	0.01235
Unstructured model $f(x)$, hidden dim = 2048	0.00129	0.01395

Table 2: The ID and OOD losses of unstructured models with different size.

The effect of regularization. We also add ℓ_1 and ℓ_2 regularization to the unstructured model, since the structured model is 1/2 sparse when viewed as a unstructured model. However, even though the regularization helps OOD loss to some extent, the OOD loss of the unstructured model is still far worse than the OOD loss of the structured model (Table 3).

Model Class	ID loss	OOD loss
Unstructured model $f(x)$, ℓ_1 regularization	0.00116	0.00899
Unstructured model $f(x)$, ℓ_2 regularization	0.00142	0.01623

Table 3: The ID and OOD losses of unstructured models with different regularizations.

D.2 JUSTIFICATION OF THE GAUSSIANTITY ASSUMPTION

In this section, we justify the Gaussianity assumption required by Theorem 3.

On the theory side, our pairwise Gaussian assumption is substantially weaker than the joint Gaussian assumption required by some prior works on the theory of domain shift (e.g., Chen et al. (2020);

⁴Because the unstructured model has more parameters, it trains slower than the structured model.

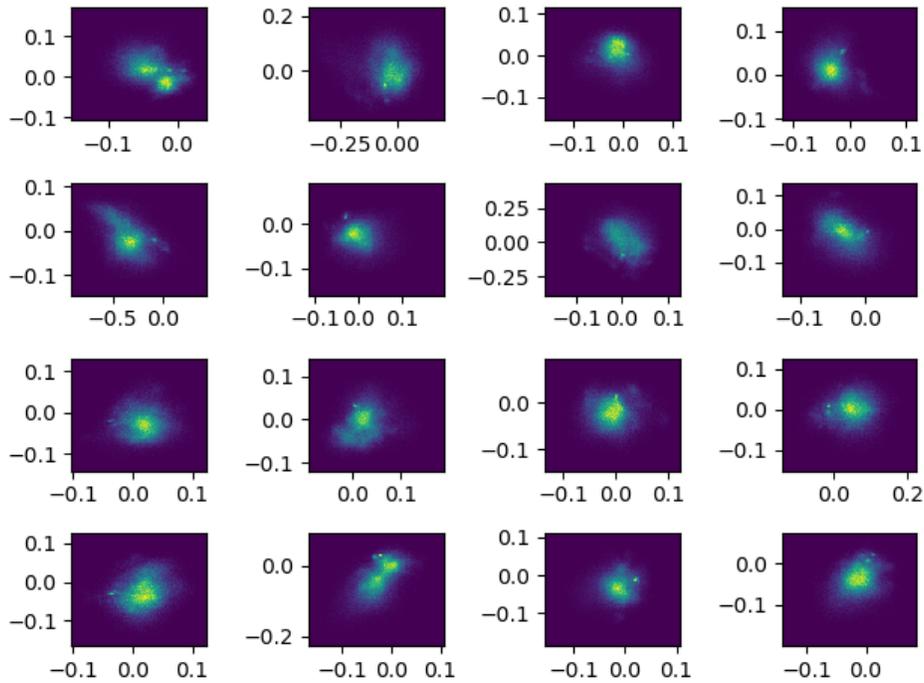


Figure 2: Visualization of the distribution of random feature pairs (x_i, x_j) .

Kumar et al. (2020)). Even with the Gaussianity assumption, existing results cannot prove any positive results for extrapolation to a distribution with different support.

On the empirical side, we can test whether the pairwise Gaussianity assumption holds approximately. On the iwildcam dataset (Beery et al., 2021), we take the output of a Resnet50 pretrained with CLIP dataset as our features x (Radford et al., 2021), and we visualize the distributions of (x_i, x_j) for randomly selected (i, j) pairs in Figure 2. The Gaussianity assumption approximately holds in this case.