

Deep Nonparametric Estimation of Operators between Infinite Dimensional Spaces ^{*}

Hao Liu, Haizhao Yang, Minshuo Chen, Tuo Zhao, and Wenjing Liao [†]

January 2, 2022

Abstract

Learning operators between infinitely dimensional spaces is an important learning task arising in wide applications in machine learning, imaging science, mathematical modeling and simulations, etc. This paper studies the nonparametric estimation of Lipschitz operators using deep neural networks. Non-asymptotic upper bounds are derived for the generalization error of the empirical risk minimizer over a properly chosen network class. Under the assumption that the target operator exhibits a low dimensional structure, our error bounds decay as the training sample size increases, with an attractive fast rate depending on the intrinsic dimension in our estimation. Our assumptions cover most scenarios in real applications and our results give rise to fast rates by exploiting low dimensional structures of data in operator estimation. We also investigate the influence of network structures (e.g., network width, depth, and sparsity) on the generalization error of the neural network estimator and propose a general suggestion on the choice of network structures to maximize the learning efficiency quantitatively.

1 Introduction

Learning nonlinear operators from a Hilbert space to another via nonparametric estimation has been an important topic with broad applications. For example, in reduced-order modeling, a data-driven approach desires to map a full model trajectory to a reduced model trajectory or vice versa [70]. In solving parametric partial differential equations (PDEs), it is desired to learn a map from the parametric function space to the PDE solution space [45, 54, 60]. In forward and inverse scattering problems [46, 92], it is interesting to learn an operator mapping the observed data function space to the parametric function space that models the underlying PDE. In density functional theory, it is desired to learn a nonlinear operator mapping a potential function to a density function [28]. In phase retrieval [22], an operator from the observed data function space

^{*}on going work.

[†]Hao Liu is affiliated with the Math department of Hong Kong Baptist University; Haizhao Yang is affiliated with the Math department of Purdue University; Minshuo Chen and Tuo Zhao are affiliated with the ISYE department at Georgia Tech; Wenjing Liao is affiliated with the Math department at Georgia Tech; Haizhao Yang and Wenjing Liao are co-corresponding authors. Email: haoliu@hkbu.edu.hk, haizhao@purdue.edu, {mchen393, tzhao80, wliao60}@gatech.edu.

to the reconstructed image function space is learned. Other image processing problems, e.g., image super-resolution [71], image denoising [89], image inpainting [72], are similar to the deep learning-based phase retrieval, where an operator from a function space to another function space is learned.

As a powerful tool of nonparametric estimation, deep learning [33] has made astonishing breakthroughs in various applications, including computer vision [50], natural language processing [34], speech recognition [39], healthcare [62], as well as nonlinear operator learning [45, 97, 28, 29, 46, 15, 60, 51, 7, 54, 66]. A typical method for operator learning is to first discretize the function spaces and represent each function by a vector using sampling. Then deep neural networks are applied to learn the map between these vector spaces [45, 97, 28, 29, 46]. Such methods are mesh dependent: if a different discretization scheme is used, the network needs to be trained again. Though empirical successes have been demonstrated in learning nonlinear operators by this approach in many applications, it is computationally expensive to train these algorithms and the training procedure has to be repeated when the dimension of vector spaces is changed. Another approach based on the theory of approximating operators by neural networks [15] can alleviate this issue to a certain extent by avoiding the discretization of the output Hilbert space of the operator. This approach was first proposed in [15] with two-layer neural networks and recently revisited with deeper neural networks in [60] with successful applications [55, 9]. However, the methods in [15, 60, 55, 9] are still mesh-dependent due to the requirement of a fixed number of sample points for the input function of the operator. More recently, a discretization-invariant (mesh-independent) operator learning method was proposed for problems with a sparsity structure in [1, 7, 54, 66] by taking the advantage of graph kernel networks, principal component analysis (PCA), and kernel integral operators, etc. With discretization-invariant approaches, the training procedure does not need to be performed again when the discretization scheme changes.

Although operator learning via deep learning-based nonparametric estimation has been successful in many applications, its statistical learning theory is still in its infancy, especially when the operator is from an infinite dimensional space to another. The successes of deep neural networks are largely due to their universal approximation power [20, 40], showing the existence of a neural network with a proper size fulfilling the approximation task for certain function classes. Quantitative function approximation theories, provably better than traditional tools, have been extensively studied with various network architectures and activation functions, e.g., for continuous functions [93, 79, 81, 82, 83, 95], for functions with certain smoothness [94, 96, 59, 87], and for functions with integral representations [3, 25, 26, 85]. In theory, deep neural networks can approximate certain high dimensional functions with a fast rate that is independent of the input dimension [3, 25, 26, 85, 81, 82, 96, 80, 15, 13, 14, 56, 44, 17, 77, 74, 23, 64]. However, in the context of operator approximation, deep learning theory is very limited. Probably the first result is the universal approximation theorem for operators in [15]. More recently, quantitative approximation results for operators between infinite dimensional spaces were given in [7, 51, 49] based on the function approximation theory in [93]. Note that the function approximation results in [93] does not give the flexibility to choose arbitrary width and depth of neural networks. In this paper, we provide a new operator approximation theory based on nearly optimal function approxima-

tion results where the width and depth of the network can be chosen flexibly. In comparison with [51], the flexibility of choosing arbitrary width and depth provides an explicit guideline to balance the approximation error and the statistical variance to achieve a better generalization error in operator learning.

We also establish a novel statistical theory for deep nonparametric estimation of Lipschitz operators between infinite dimensional Hilbert spaces. The core question to be answered is: how the generalization error scales when the number of training samples increases and whether the scaling is dimension-independent without the curse of dimensionality. In literature, the statistical theory for function regression via neural networks has been a popular research topic [38, 47, 43, 5, 75, 10, 13, 48, 65, 30, 56, 44]. These works have proved that deep nonparametric regression can achieve the optimal minimax rate of regression established in [86, 36]. When the target function has low complexity or the function domain is a low dimensional set, deep neural networks can achieve a fast rate depending on the intrinsic dimension [12, 13, 14, 56, 79, 44, 17, 77, 74, 23, 64]. In more sophisticated cases when a mathematical modeling problem is transferred to a special regression problem, e.g., solving high dimensional PDEs and identifying the governing equation of spatial-temporal data, the generalization analysis of deep learning has been proposed in [6, 84, 61, 63, 58, 57, 24, 35]. All these results focus on the regression problem when the target function is a mapping from a finite dimensional space to a finite dimensional space. Therefore, these results cannot be applied to mappings from an infinite dimensional space to another. To our best knowledge, the only work on the generalization error analysis of deep operator learning in Hilbert spaces is [51] for the algorithm in [60], which is not completely discretization-invariant. The generalization error in [51] is a posteriori depending on the properties of neural networks fitting the target operator. Recently, the posterior rates on learning linear operators by Bayesian inversion have been studied in [21].

In this paper, we establish a priori generalization error for a discretization-invariant operator learning algorithm for operators between Hilbert spaces. As we shall see later, our theory can be applied to operator learning from a finite dimensional vector space to another as a special case. Therefore, the theoretical result in this paper can facilitate the understanding of many operator learning algorithms by neural networks in the literature. Our contributions are summarized as follows:

1. We derive an upper bound on the generalization error for a general framework of learning operators between infinite dimensional spaces by deep neural networks. The framework considered here first encodes the input and output space into finite-dimensional spaces by some encoders and decoders. Then a transformation between the dimension reduced spaces is learned using deep neural networks. Our upper bound is derived for two network architectures: one has constraints on the number of nonzero weight parameters and parameter magnitude; The other network architecture does not have such constraints and allows one to flexibly choose the depth and width. Our upper bound consists of two parts: the error of learning the transformation by deep neural networks, and the dimension reduction error with encoders and decoders.
2. Our analysis is general and can be applied for a wide range of popular choices of encoders

and decoders in the numerical implementation, such as those derived from Legendre polynomials, trigonometric bases, and principal component analysis. The generalization error is given for each of these examples.

3. We discuss two scenarios to further exploit the additional low-dimensional structures of data in operator estimation motivated by practical considerations and classical numerical methods. The first scenario is when encoded vectors in the input space are on a low-dimensional manifold. In this scenario, we show that the generalization error converges as the training sample increases with a fast rate depending on the intrinsic dimension of the manifold. The second scenario is when the operator itself has low complexity: the composition of the operator with a certain encoder and decoder is a multi-index model. In this scenario, we show that the convergence rate of the generalization error depends on the intrinsic dimension of the composed operator.

We organize this paper as follows. In Section 2, we introduce our notations and the learning framework considered in this paper. Our main results with general encoders and decoders are presented in Section 3. We discuss the applications of our main results to specific encoders and decoders derived from certain function basis and PCA in Section 4 and 5, respectively. To further exploit additional low-dimensional structures of data, we discuss the application of our results to two scenarios in Section 6. The proofs of all results are given in Section 7. We conclude this paper in Section 8.

2 A general framework

2.1 Preliminaries

We first briefly introduce some definitions and notations on a Hilbert space, encoders, decoders, and feedforward neural networks used in this paper. A Hilbert space is a Banach space equipped with an inner product. It is separable if it admits a countable orthonormal basis. Let \mathcal{H} be a separable Hilbert space. An encoder for \mathcal{H} is an operator $E_{\mathcal{H}} : \mathcal{H} \rightarrow \mathbb{R}^d$, where d is a positive integer representing the encoding dimension. The associated decoder is an operator $D_{\mathcal{H}} : \mathbb{R}^d \rightarrow \mathcal{H}$. The composition $\Pi_{\mathcal{H}} = D_{\mathcal{H}} \circ E_{\mathcal{H}} : \mathcal{H} \rightarrow \mathcal{H}$ is a projection. For any $u \in \mathcal{H}$, we define the projection error as $\|\Pi_{\mathcal{H}}(u) - u\|_{\mathcal{H}}$.

In this paper, we consider the ReLU Feedforward Neural Network (FNN) in the form of

$$f(\mathbf{x}) = W_L \cdot \text{ReLU}(W_{L-1} \cdots \text{ReLU}(W_1 \mathbf{x} + \mathbf{b}_1) + \cdots + \mathbf{b}_{L-1}) + \mathbf{b}_L, \quad (1)$$

where W_l 's are weight matrices, \mathbf{b}_l 's are biases, and $\text{ReLU}(a) = \max\{a, 0\}$ is the rectified linear unit activation (ReLU) applied element-wise.

We consider two classes of network architectures whose inputs are in a compact domain of a vector space and whose outputs are vectors in \mathbb{R}^d . The dimension of the input and output spaces are to be specified later. The first class is defined as

$$\mathcal{F}_{\text{NN}}(d, L, p, K, \kappa, M) = \{\Gamma = [f_1, f_2, \dots, f_d]^{\top} : \text{for each } k = 1, \dots, d,$$

$f_k(\mathbf{x})$ is in the form of (1) with L layers, width bounded by p ,

$$\|f_k\|_\infty \leq M, \|W_l\|_{\infty, \infty} \leq \kappa, \|\mathbf{b}_l\|_\infty \leq \kappa, \sum_{l=1}^L \|W_l\|_0 + \|\mathbf{b}_l\|_0 \leq K, \quad (2)$$

where $\|f\|_\infty = \sup_{\mathbf{x}} |f(\mathbf{x})|$, $\|W\|_{\infty, \infty} = \max_{i,j} |W_{i,j}|$, $\|\mathbf{b}\|_\infty = \max_i |b_i|$ for any function f , matrix W , and vector \mathbf{b} with $\|\cdot\|_0$ denoting the number of nonzero elements of its argument. The function class given by this first network architecture has an upper bound on all weight parameters (the magnitude of all weight parameters are upper bounded by κ) and a cardinality constraint (the total number of nonzero parameters are no more than K). Each element of the output is upper bounded by M . This constraint on the output is often enforced by clipping the output in the testing procedure. Such a clipping can be realized with a two-layer network, which is fixed during training. This clipping step is common in nonparametric regression [36].

In the second class of network architecture, we drop the magnitude and cardinality constraints for practical concerns on training. The second network architecture is parameterized by L, p, M only:

$$\begin{aligned} \mathcal{F}_{\text{NN}}(d, L, p, M) = \{ \Gamma = [f_1, f_2, \dots, f_d]^\top : \text{for each } k = 1, \dots, d, \\ f_k(\mathbf{x}) \text{ is in the form of (1) with } L \text{ layers, width bounded by } p, \\ \|f_k\|_\infty \leq M \}. \end{aligned} \quad (3)$$

All theoretical results in this paper can be applied to both network architectures.

Notations: We use bold lowercase letters to denote vectors, and normal font letters to denote scalars. The notation $\mathbf{0}$ represents a zero vector. For a d dimensional vector $\mathbf{k} = [k_1, \dots, k_d]^\top$, we denote $|\mathbf{k}| = \sum_{i=1}^d k_i$. The vector norms are defined as $\|\mathbf{k}\|_\infty = \max_i |k_i|$ and $\|\mathbf{k}\|_2 = \sqrt{\sum_{i=1}^d k_i^2}$. For any scalar s , we denote $\lceil s \rceil$ as the smallest integer that is no less than s . We use \mathbb{N} to denote the set of positive integers and $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$. For a function $f : \Omega \rightarrow \mathbb{R}$ in a Hilbert space \mathcal{H} , we define the function norms as $\|f\|_\infty = \sup_{\mathbf{x} \in \Omega} |f(\mathbf{x})|$ and $\|f\|_{\mathcal{H}} = \sqrt{\langle f, f \rangle_{\mathcal{H}}}$, where $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ denotes the inner product of \mathcal{H} . For an operator $A : \mathcal{H} \rightarrow \mathcal{H}$, we denote its operator norm by $\|A\|_{\text{op}}$ and its Hilbert-Schmidt norm by $\|A\|_{\text{HS}}$. More notations used in this paper is summarized in Table 1.

2.2 Problem setup and a learning framework

Let \mathcal{X} and \mathcal{Y} be two separable Hilbert spaces and $\Psi : \mathcal{X} \rightarrow \mathcal{Y}$ be an unknown operator. Our goal is to learn the operator Ψ from a finite number of samples $\mathcal{S} = \{u_i, v_i\}_{i=1}^{2n}$ in the following setting.

Setting 1. Let \mathcal{X}, \mathcal{Y} be two separable Hilbert spaces and γ be a probability measure on \mathcal{X} . Let $\mathcal{S} = \{u_i, v_i\}_{i=1}^{2n}$ be the given data where u_i 's are i.i.d. samples from γ and the v_i 's are generated according to model:

$$v_i = \Psi(u_i) + \tilde{\epsilon}_i, \quad (4)$$

where the $\tilde{\epsilon}_i$'s are i.i.d. samples from a probability measure μ on \mathcal{Y} , independently of u_i 's. We denote the probability measure of v by ζ .

The pushforward measure of γ under Ψ is denoted by $\Psi_{\#}\gamma$, such that for any $\Omega \subset \mathcal{Y}$,

$$\Psi_{\#}\gamma(\Omega) = \gamma(\{u : \Psi(u) \in \Omega\}).$$

Without additional assumptions, the estimation error of Ψ based on a finite number of samples may not converge to zero since Ψ is an operator between infinite-dimensional spaces. In this paper, we exploit the low-dimensional structures in this estimation problem arising from practical applications, and prove a nonparametric estimation error of Ψ by deep neural networks.

Our learning framework follows the idea of model reduction [7]. It consists of encoding and decoding in both the \mathcal{X} and \mathcal{Y} spaces, and deep learning of a transformation between the encoded vectors for the elements in \mathcal{X} and \mathcal{Y} . We first encode the elements in \mathcal{X} and \mathcal{Y} to finite dimensional vectors by an encoding operator. For fixed positive integers $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$, let $E_{\mathcal{X}} : \mathcal{X} \rightarrow \mathbb{R}^{d_{\mathcal{X}}}$ and $D_{\mathcal{X}} : \mathbb{R}^{d_{\mathcal{X}}} \rightarrow \mathcal{X}$ be the encoder and decoder of \mathcal{X} , and $E_{\mathcal{Y}} : \mathcal{Y} \rightarrow \mathbb{R}^{d_{\mathcal{Y}}}$ and $D_{\mathcal{Y}} : \mathbb{R}^{d_{\mathcal{Y}}} \rightarrow \mathcal{Y}$ be the encoder and decoder of \mathcal{Y} such that

$$D_{\mathcal{X}} \circ E_{\mathcal{X}} \approx I \text{ and } D_{\mathcal{Y}} \circ E_{\mathcal{Y}} \approx I.$$

The empirical counterparts of encoders and decoders are denoted by $E_{\mathcal{X}}^n, D_{\mathcal{X}}^n, E_{\mathcal{Y}}^n$ and $D_{\mathcal{Y}}^n$, and we call them empirical encoders and decoders.

The simplest encoder in a function space is the discretization operator. When \mathcal{X} is a function space containing functions defined on a compact subset of \mathbb{R}^D , we can discretize the domain with a fixed grid, and take the encoder as the sampling operator on this grid. However, the discretization operator may not reveal the low-dimensional structures in the functions of interest, and therefore may not effectively reduce the dimension.

A popular choice of encoders in applications is the basis encoder, such as the Fourier transform with trigonometric basis, or PCA with data-driven basis, etc. Given an orthonormal basis of \mathcal{X} and a positive integer $d_{\mathcal{X}}$, the basis encoder maps an element in \mathcal{X} to $d_{\mathcal{X}}$ coefficients associated with a fixed set of $d_{\mathcal{X}}$ bases. For any coefficient vector $\mathbf{a} \in \mathbb{R}^{d_{\mathcal{X}}}$, the decoder $D_{\mathcal{X}}(\mathbf{a})$ gives rise to a linear combination of these $d_{\mathcal{X}}$ bases weighted by \mathbf{a} . See Section 4 for the details about the basis encoder. The trigonometric basis and orthogonal polynomials are commonly used bases in applications. These bases are a priori given, independently of the training data. In this case, the basis encoder can be viewed as a deterministic operator, which is given independently of the training data. The empirical encoder and decoder are the same as the oracle encoder and decoder: $E_{\mathcal{X}}^n = E_{\mathcal{X}}$ and $D_{\mathcal{X}}^n = D_{\mathcal{X}}$.

PCA [69, 41, 42] is an effective dimension reduction technique, when u_i 's exhibit a low-dimensional linear structure. The PCA encoder encodes an element in \mathcal{X} to the $d_{\mathcal{X}}$ coefficients associated with the top $d_{\mathcal{X}}$ eigenbasis of a trace operator. The PCA decoder gives a linear combination of the eigenbasis weighted by the given coefficient vector. In practice, one needs to estimate this trace operator from the training data and obtain an empirical estimation of $E_{\mathcal{X}}$ and $D_{\mathcal{X}}$, which are denoted by $E_{\mathcal{X}}^n$ and $D_{\mathcal{X}}^n$, respectively. The PCA encoder is data-driven, and we expect $E_{\mathcal{X}}^n \approx E_{\mathcal{X}}$, $D_{\mathcal{X}}^n \approx D_{\mathcal{X}}$ when the sample size n is sufficiently large. The encoding and decoding operator in \mathcal{Y} can be defined analogously.

The operator $D_{\mathcal{X}} \circ E_{\mathcal{X}}$ is the projection operator associated with the encoder $E_{\mathcal{X}}$ and decoder $D_{\mathcal{X}}$. We have the following projections and their empirical counterparts:

$$\begin{aligned}\Pi_{\mathcal{X},d_{\mathcal{X}}} &= D_{\mathcal{X}} \circ E_{\mathcal{X}}, & \Pi_{\mathcal{X},d_{\mathcal{X}}}^n &= D_{\mathcal{X}}^n \circ E_{\mathcal{X}}^n, \\ \Pi_{\mathcal{Y},d_{\mathcal{Y}}} &= D_{\mathcal{Y}} \circ E_{\mathcal{Y}}, & \Pi_{\mathcal{Y},d_{\mathcal{Y}}}^n &= D_{\mathcal{Y}}^n \circ E_{\mathcal{Y}}^n.\end{aligned}$$

After the empirical encoders $E_{\mathcal{X}}^n, E_{\mathcal{Y}}^n$ and decoders $D_{\mathcal{X}}^n, D_{\mathcal{Y}}^n$ are computed, our objective is to learn a transformation $\Gamma : \mathbb{R}^{d_{\mathcal{X}}} \rightarrow \mathbb{R}^{d_{\mathcal{Y}}}$ such that

$$D_{\mathcal{Y}}^n \circ \Gamma \circ E_{\mathcal{X}}^n \approx \Psi. \quad (5)$$

We learn Γ using a two-stage algorithm. Given the training data $\mathcal{S} = \{u_i, v_i\}_{i=1}^{2n}$, we split the data into two subsets $\mathcal{S}_1 = \{u_i, v_i\}_{i=1}^n$ and $\mathcal{S}_2 = \{u_i, v_i\}_{i=n+1}^{2n}$ ¹, where \mathcal{S}_1 is used to compute the encoders and decoders and \mathcal{S}_2 is used to learn the transformation Γ between the encoded vectors. Our two-stage algorithm follows

Stage 1: Compute the empirical encoders and decoders $E_{\mathcal{X}}^n, D_{\mathcal{X}}^n, E_{\mathcal{Y}}^n, D_{\mathcal{Y}}^n$ based on \mathcal{S}_1 . In the case of deterministic encoders, we skip Stage 1 and let $E_{\mathcal{X}}^n = E_{\mathcal{X}}, D_{\mathcal{X}}^n = D_{\mathcal{X}}, E_{\mathcal{Y}}^n = E_{\mathcal{Y}}, D_{\mathcal{Y}}^n = D_{\mathcal{Y}}$.

Stage 2: Learn Γ with \mathcal{S}_2 by solving the following optimization problem

$$\Gamma_{\text{NN}} \in \underset{\Gamma \in \mathcal{F}_{\text{NN}}}{\operatorname{argmin}} \frac{1}{n} \sum_{i=n+1}^{2n} \|\Gamma \circ E_{\mathcal{X}}^n(u_i) - E_{\mathcal{Y}}^n(v_i)\|_2^2 \quad (6)$$

for some \mathcal{F}_{NN} class with a proper choice of parameters.

Our estimator of Ψ is given as

$$\Psi_{\text{NN}} := D_{\mathcal{Y}}^n \circ \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n,$$

and the mean squared generalization error is defined as

$$\mathbb{E}_{\mathcal{S}} \mathbb{E}_{u \sim \gamma} \|\Psi_{\text{NN}}(u) - \Psi(u)\|_{\mathcal{Y}}^2. \quad (7)$$

3 Main results

The main results of this paper provide statistical guarantees on the mean squared generalization error for the estimation of Lipschitz operators.

3.1 Assumptions

We first make some assumptions on the measure γ and the operator Ψ .

Assumption 1 (Compactly supported measure). The probability distribution γ is supported on a compact set $\Omega_{\mathcal{X}} \subset \mathcal{X}$. There exists $R_{\mathcal{X}} > 0$ such that, for any $u \in \Omega_{\mathcal{X}}$, we have

$$\|u\|_{\mathcal{X}} \leq R_{\mathcal{X}}. \quad (8)$$

¹The data can be split unevenly as well.

Notation	Description	Notation	Description
\mathcal{X}	Input space	\mathcal{Y}	Output space
$\Psi : \mathcal{X} \rightarrow \mathcal{Y}$	An unknown operator	$\mathcal{S} = \{u_i, v_i\}_{i=1}^{2n}$	Given data set
γ	A probability measure on \mathcal{X}	$\Psi_{\#}\gamma$	Push forward measure of γ under Ψ
μ	The probability measure of noise $\tilde{\varepsilon}$	ζ	The probability measure of $v = \Psi(u) + \tilde{\varepsilon}$
$E_{\mathcal{X}}, D_{\mathcal{X}}$	Encoder and decoder of \mathcal{X}	$E_{\mathcal{Y}}, D_{\mathcal{Y}}$	Encoder and decoder of \mathcal{Y}
$E_{\mathcal{X}}^n, D_{\mathcal{X}}^n$	Empirical estimations of $E_{\mathcal{X}}, D_{\mathcal{X}}$ from noisy data	$E_{\mathcal{Y}}^n, D_{\mathcal{Y}}^n$	Empirical estimations of $E_{\mathcal{Y}}, D_{\mathcal{Y}}$ from noisy data
$d_{\mathcal{X}}$	Encoding dimension of \mathcal{X}	$d_{\mathcal{Y}}$	Encoding dimension of \mathcal{Y}
$\Pi_{\mathcal{X}, d_{\mathcal{X}}}$	Projection $D_{\mathcal{X}} \circ E_{\mathcal{X}}$	$\Pi_{\mathcal{Y}, d_{\mathcal{Y}}}$	Projection $D_{\mathcal{Y}} \circ E_{\mathcal{Y}}$
$\Pi_{\mathcal{X}, d_{\mathcal{X}}}^n$	Empirical projection $D_{\mathcal{X}}^n \circ E_{\mathcal{X}}^n$	$\Pi_{\mathcal{Y}, d_{\mathcal{Y}}}^n$	Empirical projection $D_{\mathcal{Y}}^n \circ E_{\mathcal{Y}}^n$
$\ \Pi_{\mathcal{X}, d_{\mathcal{X}}}(u) - u\ _{\mathcal{X}}$	Encoding error for u in \mathcal{X}	$\ \Pi_{\mathcal{Y}, d_{\mathcal{Y}}}(v) - v\ _{\mathcal{Y}}$	Encoding error for v in \mathcal{Y}
\mathcal{F}_{NN}	Neural network class	Γ_{NN}	Neural network estimator in (6)

Table 1: Notations used in this paper.

Assumption 2 (Lipschitz operator). There exists $L_{\Psi} > 0$ such that

$$\|\Psi(u_1) - \Psi(u_2)\|_{\mathcal{Y}} \leq L_{\Psi} \|u_1 - u_2\|_{\mathcal{X}}, \text{ for any } u_1, u_2 \in \Omega_{\mathcal{X}}.$$

Assumption 1 and 2 assume that γ is compactly supported and Ψ is Lipschitz continuous. We denote the image of $\Omega_{\mathcal{X}}$ under the transformation Ψ as

$$\Omega_{\mathcal{Y}} = \{v \in \mathcal{Y} : v = \Psi(u) \text{ for some } u \in \Omega_{\mathcal{X}}\}.$$

Assumption 1 and 2 imply that $\Omega_{\mathcal{Y}}$ is bounded: there exists a constant $R_{\mathcal{Y}} = L_{\Psi} R_{\mathcal{X}} > 0$ such that for any $v \in \Omega_{\mathcal{Y}}$, we have $\|v\|_{\mathcal{Y}} \leq R_{\mathcal{Y}}$.

We next make some natural assumptions on the empirical encoders and decoders:

Assumption 3 (Lipchitz encoders and decoders). The empirical encoders and decoders $E_{\mathcal{X}}^n, D_{\mathcal{X}}^n, E_{\mathcal{Y}}^n, D_{\mathcal{Y}}^n$ satisfy:

$$E_{\mathcal{X}}^n(0_{\mathcal{X}}) = \mathbf{0}, D_{\mathcal{X}}^n(\mathbf{0}) = 0_{\mathcal{X}}, E_{\mathcal{Y}}^n(0_{\mathcal{Y}}) = \mathbf{0}, D_{\mathcal{Y}}^n(\mathbf{0}) = 0_{\mathcal{Y}},$$

where $\mathbf{0}$ denotes the zero vector, $0_{\mathcal{X}}$ is the zero function in \mathcal{X} and $0_{\mathcal{Y}}$ is the zero function in \mathcal{Y} .

They are also Lipschitz: there exist $L_{E_{\mathcal{X}}^n}, L_{D_{\mathcal{X}}^n}, L_{E_{\mathcal{Y}}^n}, L_{D_{\mathcal{Y}}^n} > 0$ such that, for any $u_1, u_2 \in \mathcal{X}$ and any $\mathbf{a}_1, \mathbf{a}_2 \in \mathbb{R}^{d_{\mathcal{X}}}$, we have

$$\|E_{\mathcal{X}}^n(u_1) - E_{\mathcal{X}}^n(u_2)\|_2 \leq L_{E_{\mathcal{X}}^n} \|u_1 - u_2\|_{\mathcal{X}}, \quad \|D_{\mathcal{X}}^n(\mathbf{a}_1) - D_{\mathcal{X}}^n(\mathbf{a}_2)\|_{\mathcal{X}} \leq L_{D_{\mathcal{X}}^n} \|\mathbf{a}_1 - \mathbf{a}_2\|_2,$$

and for any $v_1, v_2 \in \mathcal{Y}$ and any $\mathbf{a}_1, \mathbf{a}_2 \in \mathbb{R}^{d_{\mathcal{Y}}}$, we have

$$\|E_{\mathcal{Y}}^n(v_1) - E_{\mathcal{Y}}^n(v_2)\|_2 \leq L_{E_{\mathcal{Y}}^n} \|v_1 - v_2\|_{\mathcal{Y}}, \quad \|D_{\mathcal{Y}}^n(\mathbf{a}_1) - D_{\mathcal{Y}}^n(\mathbf{a}_2)\|_{\mathcal{Y}} \leq L_{D_{\mathcal{Y}}^n} \|\mathbf{a}_1 - \mathbf{a}_2\|_2.$$

Remark 1. Assumption 3 is made on empirical encoders and decoders. The basis encoders and PCA encoders, which are most commonly used, satisfy Assumption 3 with the Lipschitz constants $L_{E_{\mathcal{X}}^n} = L_{D_{\mathcal{X}}^n} = L_{E_{\mathcal{Y}}^n} = L_{D_{\mathcal{Y}}^n} = 1$, independently of the training data (see Lemma 1 and Lemma 2).

Assumption 3 implies that $E_{\mathcal{X}}^n(u)$ and $E_{\mathcal{Y}}^n(v)$ are bounded for any $u \in \Omega_{\mathcal{X}}$ and $v \in \Omega_{\mathcal{Y}}$. For any $u \in \Omega_{\mathcal{X}}$, we have $\|E_{\mathcal{X}}^n(u)\|_2 \leq \|E_{\mathcal{X}}^n(u) - E_{\mathcal{X}}^n(0)\|_2 + \|E_{\mathcal{X}}^n(0)\|_2 \leq L_{E_{\mathcal{X}}^n} R_{\mathcal{X}}$. Similarly, for any $v \in \Omega_{\mathcal{Y}}$, we have $\|E_{\mathcal{Y}}^n(v)\|_2 \leq L_{E_{\mathcal{Y}}^n} R_{\mathcal{Y}}$.

Assumption 4 (Noise). The random noise $\tilde{\epsilon}$ satisfies

- (i) $\tilde{\epsilon}$ is independent of u .
- (ii) $\mathbb{E}[\tilde{\epsilon}] = 0$.
- (iii) There exists $\tilde{\sigma} > 0$ such that $\|\tilde{\epsilon}\|_{\mathcal{Y}} \leq \tilde{\sigma}$.

Assumption 4(i)-(iii) are natural assumptions on noise. Assumption 4(i) is about the independence of the input and the noise, which is commonly used in nonparametric regression. Assumption 4(iii) together with Assumption 3 imply that the perturbation of the encoded vectors are bounded: $\|E_{\mathcal{Y}}^n(\Psi(u) + \tilde{\epsilon}) - E_{\mathcal{Y}}^n(\Psi(u))\|_{\infty} \leq L_{E_{\mathcal{Y}}^n} \tilde{\sigma}$. We denote $\sigma = L_{E_{\mathcal{Y}}^n} \tilde{\sigma}$ such that

$$\|E_{\mathcal{Y}}^n(\Psi(u) + \tilde{\epsilon}) - E_{\mathcal{Y}}^n(\Psi(u))\|_{\infty} \leq \sigma \text{ for any } u \text{ and } \tilde{\epsilon}.$$

Assumption 5 (Noise and encoder). For any noise satisfying Assumption 4 and any given \mathcal{S}_1 , the conditional expectation satisfies

$$\mathbb{E}_{\tilde{\epsilon}} \left[E_{\mathcal{Y}}^n(\Psi(u) + \tilde{\epsilon}) - E_{\mathcal{Y}}^n(\Psi(u)) | \mathcal{S}_1 \right] = \mathbf{0}, \text{ for any } u \in \Omega_{\mathcal{X}},$$

where $E_{\mathcal{Y}}^n$ is the empirical encoder computed with \mathcal{S}_1 .

Assumption 5 requires that, if we condition on \mathcal{S}_1 based on which the empirical encoder $E_{\mathcal{Y}}^n$ is computed, the perturbation on the encoded vector resulted from noise has zero expectation. Assumption 5 is guaranteed for all linear encoders as long as Assumption 4(ii) holds:

$$\mathbb{E}_{\tilde{\epsilon}} \left[E_{\mathcal{Y}}^n(\Psi(u) + \tilde{\epsilon}) - E_{\mathcal{Y}}^n(\Psi(u)) | \mathcal{S}_1 \right] = \mathbb{E}_{\tilde{\epsilon}} \left[E_{\mathcal{Y}}^n(\tilde{\epsilon}) | \mathcal{S}_1 \right] = \mathbf{0}.$$

Basis encoders, including the PCA encoder, are linear encoders, so they all satisfy Assumption 5.

3.2 Generalization error with general encoders and decoders

Our main result is an upper bound of the generalization error in (7) with general encoders and decoders. Our results can be applied to both network architectures defined in (2) and (3). Our first theorem gives an upper bound of the generalization error with the network architecture defined in (2).

Theorem 1. In Setting 1, suppose Assumption 1 – 5 hold. Let Γ_{NN} be the minimizer of (6) with the network architecture $\mathcal{F}(d_{\mathcal{Y}}, L, p, K, \kappa, M)$ in (2), where

$$\begin{aligned} L &= O(\log n + \log d_{\mathcal{Y}}), \quad p = O\left(d_{\mathcal{Y}}^{-\frac{d_{\mathcal{X}}}{2+d_{\mathcal{X}}}} n^{\frac{d_{\mathcal{X}}}{2+d_{\mathcal{X}}}}\right), \quad K = O\left(d_{\mathcal{Y}}^{-\frac{d_{\mathcal{X}}}{2+d_{\mathcal{X}}}} n^{\frac{d_{\mathcal{X}}}{2+d_{\mathcal{X}}}} \log n\right), \\ M &= \sqrt{d_{\mathcal{Y}}} L_{E_{\mathcal{Y}}^n} R_{\mathcal{Y}}, \quad \kappa = \max\left\{1, \sqrt{d_{\mathcal{Y}}} L_{E_{\mathcal{Y}}^n} R_{\mathcal{Y}}, \sqrt{d_{\mathcal{X}}} L_{E_{\mathcal{X}}^n} R_{\mathcal{X}}, L_{E_{\mathcal{Y}}^n} L_{D_{\mathcal{X}}^n} L_{\Psi}\right\}. \end{aligned} \tag{9}$$

Then we have

$$\begin{aligned}
& \mathbb{E}_{\mathcal{S}} \mathbb{E}_{u \sim \gamma} \|D_{\mathcal{Y}}^n \circ \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u) - \Psi(u)\|_{\mathcal{Y}}^2 \\
& \leq C_1(\sigma^2 + R_{\mathcal{Y}}^2) d_{\mathcal{Y}}^{\frac{4+d_{\mathcal{X}}}{2+d_{\mathcal{X}}}} n^{-\frac{2}{2+d_{\mathcal{X}}}} \log^3 n + C_2(\sigma^2 + R_{\mathcal{Y}}^2) d_{\mathcal{Y}}^2 (\log d_{\mathcal{Y}}) n^{-1} \\
& \quad + C_3 \mathbb{E}_{\mathcal{S}} \mathbb{E}_{u \sim \gamma} \|\Pi_{\mathcal{X}, d_{\mathcal{X}}}^n(u) - u\|_{\mathcal{X}}^2 + 2 \mathbb{E}_{\mathcal{S}} \mathbb{E}_{w \sim \Psi_{\#} \gamma} \|\Pi_{\mathcal{Y}, d_{\mathcal{Y}}}^n(w) - w\|_{\mathcal{Y}}^2, \tag{10}
\end{aligned}$$

where C_1, C_2 are constants depending on $d_{\mathcal{X}}, R_{\mathcal{X}}, R_{\mathcal{Y}}, L_{E_{\mathcal{X}}^n}, L_{E_{\mathcal{Y}}^n}, L_{D_{\mathcal{X}}^n}, L_{D_{\mathcal{Y}}^n}, L_{\Psi}$ and $C_3 = 16L_{D_{\mathcal{Y}}^n}^2 L_{E_{\mathcal{Y}}^n}^2 L_{\Psi}^2$.

Our second theorem gives an upper bound of the generalization error with the network architecture defined in (3).

Theorem 2. In Setting 1, suppose Assumption 1 – 5 hold. Let Γ_{NN} be the minimizer of (6) with the network architecture $\mathcal{F}(d_{\mathcal{Y}}, L, p, M)$ in (3) with

$$L = O(\widetilde{L} \log \widetilde{L}), p = O(\widetilde{p} \log \widetilde{p}), M = \sqrt{d_{\mathcal{Y}} L_{E_{\mathcal{Y}}^n} R_{\mathcal{Y}}}, \tag{11}$$

where $\widetilde{L}, \widetilde{p} > 0$ are positive integers satisfying

$$\widetilde{L} \widetilde{p} = \left\lceil d_{\mathcal{Y}}^{-\frac{d_{\mathcal{X}}}{4+2d_{\mathcal{X}}}} n^{\frac{d_{\mathcal{X}}}{4+2d_{\mathcal{X}}}} \right\rceil. \tag{12}$$

Then we have

$$\begin{aligned}
& \mathbb{E}_{\mathcal{S}} \mathbb{E}_{u \sim \gamma} \|D_{\mathcal{Y}}^n \circ \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u) - \Psi(u)\|_{\mathcal{Y}}^2 \\
& \leq C_4(\sigma^2 + R_{\mathcal{Y}}^2) d_{\mathcal{Y}}^{\frac{4+d_{\mathcal{X}}}{2+d_{\mathcal{X}}}} n^{-\frac{2}{2+d_{\mathcal{X}}}} \log^6 n + C_3 \mathbb{E}_{\mathcal{S}} \mathbb{E}_{u \sim \gamma} \|\Pi_{\mathcal{X}, d_{\mathcal{X}}}^n(u) - u\|_{\mathcal{X}}^2 + 2 \mathbb{E}_{\mathcal{S}} \mathbb{E}_{w \sim \Psi_{\#} \gamma} \|\Pi_{\mathcal{Y}, d_{\mathcal{Y}}}^n(w) - w\|_{\mathcal{Y}}^2, \tag{13}
\end{aligned}$$

where C_4 is a constant depending on $d_{\mathcal{X}}, R_{\mathcal{X}}, R_{\mathcal{Y}}, L_{E_{\mathcal{X}}^n}, L_{E_{\mathcal{Y}}^n}, L_{D_{\mathcal{X}}^n}, L_{\Psi}$ and $C_3 = 16L_{D_{\mathcal{Y}}^n}^2 L_{E_{\mathcal{Y}}^n}^2 L_{\Psi}^2$ is the same one in Theorem 1.

Theorem 1 is proved in Section 7.2 and Theorem 2 is proved in Section 7.3. In Theorem 2, we have chosen the optimal $\widetilde{L} \widetilde{p}$ to balance the bias and variance term. For readers who are interested in the generalization error with arbitrary network depth L and width p , please see our proof in Section 7.3. The constants in both theorems only depend on the settings of the problem, and the choices of encoders and decoders. They do not depend on properties of Γ_{NN} . For both network architectures, the upper bound in (10) and (13) consists of a network estimation error and the projection errors in the \mathcal{X} and \mathcal{Y} space.

- The first two terms in (10) and the first term in (13) represent the **network estimation error** for the transformation $\Gamma : \mathbb{R}^{d_{\mathcal{X}}} \rightarrow \mathbb{R}^{d_{\mathcal{Y}}}$ which maps the encoded vector $E_{\mathcal{X}}^n(u)$ for u in \mathcal{X} to the encoded vector $E_{\mathcal{Y}}^n(\Phi(u))$ for $\Phi(u)$ in \mathcal{Y} . This error decays exponentially as the sample size n increases with an exponent depending on the dimension $d_{\mathcal{X}}$ of the encoded space. The dimension $d_{\mathcal{X}}$ appears in the exponent and $d_{\mathcal{Y}}$ appears as a constant factor. This is because that the transformation Γ has $d_{\mathcal{Y}}$ outputs and each output is a function from $\mathbb{R}^{d_{\mathcal{X}}}$ to \mathbb{R} . Therefore the rate is only cursed by the input dimension $d_{\mathcal{X}}$.

- The last two terms in (10) and (13) are **projection errors** in the \mathcal{X} and \mathcal{Y} space, respectively. If the measure γ is concentrated near a $d_{\mathcal{X}}$ -dimensional subspace in \mathcal{X} , both projection errors can be made small if the encoder and decoder are properly chosen as the projection onto this $d_{\mathcal{X}}$ -dimensional subspace (see Section 6).

We next compare the difference between the network architectures in Theorem 1 and Theorem 2. Denote the network architecture in Theorem 1 and Theorem 2 by \mathcal{F}_1 and \mathcal{F}_2 , respectively. The architecture \mathcal{F}_1 has the depth and width scaling properly with respect to each other, and an upper bound on all weight parameters and a cardinality constraint. The cardinality constraint is nonconvex and therefore not practical for training this neural network. The architecture \mathcal{F}_2 has more flexibility in the choice of depth and width as long as (12) is satisfied. The cardinality is removed for practical concerns. When we set $\tilde{L} = O(\log n), \tilde{p} = O(n^{\frac{d_{\mathcal{X}}}{4+2d_{\mathcal{X}}}} \log^{-1} n)$ in \mathcal{F}_2 , both networks have a depth of $O(\log n)$, while the width of \mathcal{F}_1 is the square of that of \mathcal{F}_2 , i.e., \mathcal{F}_1 is wider than \mathcal{F}_2 . The comparison between \mathcal{F}_1 and \mathcal{F}_2 is summarized in Table 2.

	\mathcal{F}_1 in (2)	\mathcal{F}_2 in (3)
General comparison		
Network architecture with a given n	Fixed L and p depending on n	One has the flexibility to choose L and p as long as (12) depending on n is satisfied
Constraints on cardinality	Yes	No
Constraints on the magnitude of weight parameters	Yes	No
Set $\tilde{L} = O(\log n), \tilde{p} = O(n^{\frac{d_{\mathcal{X}}}{4+2d_{\mathcal{X}}}} \log^{-1} n)$ in \mathcal{F}_2		
Depth L	$O(\log n)$	$O(\log n)$
Width p	$O\left(d_{\mathcal{Y}}^{-\frac{d_{\mathcal{X}}}{2+d_{\mathcal{X}}}} n^{\frac{d_{\mathcal{X}}}{2+d_{\mathcal{X}}}}\right)$	$O\left(d_{\mathcal{Y}}^{-\frac{d_{\mathcal{X}}}{4+2d_{\mathcal{X}}}} n^{\frac{d_{\mathcal{X}}}{4+2d_{\mathcal{X}}}}\right)$

Table 2: Comparison of the network architectures in Theorem 1 and 2.

In the rest of this paper, we focus on the network architecture in Theorem 2 and discuss its applications in various scenarios. Theorem 1 can also be applied in each case with a similar upper bound.

4 Generalization error with basis encoders and decoders

In this section, we discuss the application of Theorem 2 when the encoder is chosen to be a deterministic basis encoder with a given orthonormal basis of the Hilbert space. Popular choices of orthonormal bases include orthogonal polynomials (e.g., Legendre polynomials [88, 16, 18]) and trigonometric functions [68, 11, 53].

4.1 Basis encoders and decoders

Let \mathcal{H} be a separable Hilbert space equipped with an inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$, and $\{\phi_k\}_{k=1}^{\infty}$ be an orthonormal basis of \mathcal{H} such that $\langle \phi_{k_1}, \phi_{k_2} \rangle_{\mathcal{H}} = 0$ whenever $k_1 \neq k_2$ and $\|\phi_k\|_{\mathcal{H}} = 1$ for any k . For

any $u \in \mathcal{H}$, we have

$$u = \sum_{k=1}^{\infty} \langle u, \phi_k \rangle_{\mathcal{H}} \phi_k. \quad (14)$$

For a fixed positive integer d representing the encoding dimension, we define the encoder of \mathcal{H} as

$$E_{\mathcal{H},d}(u) = [\langle u, \phi_1 \rangle_{\mathcal{H}}, \dots, \langle u, \phi_d \rangle_{\mathcal{H}}]^{\top} \in \mathbb{R}^d, \text{ for any } u \in \mathcal{H}, \quad (15)$$

which gives rise to the coefficients associated with a fixed set of d basis functions in the decomposition (14). The decoder $D_{\mathcal{H},d}$ is defined as

$$D_{\mathcal{H},d}(\mathbf{a}) = \sum_{k=1}^d a_k \phi_k, \text{ for any } \mathbf{a} \in \mathbb{R}^d. \quad (16)$$

The basis encoder and decoder naturally satisfy the Lipschitz property with a Lipschitz constant 1 (see a proof of Lemma 1 in Appendix B).

Lemma 1. The encoder $E_{\mathcal{H},d}$ and decoder $D_{\mathcal{H},d}$ defined in (15) and (16) satisfy

$$\|E_{\mathcal{H},d}(u) - E_{\mathcal{H},d}(\tilde{u})\|_2 \leq \|u - \tilde{u}\|_{\mathcal{H}}, \quad (17)$$

$$\|D_{\mathcal{H},d}(\mathbf{a}) - D_{\mathcal{H},d}(\tilde{\mathbf{a}})\|_{\mathcal{H}} = \|\mathbf{a} - \tilde{\mathbf{a}}\|_2, \quad (18)$$

for any $u, \tilde{u} \in \mathcal{H}$ and $\mathbf{a}, \tilde{\mathbf{a}} \in \mathbb{R}^d$.

Remark 2. All encoders in the form of (15) are linear operators and therefore satisfy Assumption 5 as long as Assumption 4(ii) holds.

4.2 Generalization error with basis encoders

We next consider the generalization error when the elements in \mathcal{X} and \mathcal{Y} are encoded by basis encoders with the encoding dimension $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$, respectively. Substituting the Lipschitz constants of all encoders and decoders by 1 in Theorem 2, we obtain the following corollary:

Corollary 1. In Setting 1, suppose Assumption 1 – 4 hold. Let Γ_{NN} be the minimizer of (6) with the network architecture $\mathcal{F}(d_{\mathcal{Y}}, L, p, M)$ in (3) with

$$L = O(\tilde{L} \log \tilde{L}), \quad p = O(\tilde{p} \log \tilde{p}), \quad M = \sqrt{d_{\mathcal{Y}}} R_{\mathcal{Y}}, \quad (19)$$

where $\tilde{L}, \tilde{p} > 0$ are positive integers satisfying (12). Then we have

$$\begin{aligned} & \mathbb{E}_{\mathcal{S}} \mathbb{E}_{u \sim \gamma} \|D_{\mathcal{Y}}^n \circ \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u) - \Psi(u)\|_{\mathcal{Y}}^2 \\ & \leq C_4 (\sigma^2 + R_{\mathcal{Y}}^2) d_{\mathcal{Y}}^{\frac{4+d_{\mathcal{X}}}{2+d_{\mathcal{X}}}} n^{-\frac{2}{2+d_{\mathcal{X}}}} \log^6 n + 16L_{\Psi}^2 \mathbb{E}_{\mathcal{S}} \mathbb{E}_{u \sim \gamma} \|\Pi_{\mathcal{X},d_{\mathcal{X}}}^n(u) - u\|_{\mathcal{X}}^2 + 2\mathbb{E}_{\mathcal{S}} \mathbb{E}_{w \sim \Psi_{\#}\gamma} \|\Pi_{\mathcal{Y},d_{\mathcal{Y}}}^n(w) - w\|_{\mathcal{Y}}^2, \end{aligned} \quad (20)$$

where C_4 is a constant depending on $d_{\mathcal{X}}, R_{\mathcal{X}}, R_{\mathcal{Y}}, L_{\Psi}$.

Popular choices of orthonormal bases are orthogonal polynomials and trigonometric functions. We next provide an upper bound on the generalization error when Legendre polynomials or trigonometric functions are used for encoding and decoding. In the rest of this section, we assume $\mathcal{X} = \mathcal{Y} = L^2([-1, 1]^D)$ with the inner product

$$\langle u_1, u_2 \rangle = \int_{[-1, 1]^D} u_1(\mathbf{x}) \overline{u_2(\mathbf{x})} d\mathbf{x}, \quad (21)$$

where $\overline{u_2(\mathbf{x})}$ denotes the complex conjugate of $u_2(\mathbf{x})$.

4.3 Legendre polynomials

On the interval $[-1, 1]$, one-dimensional Legendre polynomials $\{\tilde{P}_k\}_{k=0}^\infty$ are defined recursively as

$$\begin{cases} \tilde{P}_0(x) = 1, \\ \tilde{P}_1(x) = x, \\ \tilde{P}_{k+1}(x) = \frac{1}{k+1} \left[(2k+1)x\tilde{P}_k(x) - k\tilde{P}_{k-1}(x) \right]. \end{cases}$$

The Legendre polynomials satisfy

$$\int_{-1}^1 \tilde{P}_k(x) \tilde{P}_l(x) dx = \frac{2}{2k+1} \delta_{kl},$$

where δ_{kl} is the Kronecker delta which equals to 1 if $k = l$ and equals to 0 otherwise. We define the normalized Legendre polynomials as

$$P_k(x) = \sqrt{\frac{2k+1}{2}} \tilde{P}_k(x). \quad (22)$$

In the Hilbert space $L^2([-1, 1]^D)$, the D -variate normalized Legendre polynomials are defined as

$$\phi_{\mathbf{k}}^L = \prod_{j=1}^D P_{k_j}(x_j), \quad (23)$$

where $\mathbf{k} = [k_1 \cdots k_D]^\top$. The orthonormal basis of Legendre polynomials in $L^2([-1, 1]^D)$ is $\{\phi_{\mathbf{k}}^L\}_{\mathbf{k} \in \mathbb{N}_0^D}$.

The encoder with Legendre polynomials can be naturally defined as the expansion coefficients associated with low-order polynomials. Specifically, when $\mathcal{X} = L^2([0, 1]^D)$, we fix a positive integer $r_{\mathcal{X}}$ representing the highest degree of the polynomials in each dimension and consider the following set of low-order polynomials

$$\Phi^{L, r_{\mathcal{X}}} := \{\phi_{\mathbf{k}}^L : \|\mathbf{k}\|_\infty \leq r_{\mathcal{X}}\}.$$

The encoder $E_{\mathcal{X}}$ and decoder $D_{\mathcal{X}}$ can be defined according to (15) and (16) using the basis functions in $\Phi^{L, r_{\mathcal{X}}}$. In the space $\mathcal{Y} = L^2([0, 1]^D)$, the encoder $E_{\mathcal{Y}}$ and decoder $D_{\mathcal{Y}}$ can be defined similarly with basis functions in $\Phi^{L, r_{\mathcal{Y}}}$ for some positive integer $r_{\mathcal{Y}}$.

When Legendre polynomials are used for encoding, the encoding error is guaranteed for regular functions, such as Hölder functions.

Definition 1 (Hölder space). Let $k \geq 0$ be an integer and $0 < \alpha \leq 1$. A function $f : [-1, 1]^D \rightarrow \mathbb{R}$ belongs to the Hölder space $\mathcal{C}^{k,\alpha}([-1, 1]^D)$ if

$$\|f\|_{\mathcal{C}^{k,\alpha}} := \max_{|\mathbf{k}| \leq k} \sup_{\mathbf{x} \in [-1, 1]^D} |\partial^{\mathbf{k}} f(\mathbf{x})| + \max_{|\mathbf{k}|=k} \sup_{\mathbf{x}_1 \neq \mathbf{x}_2 \in [-1, 1]^D} \frac{|\partial^{\mathbf{k}} f(\mathbf{x}_1) - \partial^{\mathbf{k}} f(\mathbf{x}_2)|}{\|\mathbf{x}_1 - \mathbf{x}_2\|_2^\alpha} < \infty,$$

where $\partial^{\mathbf{k}} f = \frac{\partial^{|\mathbf{k}|} f}{\partial x_1^{k_1} \partial x_2^{k_2} \dots \partial x_D^{k_D}}$.

For a given k and α , any functions in $\mathcal{C}^{k,\alpha}([-1, 1]^D)$ has continuous partial derivatives up to order k . In particular, $\mathcal{C}^{0,1}([-1, 1]^D)$ consists of all Lipschitz functions defined on $[-1, 1]^D$.

We assume that the probability measure γ in \mathcal{X} and the pushforward measure $\Psi_{\#}\gamma$ in \mathcal{Y} are supported on subsets of the Hölder space.

Assumption 6 (Hölder input and output). Let $\mathcal{X} = \mathcal{Y} = L^2([-1, 1]^D)$ with the inner product (21). For some integer $k > 0$ and $0 < \alpha \leq 1$, the support of the probability measure γ and the pushforward measure $\Psi_{\#}\gamma$ satisfies

$$\Omega_{\mathcal{X}} \subset \mathcal{C}^{k,\alpha}([-1, 1]^D), \quad \Omega_{\mathcal{Y}} \subset \mathcal{C}^{k,\alpha}([-1, 1]^D).$$

There exist $C_{\mathcal{H},\mathcal{X}} > 0$ and $C_{\mathcal{H},\mathcal{Y}} > 0$ such that, for any $u \in \Omega_{\mathcal{X}}$ and $v \in \Omega_{\mathcal{Y}}$

$$\|u\|_{\mathcal{C}^{k,\alpha}} < C_{\mathcal{H},\mathcal{X}}, \quad \|v\|_{\mathcal{C}^{k,\alpha}} < C_{\mathcal{H},\mathcal{Y}}.$$

When Legendre polynomials are used to encode Hölder functions, the generalization error for the operator is given as below:

Corollary 2. In Setting 1, suppose Assumption 1–6 hold. Denote $s = k + \alpha$. Fix positive integers $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$ such that $d_{\mathcal{X}}^{1/D}$ and $d_{\mathcal{Y}}^{1/D}$ are integers. Suppose the encoders and decoders are chosen as in (15) and (16) with basis functions $\Phi^{L,d_{\mathcal{X}}^{1/D}}$ and $\Phi^{L,d_{\mathcal{Y}}^{1/D}}$ in \mathcal{X} and \mathcal{Y} , respectively. Let Γ_{NN} be the minimizer of (6) with the network architecture $\mathcal{F}(d_{\mathcal{Y}}, L, p, M)$ in (3) where L, p, M are set as in (19). We have

$$\mathbb{E}_{\mathcal{S}} \mathbb{E}_{u \sim \gamma} \|D_{\mathcal{Y}}^n \circ \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u) - \Psi(u)\|_{\mathcal{Y}}^2 \leq C_4(\sigma^2 + R_{\mathcal{Y}}^2) d_{\mathcal{Y}}^{\frac{4+d_{\mathcal{X}}}{2+d_{\mathcal{X}}}} n^{-\frac{2}{2+d_{\mathcal{X}}}} \log^6 n + C_5 L_{\Psi}^2 d_{\mathcal{X}}^{-\frac{2s}{D}} + C_6 d_{\mathcal{Y}}^{-\frac{2s}{D}}.$$

where C_4 depends on $d_{\mathcal{X}}, R_{\mathcal{X}}, R_{\mathcal{Y}}, L_{\Psi}$, and C_5, C_6 depend on $D, C_{\mathcal{H},\mathcal{X}}, C_{\mathcal{H},\mathcal{Y}}, L_{\Psi}$.

Corollary 2 is proved in Section 7.4. In Corollary 2, the last two terms represent the projection errors in \mathcal{X} and \mathcal{Y} , respectively. When D is large, both terms decay slowly as $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$ increase. These two error terms remain the same if we choose the encoders given by finite element bases in traditional numerical PDE methods. For example, we consider learning a PDE solver where the operator Ψ represents a map from the initial condition to the PDE solution at a certain time. Assumption 6 assumes that the initial condition and the PDE solution are Hölder functions. Suppose we discretize the domain and represent the solution by finite element basis such that the diameter of all finite elements is no larger than h for some $0 < h < 1$. Let $W^{k,2}([-1, 1]^D)$ denote the Sobolev space. We say a set of basis functions are k -order if they are in $W^{k,2}([-1, 1]^D)$. If the

finite element method with k -th order basis functions is used to approximate the PDE solution, under appropriate assumptions and for any positive integer k , the squared approximation error is $O(h^{2k})$ [27, Corollary 1.109]. In this case, the total number of basis functions is $O(h^{-D})$. Taking such a finite element approximation as our encoder for $\Omega_{\mathcal{X}}$, we have $d_{\mathcal{X}} = O(h^{-D})$ and the resulting squared projection error is of $O(d_{\mathcal{X}}^{-\frac{2k}{D}})$. In particular, if sparse grids [8] are used to construct basis functions, the approximation errors for the encoder and decoder can be further reduced.

In the setting of Corollary 2, we only assume the global smoothness of input and output functions. The global approximation encoder by Legendre polynomials (or trigonometric functions in the following subsection) leads to a slow rate of convergence: In Corollary 2, if we choose $d_{\mathcal{X}} = (\log n)^{\frac{1}{2}}$ and when $n \geq \exp\left(\max\left\{100, \left(\frac{7}{2} + \frac{s}{2D}\right)^6\right\}\right)$, the squared generalization error decays in the order of $(\log n)^{-\frac{s}{D}}$ (see a derivation in Appendix A).

However, in practice, when we solve PDEs, the initial conditions and PDE solutions often exhibit low-dimensional structures. For example, the initial conditions and PDE solutions often lie on a low-dimensional subspace or manifold, or the solver itself has low complexity (see Section 6 and [37, 73] for details). Therefore, one can use a few bases (small $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$) to achieve a small projection error, leading to a fast rate of convergence in the generalization error.

4.4 Trigonometric functions

Trigonometric functions and the Fourier transform have been widely used in various applications where the computation is converted from the spacial domain to the frequency domain. Let $\{T_k(x)\}_{k=1}^{\infty}$ be one-dimensional trigonometric functions defined on $[-1, 1]$ such that

$$\begin{cases} T_1 = 1/2, \\ T_{2k} = \sin(k\pi x) \text{ for } k > 1, \\ T_{2k+1} = \cos(k\pi x) \text{ for } k > 1. \end{cases} \quad (24)$$

In the Hilbert space $L^2([-1, 1]^D)$, the trigonometric basis is given as $\{\phi_{T, \mathbf{k}}\}_{\mathbf{k} \in \mathbb{N}^D}$ with

$$\phi_{\mathbf{k}}^T(\mathbf{x}) = \prod_{j=1}^D T_{k_j}(x_j). \quad (25)$$

When $\mathcal{X} = L^2([0, 1]^D)$, we fix a positive integer $r_{\mathcal{X}}$ and define the set of low-frequency basis

$$\Phi^{T, r_{\mathcal{X}}} = \{\phi_{\mathbf{k}}^T : \|\mathbf{k}\|_{\infty} \leq r_{\mathcal{X}}\}.$$

We set the encoder $E_{\mathcal{X}}$ and decoder $D_{\mathcal{X}}$ in \mathcal{X} according to (15) and (16) using the basis functions in $\Phi^{T, r_{\mathcal{X}}}$. Similarly, we set the encoder $E_{\mathcal{Y}}$ and decoder $D_{\mathcal{Y}}$ in \mathcal{Y} using the basis functions in $\Phi^{T, r_{\mathcal{Y}}}$ for some positive integer $r_{\mathcal{Y}}$.

Let \mathcal{P} be the set of periodic functions on $[-1, 1]^D$. We assume that the input and output functions are periodic Hölder functions.

Assumption 7. Let $\mathcal{X} = \mathcal{Y} = L^2([-1, 1]^D)$ with the inner product (21). For some integer $k > 0$ and $0 < s \leq 1$, the support of the probability measure γ and the pushforward measure $\Psi_{\#}\gamma$ satisfies

$$\Omega_{\mathcal{X}} \subset \mathcal{P} \cap \mathcal{C}^{k,\alpha}([-1, 1]^D), \quad \Omega_{\mathcal{Y}} \subset \mathcal{P} \cap \mathcal{C}^{k,\alpha}([-1, 1]^D).$$

There exist $C_{\mathcal{H}_p, \mathcal{X}} > 0$ and $C_{\mathcal{H}_p, \mathcal{Y}} > 0$ such that for any $u \in \Omega_{\mathcal{X}}$ and $v \in \Omega_{\mathcal{Y}}$

$$\|u\|_{\mathcal{C}^{k,\alpha}} < C_{\mathcal{H}_p, \mathcal{X}}, \quad \|v\|_{\mathcal{C}^{k,\alpha}} < C_{\mathcal{H}_p, \mathcal{Y}}.$$

When trigonometric functions are used to encode periodic Hölder functions, the generalization error for the operator is given as below:

Corollary 3. Consider Setting 1. Suppose Assumption 1–5 and 7 hold. Denote $s = k + \alpha$. Fix positive integers $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$ such that $d_{\mathcal{X}}^{1/D}$ and $d_{\mathcal{Y}}^{1/D}$ are integers. Suppose the encoders and decoders are chosen as in (15) and (16) with basis functions $\Phi^{\mathbb{T}, d_{\mathcal{X}}^{1/D}}$ and $\Phi^{\mathbb{T}, d_{\mathcal{Y}}^{1/D}}$ for \mathcal{X} and \mathcal{Y} , respectively. Let Γ_{NN} be the minimizer of (6) with the network architecture $\mathcal{F}(d_{\mathcal{Y}}, L, p, M)$ in (3) where L, p, M are set as in (19). We have

$$\mathbb{E}_{\mathcal{S}} \mathbb{E}_{u \sim \gamma} \|D_{\mathcal{Y}}^n \circ \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u) - \Psi(u)\|_{\mathcal{Y}}^2 \leq C_4(\sigma^2 + R_{\mathcal{Y}}^2) d_{\mathcal{Y}}^{\frac{4+d_{\mathcal{X}}}{2+d_{\mathcal{X}}}} n^{-\frac{2}{2+d_{\mathcal{X}}}} \log^6 n + C_7 L_{\Psi}^2 d_{\mathcal{X}}^{-\frac{2s}{D}} + C_8 d_{\mathcal{Y}}^{-\frac{2s}{D}}.$$

where C_4 depends on $d_{\mathcal{X}}, R_{\mathcal{X}}, R_{\mathcal{Y}}, L_{\Psi}$, and C_7, C_8 depend on $D, C_{\mathcal{H}_p, \mathcal{X}}, C_{\mathcal{H}_p, \mathcal{Y}}, L_{\Psi}$.

Corollary 3 is proved in Section 7.5. The generalization error with trigonometric basis encoder in Corollary 3 is similar to the error with Legendre polynomials in Corollary 2. If only the global smoothness of input and output functions is assumed, the generalization error decays at a low rate. A faster rate can be achieved if we exploit the low-dimensional structures of the input and output functions.

5 Generalization error for PCA encoders and decoders

When the given data are concentrated near a low-dimensional subspace, PCA is an effective tool for dimension reduction. In this section, we consider the PCA encoder, where the orthonormal basis is estimated from the training data.

5.1 PCA encoders and decoders

Let ρ be a probability measure on a separable Hilbert space \mathcal{H} . Define the covariance operator with respect to ρ as

$$G_{\rho} = \mathbb{E}_{u \sim \rho} [u \otimes u], \tag{26}$$

where \otimes denotes the outer product $(f \otimes g)(h) = \langle g, h \rangle_{\mathcal{H}} f$ for any $f, g, h \in \mathcal{H}$, and $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ denotes the inner product in \mathcal{H} . Let $\{\lambda_k\}_{k=1}^{\infty}$ be the eigenvalues of G_{ρ} in a non-increasing order, and ϕ_k be the eigenfunction associated with λ_k . For any $u \in \mathcal{H}$, we have

$$u = \sum_{j=1}^{\infty} \langle u, \phi_j \rangle_{\mathcal{H}} \phi_j.$$

For a fixed positive integer d , the eigenfunctions $\{\phi_k\}_{k=1}^d$ associated with the top d eigenvalues are called the first d principal components. Fixing d , we define the encoder operator $E_{\mathcal{H},d} : \mathcal{H} \rightarrow \mathbb{R}^d$ as

$$E_{\mathcal{H},d}(u) = [\langle u, \phi_1 \rangle, \langle u, \phi_2 \rangle, \dots, \langle u, \phi_d \rangle]^\top, \text{ for any } u \in \mathcal{H}, \quad (27)$$

which gives rise to the coefficients of u associated with the first d principal components. The decoder $D_{\mathcal{H},d} : \mathbb{R}^d \rightarrow \mathcal{H}$ is defined as

$$D_{\mathcal{H},d}(\mathbf{a}) = \sum_{j=1}^d a_j \phi_j, \text{ for any } \mathbf{a} = [a_1, \dots, a_d]^\top \in \mathbb{R}^d. \quad (28)$$

Given n i.i.d samples $\{u_i\}_{i=1}^n$ from ρ , the empirical covariance operator is

$$G_\rho^n = \frac{1}{n} \sum_{i=1}^n u_i \otimes u_i. \quad (29)$$

Let $\{\lambda_k^n\}_{k=1}^\infty$ be the eigenvalues of G_ρ^n in a non-increasing order, and ϕ_k^n be the eigenfunction associated with λ_k^n . We define the empirical encoder $E_{\mathcal{H},d}^n : \mathcal{H} \rightarrow \mathbb{R}^d$ as

$$E_{\mathcal{H},d}^n(u) = [\langle u, \phi_1^n \rangle, \langle u, \phi_2^n \rangle, \dots, \langle u, \phi_d^n \rangle]^\top \text{ for any } u \in \mathcal{H}. \quad (30)$$

The empirical decoder is

$$D_{\mathcal{H},d}^n(\mathbf{a}) = \sum_{j=1}^d a_j \phi_j^n \text{ for any } \mathbf{a} \in \mathbb{R}^d. \quad (31)$$

The PCA encoders and decoders $E_{\mathcal{H},d}, D_{\mathcal{H},d}, E_{\mathcal{H},d}^n, D_{\mathcal{H},d}^n$ are Lipchitz operators with a Lipchitz constant 1.

Lemma 2. Let \mathcal{H} be a separable Hilbert space and ρ be a probability measure on \mathcal{H} . For any integer $d > 0$, let $E_{\mathcal{H},d}$ and $D_{\mathcal{H},d}$ be the PCA encoder and decoder and $E_{\mathcal{H},d}^n$ and $D_{\mathcal{H},d}^n$ be their empirical counterparts. Then we have

$$\begin{aligned} \|E_{\mathcal{H},d}^n(u) - E_{\mathcal{H},d}^n(\tilde{u})\|_2 &\leq \|u - \tilde{u}\|_{\mathcal{H}}, \text{ for any } u, \tilde{u} \in \mathcal{H}, \\ \|D_{\mathcal{H},d}^n(\mathbf{a}) - D_{\mathcal{H},d}^n(\tilde{\mathbf{a}})\|_{\mathcal{H}} &= \|\mathbf{a} - \tilde{\mathbf{a}}\|_2, \text{ for any } \mathbf{a}, \tilde{\mathbf{a}} \in \mathbb{R}^d. \end{aligned}$$

Lemma 2 can be proved in the same way as Lemma 1. The proof is omitted here.

5.2 Generalization error with PCA encoders and decoders

In this subsection, we choose PCA encoders and decoders for \mathcal{X} and \mathcal{Y} . For the \mathcal{X} space, we define the covariance operator and its empirical counterpart as

$$G_\gamma = \mathbb{E}_{u \sim \gamma} u \otimes u \quad \text{and} \quad G_\gamma^n = \frac{1}{n} \sum_{i=1}^n u_i \otimes u_i.$$

Let $\{\phi_{\gamma,k}\}_{k=1}^{d_\mathcal{X}}$ and $\{\phi_{\gamma,k}^n\}_{k=1}^{d_\mathcal{X}}$ be the first $d_\mathcal{X}$ principle components of G_γ and G_γ^n , respectively. The PCA encoder and its empirical counterpart are given as

$$E_{\mathcal{X}}(u) = \left[\langle u, \phi_{\gamma,1} \rangle, \langle u, \phi_{\gamma,2} \rangle, \dots, \langle u, \phi_{\gamma,d_\mathcal{X}} \rangle \right]^\top, \quad D_{\mathcal{X}}(\mathbf{a}) = \sum_{j=1}^d a_j \phi_{\gamma,j}, \quad (32)$$

$$E_{\mathcal{X}}^n(u) = \left[\langle u, \phi_{\gamma,1}^n \rangle, \langle u, \phi_{\gamma,2}^n \rangle, \dots, \langle u, \phi_{\gamma,d_\mathcal{X}}^n \rangle \right]^\top, \quad D_{\mathcal{X}}^n(\mathbf{a}) = \sum_{j=1}^d a_j \phi_{\gamma,j}^n \quad (33)$$

for any $u \in \mathcal{X}$ and $\mathbf{a} \in \mathbb{R}^{d_\mathcal{X}}$.

For the \mathcal{Y} space, the ideal covariance operator in the noiseless case is defined based on the pushforward measure $\Psi_\# \gamma$. In the noisy case, the samples $\{v_i\}_{i=1}^n$ are random copies of $\Psi(u) + \tilde{\epsilon}$. Denote the probability measure of v by ζ . The ideal and empirical covariance operators are defined as

$$G_{\Psi_\# \gamma} = \mathbb{E}_{w \sim \Psi_\# \gamma} w \otimes w \quad \text{and} \quad G_\zeta^n = \frac{1}{n} \sum_{i=1}^n v_i \otimes v_i.$$

Notice that G_ζ^n is the empirical counterpart of G_ζ , which is different from $G_{\Psi_\# \gamma}$ in the noisy case.

Let $\{\phi_{\Psi_\# \gamma,k}\}_{k=1}^{d_\mathcal{Y}}$ and $\{\phi_{\zeta,k}^n\}_{k=1}^{d_\mathcal{Y}}$ be the first $d_\mathcal{Y}$ principle components of $G_{\Psi_\# \gamma}$ and G_ζ^n , respectively. We choose the PCA encoder:

$$E_{\mathcal{Y}}(w) = \left[\langle w, \phi_{\Psi_\# \gamma,1} \rangle, \langle w, \phi_{\Psi_\# \gamma,2} \rangle, \dots, \langle w, \phi_{\Psi_\# \gamma,d_\mathcal{Y}} \rangle \right]^\top, \quad D_{\mathcal{Y}}(\mathbf{a}) = \sum_{j=1}^d a_j \phi_{\Psi_\# \gamma,j}, \quad (34)$$

$$E_{\mathcal{Y}}^n(w) = \left[\langle w, \phi_{\zeta,1}^n \rangle, \langle w, \phi_{\zeta,2}^n \rangle, \dots, \langle w, \phi_{\zeta,d_\mathcal{Y}}^n \rangle \right]^\top, \quad D_{\mathcal{Y}}^n(\mathbf{a}) = \sum_{j=1}^d a_j \phi_{\zeta,j}^n \quad (35)$$

for any $w \in \mathcal{Y}$ and $\mathbf{a} \in \mathbb{R}^{d_\mathcal{Y}}$.

The following theorem gives a bound on the generalization error of operator estimation with PCA encoders:

Theorem 3. In Setting 1, suppose Assumption 1–2 and 4 hold. Consider the PCA encoders and decoders defined in (32)–(35). Let $\{\lambda_k\}_{k=1}^\infty$ be the eigenvalues of the covariance operator $G_{\Psi_\# \gamma}$ in nonincreasing order. Let Γ_{NN} be the minimizer of (6) with the network architecture $\mathcal{F}(d_\mathcal{Y}, L, p, M)$ in (3), where L, p, M are set as in (19). We have

$$\begin{aligned} & \mathbb{E}_{\mathcal{S}} \mathbb{E}_{u \sim \gamma} \|D_{\mathcal{Y}}^n \circ \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u) - \Psi(u)\|_{\mathcal{Y}}^2 \\ & \leq C_4 (\sigma^2 + R_{\mathcal{Y}}^2) d_{\mathcal{Y}}^{\frac{4+d_\mathcal{X}}{2+d_\mathcal{X}}} n^{-\frac{2}{2+d_\mathcal{X}}} \log^6 n + 8 \left(4R_{\mathcal{X}}^2 L_{\Psi}^2 \sqrt{d_\mathcal{X}} + (R_{\mathcal{Y}} + \tilde{\sigma})^2 \sqrt{d_\mathcal{Y}} \right) n^{-\frac{1}{2}} + 16\tilde{\sigma}^2 \left(\frac{\tilde{\sigma}}{\lambda_{d_\mathcal{Y}} - \lambda_{d_\mathcal{Y}+1}} \right)^2 (R_{\mathcal{Y}} + \tilde{\sigma})^2 \\ & \quad + 20\tilde{\sigma}^2 + 16L_{\Psi}^2 \mathbb{E}_{u \sim \gamma} \|\Pi_{\mathcal{X},d_\mathcal{X}}(u) - u\|_2^2 + 16\mathbb{E}_{w \sim \Psi_\# \gamma} \|\Pi_{\mathcal{Y},d_\mathcal{Y}}(w) - w\|_{\mathcal{Y}}^2 \end{aligned}$$

where C_4 is a constant depending on $d_\mathcal{X}, R_{\mathcal{X}}, R_{\mathcal{Y}}, L_{\Psi}$.

Theorem 3 is proved in Section 7.6. PCA is effective when the input and output samples are concentrated near low-dimensional subspaces. In this case, an orthonormal basis of the subspace

is estimated from the samples. Since the PCA encoder and decoder are data-driven, we expect the corresponding projection errors are smaller than those by Legendre polynomials or trigonometric functions.

In the generalization error in Theorem 3, the error $16\tilde{\sigma}^2 \left(\frac{\tilde{\sigma}}{\lambda_{d_y} - \lambda_{d_y+1}} \right)^2 (R_y + \tilde{\sigma})^2 + 20\tilde{\sigma}^2$ does not decay as n increases. This is because PCA extracts the principal components from noisy data but does not denoise the data set without additional assumptions on noise. If the noise does not perturb the space spanned by the first d_y principal eigenfunctions of $G_{\Psi_{\#}\gamma}$, the constant terms can be dropped as the following corollary.

Corollary 4. Under the conditions of Theorem 3, if the eigenspace spanned by the first d_y principal eigenfunctions of G_{μ} coincides with that of $G_{\Psi_{\#}\gamma}$, then we have

$$\begin{aligned} & \mathbb{E}_{\mathcal{S}} \mathbb{E}_{u \sim \gamma} \|D_y^n \circ \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u) - \Psi(u)\|_y^2 \\ & \leq C_4(\sigma^2 + R_y^2) d_y^{\frac{4+d_{\mathcal{X}}}{2+d_{\mathcal{X}}}} n^{-\frac{2}{2+d_{\mathcal{X}}}} \log^6 n + 8 \left(4R_{\mathcal{X}}^2 L_{\Psi}^2 \sqrt{d_{\mathcal{X}}} + (R_y + \tilde{\sigma})^2 \sqrt{d_y} \right) n^{-\frac{1}{2}} \\ & \quad + 16L_{\Psi}^2 \mathbb{E}_{u \sim \gamma} \|\Pi_{\mathcal{X}, d_{\mathcal{X}}}(u) - u\|_2^2 + 16\mathbb{E}_{w \sim \Psi_{\#}\gamma} \|\Pi_{\mathcal{Y}, d_y}(w) - w\|_y^2. \end{aligned}$$

Corollary 4 is proved in Section 7.7.

6 Exploit additional low-dimensional structures

Section 4 and Section 5 are suitable for the case where the input and output samples are concentrated near a low-dimensional subspace. While in practice, the low-dimensional subspace is not a priori known. In order to capture such a subspace, we need to choose a large encoding dimension so that the low-dimensional subspace is enclosed by the encoded space, which guarantees a small projection error. However, the network estimation error (see Section 3.2 for the definition) has an exponential dependence on $d_{\mathcal{X}}$. The error decays slowly when $d_{\mathcal{X}}$ is large.

Additionally, the given data may be located on a low-dimensional manifold enclosed by the encoded space, or the operator Ψ may have low complexity. In this section, we will exploit such additional low-dimensional structures. We will show that, even though $d_{\mathcal{X}}$ and d_y are chosen to be large in order to guarantee small projection errors, the exponent in the network estimation error only depends on the intrinsic dimension of the additional low-dimensional structures of data, instead of $d_{\mathcal{X}}$. Specifically, we consider two scenarios : (1) when the collection of encoded vectors $E_{\mathcal{X}}(\Omega_{\mathcal{X}})$ is on a low-dimensional manifold and (2) when the operator Ψ only depends on a few directions in the encoded space.

6.1 When encoded vectors lie on a low-dimensional manifold

We first consider the case when the given data exhibit a nonlinear low-dimensional structure: For a given encoder $E_{\mathcal{X}} : \mathcal{X} \rightarrow \mathbb{R}^{d_{\mathcal{X}}}$, the encoded vectors $\{E_{\mathcal{X}}(u) : u \text{ is randomly sampled from } \gamma\}$ lie on a d_0 -dimensional manifold with $d_0 \ll d_{\mathcal{X}}$. This scenario is observed in many applications. For example, the solutions of most PDEs are in an infinite-dimensional function space. After uniform discretization, the solutions are encoded to vectors in a very high dimensional space. For many

PDEs, it is commonly observed that the solutions actually lie on a low-dimensional manifold enclosed by the discretized high-dimensional space. Therefore the solution manifold can be well-approximated using much fewer bases than those used in the discretization. This observation leads to the success of the reduced basis method [37, 73]. Another concrete example is described as follows:

Example 1. Let $\mathcal{X} = L^2([-1, 1])$ and $d_0, d_{\mathcal{X}}$ be positive integers such that $d_0 < d_{\mathcal{X}}$. Let $\{T_k\}_{k=1}^{\infty}$ be the trigonometric functions defined in (24) and $\{g_k\}_{k=d_0+1}^{d_{\mathcal{X}}}$ be some real valued functions. Suppose the probability measure γ is supported on

$$\Omega_{\mathcal{X}} = \left\{ u : u = \sum_{k=1}^{d_{\mathcal{X}}} a_k T_k \text{ with } a_k \in \mathbb{R} \text{ for } k = 1, \dots, d_0, \text{ and } a_k = g_k(a_1, \dots, a_{d_0}) \text{ for } k = d_0 + 1, \dots, d_{\mathcal{X}} \right\}.$$

The support set $\Omega_{\mathcal{X}}$ has an intrinsic dimension d_0 . If we choose the basis encoder $E_{\mathcal{X}} : \mathcal{X} \rightarrow \mathbb{R}^{d_{\mathcal{X}}}$ using the trigonometric functions $\{T_k\}_{k=1}^{d_{\mathcal{X}}}$, then the encoded vectors $\{E_{\mathcal{X}}(u) : u \text{ is randomly sampled from } \gamma\}$ lie on a d_0 -dimensional manifold embedded in $\mathbb{R}^{d_{\mathcal{X}}}$. Figure 1 shows this manifold when $d_{\mathcal{X}} = 3, d_0 = 2$ and $g_3 = a_1^2 + a_2$.

This nonlinear low-dimensional structure of data can be described as follows:

Assumption 8. Let $d_0, d_{\mathcal{X}}$ be positive integers such that $d_0 < d_{\mathcal{X}}$. In Setting 1, there exists an encoder $E_{\mathcal{X}} : \mathcal{X} \rightarrow \mathbb{R}^{d_{\mathcal{X}}}$ such that the encoded vectors $\{E_{\mathcal{X}}(u) : u \text{ is randomly sampled from } \gamma\}$ is on a d_0 -dimensional compact smooth Riemannian manifold \mathcal{M} isometrically embedded in $\mathbb{R}^{d_{\mathcal{X}}}$. The reach of \mathcal{M} [31, 67] is $\tau > 0$.

Under Assumption 8 and Setting 1, the output $\Psi(u)$ is perturbed by noise, while the input u is clean and its encoded vector is located on \mathcal{M} . Such a setting is common in practice when

a series of experiments is conducted to simulate a scientific phenomenon. In experiments, one designs the inputs and takes measurements of the outputs. Usually, the inputs are generated according to some physical laws that lead to low-dimensional structures. Due to the limitations of sensors and equipment, the measured outputs are perturbed by noise.

Approximation and statistical estimation theories of deep neural networks for functions on a low-dimensional manifold have been studied in [12, 13, 14, 56, 79, 44, 17, 77, 74, 23, 64]. In this subsection, we show that deep neural networks can automatically adapt to nonlinear low-dimensional structures of data, and give rise to a sample complexity depending on the intrinsic dimension d_0 . The following theorem gives a generalization error in this scenario.




Figure 1: An illustration of Example 1 with $d_{\mathcal{X}} = 3, d_0 = 2$ and $g_3 = a_1^2 + a_2$.

Theorem 4. In Setting 1, suppose Assumption 1–5 and 8 hold, and the encoder $E_{\mathcal{X}}$ in Assumption 8 is given. Let Γ_{NN} be the minimizer of (6) with the network architecture $\mathcal{F}(d_y, L, p, M)$ in (3) with

$$L = O(\tilde{L} \log \tilde{L}), \quad p = O(d_{\mathcal{X}} \tilde{p} \log \tilde{p}), \quad M = \sqrt{d_y} L_{E_y^n} R_y, \quad (36)$$

where $\tilde{L}, \tilde{p} > 0$ are positive integers satisfying

$$\tilde{L} \tilde{p} = \left\lceil d_y^{-\frac{d_0}{4+2d_0}} n^{\frac{d_0}{4+2d_0}} \right\rceil. \quad (37)$$

Then we have

$$\begin{aligned} & \mathbb{E}_{\mathcal{S}} \mathbb{E}_{u \sim \gamma} \|D_y^n \circ \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u) - \Psi(u)\|_y^2 \\ & \leq C_5(\sigma^2 + R_y^2) d_y^{\frac{4+d_0}{2+d_0}} d_{\mathcal{X}}^2 n^{-\frac{2}{2+d_0}} \log^6 n + C_3 \mathbb{E}_{\mathcal{S}} \mathbb{E}_{u \sim \gamma} \|\Pi_{\mathcal{X}, d_{\mathcal{X}}}(u) - u\|_{\mathcal{X}}^2 + 2 \mathbb{E}_{\mathcal{S}} \mathbb{E}_{w \sim \Psi_{\#} \gamma} \|\Pi_{\mathcal{Y}, d_y}^n(w) - w\|_y^2, \end{aligned} \quad (38)$$

where C_5 depends on $d_0, \log d_{\mathcal{X}}, R_{\mathcal{X}}, R_y, L_{E_{\mathcal{X}}^n}, L_{E_y^n}, L_{D_{\mathcal{X}}}, L_{D_y^n}, L_{\Psi}, \tau$, the surface area of \mathcal{M} , and $C_3 = 16L_{D_y^n}^2 L_{E_y^n}^2 L_{\Psi}^2$.

Theorem 4 is proved in Section 7.8. The convergence rate in Theorem 4 has an exponential dependence on d_0 , instead of $d_{\mathcal{X}}$. Theorem 4 shows that when the encoded vectors are located on a low-dimensional manifold, deep neural networks are adaptive to such nonlinear geometric structures of data.

6.2 When the operator Ψ has low complexity

In our framework, learning Ψ is converted to learning the transformation $\Gamma : \mathbb{R}^{d_{\mathcal{X}}} \rightarrow \mathbb{R}^{d_y}$, as defined in (5). The second scenario we consider in this subsection is that, even though the u_i 's and v_i 's are in infinite-dimensional spaces, the operator Ψ has low complexity: its corresponding transformation Γ can be approximated by some low-dimensional functions that only depend on few directions in $\mathbb{R}^{d_{\mathcal{X}}}$. For example, consider solving a linear PDE with constant coefficients by the Fourier spectral method. In this case, the operator Ψ is the PDE solver that maps initial conditions to solutions at certain time. By taking the Fourier transform on both sides of the PDE, solving the PDEs is converted to solving a series of independent ODEs, each of which controls the evolution of a Fourier coefficient of the solution [78, Chapter 2]. The operator Ψ can be fully characterized by a system of one-dimensional ODEs. We next adapt this setting to our framework in order to learn Ψ . We use trigonometric functions as our encoders and decoders: the initial conditions and solutions are approximated by the first $d_{\mathcal{X}} = d_y$ terms of their Fourier series expansion. Then learning Ψ reduces to learning d_y one-dimensional functions, each of which corresponds to an ODE of a Fourier coefficient, instead of learning d_y $d_{\mathcal{X}}$ -dimensional functions.

In this subsection, we show that we can get a faster rate by exploiting the low complexity of Ψ . We first make an assumption on Ψ :

Assumption 9. Let $0 < d_0 \leq d_{\mathcal{X}}$ be integers. Assume there exist $E_{\mathcal{X}}, D_{\mathcal{X}}, E_{\mathcal{Y}}, D_{\mathcal{Y}}$ such that for any $u \in \Omega_{\mathcal{X}}$, we have

$$\Pi_{\mathcal{Y}, d_{\mathcal{Y}}} \circ \Psi(u) = D_{\mathcal{Y}} \circ \mathbf{g} \circ E_{\mathcal{X}}(u) \quad (39)$$

with $\mathbf{g} : \mathbb{R}^{d_{\mathcal{X}}} \rightarrow \mathbb{R}^{d_{\mathcal{Y}}}$ in the form:

$$\mathbf{g}(\mathbf{a}) = \left[g_1(V_1^{\top} \mathbf{a}) \quad \cdots \quad g_{d_{\mathcal{Y}}}(V_{d_{\mathcal{Y}}}^{\top} \mathbf{a}) \right]^{\top}, \quad (40)$$

for some unknown matrix $V_k \in \mathbb{R}^{d_{\mathcal{X}} \times d_0}$, and some unknown real valued function $g_k : \mathbb{R}^{d_0} \rightarrow \mathbb{R}$ where $k = 1, \dots, d_{\mathcal{Y}}$.

In statistics, the functions g_k 's in Assumption 9 are known as single-index models for $d_0 = 1$, and are known as multi-index models for $d_0 > 1$. For any given $u \sim \gamma$, we decompose $\Psi(u)$ into two parts: the first part is its projection to the set of encoded vectors $E_{\mathcal{Y}}(\Omega_{\mathcal{Y}})$; the second part is the rest orthogonal to the first part. Assumption 9 assumes that the operator mapping u to the first part follows a multi-index model. When $d_{\mathcal{X}}$ is large enough, the second part has a small magnitude and is included in the projection error. In the following example, we give a simple illustration when the second part vanishes.

Example 2. Let $\mathcal{X} = L^2([-1, 1])$, $\Omega_{\mathcal{X}} \subset \mathcal{X}$ be a compact set in \mathcal{X} and $0 < d_0 < d_{\mathcal{X}}$ be integers. Let $\{T_k\}_{k=1}^{\infty}$ be trigonometric functions defined in (24). Any $u \in \Omega_{\mathcal{X}}$ can be written as $u = \sum_{k=1}^{\infty} a_k T_k$ for some a_k 's. Denote $\mathbf{a}_u = [a_1 \quad \cdots \quad a_{d_{\mathcal{X}}}]^{\top}$. Suppose the operator we want to learn has the following form

$$\Psi(u) = \sum_{k=1}^{d_{\mathcal{Y}}} g_k(V_k^{\top} \mathbf{a}_u) T_k, \quad (41)$$

with $V_k \in \mathbb{R}^{d_{\mathcal{X}} \times d_0}$ and $g_k : \mathbb{R}^{d_0} \rightarrow \mathbb{R}$ for $k = 1, \dots, d_{\mathcal{Y}}$. We set $E_{\mathcal{X}}, D_{\mathcal{X}}$ as the basis encoder and decoder using the basis functions $\{T_k\}_{k=1}^{d_{\mathcal{X}}}$, and $E_{\mathcal{Y}}, D_{\mathcal{Y}}$ as encoder and decoder derived using basis $\{T_k\}_{k=1}^{d_{\mathcal{Y}}}$. In this example, $\Pi_{\mathcal{Y}, d_{\mathcal{Y}}} \circ \Psi(u) = \Psi(u)$ for any $u \sim \gamma$. Then learning Ψ reduces to learning the g_k 's and the V_k 's. An illustration of the estimator is shown in Figure 2. In neural networks, the V_k 's can be realized by a single layer. Therefore, our major task is to learn good approximations of the g_k 's. Note that each g_k is a d_0 -dimensional function. By exploiting such low complexity of the operator, we can convert the learning task from learning $d_{\mathcal{Y}} d_{\mathcal{X}}$ -dimensional functions to learning $d_{\mathcal{Y}} d_0$ -dimensional functions.

With Assumption 9, the following theorem gives a faster rate on the generalization error:

Theorem 5. In Setting 1, suppose Assumption 1–5 and 9 hold. Assume that the encoders and decoders $E_{\mathcal{X}}, D_{\mathcal{X}}, E_{\mathcal{Y}}, D_{\mathcal{Y}}$ in Assumption 9 are given. Let Γ_{NN} be the minimizer of (6) with the network architecture $\mathcal{F}(d_{\mathcal{Y}}, L, p, M)$ in (3), where

$$L = O(\tilde{L} \log \tilde{L}), \quad p = O(\tilde{p} \log \tilde{p}), \quad M = \sqrt{d_{\mathcal{Y}}} L_{E_{\mathcal{Y}}} R_{\mathcal{Y}} \quad (42)$$

and $\tilde{L}, \tilde{p} > 0$ are integers and satisfy (37).

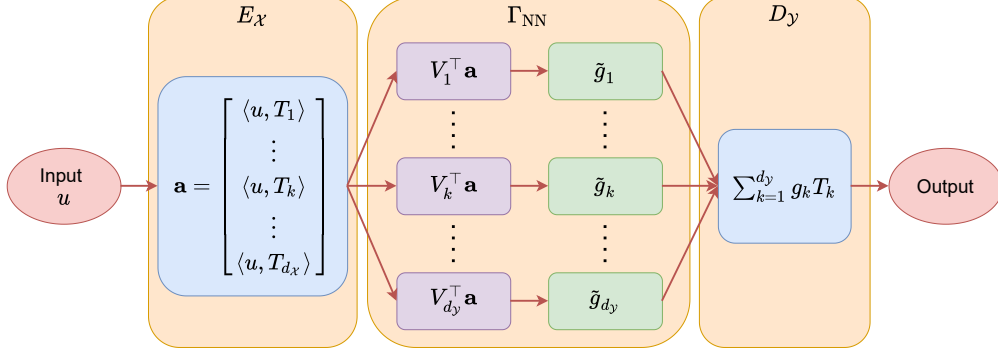


Figure 2: An illustration of Example 2, where the \tilde{g}_k 's represent network approximations of the g_k 's in (41).

We have

$$\begin{aligned}
& \mathbb{E}_{\mathcal{S}} \mathbb{E}_{u \sim \gamma} \|D_Y^n \circ \Gamma_{\text{NN}} \circ E_X^n(u) - \Psi(u)\|_Y^2 \\
& \leq C_6 (\sigma^2 + R_Y^2) d_Y^{\frac{4+d_0}{2+d_0}} \max \left\{ n^{-\frac{2}{2+d_0}}, d_X n^{-\frac{4+d_0}{4+2d_0}} \right\} \log^6 n \\
& \quad + C_3 \mathbb{E}_{\mathcal{S}} \mathbb{E}_{u \sim \gamma} \|\Pi_{\mathcal{X}, d_X}(u) - u\|_{\mathcal{X}}^2 + 2 \mathbb{E}_{\mathcal{S}} \mathbb{E}_{w \sim \Psi_{\#} \gamma} \|\Pi_{Y, d_Y}(w) - w\|_Y^2.
\end{aligned} \tag{43}$$

where C_6 depends on $d_0, \log d_X, R_X, R_Y, L_{E_X^n}, L_{E_Y^n}, L_{D_X^n}, L_{D_Y^n}, L_{\Psi}$, and $C_3 = 16L_{D_Y^n}^2 L_{E_Y^n}^2 L_{\Psi}^2$.

Theorem 5 is proved in Section 7.9. In Assumption 9, each V_k is a linear transformation that can be realized by a singly layer. In our network construction, the first layer is used to learn these transformations and the rest is used to learn the functions g_k 's.

7 Proof of main results

7.1 Preliminaries

In this section, we define several quantities that will be used in the proof. We first define two types of covering number of function classes. The first type is independent of data and will be used to prove Theorem 1.

Definition 2 (Cover). Let \mathcal{F} be a class of functions. A set of functions \mathcal{S} is a δ -cover of \mathcal{F} with respect to a norm $\|\cdot\|$ if for any $f \in \mathcal{F}$, one has

$$\inf_{f^* \in \mathcal{S}} \|f - f^*\| \leq \delta.$$

Definition 3 (Covering number, Definition 2.1.5 of [91]). Let \mathcal{F} be a class of functions. For any $\delta > 0$, the covering number of \mathcal{F} is defined as

$$\mathcal{N}(\delta, \mathcal{F}, \|\cdot\|) = \min\{|\mathcal{S}_f| : \mathcal{S}_f \text{ is a } \delta\text{-cover of } \mathcal{F}\},$$

where $|\mathcal{S}_f|$ denotes the cardinality of \mathcal{S}_f .

Definition 2 and 3 depend on the norm $\|\cdot\|$. In the following, we choose $\|\cdot\|$ as a sample dependent norm and define the so-called uniform covering number. We first define the cover with respect to samples:

Definition 4 (Cover with respect to samples). Let \mathcal{F} be a class of functions from \mathbb{R}^{d_1} to \mathbb{R}^{d_2} . Given a set of samples $X = \{\mathbf{x}_k\}_{k=1}^m \subset \mathbb{R}^{d_1}$, for any $\delta > 0$, a function set $\mathcal{S}_f(X)$ is a δ -cover of F with respect to X if for any $f \in \mathcal{F}$, there exists $f^* \in \mathcal{S}_f(X)$ such that

$$\|f(\mathbf{x}_k) - f^*(\mathbf{x}_k)\|_\infty \leq \delta, \quad \forall 1 \leq k \leq m.$$

Definition 4 is a special case of Definition 2 in which the norm $\|\cdot\|$ is chosen as the ℓ^∞ norm of the collection of its argument's values over samples X . Based on Definition 4, we define the uniform covering number as follows:

Definition 5 (Uniform covering number, Section 10.2 of [2]). Let \mathcal{F} be a class of functions from \mathbb{R}^d to \mathbb{R} . For any set of samples $X = \{\mathbf{x}_k\}_{k=1}^m \subset \mathbb{R}^d$, denote

$$\mathcal{F}|_X = \{(f(\mathbf{x}_1), \dots, f(\mathbf{x}_m)) : f \in \mathcal{F}\}.$$

For any $\delta > 0$, the uniform covering number of \mathcal{F} with m samples is defined as

$$\mathcal{N}(\delta, \mathcal{F}, m) = \max_{X \subset \mathbb{R}^d, |X|=m} \min_{\mathcal{S}_f(X)} \{|\mathcal{S}_f(X)| : \mathcal{S}_f(X) \text{ is a } \delta\text{-cover of } \mathcal{F} \text{ with respect to } X\}. \quad (44)$$

This covering number is used to prove Theorem 2.

7.2 Proof of Theorem 1

To prove Theorem 1, we first decompose the squared L^2 error $\mathbb{E}_{\mathcal{S}} \mathbb{E}_{u \sim \gamma} \|D_{\mathcal{Y}}^n \circ \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u) - \Psi(u)\|_{\mathcal{Y}}^2$ into a network estimation error and a projection error. The network estimation error can be further decomposed into a bias term and a variance term. The bias term heavily depends on the approximation error of the network class (2). The variance term is upper bounded in terms of the covering number of the network class.

Proof of Theorem 1. We first decompose the squared L^2 error as

$$\begin{aligned} & \mathbb{E}_{\mathcal{S}} \mathbb{E}_{u \sim \gamma} \left[\left\| D_{\mathcal{Y}}^n \circ \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u) - \Psi(u) \right\|_{\mathcal{Y}}^2 \right] \\ & \leq \underbrace{2\mathbb{E}_{\mathcal{S}} \mathbb{E}_{u \sim \gamma} \left[\left\| D_{\mathcal{Y}}^n \circ \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u) - D_{\mathcal{Y}}^n \circ E_{\mathcal{Y}}^n \circ \Psi(u) \right\|_{\mathcal{Y}}^2 \right]}_{\text{I}} + \underbrace{2\mathbb{E}_{\mathcal{S}} \mathbb{E}_{u \sim \gamma} \left[\left\| D_{\mathcal{Y}}^n \circ E_{\mathcal{Y}}^n \circ \Psi(u) - \Psi(u) \right\|_{\mathcal{Y}}^2 \right]}_{\text{II}}. \end{aligned} \quad (45)$$

Here I is the network estimation error in the \mathcal{Y} space, II is the empirical projection error, which can be rewritten as

$$\text{II} = 2\mathbb{E}_{\mathcal{S}} \mathbb{E}_{w \sim \Psi_{\#} \gamma} \left[\left\| \Pi_{\mathcal{Y}, d_{\mathcal{Y}}}^n(w) - w \right\|_{\mathcal{Y}}^2 \right]. \quad (46)$$

In the remaining of this subsection, we derive an upper bound of I. Note that I can be bounded as

$$\begin{aligned} I &= 2\mathbb{E}_{\mathcal{S}}\mathbb{E}_{u\sim\gamma}\left[\left\|D_{\mathcal{Y}}^n\circ\Gamma_{\text{NN}}\circ E_{\mathcal{X}}^n(u)-D_{\mathcal{Y}}^n\circ E_{\mathcal{Y}}^n\circ\Psi(u)\right\|_{\mathcal{Y}}^2\right] \\ &\leq 2L_{D_{\mathcal{Y}}^n}^2\mathbb{E}_{\mathcal{S}}\mathbb{E}_{u\sim\gamma}\left[\left\|\Gamma_{\text{NN}}\circ E_{\mathcal{X}}^n(u)-E_{\mathcal{Y}}^n\circ\Psi(u)\right\|_2^2\right]. \end{aligned} \quad (47)$$

If the training samples in \mathcal{S}_1 are fixed, we have the following conditioned on \mathcal{S}_1 :

$$\begin{aligned} &\mathbb{E}_{\mathcal{S}_2}\mathbb{E}_{u\sim\gamma}\left[\left\|\Gamma_{\text{NN}}\circ E_{\mathcal{X}}^n(u)-E_{\mathcal{Y}}^n\circ\Psi(u)\right\|_2^2\right] \\ &= 2\mathbb{E}_{\mathcal{S}_2}\underbrace{\left[\frac{1}{n}\sum_{i=n+1}^{2n}\left\|\Gamma_{\text{NN}}\circ E_{\mathcal{X}}^n(u_i)-E_{\mathcal{Y}}^n\circ\Psi(u_i)\right\|_2^2\right]}_{T_1} \\ &\quad + \underbrace{\mathbb{E}_{\mathcal{S}_2}\mathbb{E}_{u\sim\gamma}\left[\left\|\Gamma_{\text{NN}}\circ E_{\mathcal{X}}^n(u)-E_{\mathcal{Y}}^n\circ\Psi(u)\right\|_2^2\right]-\mathbb{E}_{\mathcal{S}_2}\left[\frac{2}{n}\sum_{i=n+1}^{2n}\left\|\Gamma_{\text{NN}}\circ E_{\mathcal{X}}^n(u_i)-E_{\mathcal{Y}}^n\circ\Psi(u_i)\right\|_2^2\right]}_{T_2}. \end{aligned} \quad (48)$$

In the decomposition of (48), the term T_1 consists of the bias of using neural network to approximate the transformation Γ and the projection error of $\Pi_{\mathcal{X},d_{\mathcal{X}}}^n$ in the \mathcal{X} space. The term T_2 captures the variance. We next derive bounds for T_1 and T_2 respectively.

Upper bound of T_1 . The term T_1 is the expected mean squared error of the learned transformation Γ_{NN} with respect to \mathcal{S}_2 . We will derive an upper bound using the network approximation error and network architecture's covering number. The network approximation error is the bias. We use network architecture's covering number to bound the stochastic error.

Define the transformation $\Gamma_d^n : \mathbb{R}^{d_{\mathcal{X}}} \rightarrow \mathbb{R}^{d_{\mathcal{Y}}}$

$$\Gamma_d^n = E_{\mathcal{Y}}^n \circ \Psi \circ D_{\mathcal{X}}^n, \quad (49)$$

which maps the encoded vector $E_{\mathcal{X}}^n(u)$ in \mathcal{X} to the encoded vector $E_{\mathcal{Y}}^n(v)$ in \mathcal{Y} . The transformation Γ_d^n is the target transformation to be estimated by Γ_{NN} . It is straightforward to show that Γ_d^n is a Lipschitz transformation (see a proof of Lemma 3 in Appendix C).

Lemma 3. Assume Assumption 2 and 3. Γ_d^n is Lipschitz with a Lipschitz constant $L_{E_{\mathcal{Y}}^n}L_{D_{\mathcal{X}}^n}L_{\Psi}$.

Denote

$$\epsilon_i = E_{\mathcal{Y}}^n(v_i) - E_{\mathcal{Y}}^n(\Psi(u_i)). \quad (50)$$

According to Assumption 3 and Assumption 4(iii)–(iv), we have

$$\mathbb{E}[\epsilon_i] = \mathbf{0}, \text{ and } \|\epsilon_i\|_{\infty} < \sigma.$$

We decompose T_1 as

$$\begin{aligned}
T_1 &= 2\mathbb{E}_{\mathcal{S}_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} \left\| \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i) - E_{\mathcal{Y}}^n \circ \Psi(u_i) \right\|_2^2 \right] \\
&= 2\mathbb{E}_{\mathcal{S}_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} \left\| \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i) - E_{\mathcal{Y}}^n \circ \Psi(u_i) - \epsilon_i + \epsilon_i \right\|_2^2 \right] \\
&= 2\mathbb{E}_{\mathcal{S}_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} \left\| \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i) - E_{\mathcal{Y}}^n \circ \Psi(u_i) - \epsilon_i \right\|_2^2 \right] \\
&\quad + 4\mathbb{E}_{\mathcal{S}_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} \langle \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i) - E_{\mathcal{Y}}^n \circ \Psi(u_i) - \epsilon_i, \epsilon_i \rangle \right] + 2\mathbb{E}_{\mathcal{S}_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} \|\epsilon_i\|_2^2 \right] \\
&= 2\mathbb{E}_{\mathcal{S}_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} \left\| \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i) - E_{\mathcal{Y}}^n(v_i) \right\|_2^2 \right] + 4\mathbb{E}_{\mathcal{S}_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} \langle \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i), \epsilon_i \rangle \right] - 2\mathbb{E}_{\mathcal{S}_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} \|\epsilon_i\|_2^2 \right] \\
&= 2\mathbb{E}_{\mathcal{S}_2} \left[\inf_{\Gamma \in \mathcal{F}_{\text{NN}}} \frac{1}{n} \sum_{i=n+1}^{2n} \left\| \Gamma \circ E_{\mathcal{X}}^n(u_i) - E_{\mathcal{Y}}^n(v_i) \right\|_2^2 \right] \\
&\quad + 4\mathbb{E}_{\mathcal{S}_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} \langle \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i), \epsilon_i \rangle \right] - 2\mathbb{E}_{\mathcal{S}_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} \|\epsilon_i\|_2^2 \right] \quad \text{by the definition of } \Gamma_{\text{NN}} \text{ in (6)} \\
&\leq 2 \inf_{\Gamma \in \mathcal{F}_{\text{NN}}} \mathbb{E}_{\mathcal{S}_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} \left\| \Gamma \circ E_{\mathcal{X}}^n(u_i) - E_{\mathcal{Y}}^n(v_i) \right\|_2^2 \right] + 4\mathbb{E}_{\mathcal{S}_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} \langle \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i), \epsilon_i \rangle \right] \\
&\quad - 2\mathbb{E}_{\mathcal{S}_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} \|\epsilon_i\|_2^2 \right] \\
&= 2 \inf_{\Gamma \in \mathcal{F}_{\text{NN}}} \mathbb{E}_{\mathcal{S}_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} \left[\left\| \Gamma \circ E_{\mathcal{X}}^n(u_i) - E_{\mathcal{Y}}^n \circ \Psi(u_i) - \epsilon_i \right\|_2^2 - \|\epsilon_i\|_2^2 \right] \right] \\
&\quad + 4\mathbb{E}_{\mathcal{S}_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} \langle \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i), \epsilon_i \rangle \right] \\
&= 2 \inf_{\Gamma \in \mathcal{F}_{\text{NN}}} \mathbb{E}_{u \sim \gamma} \left[\left\| \Gamma \circ E_{\mathcal{X}}^n(u) - E_{\mathcal{Y}}^n \circ \Psi(u) \right\|_2^2 \right] + 4\mathbb{E}_{\mathcal{S}_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} \langle \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i), \epsilon_i \rangle \right]. \tag{51}
\end{aligned}$$

In (51), the first term is the neural network approximation error, and the second term is the stochastic error from noise. To derive an upper bound of the first term, we use the following lemma which shows that for any function f in the Sobolev space $W^{k,\infty}$, when the network architecture is properly set, FNN can approximate f with arbitrary accuracy:

Lemma 4 (Theorem 1 of [93]). Let $k \geq 0$ be a positive integer. There exists an FNN architecture $\mathcal{F}_{\text{NN}}(1, L, p, K, \kappa, M)$ capable of approximating any function in $W^{k,\infty}([-B, B]^d)$, i.e., for any given $\epsilon \in (0, 1)$ and if $f \in W^{k,\infty}([-B, B]^d)$, the network architecture gives rise to a function \tilde{f} satisfying

$$\|\tilde{f} - f\|_{\infty} \leq \epsilon.$$

The hyperparameters in \mathcal{F}_{NN} are chosen as

$$L = O\left(\log \frac{1}{\varepsilon}\right), p = O\left(\varepsilon^{-\frac{d}{k}}\right), K = O\left(\varepsilon^{-\frac{d}{k}} \log \frac{1}{\varepsilon}\right), \kappa = \max\{1, B, R\}, M = R.$$

The constant hidden in $O(\cdot)$ depends on k, α, d, B, R .

Since Γ_d^n is Lipschitz by Lemma 3, according to Lemma 4 with $k = 1$, for any $\varepsilon_1 > 0$, there is a network architecture $\mathcal{F}_{\text{NN}}(d_y, L, p, K, \kappa, M)$, such that for any Γ_d^n defined in (49), there exists a $\tilde{\Gamma}_d^n \in \mathcal{F}_{\text{NN}}(d_y, L, p, K, \kappa, M)$ with

$$\|\tilde{\Gamma}_d^n - \Gamma_d^n\|_\infty \leq \varepsilon_1. \quad (52)$$

Such a network architecture has

$$\begin{aligned} L &= O(\log \varepsilon_1), p = O(\varepsilon_1^{-d_\chi}), K = O(\varepsilon_1^{-d_\chi} \log \varepsilon_1), \\ \kappa &= \max\left\{1, \sqrt{d_y} L_{E_y^n} R_y, \sqrt{d_\chi} L_{E_\chi^n} R_\chi, L_{E_y^n} L_{D_\chi^n} L_\Psi\right\}, M = \sqrt{d_y} L_{E_y^n} R_y. \end{aligned} \quad (53)$$

We bound the first term in (51) as

$$\begin{aligned} & \inf_{\Gamma \in \mathcal{F}_{\text{NN}}} \mathbb{E}_{u \sim \gamma} \left[\left\| \Gamma \circ E_\chi^n(u) - E_y^n \circ \Psi(u) \right\|_2^2 \right] \\ & \leq \mathbb{E}_{u \sim \gamma} \left[\left\| \tilde{\Gamma}_d^n \circ E_\chi^n(u) - E_y^n \circ \Psi(u) \right\|_2^2 \right] \\ & \leq 2\mathbb{E}_{u \sim \gamma} \left[\left\| \tilde{\Gamma}_d^n \circ E_\chi^n(u) - \Gamma_d \circ E_\chi^n(u) \right\|_2^2 \right] + 2\mathbb{E}_{u \sim \gamma} \left[\left\| \Gamma_d \circ E_\chi^n(u) - E_y^n \circ \Psi(u) \right\|_2^2 \right] \\ & \leq 2d_y \varepsilon_1^2 + 2\mathbb{E}_{u \sim \gamma} \left[\left\| \Gamma_d \circ E_\chi^n(u) - E_y^n \circ \Psi(u) \right\|_2^2 \right] \\ & = 2d_y \varepsilon_1^2 + 2\mathbb{E}_{u \sim \gamma} \left[\left\| E_y^n \circ \Psi \circ D_\chi^n \circ E_\chi^n(u) - E_y^n \circ \Psi(u) \right\|_2^2 \right] \quad \text{by the definition of } \Gamma_d \text{ in (49)} \\ & \leq 2d_y \varepsilon_1^2 + 2L_{E_y^n}^2 L_\Psi^2 \mathbb{E}_{u \sim \gamma} \left[\left\| D_\chi^n \circ E_\chi^n(u) - u \right\|_\chi^2 \right] \\ & = 2d_y \varepsilon_1^2 + 2L_{E_y^n}^2 L_\Psi^2 \mathbb{E}_{u \sim \gamma} \left[\left\| \Gamma_{\chi, d_\chi}^n(u) - u \right\|_\chi^2 \right]. \end{aligned} \quad (54)$$

An upper bound of the second term in (51) is provided by the following lemma (see a proof in Appendix D):

Lemma 5. Under the conditions of Theorem 1, for any $\delta \in (0, 1)$, we have

$$\begin{aligned} & \mathbb{E}_{\mathcal{S}_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} \langle \Gamma_{\text{NN}} \circ E_\chi^n(u_i), \epsilon_i \rangle \right] \\ & \leq 2\sqrt{2d_y} \sigma \left(\sqrt{\mathbb{E}_{\mathcal{S}_2} \|\Gamma_{\text{NN}} \circ E_\chi^n(u_i) - \Gamma_d^n \circ E_\chi^n(u_i)\|_n^2} + \sqrt{d_y} \delta \right) \sqrt{\frac{\log \mathcal{N}(\delta, \mathcal{F}_{\text{NN}}, \|\cdot\|_\infty) + 2}{n}} + d_y \sigma \delta. \end{aligned} \quad (55)$$

Let \mathcal{F}_{NN} be the network architecture specified in (53). Substituting (54) and (55) into (51), we have

$$T_1 = 2\mathbb{E}_{\mathcal{S}_2} \left[\left\| \Gamma_{\text{NN}} \circ E_\chi^n(u_i) - E_y^n \circ \Psi(u_i) \right\|_n^2 \right]$$

$$\begin{aligned}
&\leq 4d_y \varepsilon_1^2 + 8\sqrt{2d_y} \sigma \left(\sqrt{\mathbb{E}_{\mathcal{S}_2} \left\| \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i) - \Gamma_d^n \circ E_{\mathcal{X}}^n(u_i) \right\|_n^2} + \sqrt{d_y} \delta \right) \sqrt{\frac{\log \mathcal{N}(\delta, \mathcal{F}_{\text{NN}}, \|\cdot\|_\infty) + 2}{n}} \\
&\quad + 4d_y \sigma \delta + 4L_{E_y}^2 L_\Psi^2 \mathbb{E}_{u \sim \gamma} \left[\left\| \Pi_{\mathcal{X}, d_{\mathcal{X}}}^n(u) - u \right\|_{\mathcal{X}}^2 \right]. \tag{56}
\end{aligned}$$

Denote

$$\begin{aligned}
\rho &= \sqrt{\mathbb{E}_{\mathcal{S}_2} \left[\left\| \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i) - E_y^n \circ \Psi(u_i) \right\|_n^2 \right]}, \\
a &= 2d_y \varepsilon_1^2 + 2d_y \sigma \delta + 2L_{E_y}^2 L_\Psi^2 \mathbb{E}_{u \sim \gamma} \left[\left\| \Pi_{\mathcal{X}, d_{\mathcal{X}}}^n(u) - u \right\|_{\mathcal{X}}^2 \right] + 4\sqrt{2} d_y \sigma \delta \sqrt{\frac{\log \mathcal{N}(\delta, \mathcal{F}_{\text{NN}}, \|\cdot\|_\infty) + 2}{n}}, \\
b &= 2\sqrt{2d_y} \sigma \sqrt{\frac{\log \mathcal{N}(\delta, \mathcal{F}_{\text{NN}}, \|\cdot\|_\infty) + 2}{n}}.
\end{aligned}$$

Inequality (56) can be rewritten as

$$\rho^2 \leq a + 2b\rho,$$

from which we deduce that

$$(\rho - b)^2 \leq a + b^2 \Rightarrow \rho^2 \leq 2a + 4b^2.$$

Therefore,

$$\begin{aligned}
T_1 &= 2\rho^2 \leq 8d_y \varepsilon_1^2 + 64d_y \sigma^2 \frac{\log \mathcal{N}(\delta, \mathcal{F}_{\text{NN}}, \|\cdot\|_\infty) + 2}{n} + 16\sqrt{2} d_y \sigma \delta \sqrt{\frac{\log \mathcal{N}(\delta, \mathcal{F}_{\text{NN}}, \|\cdot\|_\infty) + 2}{n}} \\
&\quad + 8d_y \sigma \delta + 8L_{E_y}^2 L_\Psi^2 \mathbb{E}_{u \sim \gamma} \left[\left\| \Pi_{\mathcal{X}, d_{\mathcal{X}}}^n(u) - u \right\|_{\mathcal{X}}^2 \right]. \tag{57}
\end{aligned}$$

Upper bound of T_2 . The term T_2 is the difference between the population risk and the empirical risk of the network estimator Γ_{NN} , while there is a factor 2 ahead of the empirical risk. Utilizing a covering of $\mathcal{F}_{\text{NN}}(d_y, L, p, K, \kappa, M)$ and Bernstein-type inequalities, we establish a fast convergence of T_2 . The upper bound is presented in the following lemma (see a proof in Appendix E).

Lemma 6. Under the conditions of Theorem 1, we have

$$T_2 \leq \frac{35d_y L_{E_y}^2 R_{E_y}^2}{n} \log \mathcal{N} \left(\frac{\delta}{4d_y L_{E_y} R_y}, \mathcal{F}_{\text{NN}}, \|\cdot\|_\infty \right) + 6\delta. \tag{58}$$

Substituting (57) and (58) into (47) gives rise to

$$\begin{aligned}
\text{I} &\leq 2L_{D_y}^2 \mathbb{E}_{\mathcal{S}_1} \mathbb{E}_{u \sim \gamma} \left[\left\| \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u) - E_y^n \circ \Psi(u) \right\|_2^2 \right] \\
&= 2L_{D_y}^2 \mathbb{E}_{\mathcal{S}_1} [T_1] + 2L_{D_y}^2 \mathbb{E}_{\mathcal{S}_1} [T_2] \\
&\leq 16d_y L_{D_y}^2 \varepsilon_1^2 + 128d_y \sigma^2 L_{D_y}^2 \frac{\log \mathcal{N}(\delta, \mathcal{F}_{\text{NN}}, \|\cdot\|_\infty) + 2}{n} + 32\sqrt{2} d_y \sigma L_{D_y}^2 \delta \sqrt{\frac{\log \mathcal{N}(\delta, \mathcal{F}_{\text{NN}}, \|\cdot\|_\infty) + 2}{n}} \\
&\quad + 16d_y \sigma L_{D_y}^2 \delta + 16L_{D_y}^2 L_{E_y}^2 L_\Psi^2 \mathbb{E}_{u \sim \gamma} \left\| \Pi_{\mathcal{X}, d_{\mathcal{X}}}^n(u) - u \right\|_{\mathcal{X}}^2
\end{aligned}$$

$$\begin{aligned}
& + \frac{70d_y L_{D_y}^2 L_{E_y}^2 R_y^2}{n} \log \mathcal{N} \left(\frac{\delta}{4d_y L_{E_y} R_y}, \mathcal{F}_{\text{NN}}, \|\cdot\|_\infty \right) + 12L_{D_y}^2 \delta \\
\leq & 16d_y L_{D_y}^2 \varepsilon_1^2 + \frac{128d_y \sigma^2 L_{D_y}^2 + 70d_y L_{D_y}^2 L_{E_y}^2 R_y^2}{n} \log \mathcal{N} \left(\frac{\delta}{4d_y L_{E_y} R_y}, \mathcal{F}_{\text{NN}}, \|\cdot\|_\infty \right) \\
& + 64d_y \sigma L_{D_y}^2 \delta \sqrt{\frac{\log \mathcal{N}(\delta, \mathcal{F}_{\text{NN}}, \|\cdot\|_\infty)}{n}} + (16d_y \sigma + 12)L_{D_y}^2 \delta + 16L_{D_y}^2 L_{E_y}^2 L_\Psi^2 \mathbb{E}_{u \sim \gamma} \|\Pi_{\mathcal{X}, d_{\mathcal{X}}}^n(u) - u\|_{\mathcal{X}}^2,
\end{aligned} \tag{59}$$

when $\delta < 1$. The covering number of $\mathcal{F}_{\text{NN}}(d_y, L, p, K, \kappa, M)$ can be bounded in terms of its parameters, which is summarized in the following lemma:

Lemma 7 (Lemma 6 of [13]). Let $\mathcal{F}_{\text{NN}}(d_y, L, p, K, \kappa, M)$ be a class of network: $[-B, B]^{d_{\mathcal{X}}} \rightarrow [-M, M]^{d_y}$. For any $\delta > 0$, the δ -covering number of $\mathcal{F}_{\text{NN}}(L, p, K, \kappa, M)$ is bounded by

$$\mathcal{N}(\delta, \mathcal{F}_{\text{NN}}(d_y, L, p, K, \kappa, M), \|\cdot\|_\infty) \leq \left(\frac{2L^2(pB+2)\kappa^L p^{L+1}}{\delta} \right)^{d_y K}. \tag{60}$$

Combining (53) and (60) gives

$$\log \mathcal{N}(\delta, \mathcal{F}_{\text{NN}}(d_y, L, p, K, \kappa, M), \|\cdot\|_\infty) \leq C_7 d_y \left(\varepsilon_1^{-d_{\mathcal{X}}} \log^3 \varepsilon_1^{-1} + \log \delta + \log d_y \right), \tag{61}$$

where C_7 is a constant depending on $d_{\mathcal{X}}, R_{\mathcal{X}}, R_y, L_{E_{\mathcal{X}}}, L_{E_y}, L_{D_{\mathcal{X}}}$ and L_Ψ . Substituting (61) into (59) yields

$$\begin{aligned}
\text{I} \leq & 16d_y L_{D_y}^2 \varepsilon_1^2 + C_7 d_y^2 L_{D_y}^2 \frac{128\sigma^2 + 70L_{E_y}^2 R_y^2}{n} \left(\varepsilon_1^{-d_{\mathcal{X}}} \log^3 \varepsilon_1^{-1} + \log \delta + \log d_y \right) \\
& + 64d_y \sigma L_{D_y}^2 \delta \sqrt{\frac{C_7 d_y \left(\varepsilon_1^{-d_{\mathcal{X}}} \log^3 \varepsilon_1^{-1} + \log \delta + \log d_y \right)}{n}} \\
& + (16d_y \sigma + 12)L_{D_y}^2 \delta + 16L_{D_y}^2 L_{E_y}^2 L_\Psi^2 \mathbb{E}_{u \sim \gamma} \left[\left\| \Pi_{\mathcal{X}, d_{\mathcal{X}}}^n(u) - u \right\|_{\mathcal{X}}^2 \right].
\end{aligned} \tag{62}$$

Setting

$$\varepsilon_1 = d_y^{\frac{1}{2+d_{\mathcal{X}}}} n^{-\frac{1}{2+d_{\mathcal{X}}}}, \delta = n^{-1},$$

we get an upper bound of I

$$\begin{aligned}
\text{I} \leq & C_1 (\sigma^2 + R_y^2) d_y^{\frac{4+d_{\mathcal{X}}}{2+d_{\mathcal{X}}}} n^{-\frac{2}{2+d_{\mathcal{X}}}} \log^3 n + C_2 (\sigma^2 + R_y^2) d_y^2 (\log d_y) n^{-1} \\
& + 16L_{D_y}^2 L_{E_y}^2 L_\Psi^2 \mathbb{E}_{\mathcal{S}_1} \mathbb{E}_{u \sim \gamma} \left[\left\| \Pi_{\mathcal{X}, d_{\mathcal{X}}}^n(u) - u \right\|_2^2 \right]
\end{aligned} \tag{63}$$

for some constants C_1, C_2 depending on $d_{\mathcal{X}}, R_{\mathcal{X}}, R_y, L_{E_{\mathcal{X}}}, L_{E_y}, L_{D_{\mathcal{X}}}, L_\Psi$. The constants C_1, C_2 are the same ones as in Theorem 1. The resulting network architecture $\mathcal{F}(d_y, L, p, K, \kappa, M)$ has

$$\begin{aligned}
L & = O(\log n + \log d_y), \quad p = O\left(d_y^{-\frac{d_{\mathcal{X}}}{2+d_{\mathcal{X}}}} n^{\frac{d_{\mathcal{X}}}{2+d_{\mathcal{X}}}}\right), \quad K = O\left(d_y^{-\frac{d_{\mathcal{X}}}{2+d_{\mathcal{X}}}} n^{\frac{d_{\mathcal{X}}}{2+d_{\mathcal{X}}}} \log n\right), \\
\kappa & = \max\left\{1, \sqrt{d_y L_{E_y} R_y}, \sqrt{d_{\mathcal{X}}} L_{E_{\mathcal{X}}} R_{\mathcal{X}}, \sqrt{d_{\mathcal{X}}} L_{E_{\mathcal{X}}} L_{E_y} L_{D_{\mathcal{X}}} L_\Psi R_{\mathcal{X}}}\right\}, \quad M = \sqrt{d_y L_{E_y} R_y}.
\end{aligned} \tag{64}$$

Combining the bounds of I and II. Putting (62) and (46) together gives rise to

$$\begin{aligned}
& \mathbb{E}_{\mathcal{S}} \mathbb{E}_{u \sim \gamma} \left[\left\| D_{\mathcal{Y}}^n \circ \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u) - \Psi(u) \right\|_{\mathcal{Y}}^2 \right] \\
& \leq \text{I} + \text{II} \\
& \leq C_1(\sigma^2 + R_{\mathcal{Y}}^2) d_{\mathcal{Y}}^{\frac{4+d_{\mathcal{X}}}{2+d_{\mathcal{X}}}} n^{-\frac{2}{2+d_{\mathcal{X}}}} \log^3 n + C_2(\sigma^2 + R_{\mathcal{Y}}^2) d_{\mathcal{Y}}^2 (\log d_{\mathcal{Y}}) n^{-1} \\
& \quad + 16L_{D_{\mathcal{Y}}}^2 L_{E_{\mathcal{Y}}}^2 L_{\Psi}^2 \mathbb{E}_{\mathcal{S}_1} \mathbb{E}_{u \sim \gamma} \left[\left\| \Pi_{\mathcal{X}, d_{\mathcal{X}}}^n(u) - u \right\|_2^2 \right] + 2\mathbb{E}_{\mathcal{S}_1} \mathbb{E}_{v^* \sim \Psi_{\#} \gamma} \left[\left\| \Pi_{\mathcal{Y}, d_{\mathcal{Y}}}^n(v^*) - v^* \right\|_{\mathcal{Y}}^2 \right], \tag{65}
\end{aligned}$$

where C_1, C_2 are constants depending on $d_{\mathcal{X}}, R_{\mathcal{X}}, R_{\mathcal{Y}}, L_{E_{\mathcal{X}}}, L_{E_{\mathcal{Y}}}, L_{D_{\mathcal{X}}}, L_{\Psi}$.

□

7.3 Proof of Theorem 2

The main framework of the proof of Theorem 2 is the same as that of Theorem 1, except special attentions need to be paid on bounding T_1 and T_2 in (48):

- For T_1 , we establish a new result on the approximation error of deep neural networks with architecture $\mathcal{F}_{\text{NN}}(d_{\mathcal{Y}}, L, p, M)$.
- For T_2 , we derive an upper bound using the uniform covering numbers. The motivation to use $\mathcal{F}_{\text{NN}}(d_{\mathcal{Y}}, L, p, M)$ is that it removes parameter upper bound, which is appealing to practical training. However, removing parameter upper bound leads to technical issues in bounding T_2 . We address these issues using the uniform covering numbers thanks to the boundedness of network outputs inspired by [44].

The first part of our proof is the same as that of Theorem 1 up to (48), which is omitted here. In the following, we bound T_1 and T_2 in order.

Upper bound of T_1 . The upper bound of T_1 can be derived similarly as that in Section 7.2, except we make two changes:

- Replace Lemma 4 by the following one

Lemma 8 (Theorem 1.1 of [79]). Let $0 < \alpha \leq 1$ be a real number. There exists a FNN architecture $\mathcal{F}_{\text{NN}}(1, L, p, M)$ with $d_{\mathcal{Y}} = 1$ such that for any integers $\tilde{L}, \tilde{p} > 0$ and $f \in \mathcal{C}^{0, \alpha}([-B, B]^d)$ with $\|f\|_{\mathcal{C}^{0, \alpha}} \leq R$, such an architecture gives rise to an FNN \tilde{f} with

$$\|\tilde{f} - f\|_{\infty} \leq C \tilde{L}^{-\frac{2\alpha}{d}} \tilde{p}^{-\frac{2\alpha}{d}}$$

for some constant C depending on α, d, B, R . This architecture has

$$L = O(\tilde{L} \log \tilde{L}), \quad p = O(\tilde{p} \log \tilde{p}), \quad M = R.$$

The constant hidden in $O(\cdot)$ depends on α, d, B, R .

According to Lemma 8 with $\alpha = 1$, for any $\varepsilon_1 > 0$, there is a network architecture $\mathcal{F}_{\text{NN}}(d_y, L, p, M)$, such that for any Γ_{NN}^n defined in (49), there exists a $\tilde{\Gamma}_d^n \in \mathcal{F}_{\text{NN}}(d_y, L, p, M)$ with

$$\|\tilde{\Gamma}_d^n - \Gamma_d^n\|_\infty \leq \varepsilon_1.$$

Such a network architecture has

$$L = O(\tilde{L} \log \tilde{L}), p = O(\tilde{p} \log \tilde{p}), M = \sqrt{d_y L_{E_y^n} R_y}, \quad (66)$$

where $\tilde{L}, \tilde{p} > 0$ are integers satisfying $\tilde{L}\tilde{p} = \lceil \varepsilon_1^{-d_{\mathcal{X}}/2} \rceil$. The constant hidden in $O(\cdot)$ depends on $d_{\mathcal{X}}, L_{E_y^n}, L_{D_{\mathcal{X}}^n}, L_{\Psi}, B$ and M .

- Replace Lemma 5 by

Lemma 9. Under the conditions of Theorem 2, for any $\delta \in (0, 1)$, we have

$$\begin{aligned} & \mathbb{E}_{\mathcal{S}_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} \langle \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i), \epsilon_i \rangle \right] \\ & \leq 2\sqrt{2d_y\sigma} \left(\sqrt{\mathbb{E}_{\mathcal{S}_2} \|\Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i) - \Gamma_d^n \circ E_{\mathcal{X}}^n(u_i)\|_n^2} + \sqrt{d_y\sigma} \delta \right) \sqrt{\frac{\log \mathcal{N}(\delta, \mathcal{F}_{\text{NN}}, n) + 2}{n}} + d_y\sigma\delta. \end{aligned} \quad (67)$$

Lemma 9 can be proved similarly as Lemma 5. We need to replace the δ -cover $\mathcal{F}^* = \{\Gamma_j^*\}_{j=1}^{\mathcal{N}(\delta, \mathcal{F}_{\text{NN}}, \|\cdot\|_\infty)}$ by a δ -cover of \mathcal{F}_{NN} with respect to \mathcal{S}_2 : $\mathcal{F}^* = \{\Gamma_j^*\}_{j=1}^{\mathcal{N}(\delta, \mathcal{F}_{\text{NN}}, n)}$, where $\mathcal{N}(\delta, \mathcal{F}_{\text{NN}}, n)$ is the uniform covering number. Here the cover \mathcal{F}^* depends on the samples $\{E_{\mathcal{X}}^n(u_i)\}_{i=n+1}^{2n}$. Then there exists $\Gamma^* \in \mathcal{F}^*$ satisfying $\|\Gamma^* \circ E_{\mathcal{X}}^n(u_i) - \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i)\|_\infty \leq \delta$ for any $n+1 \leq i \leq 2n$. The proof is omitted here.

Following the rest of the proof for T_1 in Section 7.2, we can derive that

$$\begin{aligned} T_1 & \leq 8d_y\varepsilon_1^2 + 64d_y\sigma^2 \frac{\log \mathcal{N}(\delta, \mathcal{F}_{\text{NN}}, n) + 2}{n} + 16\sqrt{2}d_y\sigma\delta \sqrt{\frac{\log \mathcal{N}(\delta, \mathcal{F}_{\text{NN}}, n) + 2}{n}} \\ & \quad + 8d_y\sigma\delta + 8L_{E_y^n}^2 L_{\Psi}^2 \mathbb{E}_{u \sim \gamma} \left[\left\| \Pi_{\mathcal{X}, d_{\mathcal{X}}}^n(u) - u \right\|_{\mathcal{X}}^2 \right]. \end{aligned} \quad (68)$$

The network architecture of $\mathcal{F}_{\text{NN}}(d_y, L, p, M)$ is specified in (66).

Upper bound of T_2 . Using the covering number defined in Definition 5, we have the following bound of T_2 .

Lemma 10. Under the conditions of Theorem 2, we have

$$T_2 \leq \frac{35d_y R_y^2}{n} \log \mathcal{N} \left(\frac{\delta}{4d_y L_{E_y^n} R_y}, \mathcal{F}_{\text{NN}}, 2n \right) + 6\delta. \quad (69)$$

Lemma 10 is proved in Appendix F using techniques similar to those in the proof of Lemma 6. Substituting (68) and (69) into (47) gives rise to

$$\begin{aligned}
\text{I} &\leq 2L_{D_y}^2 \mathbb{E}_{\mathcal{S}_1} \mathbb{E}_{u \sim \gamma} \left[\left\| \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u) - E_{\mathcal{Y}}^n \circ \Psi(u) \right\|_2^2 \right] \\
&= 2L_{D_y}^2 \mathbb{E}_{\mathcal{S}_1} [\text{T}_1] + 2L_{D_y}^2 \mathbb{E}_{\mathcal{S}_1} [\text{T}_2] \\
&\leq 16d_y L_{D_y}^2 \varepsilon_1^2 + 128d_y \sigma^2 L_{D_y}^2 \frac{\log \mathcal{N}(\delta, \mathcal{F}_{\text{NN}}, n) + 2}{n} + 32\sqrt{2}d_y \sigma L_{D_y}^2 \delta \sqrt{\frac{\log \mathcal{N}(\delta, \mathcal{F}_{\text{NN}}, n) + 2}{n}} \\
&\quad + 16d_y \sigma L_{D_y}^2 \delta + 16L_{D_y}^2 L_{E_y}^2 L_{\Psi}^2 \mathbb{E}_{u \sim \gamma} \left[\left\| \Pi_{\mathcal{X}, d_{\mathcal{X}}}^n(u) - u \right\|_{\mathcal{X}}^2 \right] \\
&\quad + \frac{70d_y L_{D_y}^2 L_{E_y}^2 R_y^2}{n} \log \mathcal{N} \left(\frac{\delta}{4d_y L_{E_y} R_y}, \mathcal{F}_{\text{NN}}, 2n \right) + 12L_{D_y}^2 \delta \\
&\leq 16d_y L_{D_y}^2 \varepsilon_1^2 + \frac{128d_y \sigma^2 L_{D_y}^2 + 70d_y L_{D_y}^2 L_{E_y}^2 R_y^2}{3n} \log \mathcal{N} \left(\frac{\delta}{4d_y L_{E_y} R_y}, \mathcal{F}_{\text{NN}}, 2n \right) \\
&\quad + 64d_y \sigma L_{D_y}^2 \delta \sqrt{\frac{\log \mathcal{N}(\delta, \mathcal{F}_{\text{NN}}, n)}{n}} + (16d_y \sigma + 12)L_{D_y}^2 \delta + 16L_{D_y}^2 L_{E_y}^2 L_{\Psi}^2 \mathbb{E}_{u \sim \gamma} \left[\left\| \Pi_{\mathcal{X}, d_{\mathcal{X}}}^n(u) - u \right\|_{\mathcal{X}}^2 \right]
\end{aligned} \tag{70}$$

The covering number in (71) can be bounded using the pseudo-dimension of the network class:

Lemma 11 (Theorem 12.2 of [2]). Let F be a class of functions from some domain Ω to $[-M, M]$. Denote the pseudo-dimension of F by $\text{Pdim}(F)$. For any $\delta > 0$, we have

$$\mathcal{N}(\delta, F, m) \leq \left(\frac{2eMm}{\delta \text{Pdim}(F)} \right)^{\text{Pdim}(F)} \tag{72}$$

for $m > \text{Pdim}(F)$.

The next lemma shows that the pseudo-dimension of $\mathcal{F}_{\text{NN}}(1, L, p, M)$ can be bounded using its parameters:

Lemma 12 (Theorem 7 of [4]). For any network architecture \mathcal{F}_{NN} with L layers and U parameters, there exists a universal constant C such that

$$\text{Pdim}(\mathcal{F}_{\text{NN}}) \leq CLU \log(U). \tag{73}$$

Now consider the network architecture $\mathcal{F}_{\text{NN}}(1, L, p, M)$, the number of parameters is bounded by $U = Lp^2$. Combing Lemma 11 and 12, we have

$$\log \mathcal{N} \left(\frac{\delta}{4d_y L_{E_y} R_y}, \mathcal{F}_{\text{NN}}(d_y, L, p, M), 2n \right) \leq C_8 d_y p^2 L^2 \log(p^2 L) (\log M + \log \delta^{-1} + \log n) \tag{74}$$

when $2n > C_9 p^2 L^2 \log(p^2 L)$ for some universal constant C_8, C_9 . Substituting (66) into (74) gives rise to

$$\log \mathcal{N} \left(\frac{\delta}{4d_y L_{E_y} R_y}, \mathcal{F}_{\text{NN}}, 2n \right) \leq C_8 d_y \varepsilon_1^{-d_{\mathcal{X}}} \log^5(\varepsilon_1^{-1}) (\log \delta^{-1} + \log n). \tag{75}$$

Substituting (75) into (71) yields

$$\begin{aligned}
\text{I} &\leq 16d_y L_{D_y^n}^2 \varepsilon_1^2 + C_8 d_y^2 L_{D_y^n}^2 \frac{128\sigma^2 + 70L_{E_y^n}^2 R_y^2}{n} \varepsilon_1^{-d_x} \log^5(\varepsilon_1^{-1}) (\log \delta^{-1} + \log n) \\
&\quad + 64d_y \sigma L_{D_y^n}^2 \delta \sqrt{\frac{C_8 d_y \varepsilon_1^{-d_x} \log^5(\varepsilon_1^{-1}) (\log \delta^{-1} + \log n)}{n}} \\
&\quad + (16d_y \sigma + 12) L_{D_y^n}^2 \delta + 16L_{D_y^n}^2 L_{E_y^n}^2 L_{\Psi}^2 \mathbb{E}_{u \sim \gamma} \left[\left\| \Pi_{\mathcal{X}, d_x}^n(u) - u \right\|_{\mathcal{X}}^2 \right]. \tag{76}
\end{aligned}$$

Setting

$$\varepsilon_1 = d_y^{\frac{1}{2+d_x}} n^{-\frac{1}{2+d_x}}, \delta = n^{-1},$$

we have

$$\text{I} \leq C_4 (\sigma^2 + R_y^2) d_y^{\frac{4+d_x}{2+d_x}} n^{-\frac{2}{2+d_x}} \log^6 n + 16L_{D_y^n}^2 L_{E_y^n}^2 L_{\Psi}^2 \mathbb{E}_{u \sim \gamma} \left[\left\| \Pi_{\mathcal{X}, d_x}^n(u) - u \right\|_{\mathcal{X}}^2 \right], \tag{77}$$

where C_4 is a constant depending on $d_x, R_x, R_y, L_{E_x^n}, L_{E_y^n}, L_{D_x^n}, L_{\Psi}$, the same constant in Theorem 2. The resulting network architecture $\mathcal{F}(L, p, M)$ has

$$L = O(\tilde{L} \log \tilde{L}), p = O(\tilde{p} \log \tilde{p}), M = \sqrt{d_y} L_{E_y^n} R_y, \tag{78}$$

where $\tilde{L} \tilde{p} = d_y^{-\frac{d_x}{4+2d_x}} n^{\frac{d_x}{4+2d_x}}$. Now we check the condition in Lemma 11. Under the choice of L and p above, we have

$$L^2 p^2 \log(p^2 L) = O\left(n^{\frac{2d_x}{4+2d_x}} \log^5 n\right) < 2n$$

when n is large enough. The condition is satisfied.

Combining the bounds of I and II. Putting (77) and (46) together gives rise to

$$\begin{aligned}
&\mathbb{E}_{\mathcal{S}} \mathbb{E}_{u \sim \gamma} \left[\left\| D_y^n \circ \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u) - \Psi(u) \right\|_y^2 \right] \\
&\leq \text{I} + \text{II} \\
&\leq C_4 (\sigma^2 + R_y^2) d_y^{\frac{4+d_x}{2+d_x}} n^{-\frac{2}{2+d_x}} \log^6 n \\
&\quad + 16L_{D_y^n}^2 L_{E_y^n}^2 L_{\Psi}^2 \mathbb{E}_{u \sim \gamma} \left[\left\| \Pi_{\mathcal{X}, d_x}^n(u) - u \right\|_{\mathcal{X}}^2 \right] + 2\mathbb{E}_{\mathcal{S}_1} \mathbb{E}_{v^* \sim \Psi_{\#} \gamma} \left[\left\| \Pi_{y, d_y}^n(v^*) - v^* \right\|_y^2 \right], \tag{79}
\end{aligned}$$

which finishes the proof.

7.4 Proof of Corollary 2

Proof of Corollary 2. We only need to derive upper bounds of

$$\mathbb{E}_{u \sim \gamma} \left[\left\| \Pi_{\mathcal{X}, d_x}(u) - u \right\|_2^2 \right] \text{ and } \mathbb{E}_{v \sim \Psi_{\#} \gamma} \left[\left\| \Pi_{y, d_y}(v) - v \right\|_y^2 \right].$$

Then Corollary 2 is a direct result of Corollary 1. Our proof relies on the following lemma which gives an approximation error of Legendre polynomials for Hölder functions:

Lemma 13 (Theorem 4.5(ii) of [76]). Let $k \geq 0$ be an integer and $\alpha > 0$. For any $f \in \mathcal{C}^{k,\alpha}([0,1]^D)$ with $\|f\|_{\mathcal{C}^{k,\alpha}} < \infty$, there exists $\tilde{f} \in \text{span}(\Phi^{L,r})$ such that

$$\|f - \tilde{f}\|_{\infty} \leq \frac{C}{r^{k+\alpha}}, \quad (80)$$

where C is a constant depending on D and $\|f\|_{\mathcal{C}^{k,\alpha}}$.

We first derive an upper bound of $\mathbb{E}_{u \sim \gamma} \left[\|\Pi_{\mathcal{X},d_{\mathcal{X}}}(u) - u\|_2^2 \right]$. For any $u \in \Omega_{\mathcal{X}}$, according to Lemma 13, there exists $\tilde{u} \in \text{span}(\Phi^{L,r_{\mathcal{X}}})$ such that

$$\|u - \tilde{u}\|_{\infty} \leq C_{10} r_{\mathcal{X}}^{-s},$$

where $s = k + \alpha$, C_{10} is a constant depending on D and $C_{\mathcal{H}_p, \mathcal{X}}$. We deduce that

$$\begin{aligned} \|\Pi_{\mathcal{X},d_{\mathcal{X}}}(u) - u\|_{\mathcal{X}}^2 &= \min_{\tilde{u} \in \text{span}(\Phi^{L,r_{\mathcal{X}}})} \|\tilde{u} - u\|_{\mathcal{X}}^2 \\ &\leq \|\tilde{u} - u\|_{\mathcal{X}}^2 \\ &\leq \int_{[-1,1]^D} |\tilde{u} - u|^2 dx \\ &\leq 2^D C_{10} r_{\mathcal{X}}^{-2s} \\ &= 2^D C_{10} d_{\mathcal{X}}^{-\frac{2s}{D}}, \end{aligned}$$

where in the last equality $d_{\mathcal{X}} = r_{\mathcal{X}}^D$ is used. Therefore

$$\mathbb{E}_{u \sim \gamma} \left[\|\Pi_{\mathcal{X},d_{\mathcal{X}}}(u) - u\|_2^2 \right] \leq C_5 d_{\mathcal{X}}^{-\frac{2s}{D}}.$$

where C_5 is a constant depending on D and $C_{\mathcal{H}_p, \mathcal{X}}$. Similarly, one can show

$$\mathbb{E}_{v \sim \Psi_{\#} \gamma} \left[\|\Pi_{\mathcal{Y},d_{\mathcal{Y}}}(v) - v\|_{\mathcal{Y}}^2 \right] \leq C_6 d_{\mathcal{Y}}^{-\frac{2s}{D}},$$

where C_6 is a constant depending on D and $C_{\mathcal{H}_p, \mathcal{Y}}$. The theorem is proved. \square

7.5 Proof of Corollary 3

Proof of Corollary 3. Our proof relies on the following lemma which gives an approximation error of trigonometric bases for periodic Hölder functions

Lemma 14 (Theorem 4.3(ii) of [76]). Let $k \geq 0$ be an integer and $0 < \alpha \leq 1$. For any $f \in \mathcal{P} \cap \mathcal{C}^{k,\alpha}([0,1]^D)$ with $\|f\|_{\mathcal{C}^{k,\alpha}} < \infty$, there exists $\tilde{f} \in \text{span}(\Phi^{T,r})$ such that

$$\|f - \tilde{f}\|_{\infty} \leq \frac{C}{r^s}, \quad (81)$$

where C is a constant depending on D and $\|f\|_{\mathcal{C}^{k,\alpha}}$.

Corollary 3 can be proved by following the proof of Corollary 3 in which Lemma 13 is replaced by Lemma 14. \square

7.6 Proof of Theorem 3

Proof of Theorem 3. Lemma 2 implies that $E_{\mathcal{X}}^n, D_{\mathcal{X}}^n, E_{\mathcal{Y}}^n, D_{\mathcal{Y}}^n$ are Lipschitz with a Lipschitz constant 1. Therefore Corollary 1 can be applied. We only need to bound $\mathbb{E}_{\mathcal{S}}\mathbb{E}_{u\sim\gamma} \left[\left\| \Pi_{\mathcal{X},d_{\mathcal{X}}}^n(u) - u \right\|_{\mathcal{X}}^2 \right]$ and $\mathbb{E}_{\mathcal{S}}\mathbb{E}_{w\sim\Psi_{\#}\gamma} \left[\left\| \Pi_{\mathcal{Y},d_{\mathcal{Y}}}^n(w) - w \right\|_{\mathcal{Y}}^2 \right]$ in (20). We use the following lemma:

Lemma 15 (Theorem 3.4 of [7]). Let \mathcal{H} be a separable Hilbert space and ρ be a probability measure defined on it. Define the covariance operator $G_{\rho} = \mathbb{E}_{u\sim\rho} u \otimes u$ and its empirical estimation from n samples by $G_{\rho}^n = \frac{1}{n} \sum_{i=1}^n u_i \otimes u_i$ where $\{u_i\}_{i=1}^n$ are i.i.d. samples sampled from ρ . For some integer $d > 0$, let $\Pi_{\mathcal{H},d}$ and $\Pi_{\mathcal{H},d}^n$ be the projectors that project any $u \in \mathcal{H}$ to the space spanned by the eigenfunctions corresponding to the largest d eigenvalues of G_{ρ} and G_{ρ}^n , respectively. We have

$$\mathbb{E}_{\{u_k\}_{k=1}^n \sim \rho} \mathbb{E}_{u\sim\rho} \left[\left\| \Pi_{\mathcal{H},d}^n(u) - u \right\|_{\mathcal{H}}^2 \right] \leq \sqrt{\frac{Cd}{n}} + \mathbb{E}_{u\sim\rho} \left[\left\| \Pi_{\mathcal{H},d}(u) - u \right\|_{\mathcal{H}}^2 \right] \quad (82)$$

with $C = \mathbb{E}_{\{u_i\}_{i=1}^n \sim \rho} \left[\left\| G^n - G \right\|_{\text{HS}}^2 \right]$, where $\|\cdot\|_{\text{HS}}$ is the Hilbert-Schmidt norm.

We first bound $\mathbb{E}_{\mathcal{S}}\mathbb{E}_{u\sim\gamma} \left[\left\| \Pi_{\mathcal{X},d_{\mathcal{X}}}^n(u) - u \right\|_{\mathcal{X}}^2 \right]$. For any $u \sim \gamma$, we have $\|u\|_{\mathcal{X}} \leq R_{\mathcal{X}}$. Therefore

$$\mathbb{E}_{u\sim\gamma} \left[\left\| G_{\mathcal{X}}^n - G_{\mathcal{X}} \right\|_{\text{HS}}^2 \right] \leq 4\mathbb{E}_{u\sim\gamma} \left[\|u\|_{\mathcal{X}}^4 \right] \leq 4R_{\mathcal{X}}^4$$

and Lemma 15 gives

$$\mathbb{E}_{\mathcal{S}}\mathbb{E}_{u\sim\gamma} \left[\left\| \Pi_{\mathcal{X},d_{\mathcal{X}}}^n(u) - u \right\|_{\mathcal{X}}^2 \right] \leq \sqrt{\frac{4R_{\mathcal{X}}^4 d_{\mathcal{X}}}{n}} + \mathbb{E}_{u\sim\gamma} \left[\left\| \Pi_{\mathcal{X},d_{\mathcal{X}}}(u) - u \right\|_{\mathcal{X}}^2 \right]. \quad (83)$$

An upper bound of $\mathbb{E}_{\mathcal{S}}\mathbb{E}_{w\sim\Psi_{\#}\gamma} \left[\left\| \Pi_{\mathcal{Y},d_{\mathcal{Y}}}^n(w) - w \right\|_{\mathcal{Y}}^2 \right]$ is given by the following lemma (see a proof in Appendix G):

Lemma 16. Under the conditions of Theorem 3, we have

$$\begin{aligned} \mathbb{E}_{\mathcal{S}}\mathbb{E}_{w\sim\Psi_{\#}\gamma} \left[\left\| \Pi_{\mathcal{Y},d_{\mathcal{Y}}}^n(w) - w \right\|_{\mathcal{Y}}^2 \right] &\leq 4\sqrt{\frac{(R_{\mathcal{Y}} + \tilde{\sigma})^4 d_{\mathcal{Y}}}{n}} + 8 \left(\frac{\tilde{\sigma}}{\lambda_{d_{\mathcal{Y}}} - \lambda_{d_{\mathcal{Y}+1}}} \right)^2 \tilde{\sigma}^2 (R_{\mathcal{Y}} + \tilde{\sigma})^2 \\ &\quad + 10\tilde{\sigma}^2 + 8\mathbb{E}_{w\sim\Psi_{\#}\gamma} \left[\left\| \Pi_{\mathcal{Y},d_{\mathcal{Y}}}(w) - w \right\|_{\mathcal{Y}}^2 \right]. \end{aligned} \quad (84)$$

□

7.7 Proof of Corollary 4

Proof of Corollary 4. We only need to show that the eigenspace spanned by the first $d_{\mathcal{Y}}$ principal eigenfunctions of $G_{\Psi_{\#}\gamma}$ is the same as that of G_{ζ} . Then Corollary 4 can be proved by following the proof of Theorem 3 in which the upper bound of $\mathbb{E}_{\mathcal{S}}\mathbb{E}_{w\sim\Psi_{\#}\gamma} \left[\left\| \Pi_{\mathcal{Y},d_{\mathcal{Y}}}^n(w) - w \right\|_{\mathcal{Y}}^2 \right]$ can be derived in the same manner as that of $\mathbb{E}_{\mathcal{S}}\mathbb{E}_{u\sim\gamma} \left[\left\| \Pi_{\mathcal{X},d_{\mathcal{X}}}^n(u) - u \right\|_{\mathcal{X}}^2 \right]$.

Denote the eigenvalues of $G_{\Psi_\# \gamma}$ in non-increasing order by $\{\lambda_{\Psi_\# \gamma, k}\}_{k=1}^\infty$. Denote the eigenspace spanned by the first d_γ principal eigenfunctions of $G_{\Psi_\# \gamma}$ by \mathcal{K} , and its compliment by \mathcal{K}^\top . Similarly, we define \mathcal{K}_ζ and \mathcal{K}_ζ^\top for G_ζ . From our assumption, \mathcal{K} is also the eigenspace spanned by the first d_γ eigenfunctions of G_μ . We denote the eigenvalues of G_μ in non-increasing order by $\{\lambda_{\mu, k}\}_{k=1}^\infty$. We are going to show that $\mathcal{K} = \mathcal{K}_\zeta$. From (110), we have $G_\zeta = G_{\Psi_\# \gamma} + G_\mu$. Note that for any $\phi \in \mathcal{K}$ and $\tilde{\phi} \in \mathcal{K}^\top$ with unit length, we have

$$\begin{aligned} \langle G_\zeta \phi, \phi \rangle_\gamma &= \langle G_{\Psi_\# \gamma} \phi, \phi \rangle_\gamma + \langle G_\mu \phi, \phi \rangle_\gamma \\ &\geq \lambda_{\Psi_\# \gamma, d_\gamma} + \lambda_{\mu, d_\gamma+1} \\ &\geq \lambda_{\Psi_\# \gamma, d_\gamma+1} + \lambda_{\mu, d_\gamma+1} \\ &\geq \langle G_{\Psi_\# \gamma} \tilde{\phi}, \tilde{\phi} \rangle_\gamma + \langle G_\mu \tilde{\phi}, \tilde{\phi} \rangle_\gamma \\ &= \langle G_\zeta \tilde{\phi}, \tilde{\phi} \rangle_\gamma. \end{aligned}$$

Since both \mathcal{K} and \mathcal{K}_ζ have dimension d_γ , we have $\mathcal{K} = \mathcal{K}_\zeta$. The proof is finished. \square

7.8 Proof of Theorem 4

Proof of Theorem 4. Theorem 4 can be proved by following the proof of Theorem 2 with the following changes:

- Replace E_χ^n by E_χ .
- Under Assumption 2 and 8, our target function $E_\gamma \circ \Psi \circ D_\chi$ is a Lipschitz function on \mathcal{M} . We replace Lemma 8 by the following one (see a proof in Appendix H):

Lemma 17. Suppose Assumption 8 holds. Assume for any $\mathbf{a} \in \mathcal{M}$, $\|\mathbf{a}\|_\infty \leq B$ for some $B > 0$. There exists a FNN architecture $\mathcal{F}_{\text{NN}}(1, L, p, M)$ such that for any integers $\tilde{L}, \tilde{p} > 0$ and $f \in C^{0,1}(\mathcal{M})$ with $\|f\|_{C^{0,1}} \leq R$, such an architecture gives rise to a FNN \tilde{f} with

$$\|\tilde{f} - f\|_\infty \leq C \tilde{L}^{-\frac{2}{d_0}} \tilde{p}^{-\frac{2}{d_0}}$$

for some constant C depending on d_0, B, R, τ and the surface area of \mathcal{M} . This architecture has

$$L = O(\tilde{L} \log \tilde{L}), \quad p = O(d_\chi \tilde{p} \log \tilde{p}), \quad M = R. \quad (85)$$

The constant hidden in $O(\cdot)$ depends on d_0, B, R, τ and the surface area of \mathcal{M} . \square

7.9 Proof of Theorem 5

Proof of Theorem 5. Theorem 5 can be proved similarly as Theorem 4 while special attention needs to be paid on bounding $\log \mathcal{N}(\delta, \mathcal{F}_{\text{NN}}(d_\gamma, L, p, M), n)$. Note that the total number of parameters of $\mathcal{F}_{\text{NN}}(d_\gamma, L, p, M)$ is bounded by $U = Lp + d_\chi p$. Combing Lemma 11 and 12, we have

$$\log \mathcal{N}\left(\frac{\delta}{4d_\gamma L E_\gamma^n R_\gamma}, \mathcal{F}_{\text{NN}}(d_\gamma, L, p, M), 2n\right)$$

$$\leq C_{11} d_Y (p^2 L^2 + d_X p L) \log(p^2 L + d_X p L) (\log M + \log \delta^{-1} + \log n), \quad (86)$$

where C_{11} is a universal constant. According to (66), one has $Lp = O(\varepsilon_1^{-d_X/2} \log^2(\varepsilon^{-1}))$. Using this relation and substituting the choice of L, p in (66) to (86) gives rise to

$$\begin{aligned} & \log \mathcal{N} \left(\frac{\delta}{4d_Y L E_Y^n R_Y}, \mathcal{F}_{\text{NN}}(d_Y, L, p_1, p_2, M), 2n \right) \\ & \leq C_{11} d_Y (\varepsilon_1^{-d_X} + d_X \varepsilon_1^{-d_X/2}) \log^5(\varepsilon_1^{-1}) (\log \delta^{-1} + \log n). \end{aligned} \quad (87)$$

The proof can be finished by following the rest of the proof of Theorem 4. □

8 Conclusion

We study the generalization error of a general framework on learning operators between infinite-dimensional spaces by two types of deep neural networks. Our upper bound consists of a network estimation error and a projections error, and holds for general encoders and decoders under mild assumptions. The application of our results on some popular encoders and decoders are discussed, such as those using Legendre polynomials, trigonometric functions, and PCA. We also consider two scenarios where additional low dimensional structures of data can be exploited. The two scenarios are: (1) the input data can be encoded to vectors on a low dimensional manifold; (2) the operator has low complexity. In both scenarios, we show that the generalization error converges at a fast rate depending on the intrinsic dimension. Our results show that deep neural networks are adaptive to low dimensional structures of data in operator estimation. In general, our results provide a theoretical justification on the successes of deep neural networks for learning operators between infinite dimensional spaces.

References

- [1] A. Anandkumar, K. Azizzadenesheli, K. Bhattacharya, N. Kovachki, Z. Li, B. Liu, and A. Stuart. Neural operator: Graph kernel network for partial differential equations. In *ICLR 2020 Workshop on Integration of Deep Neural Models and Differential Equations*, 2020.
- [2] M. Anthony and P. Bartlett. *Neural network learning: theoretical foundations*, 1999.
- [3] A. R. Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information Theory*, 39(3):930–945, May 1993.
- [4] P. L. Bartlett, N. Harvey, C. Liaw, and A. Mehrabian. Nearly-tight vc-dimension and pseudodimension bounds for piecewise linear neural networks. *The Journal of Machine Learning Research*, 20(1):2285–2301, 2019.
- [5] B. Bauer and M. Kohler. On deep learning as a remedy for the curse of dimensionality in nonparametric regression. *The Annals of Statistics*, 47(4):2261 – 2285, 2019.

- [6] J. Berner, P. Grohs, and A. Jentzen. Analysis of the generalization error: Empirical risk minimization over deep artificial neural networks overcomes the curse of dimensionality in the numerical approximation of black-scholes partial differential equations. *CoRR*, abs/1809.03062, 2018.
- [7] K. Bhattacharya, B. Hosseini, N. B. Kovachki, and A. M. Stuart. Model reduction and neural networks for parametric pdes. *arXiv preprint arXiv:2005.03180*, 2020.
- [8] H.-J. Bungartz and M. Griebel. Sparse grids. *Acta numerica*, 13:147–269, 2004.
- [9] S. Cai, Z. Wang, L. Lu, T. A. Zaki, and G. E. Karniadakis. DeepM&Mnet: Inferring the electroconvection multiphysics fields based on operator approximation by neural networks. *Journal of Computational Physics*, 436:110296, 2021.
- [10] Y. Cao and Q. Gu. Generalization bounds of stochastic gradient descent for wide and deep neural networks. *CoRR*, abs/1905.13210, 2019.
- [11] L. Q. Chen and J. Shen. Applications of semi-implicit fourier-spectral method to phase field equations. *Computer Physics Communications*, 108(2-3):147–158, 1998.
- [12] M. Chen, H. Jiang, W. Liao, and T. Zhao. Efficient approximation of deep relu networks for functions on low dimensional manifolds. *Advances in neural information processing systems*, 32:8174–8184, 2019.
- [13] M. Chen, H. Jiang, W. Liao, and T. Zhao. Nonparametric regression on low-dimensional manifolds using deep relu networks. *arXiv preprint arXiv:1908.01842*, 2019.
- [14] M. Chen, H. Liu, W. Liao, and T. Zhao. Doubly robust off-policy learning on low-dimensional manifolds by deep neural networks. *arXiv preprint arXiv:2011.01797*, 2020.
- [15] T. Chen and H. Chen. Universal approximation to nonlinear operators by neural networks with arbitrary activation functions and its application to dynamical systems. *IEEE Transactions on Neural Networks*, 6(4):911–917, 1995.
- [16] A. Chkifa, A. Cohen, and C. Schwab. Breaking the curse of dimensionality in sparse polynomial approximation of parametric pdes. *Journal de Mathématiques Pures et Appliquées*, 103(2):400–428, 2015.
- [17] A. Cloninger and T. Klock. Relu nets adapt to intrinsic dimensionality beyond the target domain. *arXiv e-prints*, pages arXiv–2008, 2020.
- [18] A. Cohen and R. DeVore. Approximation of high-dimensional parametric pdes. *Acta Numerica*, 24:1–159, 2015.
- [19] J. H. Conway and N. J. A. Sloane. *Sphere packings, lattices and groups*, volume 290. Springer Science & Business Media, 2013.

- [20] G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.
- [21] M. V. de Hoop, N. B. Kovachki, N. H. Nelsen, and A. M. Stuart. Convergence rates for learning linear operators from noisy data. *arXiv preprint arXiv:2108.12515*, 2021.
- [22] M. Deng, S. Li, A. Goy, I. Kang, and G. Barbastathis. Learning to synthesize: robust phase retrieval at low photon counts. *Light: Science & Applications*, 9(1):36, 2020.
- [23] Q. Du, Y. Gu, H. Yang, and C. Zhou. The discovery of dynamics via linear multistep methods and deep learning: Error estimation. *arXiv preprint arXiv:2103.11488*, 2021.
- [24] C. Duan, Y. Jiao, Y. Lai, X. Lu, and Z. Yang. Convergence rate analysis for deep ritz method. *arxiv:2103.13330*, 2021.
- [25] W. E, C. Ma, and L. Wu. A priori estimates of the population risk for two-layer neural networks. *Communications in Mathematical Sciences*, 17(5):1407–1425, 2019.
- [26] W. E, C. Ma, and L. Wu. The barron space and the flow-induced function spaces for neural network models. *Constructive Approximation*, 2021.
- [27] A. Ern and J.-L. Guermond. *Theory and practice of finite elements*, volume 159. Springer, 2004.
- [28] Y. Fan, J. Feliu-Fabà, L. Lin, L. Ying, and L. Zepeda-Núñez. A multiscale neural network based on hierarchical nested bases. *Research in the Mathematical Sciences*, 6(2):21, 2019.
- [29] Y. Fan, C. Orozco Bohorquez, and L. Ying. Bcr-net: A neural network based on the nonstandard wavelet form. *Journal of Computational Physics*, 384:1–15, 2019.
- [30] M. H. Farrell, T. Liang, and S. Misra. Deep neural networks for estimation and inference. *Econometrica*, 89(1):181–213, 2021.
- [31] H. Federer. Curvature measures. *Transactions of the American Mathematical Society*, 93(3):418–491, 1959.
- [32] I. Giulini. Robust pca and pairs of projections in a hilbert space. *Electronic Journal of Statistics*, 11(2):3903–3926, 2017.
- [33] I. Goodfellow, Y. Bengio, and A. Courville. *Deep learning*. MIT press, 2016.
- [34] A. Graves, A.-r. Mohamed, and G. Hinton. Speech recognition with deep recurrent neural networks. In *2013 IEEE international conference on acoustics, speech and signal processing*, pages 6645–6649. IEEE, 2013.
- [35] Y. Gu, J. Harlim, S. Liang, and H. Yang. Stationary density estimation of itô diffusions using deep learning. *arxiv:2109.03992*, 2021.
- [36] L. Györfi, M. Kohler, A. Krzyżak, and H. Walk. *A distribution-free theory of nonparametric regression*, volume 1. Springer, 2002.

- [37] B. Haasdonk. Reduced basis methods for parametrized pdes—a tutorial introduction for stationary and instationary problems. *Model reduction and approximation: theory and algorithms*, 15:65, 2017.
- [38] M. Hamers and M. Kohler. Nonasymptotic bounds on the l2 error of neural network regression estimates. *Annals of the Institute of Statistical Mathematics*, 58(1):131–151, 2006.
- [39] G. Hinton, L. Deng, D. Yu, G. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, and B. Kingsbury. Deep neural networks for acoustic modeling in speech recognition. *IEEE Signal processing magazine*, 29, 2012.
- [40] K. Hornik. Approximation capabilities of multilayer feedforward networks. *Neural networks*, 4(2):251–257, 1991.
- [41] H. Hotelling. Analysis of a complex of statistical variables into principal components. *Journal of educational psychology*, 24(6):417, 1933.
- [42] H. Hotelling. Relations between two sets of variates. In *Breakthroughs in statistics*, pages 162–190. Springer, 1992.
- [43] A. Jacot, F. Gabriel, and C. Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *CoRR*, abs/1806.07572, 2018.
- [44] Y. Jiao, G. Shen, Y. Lin, and J. Huang. Deep nonparametric regression on approximately low-dimensional manifolds. *arXiv: Statistics Theory*, 2021.
- [45] Y. Khoo, J. Lu, and L. Ying. Solving parametric pde problems with artificial neural networks. *European Journal of Applied Mathematics*, 32(3):421–435, 2021.
- [46] Y. Khoo and L. Ying. Switchnet: A neural network model for forward and inverse scattering problems. *SIAM Journal on Scientific Computing*, 41(5):A3182–A3201, 2019.
- [47] M. Kohler and A. Krzyżak. Adaptive regression estimation with multilayer feedforward neural networks. *Nonparametric Statistics*, 17(8):891–913, 2005.
- [48] M. Kohler, A. Krzyżak, and S. Langer. Estimation of a function of low local dimensionality by deep neural networks. *arxiv:1908.11140*, 2020.
- [49] N. Kovachki, S. Lanthaler, and S. Mishra. On universal approximation and error bounds for fourier neural operators. *Journal of Machine Learning Research*, 22(290):1–76, 2021.
- [50] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [51] S. Lanthaler, S. Mishra, and G. E. Karniadakis. Error estimates for deep neural networks: A deep learning framework in infinite dimensions. *arXiv preprint arXiv:2102.09618*, 2021.

- [52] J. M. Lee. *Riemannian manifolds: an introduction to curvature*, volume 176. Springer Science & Business Media, 2006.
- [53] D. Li, Z. Qiao, and T. Tang. Characterizing the stabilization size for semi-implicit fourier-spectral method to phase field equations. *SIAM Journal on Numerical Analysis*, 54(3):1653–1681, 2016.
- [54] Z. Li, N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, and A. Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020.
- [55] C. Lin, Z. Li, L. Lu, S. Cai, M. Maxey, and G. E. Karniadakis. Operator learning for predicting multiscale bubble growth dynamics. *The Journal of Chemical Physics*, 154(10):104118, 2021.
- [56] H. Liu, M. Chen, T. Zhao, and W. Liao. Besov function approximation and binary classification on low-dimensional manifolds using convolutional residual networks. In *International Conference on Machine Learning*, 2021.
- [57] J. Lu and Y. Lu. A priori generalization error analysis of two-layer neural networks for solving high dimensional schrödinger eigenvalue problems. *arxiv:2105.01228*, 2021.
- [58] J. Lu, Y. Lu, and M. Wang. A priori generalization analysis of the deep ritz method for solving high dimensional elliptic equations. *arxiv:2101.01708*, 2021.
- [59] J. Lu, Z. Shen, H. Yang, and S. Zhang. Deep network approximation for smooth functions. *SIAM Journal on Mathematical Analysis*, to appear.
- [60] L. Lu, P. Jin, G. Pang, Z. Zhang, and G. E. Karniadakis. Learning nonlinear operators via deeponet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3):218–229, 2021.
- [61] T. Luo and H. Yang. Two-layer neural networks for partial differential equations: Optimization and generalization theory. *ArXiv*, abs/2006.15733, 2020.
- [62] R. Miotto, F. Wang, S. Wang, X. Jiang, and J. T. Dudley. Deep learning for healthcare: review, opportunities and challenges. *Briefings in bioinformatics*, 19(6):1236–1246, 2017.
- [63] S. Mishra and R. Molinaro. Estimates on the generalization error of physics informed neural networks (pinns) for approximating pdes. *arxiv:2006.16144*, 2020.
- [64] R. Nakada and M. Imaizumi. Adaptive approximation and generalization of deep neural network with intrinsic dimensionality. *J. Mach. Learn. Res.*, 21:174–1, 2020.
- [65] R. Nakada and M. Imaizumi. Adaptive approximation and generalization of deep neural network with intrinsic dimensionality. *Journal of Machine Learning Research*, 21(174):1–38, 2020.

- [66] N. H. Nelsen and A. M. Stuart. The random feature model for input-output maps between banach spaces. *arXiv preprint arXiv:2005.10224*, 2020.
- [67] P. Niyogi, S. Smale, and S. Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Discrete & Computational Geometry*, 39(1-3):419–441, 2008.
- [68] S. A. Orszag. Accurate solution of the orr–sommerfeld stability equation. *Journal of Fluid Mechanics*, 50(4):689–703, 1971.
- [69] K. Pearson. Liii. on lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin philosophical magazine and journal of science*, 2(11):559–572, 1901.
- [70] B. Peherstorfer and K. Willcox. Data-driven operator inference for nonintrusive projection-based model reduction. *Computer Methods in Applied Mechanics and Engineering*, 306:196–215, 2016.
- [71] C. Qiao, D. Li, Y. Guo, C. Liu, T. Jiang, Q. Dai, and D. Li. Evaluation and development of deep neural networks for image super-resolution in optical microscopy. *Nature Methods*, 18(2):194–202, 2021.
- [72] Z. Qin, Q. Zeng, Y. Zong, and F. Xu. Image inpainting based on deep learning: A review. *Displays*, 69:102028, 2021.
- [73] G. Rozza. Fundamentals of reduced basis method for problems governed by parametrized pdes and applications. In *Separated representations and PGD-based model reduction*, pages 153–227. Springer, 2014.
- [74] J. Schmidt-Hieber. Deep relu network approximation of functions on a manifold. *arXiv preprint arXiv:1908.00695*, 2019.
- [75] J. Schmidt-Hieber. Nonparametric regression using deep neural networks with relu activation function. *The Annals of Statistics*, 48(4):1875–1897, 2020.
- [76] M. H. Schultz. L^2 -multivariate approximation theory. *SIAM Journal on Numerical Analysis*, 6(2):161–183, 1969.
- [77] U. Shaham, A. Cloninger, and R. R. Coifman. Provable approximation properties for deep neural networks. *Applied and Computational Harmonic Analysis*, 44(3):537–557, 2018.
- [78] J. Shen, T. Tang, and L.-L. Wang. *Spectral methods: algorithms, analysis and applications*, volume 41. Springer Science & Business Media, 2011.
- [79] Z. Shen, H. Yang, and S. Zhang. Deep network approximation characterized by number of neurons. *Communications in Computational Physics*, 28(5):1768–1811, 2020.
- [80] Z. Shen, H. Yang, and S. Zhang. Deep network approximation: Achieving arbitrary accuracy with fixed number of neurons. *arxiv:2107.02397*, 2021.

- [81] Z. Shen, H. Yang, and S. Zhang. Deep network with approximation error being reciprocal of width to power of square root of depth. *Neural Computation*, 33(4):1005–1036, 03 2021.
- [82] Z. Shen, H. Yang, and S. Zhang. Neural network approximation: Three hidden layers are enough. *Neural Networks*, 141:160–173, 2021.
- [83] Z. Shen, H. Yang, and S. Zhang. Optimal approximation rate of ReLU networks in terms of width and depth. *Journal de Mathématiques Pures et Appliquées*, to appear.
- [84] Y. Shin, J. Darbon, and G. E. Karniadakis. On the convergence of physics informed neural networks for linear second-order elliptic and parabolic type pdes. *arxiv:2004.01806*, 2020.
- [85] J. W. Siegel and J. Xu. Sharp bounds on the approximation rates, metric entropy, and n -widths of shallow neural networks. *arxiv:2101.12365*, 2021.
- [86] C. J. Stone. Optimal Global Rates of Convergence for Nonparametric Regression. *The Annals of Statistics*, 10(4):1040 – 1053, 1982.
- [87] T. Suzuki. Adaptivity of deep relu network for learning in besov and mixed smooth besov spaces: optimal rate and curse of dimensionality. *arXiv preprint arXiv:1810.08033*, 2018.
- [88] G. Szeg. *Orthogonal polynomials*, volume 23. American Mathematical Soc., 1939.
- [89] C. Tian, L. Fei, W. Zheng, Y. Xu, W. Zuo, and C.-W. Lin. Deep learning on image denoising: An overview. *Neural Networks*, 131:251–275, Nov 2020.
- [90] L. W. Tu. *An introduction to manifolds*. Springer., 2011.
- [91] A. W. Van Der Vaart, A. W. van der Vaart, A. van der Vaart, and J. Wellner. *Weak convergence and empirical processes: with applications to statistics*. Springer Science & Business Media, 1996.
- [92] Z. Wei and X. Chen. Physics-inspired convolutional neural network for solving full-wave inverse scattering problems. *IEEE Transactions on Antennas and Propagation*, 67(9):6138–6148, 2019.
- [93] D. Yarotsky. Error bounds for approximations with deep relu networks. *Neural Networks*, 94:103–114, 2017.
- [94] D. Yarotsky. Optimal approximation of continuous functions by very deep ReLU networks. In S. Bubeck, V. Perchet, and P. Rigollet, editors, *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 639–649. PMLR, 06–09 Jul 2018.
- [95] D. Yarotsky. Elementary superexpressive activations. *arXiv e-prints*, 2021.

- [96] D. Yarotsky and A. Zhevnerchuk. The phase diagram of approximation rates for deep neural networks. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 13005–13015. Curran Associates, Inc., 2020.
- [97] Y. Zhu and N. Zabarar. Bayesian deep convolutional encoder–decoder networks for surrogate modeling and uncertainty quantification. *Journal of Computational Physics*, 366:415–447, 2018.

Appendix

A The derivation for the error bound in Corollary 2 when $d_{\mathcal{X}} = d_{\mathcal{Y}} = \log^{\frac{1}{2}} n$

From Corollary 2, we need to balance the two terms $d_{\mathcal{Y}}^{\frac{4+d_{\mathcal{X}}}{2+d_{\mathcal{X}}}} n^{-\frac{2}{2+d_{\mathcal{X}}}} \log^6 n$ and $d_{\mathcal{X}}^{-\frac{2s}{D}}$. By setting $d_{\mathcal{X}} = d_{\mathcal{Y}} = \log^{\frac{1}{2}} n$, the first term decays faster than the second term as n increases. We want to find a lower bound of n , denoted by n_0 , so that when $n > n_0$, the error is dominated by the second term. Note that n_0 should satisfy

$$d_{\mathcal{Y}}^{\frac{4+d_{\mathcal{X}}}{2+d_{\mathcal{X}}}} n^{-\frac{2}{2+d_{\mathcal{X}}}} \log^6 n \leq d_{\mathcal{X}}^{-\frac{2s}{D}}. \quad (88)$$

Since

$$d_{\mathcal{Y}}^{\frac{4+d_{\mathcal{X}}}{2+d_{\mathcal{X}}}} n^{-\frac{2}{2+d_{\mathcal{X}}}} \log^6 n \leq d_{\mathcal{Y}}^2 n^{-\frac{2}{2+d_{\mathcal{X}}}} \log^6 n \leq d_{\mathcal{X}}^{-\frac{2s}{D}},$$

in the following, we consider solving

$$d_{\mathcal{Y}}^2 n^{-\frac{2}{2+d_{\mathcal{X}}}} \log^6 n \leq d_{\mathcal{X}}^{-\frac{2s}{D}}.$$

Substituting the expression of $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$, we deduce

$$n^{-\frac{2}{2+\log^{1/2} n}} \log^7 n \leq \log^{-\frac{s}{D}} n \Rightarrow -\frac{2}{2+\log^{1/2} n} \log n + 7 \log \log n \leq -\frac{s}{D} \log \log n.$$

Denote $a = \log n$. We have

$$\frac{2}{2+a^{1/2}} a \geq \left(7 + \frac{s}{D}\right) \log a. \quad (89)$$

A sufficient condition of (89) is

$$\frac{2}{a^{1/2}} a \geq \left(7 + \frac{s}{D}\right) \log a \Rightarrow a \geq \frac{1}{4} \left(7 + \frac{s}{D}\right)^2 \log^2 a. \quad (90)$$

Note that $\log a < a^{1/3}$ for $a > 100$. Therefore, it is sufficient to solve

$$a \geq \frac{1}{4} \left(7 + \frac{s}{D}\right)^2 a^{\frac{2}{3}} \Rightarrow a \geq \left(\frac{7}{2} + \frac{s}{2D}\right)^6.$$

Substituting a by $\log n$, one has

$$n \geq \exp\left(\max\left\{100, \left(\frac{7}{2} + \frac{s}{2D}\right)^6\right\}\right).$$

B Proof of Lemma 1

Proof of Lemma 1. We first prove (17):

$$\begin{aligned} \|E_{\mathcal{H},d}(u) - E_{\mathcal{H},d}(\tilde{u})\|_2^2 &= \left\| [\langle u - \tilde{u}, \phi_1 \rangle_{\mathcal{H}}, \dots, \langle u - \tilde{u}, \phi_d \rangle_{\mathcal{H}}]^\top \right\|_2^2 \\ &= \sum_{k=1}^d |\langle u - \tilde{u}, \phi_k \rangle_{\mathcal{H}}|^2 \\ &\leq \sum_{k=1}^{\infty} |\langle u - \tilde{u}, \phi_k \rangle_{\mathcal{H}}|^2 \\ &= \|u - \tilde{u}\|_{\mathcal{H}}^2. \end{aligned}$$

For (18), we have

$$\|D_{\mathcal{H},d}(\mathbf{a}) - D_{\mathcal{H},d}(\tilde{\mathbf{a}})\|_{\mathcal{H}}^2 = \left\| \sum_{k=1}^d (a_k - \tilde{a}_k) \phi_k \right\|_{\mathcal{H}}^2 = \|\mathbf{a} - \tilde{\mathbf{a}}\|_2^2,$$

since $\{\phi_k\}_{k=1}^d$ is an orthonormal set. □

C Proof of Lemma 3

Proof of Lemma 3. Let $\mathbf{a}, \tilde{\mathbf{a}} \in \mathbb{R}^{d_\chi}$. We have

$$\begin{aligned} \|\Gamma_d^n(\mathbf{a}) - \Gamma_d^n(\tilde{\mathbf{a}})\|_2 &= \|E_y^n \circ \Psi \circ D_\chi^n(\mathbf{a}) - E_y^n \circ \Psi \circ D_\chi^n(\tilde{\mathbf{a}})\|_2 \\ &\leq L_{E_y^n} \|\Psi \circ D_\chi^n(\mathbf{a}) - \Psi \circ D_\chi^n(\tilde{\mathbf{a}})\|_2 \\ &\leq L_{E_y^n} L_\Psi \|D_\chi^n(\mathbf{a}) - D_\chi^n(\tilde{\mathbf{a}})\|_y \\ &\leq L_{E_y^n} L_{D_\chi^n} L_\Psi \|\mathbf{a} - \tilde{\mathbf{a}}\|_2. \end{aligned} \tag{91}$$

□

D Proof of Lemma 5

Proof of Lemma 5. We prove Lemma 5 using the covering number of \mathcal{F}_{NN} . Let $\mathcal{F}^* = \{\Gamma_j^*\}_{j=1}^{\mathcal{N}(\delta, \mathcal{F}_{\text{NN}}, \|\cdot\|_\infty)}$ be a δ -cover of \mathcal{F}_{NN} , where $\mathcal{N}(\delta, \mathcal{F}_{\text{NN}}, \|\cdot\|_\infty)$ is the covering number. Then there exists $\Gamma^* \in \mathcal{F}^*$ satisfying $\|\Gamma^* - \Gamma_{\text{NN}}\|_\infty \leq \delta$, where Γ_{NN} is our estimator in (6). Denote $\|\Gamma \circ E_\chi^n\|_n^2 = \frac{1}{n} \sum_{i=n+1}^{2n} \|\Gamma \circ E_\chi^n(u_i)\|_2^2$.

We have

$$\begin{aligned}
& \mathbb{E}_{\mathcal{S}_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} \left\langle \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i), \epsilon_i \right\rangle \right] \\
&= \mathbb{E}_{\mathcal{S}_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} \left\langle \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i) - \Gamma^* \circ E_{\mathcal{X}}^n(u_i) + \Gamma^* \circ E_{\mathcal{X}}^n(u_i) - \Gamma_d^n \circ E_{\mathcal{X}}^n(u_i), \epsilon_i \right\rangle \right] \\
&\leq \mathbb{E}_{\mathcal{S}_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} \left\langle \Gamma^* \circ E_{\mathcal{X}}^n(u_i) - \Gamma_d^n \circ E_{\mathcal{X}}^n(u_i), \epsilon_i \right\rangle \right] + \mathbb{E}_{\mathcal{S}_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} \left\| \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i) - \Gamma^* \circ E_{\mathcal{X}}^n(u_i) \right\|_2 \|\epsilon_i\|_2 \right] \\
&\leq \mathbb{E}_{\mathcal{S}_2} \left[\frac{\left\| \Gamma^* \circ E_{\mathcal{X}}^n - \Gamma_d^n \circ E_{\mathcal{X}}^n \right\|_n \sum_{i=n+1}^{2n} \left\langle \Gamma^* \circ E_{\mathcal{X}}^n(u_i) - \Gamma_d^n \circ E_{\mathcal{X}}^n(u_i), \epsilon_i \right\rangle}{\sqrt{n}} \right] + d_y \sigma \delta \\
&\leq \sqrt{2} \mathbb{E}_{\mathcal{S}_2} \left[\frac{\left\| \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i) - \Gamma_d^n \circ E_{\mathcal{X}}^n(u_i) \right\|_n + \sqrt{d_y} \delta}{\sqrt{n}} \left| \frac{\sum_{i=n+1}^{2n} \left\langle \Gamma^* \circ E_{\mathcal{X}}^n(u_i) - \Gamma_d^n \circ E_{\mathcal{X}}^n(u_i), \epsilon_i \right\rangle}{\sqrt{n} \left\| \Gamma^* \circ E_{\mathcal{X}}^n - \Gamma_d^n \circ E_{\mathcal{X}}^n \right\|_n} \right| \right] + d_y \sigma \delta,
\end{aligned} \tag{92}$$

where the first inequality follows from Cauchy–Schwarz inequality, the third inequality holds since

$$\begin{aligned}
& \left\| \Gamma^* \circ E_{\mathcal{X}}^n - \Gamma_d^n \circ E_{\mathcal{X}}^n \right\|_n \\
&= \sqrt{\frac{1}{n} \sum_{i=n+1}^{2n} \left\| \Gamma^* \circ E_{\mathcal{X}}^n(u_i) - \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i) + \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i) - \Gamma_d^n \circ E_{\mathcal{X}}^n(u_i) \right\|_2^2} \\
&\leq \sqrt{\frac{2}{n} \sum_{i=n+1}^{2n} \left\| \Gamma^* \circ E_{\mathcal{X}}^n(u_i) - \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i) \right\|_2^2 + \left\| \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i) - \Gamma_d^n \circ E_{\mathcal{X}}^n(u_i) \right\|_2^2} \\
&\leq \sqrt{\frac{2}{n} \sum_{i=n+1}^{2n} d_y \delta^2 + \left\| \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i) - \Gamma_d^n \circ E_{\mathcal{X}}^n(u_i) \right\|_2^2} \\
&\leq \sqrt{2} \left\| \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i) - \Gamma_d^n \circ E_{\mathcal{X}}^n(u_i) \right\|_n + \sqrt{2d_y} \delta.
\end{aligned} \tag{93}$$

Denote $z_j = \frac{\sum_{i=n+1}^{2n} \left\langle \Gamma_j^* \circ E_{\mathcal{X}}^n(u_i) - \Gamma_d^n \circ E_{\mathcal{X}}^n(u_i), \epsilon_i \right\rangle}{\sqrt{n} \left\| \Gamma_j^* \circ E_{\mathcal{X}}^n - \Gamma_d^n \circ E_{\mathcal{X}}^n \right\|_n}$. The expectation term in (92) can be bounded as

$$\begin{aligned}
& \mathbb{E}_{\mathcal{S}_2} \left[\frac{\left\| \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i) - \Gamma_d^n \circ E_{\mathcal{X}}^n(u_i) \right\|_n + \sqrt{d_y} \delta}{\sqrt{n}} \left| \frac{\sum_{i=n+1}^{2n} \left\langle \Gamma^* \circ E_{\mathcal{X}}^n(u_i) - \Gamma_d^n \circ E_{\mathcal{X}}^n(u_i), \epsilon_i \right\rangle}{\sqrt{n} \left\| \Gamma^* \circ E_{\mathcal{X}}^n - \Gamma_d^n \circ E_{\mathcal{X}}^n \right\|_n} \right| \right] \\
&\leq \mathbb{E}_{\mathcal{S}_2} \left[\frac{\left\| \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i) - \Gamma_d^n \circ E_{\mathcal{X}}^n(u_i) \right\|_n + \sqrt{d_y} \delta}{\sqrt{n}} \max_j |z_j| \right] \\
&= \mathbb{E}_{\mathcal{S}_2} \left[\frac{\left\| \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i) - \Gamma_d^n \circ E_{\mathcal{X}}^n(u_i) \right\|_n}{\sqrt{n}} \max_j |z_j| + \frac{\sqrt{d_y} \delta}{\sqrt{n}} \max_j |z_j| \right] \\
&\leq \mathbb{E}_{\mathcal{S}_2} \left[\sqrt{\frac{1}{n} \left\| \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i) - \Gamma_d^n \circ E_{\mathcal{X}}^n(u_i) \right\|_n^2} \sqrt{\max_j |z_j|^2} + \frac{\sqrt{d_y} \delta}{\sqrt{n}} \sqrt{\max_j |z_j|^2} \right]
\end{aligned}$$

$$\begin{aligned}
&\leq \sqrt{\frac{1}{n} \mathbb{E}_{\mathcal{S}_2} \left[\left\| \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i) - \Gamma_d^n \circ E_{\mathcal{X}}^n(u_i) \right\|_n^2 \right]} \sqrt{\mathbb{E}_{\mathcal{S}_2} \left[\max_j |z_j|^2 \right]} + \frac{\sqrt{d_y} \delta}{\sqrt{n}} \sqrt{\mathbb{E}_{\mathcal{S}_2} \left[\max_j |z_j|^2 \right]} \\
&= \left(\sqrt{\frac{1}{n} \mathbb{E}_{\mathcal{S}_2} \left[\left\| \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u_i) - \Gamma_d^n \circ E_{\mathcal{X}}^n(u_i) \right\|_n^2 \right]} + \frac{\sqrt{d_y} \delta}{\sqrt{n}} \right) \sqrt{\mathbb{E}_{\mathcal{S}_2} \left[\max_j |z_j|^2 \right]}. \tag{94}
\end{aligned}$$

where the second inequality comes from Cauchy–Schwarz inequality, the third inequality comes from Jensen’s inequality.

Since $\epsilon_i \in [-\sigma, \sigma]^{d_y}$, each component of ϵ_i is a sub-Gaussian variable with parameter σ . Therefore for given u_{n+1}, \dots, u_{2n} , each z_j is a sub-gaussian variable with parameter $\sqrt{d_y} \sigma$. The last term is the maximum of a collection of squared sub-Gaussian variables and is bounded as

$$\begin{aligned}
\mathbb{E}_{\mathcal{S}_2} \left[\max_j |z_j|^2 | u_{n+1}, \dots, u_{2n} \right] &= \frac{1}{t} \log \exp \left(t \mathbb{E}_{\mathcal{S}_2} \left[\max_j |z_j|^2 | u_{n+1}, \dots, u_{2n} \right] \right) \\
&\leq \frac{1}{t} \log \mathbb{E}_{\mathcal{S}_2} \left[\exp \left(t \max_j |z_j|^2 | u_{n+1}, \dots, u_{2n} \right) \right] \\
&\leq \frac{1}{t} \log \mathbb{E}_{\mathcal{S}_2} \left[\sum_j \exp \left(t |z_j|^2 | u_{n+1}, \dots, u_{2n} \right) \right] \\
&\leq \frac{1}{t} \log \mathcal{N}(\delta, \mathcal{F}_{\text{NN}}, \|\cdot\|_\infty) + \frac{1}{t} \log \mathbb{E}_{\mathcal{S}_2} \left[\exp \left(t |z_1|^2 | u_{n+1}, \dots, u_{2n} \right) \right]. \tag{95}
\end{aligned}$$

Since z_1 is sub-Gaussian with parameter σ^2 , we have

$$\begin{aligned}
\mathbb{E}_{\mathcal{S}_2} \left[\exp \left(t |z_1|^2 | u_{n+1}, \dots, u_{2n} \right) \right] &= 1 + \sum_{k=1}^{\infty} \frac{t^k \mathbb{E}_{\mathcal{S}_2} \left[z_1^{2k} | u_{n+1}, \dots, u_{2n} \right]}{k!} \\
&= 1 + \sum_{k=1}^{\infty} \frac{t^k}{k!} \int_0^{\infty} \mathbb{P} \left(|z_1| \geq \tau^{\frac{1}{2k}} | u_{n+1}, \dots, u_{2n} \right) d\tau \\
&\leq 1 + 2 \sum_{k=1}^{\infty} \frac{t^k}{k!} \int_0^{\infty} \exp \left(-\frac{\tau^{1/k}}{2d_y \sigma^2} \right) d\tau \\
&= 1 + \sum_{k=1}^{\infty} \frac{2k(2td_y \sigma^2)^k}{k!} \Gamma_G(k) \\
&= 1 + 2 \sum_{k=1}^{\infty} (2td_y \sigma^2)^k, \tag{96}
\end{aligned}$$

where Γ_G represents the Gamma function. Setting $t = (4d_y \sigma^2)^{-1}$ gives rise to

$$\begin{aligned}
\mathbb{E}_{\mathcal{S}_2} \left[\max_j |z_j|^2 | u_{n+1}, \dots, u_{2n} \right] &\leq 4d_y \sigma^2 \log \mathcal{N}(\delta, \mathcal{F}_{\text{NN}}, \|\cdot\|_\infty) + 4d_y \sigma^2 \log 3 \\
&\leq 4d_y \sigma^2 \log \mathcal{N}(\delta, \mathcal{F}_{\text{NN}}, \|\cdot\|_\infty) + 6d_y \sigma^2. \tag{97}
\end{aligned}$$

Combining (97), (94), (92) finishes the proof. □

E Proof of Lemma 6

Proof of Lemma 6. Our proof follows the proof of [13, Lemma 4.2]. Denote $g(u) = \|\Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u) - E_{\mathcal{Y}}^n \circ \Psi(u)\|_2^2$. We have $\|g\|_{\infty} \leq 4d_{\mathcal{Y}}L_{E_{\mathcal{Y}}^n}^2R_{\mathcal{Y}}^2$. Then

$$\begin{aligned} T_2 &= \mathbb{E}_{\mathcal{S}_2} \left[\mathbb{E}_{u \sim \gamma} [g(u)|\mathcal{S}_1] - \frac{2}{n} \sum_{i=n+1}^{2n} g(u_i) \right] \\ &= 2\mathbb{E}_{\mathcal{S}_2} \left[\frac{1}{2} \mathbb{E}_{u \sim \gamma} [g(u)|\mathcal{S}_1] - \frac{1}{n} \sum_{i=n+1}^{2n} g(u_i) \right] \\ &= 2\mathbb{E}_{\mathcal{S}_2} \left[\mathbb{E}_{u \sim \gamma} [g(u)|\mathcal{S}_1] - \frac{1}{n} \sum_{i=n+1}^{2n} g(u_i) - \frac{1}{2} \mathbb{E}_{u \sim \gamma} [g(u)|\mathcal{S}_1] \right]. \end{aligned} \quad (98)$$

A lower bound of $\frac{1}{2} \mathbb{E}_{u \sim \gamma} [g(u)|\mathcal{S}_1]$ can be derived as

$$\mathbb{E}_{u \sim \gamma} [g(u)|\mathcal{S}_1] = \mathbb{E}_{u \sim \gamma} \left[\frac{4d_{\mathcal{Y}}L_{E_{\mathcal{Y}}^n}^2R_{\mathcal{Y}}^2}{4d_{\mathcal{Y}}L_{E_{\mathcal{Y}}^n}^2R_{\mathcal{Y}}^2} g(u)|\mathcal{S}_1 \right] \geq \frac{1}{4d_{\mathcal{Y}}L_{E_{\mathcal{Y}}^n}^2R_{\mathcal{Y}}^2} \mathbb{E}_{u \sim \gamma} [g^2(u)|\mathcal{S}_1] \quad (99)$$

Substituting (99) into (98) gives

$$T_2 \leq 2\mathbb{E}_{\mathcal{S}_2} \left[\mathbb{E}_{u \sim \gamma} [g(u)|\mathcal{S}_1] - \frac{1}{n} \sum_{i=n+1}^{2n} g(u_i) - \frac{1}{8d_{\mathcal{Y}}L_{E_{\mathcal{Y}}^n}^2R_{\mathcal{Y}}^2} \mathbb{E}_{u \sim \gamma} [g^2(u)|\mathcal{S}_1] \right]. \quad (100)$$

Define the set

$$\mathcal{R} = \left\{ g(u) = \|\Gamma \circ E_{\mathcal{X}}^n(u) - E_{\mathcal{Y}}^n \circ \Psi(u)\|_2^2 : \Gamma \in \mathcal{F}_{\text{NN}} \right\}. \quad (101)$$

Denote $\mathcal{S}'_2 = \{u'_i\}_{i=n+1}^{2n}$ as an independent copy of \mathcal{S}_2 . We rewrite T_2 as

$$\begin{aligned} T_2 &\leq 2\mathbb{E}_{\mathcal{S}_2} \left[\sup_{g \in \mathcal{R}} \left(\mathbb{E}_{\mathcal{S}'_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} g(u'_i) \right] \right) - \frac{1}{n} \sum_{i=n+1}^{2n} g(u_i) - \frac{1}{8d_{\mathcal{Y}}L_{E_{\mathcal{Y}}^n}^2R_{\mathcal{Y}}^2} \left(\mathbb{E}_{\mathcal{S}'_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} g^2(u'_i) \right] \right) \right] \\ &\leq 2\mathbb{E}_{\mathcal{S}_2} \left[\sup_{g \in \mathcal{R}} \left(\mathbb{E}_{\mathcal{S}'_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} (g(u'_i) - g(u_i)) \right] \right) - \frac{1}{16d_{\mathcal{Y}}L_{E_{\mathcal{Y}}^n}^2R_{\mathcal{Y}}^2} \mathbb{E}_{\mathcal{S}_2, \mathcal{S}'_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} (g^2(u_i) + g^2(u'_i)) \right] \right] \\ &\leq 2\mathbb{E}_{\mathcal{S}_2, \mathcal{S}'_2} \left[\sup_{g \in \mathcal{R}} \left(\frac{1}{n} \sum_{i=n+1}^{2n} \left((g(u_i) - g(u'_i)) - \frac{1}{16d_{\mathcal{Y}}L_{E_{\mathcal{Y}}^n}^2R_{\mathcal{Y}}^2} [g^2(u_i) + g^2(u'_i)] \right) \right) \right]. \end{aligned} \quad (102)$$

Let $\mathcal{R}^* = \{g_i^*\}_{i=1}^{\mathcal{N}(\delta, \mathcal{R}, \|\cdot\|_{\infty})}$ be a δ -cover of \mathcal{R} . Then for any $g \in \mathcal{R}$, there exists $g^* \in \mathcal{R}^*$ such that $\|g - g^*\|_{\infty} \leq \delta$.

We next bound (102) using g^* 's. For the first term in (102), we have

$$\begin{aligned} g(u_i) - g(u'_i) &= g(u_i) - g^*(u_i) + g^*(u_i) - g^*(u'_i) + g^*(u'_i) - g(u'_i) \\ &= (g(u_i) - g^*(u_i)) + (g^*(u_i) - g^*(u'_i)) + (g^*(u'_i) - g(u'_i)) \end{aligned}$$

$$\leq (g^*(u_i) - g^*(u'_i)) + 2\delta. \quad (103)$$

We lower bound $g^2(u_i) + g^2(u'_i)$ as

$$\begin{aligned} g^2(u_i) + g^2(u'_i) &= (g^2(u_i) - (g^*)^2(u_i)) + ((g^*)^2(u_i) + (g^*)^2(u'_i)) - ((g^*)^2(u'_i) - g^2(u'_i)) \\ &\geq (g^*)^2(u_i) + (g^*)^2(u'_i) - |g(u_i) - g^*(u_i)| |g(u_i) + g^*(u_i)| - |g^*(u'_i) - g(u'_i)| |g^*(u'_i) + g(u'_i)| \\ &\geq (g^*)^2(u_i) + (g^*)^2(u'_i) - 16d_y L_{E_y}^2 R_y^2 \delta. \end{aligned} \quad (104)$$

Substituting (103) and (104) into (102) gives rise to

$$\begin{aligned} T_2 &\leq 2\mathbb{E}_{\mathcal{S}_2, \mathcal{S}'_2} \left[\sup_{g^* \in \mathcal{R}^*} \left(\frac{1}{n} \sum_{i=n+1}^{2n} \left((g^*(u_i) - g^*(u'_i)) - \frac{1}{16d_y L_{E_y}^2 R_y^2} [(g^*)^2(u_i) + (g^*)^2(u'_i)] \right) \right) \right] + 6\delta \\ &= 2\mathbb{E}_{\mathcal{S}_2, \mathcal{S}'_2} \left[\max_j \left(\frac{1}{n} \sum_{i=n+1}^{2n} \left((g_j^*(u_i) - g_j^*(u'_i)) - \frac{1}{16d_y L_{E_y}^2 R_y^2} [(g_j^*)^2(u_i) + (g_j^*)^2(u'_i)] \right) \right) \right] + 6\delta. \end{aligned} \quad (105)$$

Denote $h_j = (u_i, u'_i, \xi_i) = (g_j^*(u_i) - g_j^*(u'_i))$. We have

$$\begin{aligned} \mathbb{E}_{\mathcal{S}_2, \mathcal{S}'_2} [h_j(u_i, u'_i)] &= 0, \\ \text{Var}[h_j(u_i, u'_i)] &= \mathbb{E}[h_j^2(u_i, u'_i)] \\ &= \mathbb{E}_{\mathcal{S}_2, \mathcal{S}'_2} [(g_j^*(u_i) - g_j^*(u'_i))^2] \\ &\leq 2\mathbb{E}_{\mathcal{S}_2, \mathcal{S}'_2} [(g_j^*)^2(u_i) + (g_j^*)^2(u'_i)]. \end{aligned}$$

Thus T_2 can be bounded as

$$\begin{aligned} T_2 &\leq \tilde{T}_2 + 6\delta \\ \text{with } \tilde{T}_2 &= 2\mathbb{E}_{\mathcal{S}_2, \mathcal{S}'_2} \left[\max_j \left(\frac{1}{n} \sum_{i=n+1}^{2n} \left(h_j(u_i, u'_i) - \frac{1}{32d_y L_{E_y}^2 R_y^2} \text{Var}[h_j(u_i, u'_i)] \right) \right) \right]. \end{aligned}$$

Note that $\|h_j\|_\infty \leq 4d_y L_{E_y}^2 R_y^2$. We next derive the moment generating function of h_j . For any $0 < t < \frac{3}{4d_y L_{E_y}^2 R_y^2}$, we have

$$\begin{aligned} \mathbb{E}_{\mathcal{S}_2, \mathcal{S}'_2} [\exp(th_j(u_i, u'_i))] &= \mathbb{E}_{\mathcal{S}_2, \mathcal{S}'_2} \left[1 + th_j(u_i, u'_i) + \sum_{k=2}^{\infty} \frac{t^k h_j^k(u_i, u'_i)}{k!} \right] \\ &\leq \mathbb{E}_{\mathcal{S}_2, \mathcal{S}'_2} \left[1 + th_j(u_i, u'_i) + \sum_{k=2}^{\infty} \frac{(4d_y L_{E_y}^2 R_y^2)^{k-2} t^k h_j^2(u_i, u'_i)}{2 \times 3^{k-2}} \right] \\ &= \mathbb{E}_{\mathcal{S}_2, \mathcal{S}'_2} \left[1 + th_j(u_i, u'_i) + \frac{t^2 h_j^2(u_i, u'_i)}{2} \sum_{k=2}^{\infty} \frac{(4d_y L_{E_y}^2 R_y^2)^{k-2} t^{k-2}}{3^{k-2}} \right] \\ &= \mathbb{E}_{\mathcal{S}_2, \mathcal{S}'_2} \left[1 + th_j(u_i, u'_i) + \frac{t^2 h_j^2(u_i, u'_i)}{2} \frac{1}{1 - 4d_y L_{E_y}^2 R_y^2 t/3} \right] \end{aligned}$$

$$\begin{aligned}
&= 1 + t^2 \text{Var}[h_j(u_i, u'_i)] \frac{1}{2 - 8d_y L_{E_y^n}^2 R_y^2 t/3} \\
&\leq \exp\left(\text{Var}[h_j(u_i, u'_i)] \frac{3t^2}{6 - 8d_y L_{E_y^n}^2 R_y^2 t}\right),
\end{aligned} \tag{106}$$

where the last inequality comes from $1 + x \leq \exp(x)$ for $x \geq 0$.

Then for $0 < t/n < \frac{3}{4d_y L_{E_y^n}^2 R_y^2}$, we have

$$\begin{aligned}
&\exp\left(\frac{t\tilde{\Gamma}_2}{2}\right) \\
&= \exp\left(t \mathbb{E}_{S_2, S'_2} \left[\max_j \left(\frac{1}{n} \sum_{i=n+1}^{2n} h_j(u_i, u'_i) - \frac{1}{32d_y L_{E_y^n}^2 R_y^2} \frac{1}{n} \sum_{i=n+1}^{2n} \text{Var}[h_j(u_i, u'_i)] \right) \right]\right) \\
&\leq \mathbb{E}_{S_2, S'_2} \left[\exp\left(t \max_j \left(\frac{1}{n} \sum_{i=n+1}^{2n} h_j(u_i, u'_i) - \frac{1}{32d_y L_{E_y^n}^2 R_y^2} \frac{1}{n} \sum_{i=n+1}^{2n} \text{Var}[h_j(u_i, u'_i)] \right) \right)\right] \\
&\leq \mathbb{E}_{S_2, S'_2} \left[\sum_j \exp\left(\frac{t}{n} \sum_{i=n+1}^{2n} h_j(u_i, u'_i) - \frac{t}{32d_y L_{E_y^n}^2 R_y^2} \frac{1}{n} \sum_{i=n+1}^{2n} \text{Var}[h_j(u_i, u'_i)]\right) \right] \\
&\leq \left[\sum_j \exp\left(\sum_{i=n+1}^{2n} \text{Var}[h_j(u_i, u'_i)] \frac{3t^2/n^2}{6 - 8d_y L_{E_y^n}^2 R_y^2 t/n} - \frac{1}{32d_y L_{E_y^n}^2 R_y^2} \frac{t}{n} \text{Var}[h_j(u_i, u'_i)]\right) \right] \\
&= \left[\sum_j \exp\left(\sum_{i=n+1}^{2n} \frac{t}{n} \text{Var}[h_j(u_i, u'_i)] \left(\frac{3t/n}{6 - 8d_y L_{E_y^n}^2 R_y^2 t/n} - \frac{1}{32d_y L_{E_y^n}^2 R_y^2}\right)\right) \right],
\end{aligned} \tag{107}$$

where the first inequality follows from Jensen's inequality and the third inequality uses (106).

Setting

$$\frac{3t/n}{6 - 8d_y L_{E_y^n}^2 R_y^2 t/n} - \frac{1}{32d_y L_{E_y^n}^2 R_y^2} = 0$$

gives $t = \frac{3n}{52d_y L_{E_y^n}^2 R_y^2} < \frac{3n}{4d_y L_{E_y^n}^2 R_y^2}$. Substituting our choice of t into (107) gives

$$\frac{t\tilde{\Gamma}_2}{2} \leq \log \sum_j \exp(0).$$

Therefore

$$\tilde{\Gamma}_2 \leq \frac{2}{t} \log \mathcal{N}(\delta, \mathcal{R}, \|\cdot\|_\infty) = \frac{104d_y L_{E_y^n}^2 R_y^2}{3n} \log \mathcal{N}(\delta, \mathcal{R}, \|\cdot\|_\infty)$$

and

$$\Gamma_2 \leq \frac{104d_y L_{E_y^n}^2 R_y^2}{3n} \log \mathcal{N}(\delta, \mathcal{R}, \|\cdot\|_\infty) + 6\delta \leq \frac{35d_y L_{E_y^n}^2 R_y^2}{n} \log \mathcal{N}(\delta, \mathcal{R}, \|\cdot\|_\infty) + 6\delta.$$

We next derive a relation between the covering number of \mathcal{F}_{NN} and \mathcal{R} . For any $g, \tilde{g} \in \mathcal{R}$, we have

$$g(u) = \|\Gamma \circ E_{\mathcal{X}}^n(u) - E_y^n \circ \Psi(u)\|_2^2, \quad \tilde{g}(u) = \|\tilde{\Gamma} \circ E_{\mathcal{X}}^n(u) - E_y^n \circ \Psi(u)\|_2^2$$

for some $\Gamma, \tilde{\Gamma} \in \mathcal{F}_{\text{NN}}$. We have

$$\begin{aligned}
\|g - \tilde{g}\|_\infty &= \sup_u \left| \left\| \Gamma \circ E_{\mathcal{X}}^n(u) - E_{\mathcal{Y}}^n \circ \Psi(u) \right\|_2^2 - \left\| \tilde{\Gamma} \circ E_{\mathcal{X}}^n(u) - E_{\mathcal{Y}}^n \circ \Psi(u) \right\|_2^2 \right| \\
&= \sup_u \left| \left\langle \Gamma \circ E_{\mathcal{X}}^n(u) - \tilde{\Gamma} \circ E_{\mathcal{X}}^n(u), \Gamma \circ E_{\mathcal{X}}^n(u) + \tilde{\Gamma} \circ E_{\mathcal{X}}^n(u) - 2E_{\mathcal{Y}}^n \circ \Psi(u) \right\rangle \right| \\
&\leq \sup_u \left\| \Gamma \circ E_{\mathcal{X}}^n(u) - \tilde{\Gamma} \circ E_{\mathcal{X}}^n(u) \right\|_2 \left\| \Gamma \circ E_{\mathcal{X}}^n(u) + \tilde{\Gamma} \circ E_{\mathcal{X}}^n(u) - 2E_{\mathcal{Y}}^n \circ \Psi(u) \right\|_2 \\
&\leq 4d_{\mathcal{Y}} L_{E_{\mathcal{Y}}^n} R_{\mathcal{Y}} \left\| \Gamma - \tilde{\Gamma} \right\|_\infty.
\end{aligned}$$

As a result, we have

$$\mathcal{N}(\delta, \mathcal{R}, \|\cdot\|_\infty) \leq \mathcal{N}\left(\frac{\delta}{4d_{\mathcal{Y}} L_{E_{\mathcal{Y}}^n} R_{\mathcal{Y}}}, \mathcal{F}_{\text{NN}}, \|\cdot\|_\infty\right).$$

and Lemma 6 is proved. \square

F Proof of Lemma 10

Lemma 10 can be proved similarly to Lemma 6. Denote $g(u) = \left\| \Gamma_{\text{NN}} \circ E_{\mathcal{X}}^n(u) - E_{\mathcal{Y}}^n \circ \Psi(u) \right\|_2^2$ and let $\mathcal{S}'_2 = \{u'_i\}_{i=n+1}^{2n}$ be an independent copy of \mathcal{S}_2 . Following the proof of Lemma 6 up to (102) and replacing $\mathbb{E}_{u \sim \gamma}[g(u)|\mathcal{S}_1]$ by $\mathbb{E}_{\mathcal{S}'_2} \left[\frac{1}{n} \sum_{i=n+1}^{2n} g(u'_i) \right]$, we can derive

$$T_2 \leq 2\mathbb{E}_{\mathcal{S}_2, \mathcal{S}'_2} \left[\sup_{g \in \mathcal{R}} \left(\frac{1}{n} \sum_{i=n+1}^{2n} (g(u_i) - g(u'_i)) - \frac{1}{16d_{\mathcal{Y}} L_{E_{\mathcal{Y}}^n}^2 R_{\mathcal{Y}}^2} \frac{1}{n} \sum_{i=n+1}^{2n} (g^2(u_i) + g^2(u'_i)) \right) \right], \quad (108)$$

Let $\mathcal{R}^* = \{g_i^*\}_{i=1}^{\mathcal{N}(\delta, \mathcal{R}, 2n)}$ be a δ -cover of \mathcal{R} with respect to the data set $\tilde{\mathcal{S}} = \{u_i\}_{i=1}^n \cup \{u'_i\}_{i=1}^n$. Then for any $g \in \mathcal{R}$, there exists $g^* \in \mathcal{R}^*$ such that $|g(u) - g^*(u)| \leq \delta, \forall u \in \tilde{\mathcal{S}}$.

Lemma 10 can be proved by following the rest proof of Lemma 6.

G Proof of Lemma 16

The proof of Lemma 16 relies on the perturbation theory of operators on separable Hilbert spaces, which is stated in the following lemma:

Lemma 18 (Proposition 2.1 of [32]). Let A, \tilde{A} be two compact self-adjoint nonnegative operators on the separable real Hilbert space \mathcal{H} . Denote the eigenvalues of A and \tilde{A} in non-increasing order by $\{\lambda_1, \lambda_2, \dots\}$ and $\{\tilde{\lambda}_1, \tilde{\lambda}_2, \dots\}$, respectively. For some integer $d > 0$, let $\Pi_{\mathcal{H}, d}$ and $\tilde{\Pi}_{\mathcal{H}, d}$ be the projectors that project any $u \in \mathcal{H}$ to the space spanned by the eigenfunctions corresponding to the largest d eigenvalues of A and \tilde{A} , respectively. We have

$$\left\| \Pi_{\mathcal{H}, d} - \tilde{\Pi}_{\mathcal{H}, d} \right\|_{\text{HS}} \leq \frac{\sqrt{2} \|A - \tilde{A}\|_{\text{HS}}}{\max\{\lambda_d - \lambda_{d+1}, \tilde{\lambda}_d - \tilde{\lambda}_{d+1}\}}. \quad (109)$$

Proof of Lemma 16. Denote $w = \Psi(u)$. Recall that ζ is the probability measure of $v = \Psi(u) + \tilde{\epsilon}$. We have

$$G_\zeta = \mathbb{E}_{\{v_i\}_{i=1}^n \sim \zeta} [G_\zeta^n] = \mathbb{E}_{v \sim \zeta} [v \otimes v]$$

$$\begin{aligned}
&= \mathbb{E}_{w \sim \Psi_{\# \gamma}, \tilde{\epsilon} \sim \mu} [(w + \tilde{\epsilon}) \otimes (w + \tilde{\epsilon})] \\
&= \mathbb{E}_{w \sim \Psi_{\# \gamma}} [w \otimes w] + \mathbb{E}_{\tilde{\epsilon} \sim \mu} [\tilde{\epsilon} \otimes \tilde{\epsilon}] \\
&= G_{\Psi_{\# \gamma}} + G_{\mu},
\end{aligned} \tag{110}$$

where the third equality holds since w and $\tilde{\epsilon}$ are independent and $\mathbb{E}\tilde{\epsilon} = 0$. Recall that $\Pi_{\mathcal{Y}, d_{\mathcal{Y}}}$ (resp. $\Pi_{\mathcal{Y}, d_{\mathcal{Y}}}^n$) projects any $w \in \mathcal{Y}$ to the space spanned by the first $d_{\mathcal{Y}}$ principal eigenfunctions of $G_{\Psi_{\# \gamma}}$ (resp. G_{ζ}^n). We denote by $\tilde{\Pi}_{\mathcal{Y}, d_{\mathcal{Y}}}$ as the projection that projects any $w \in \mathcal{Y}$ to the space spanned by the first $d_{\mathcal{Y}}$ principal eigenfunctions of G_{ζ} . Relation (110) implies that

$$\mathbb{E}_{\{v_i\}_{i=1}^n \sim \zeta} [\Pi_{\mathcal{Y}, d_{\mathcal{Y}}}^n] = \tilde{\Pi}_{\mathcal{Y}, d_{\mathcal{Y}}}.$$

We have

$$\mathbb{E}_{v \sim \zeta} \left[\left\| G_{\zeta}^n - \mathbb{E}_{\{v_i\}_{i=1}^n \sim \zeta} [G_{\zeta}^n] \right\|_{\text{HS}}^2 \right] \leq 4 \mathbb{E}_{v \sim \zeta} [\|v\|_{\mathcal{Y}}^4] \leq 4(R_{\mathcal{Y}} + \tilde{\sigma})^4.$$

We deduce that

$$\begin{aligned}
&\mathbb{E}_{\mathcal{S}} \mathbb{E}_{w \sim \Psi_{\# \gamma}} \left[\left\| \Pi_{\mathcal{Y}, d_{\mathcal{Y}}}^n(w) - w \right\|_{\mathcal{Y}}^2 \right] \\
&= \mathbb{E}_{\mathcal{S}} \mathbb{E}_{\tilde{\epsilon} \sim \mu} \mathbb{E}_{w \sim \Psi_{\# \gamma}} \left[\left\| \Pi_{\mathcal{Y}, d_{\mathcal{Y}}}^n(w + \tilde{\epsilon}) - (w + \tilde{\epsilon}) - \left[\Pi_{\mathcal{Y}, d_{\mathcal{Y}}}^n(\tilde{\epsilon}) - \tilde{\epsilon} \right] \right\|_{\mathcal{Y}}^2 \right] \\
&\leq 2 \mathbb{E}_{\mathcal{S}} \mathbb{E}_{\tilde{\epsilon} \sim \mu} \mathbb{E}_{w \sim \Psi_{\# \gamma}} \left[\left\| \Pi_{\mathcal{Y}, d_{\mathcal{Y}}}^n(w + \tilde{\epsilon}) - (w + \tilde{\epsilon}) \right\|_{\mathcal{Y}}^2 \right] + 2 \mathbb{E}_{\mathcal{S}} \mathbb{E}_{\tilde{\epsilon} \sim \mu} \left[\left\| \Pi_{\mathcal{Y}, d_{\mathcal{Y}}}^n(\tilde{\epsilon}) - \tilde{\epsilon} \right\|_{\mathcal{Y}}^2 \right] \\
&\leq 2 \mathbb{E}_{\mathcal{S}} \mathbb{E}_{v \sim \zeta} \left[\left\| \Pi_{\mathcal{Y}, d_{\mathcal{Y}}}^n(v) - v \right\|_{\mathcal{Y}}^2 \right] + 2 \mathbb{E}_{\tilde{\epsilon} \sim \mu} [\|\tilde{\epsilon}\|_{\mathcal{Y}}^2] \\
&\leq 2 \sqrt{\frac{4(R_{\mathcal{Y}} + \tilde{\sigma})^4 d_{\mathcal{Y}}}{n}} + 2 \mathbb{E}_{v \sim \zeta} \left[\left\| \tilde{\Pi}_{\mathcal{Y}, d_{\mathcal{Y}}}(v) - v \right\|_{\mathcal{Y}}^2 \right] + 2\tilde{\sigma}^2,
\end{aligned} \tag{111}$$

where the last inequality comes from Lemma 15 and $\tilde{\Pi}_{\mathcal{Y}, d_{\mathcal{Y}}} = \mathbb{E}_{v \sim \zeta} [v \otimes v] = \mathbb{E}_{\{v_i\}_{i=1}^n \sim \zeta} [\Pi_{\mathcal{Y}, d_{\mathcal{Y}}}^n]$.

We bound the second term on the right-hand side as

$$\begin{aligned}
&\mathbb{E}_{v \sim \zeta} \left[\left\| \tilde{\Pi}_{\mathcal{Y}, d_{\mathcal{Y}}}(v) - v \right\|_{\mathcal{Y}}^2 \right] \\
&\leq 2 \mathbb{E}_{v \sim \zeta} \left[\left\| \tilde{\Pi}_{\mathcal{Y}, d_{\mathcal{Y}}}(v) - \Pi_{\mathcal{Y}, d_{\mathcal{Y}}}(v) \right\|_{\mathcal{Y}}^2 \right] + 2 \mathbb{E}_{v \sim \zeta} \left[\left\| \Pi_{\mathcal{Y}, d_{\mathcal{Y}}}(v) - v \right\|_{\mathcal{Y}}^2 \right] \\
&\leq 2 \mathbb{E}_{v \sim \zeta} \left[\left\| (\tilde{\Pi}_{\mathcal{Y}, d_{\mathcal{Y}}} - \Pi_{\mathcal{Y}, d_{\mathcal{Y}}})(v) \right\|_{\mathcal{Y}}^2 \right] + 2 \mathbb{E}_{\tilde{\epsilon} \sim \mu} \mathbb{E}_{w \sim \Psi_{\# \gamma}} \left[\left\| \Pi_{\mathcal{Y}, d_{\mathcal{Y}}}(w + \tilde{\epsilon}) - (w + \tilde{\epsilon}) \right\|_{\mathcal{Y}}^2 \right] \\
&\leq 2 \mathbb{E}_{v \sim \zeta} \left[\left\| \tilde{\Pi}_{\mathcal{Y}, d_{\mathcal{Y}}} - \Pi_{\mathcal{Y}, d_{\mathcal{Y}}} \right\|_{\text{op}}^2 \|v\|_{\mathcal{Y}}^2 \right] + 4 \mathbb{E}_{w \sim \Psi_{\# \gamma}} \left[\left\| \Pi_{\mathcal{Y}, d_{\mathcal{Y}}}(w) - w \right\|_{\mathcal{Y}}^2 \right] + 4 \mathbb{E}_{\tilde{\epsilon} \sim \mu} \left[\left\| \Pi_{\mathcal{Y}, d_{\mathcal{Y}}}(\tilde{\epsilon}) - \tilde{\epsilon} \right\|_{\mathcal{Y}}^2 \right] \\
&\leq 2 \mathbb{E}_{v \sim \zeta} \left[\left\| \tilde{\Pi}_{\mathcal{Y}, d_{\mathcal{Y}}} - \Pi_{\mathcal{Y}, d_{\mathcal{Y}}} \right\|_{\text{HS}}^2 \|v\|_{\mathcal{Y}}^2 \right] + 4 \mathbb{E}_{w \sim \Psi_{\# \gamma}} \left[\left\| \Pi_{\mathcal{Y}, d_{\mathcal{Y}}}(w) - w \right\|_{\mathcal{Y}}^2 \right] + 4 \mathbb{E}_{\tilde{\epsilon} \sim \mu} \left[\left\| \Pi_{\mathcal{Y}, d_{\mathcal{Y}}}(\tilde{\epsilon}) - \tilde{\epsilon} \right\|_{\mathcal{Y}}^2 \right] \\
&\leq 2 \left(\frac{\sqrt{2} \|G_{\mu}\|_{\text{HS}}}{\lambda_{d_{\mathcal{Y}}} - \lambda_{d_{\mathcal{Y}+1}}} \right)^2 (R_{\mathcal{Y}} + \tilde{\sigma})^2 + 4 \mathbb{E}_{w \sim \Psi_{\# \gamma}} \left[\left\| \Pi_{\mathcal{Y}, d_{\mathcal{Y}}}(w) - w \right\|_{\mathcal{Y}}^2 \right] + 4\tilde{\sigma}^2 \\
&\leq 4 \left(\frac{\tilde{\sigma}}{\lambda_{d_{\mathcal{Y}}} - \lambda_{d_{\mathcal{Y}+1}}} \right)^2 \tilde{\sigma}^2 (R_{\mathcal{Y}} + \tilde{\sigma})^2 + 4 \mathbb{E}_{w \sim \Psi_{\# \gamma}} \left[\left\| \Pi_{\mathcal{Y}, d_{\mathcal{Y}}}(w) - w \right\|_{\mathcal{Y}}^2 \right] + 4\tilde{\sigma}^2,
\end{aligned} \tag{112}$$

where the fourth inequality follows from Lemma 18.

Substituting (112) into (111) gives rise to (84). □

H Proof of Lemma 17

Proof of Lemma 17. Our proof relies on concepts related to functions on manifolds, such as charts, atlas, the partition of unity, and functions on manifolds. We refer the readers to [90, 52, 13, 56] for details. Following [15, Proof of Theorem 1], we first construct an atlas of \mathcal{M} in which all projections projects any point on \mathcal{M} to a tangent space of \mathcal{M} . These projections are linear functions that can be realized by a subnetwork. Then the function f is decomposed using a partition of unity subordinates to the atlas we constructed. For each chart (U, ϕ) , we use a subnetwork to approximate an indicator function that determines whether the input $\mathbf{x} \in \mathcal{M}$ belongs to U . Another subnetwork is used to approximate $f(\mathbf{x}) \circ \phi^{-1}$ on its tangent space. Finally, we multiply both subnetworks together and sum over all charts. The multiplication is approximated by another subnetwork. We prove Lemma 17 in four steps.

Step 1. In the first step, we show that there exists an atlas of \mathcal{M} , denoted by $\{U_k, \phi_k\}_{k=1}^{C_{\mathcal{M}}}$, such that ϕ_k 's are linear projections. Denote $B_r(\mathbf{c})$ as the Euclidean ball in \mathbb{R}^{d_x} centered at \mathbf{c} with radius r . For any given $r > 0$, since \mathcal{M} is compact, there exists a set of points $\{\mathbf{c}_k\}_{k=1}^{C_{\mathcal{M}}}$ such that $\mathcal{M} \subseteq \cup_k B_r(\mathbf{c}_k)$. For each $B_r(\mathbf{c}_k)$, denote $U_i = \mathcal{M} \cap B_r(\mathbf{c}_k)$. By setting $r < \tau/2$, we have that U_i is diffeomorphic to a ball in \mathbb{R}^{d_0} [67]. The minimal number of balls is upper bounded by

$$C_{\mathcal{M}} \leq \lceil \text{Area}(\mathcal{M}) T_d / r^d \rceil,$$

where $\text{Area}(\mathcal{M})$ is the area of \mathcal{M} and T_d is the thickness of U_k 's (see Chapter 2 of [19]).

We next define ϕ_k 's. For each \mathbf{c}_k , let $\{\mathbf{v}_j^k\}_{j=1}^{d_0}$ be an orthonormal basis of the tangent space of \mathcal{M} at \mathbf{c}_k . Define the matrix $V_k = [\mathbf{v}_1^k, \dots, \mathbf{v}_d^k]$. We set

$$\phi_k(\mathbf{x}) = V_k^\top (\mathbf{x} - \mathbf{c}_k).$$

Note that ϕ_k is a linear function which can be realized by a single layer. Then $\{(U_k, \phi_k)\}_{k=1}^{C_{\mathcal{M}}}$ form an atlas of \mathcal{M} .

Step 2. In the second step, we design a subnetwork that determines the chart that the input \mathbf{x} belongs to. To determine whether $\mathbf{x} \in U_k$, it is equivalent to check whether the squared distance between \mathbf{x} and \mathbf{c}_k is less than r^2 . It can be done by $\mathbb{1}_{[0, r^2]} \circ d_k^2(\mathbf{x})$ where $\mathbb{1}_{[0, r^2]}(a)$ is an indicator function that outputs 1 if $a \in [0, r^2]$, and outputs 0 otherwise. Here $d_k^2(\mathbf{x})$ is the squared distance function defined as

$$d_k^2(\mathbf{x}) = \|\mathbf{x} - \mathbf{c}_k\|_2^2 = \sum_{j=1}^{d_x} (x_j - c_{k,j})^2,$$

where the notations $\mathbf{x} = [x_1, \dots, x_{d_x}]^\top$ and $\mathbf{c}_k = [c_{k,1}, \dots, c_{k,d_x}]$ are used.

We next approximate both functions by neural networks. To approximate d_k^2 , the key issue is to approximate the square function by neural networks, for which we use the following lemma:

Lemma 19 (Lemma 4.2 of [59]). For any $B > 0$ and integers $L, p > 0$, there exists a network $\tilde{\chi}$ in $\mathcal{F}_{\text{NN}}(1, L, 9p + 1, B^2)$ with $d_y = 1$ such that for any $x, y \in [-B, B]$, we have

$$|\tilde{\chi}(x, y) - xy| \leq 24B^2 p^{-L}. \quad (113)$$

According to Lemma 19, we approximate $d_k^2(\mathbf{x})$ by

$$\tilde{d}_k^2(\mathbf{x}) = \sum_{j=1}^{d_\chi} \tilde{\chi}(x_j - c_{k,j}, x_j - c_{k,j}),$$

where $\tilde{\chi} \in \mathcal{F}_{\text{NN}}(1, 4sL_1, 9p_1 + 1, B^2)$. The approximation error is $\|\tilde{d}_k - d_k\|_\infty \leq 24d_\chi B^2 p_1^{-4sL_1}$.

For $\mathbb{1}_{[0, r]^2}$, we use the following function to approximate it

$$\tilde{\mathbb{I}}_\Delta(a) = \begin{cases} 1 & a \leq r^2 - \Delta + 24d_\chi B^2 p_1^{-4sL_1}, \\ -\frac{1}{\Delta - 48d_\chi B^2 p_1^{-4sL_1}} a + \frac{r^2 - 24d_\chi B^2 p_1^{-4sL_1}}{\Delta - 48d_\chi B^2 p_1^{-4sL_1}} & a \in [r^2 - \Delta + 24d_\chi B^2 p_1^{-4sL_1}, r^2 - 24d_\chi B^2 p_1^{-4sL_1}], \\ 0 & a \geq r^2 - 24d_\chi B^2 p_1^{-4sL_1}, \end{cases}$$

where $\Delta \geq 24d_\chi B^2 p_1^{-4sL_1}$ will be chosen later. We approximate $\tilde{\mathbb{I}}_\Delta \circ d_k^2(\mathbf{x})$ by $\tilde{\mathbb{I}}_\Delta \circ \tilde{d}_k^2(\mathbf{x})$ in which the parameter Δ is the 'width' of the error region: when $\mathbf{x} \notin U_k$, we have $d_k^2(\mathbf{x}) \geq r^2$ and $\tilde{\mathbb{I}}_\Delta \circ \tilde{d}_k^2(\mathbf{x}) = 0$; when $\mathbf{x} \in U_k$ and $d_k^2(\mathbf{x}) \leq r^2 - \Delta$, we have $\tilde{\mathbb{I}}_\Delta \circ \tilde{d}_k^2(\mathbf{x}) = 1$.

We then realize $\tilde{\mathbb{I}}_\Delta(a)$ by a subnetwork. Denoting $m_0 = \frac{1}{\Delta - 48d_\chi B^2 p_1^{-4sL_1}}$, $m_1 = r^2 - \Delta + 24d_\chi B^2 p_1^{-4sL_1}$, $m_2 = r^2 - 24d_\chi B^2 p_1^{-4sL_1}$, we rewrite $\tilde{\mathbb{I}}_\Delta(a)$ as

$$\tilde{\mathbb{I}}_\Delta(a) = -m_0(\min\{\max\{a, m_1\}, m_2\}) + m_2 m_0.$$

The function above can be realized by a network with one hidden layer:

$$\tilde{\mathbb{I}}_\Delta(a) = -m_0(m_2 - \text{ReLU}[m_2 - (\text{ReLU}(a - m_1) + m_1)]) + m_2 m_0.$$

Step 3. In this step, we decompose f using a partition of unity of \mathcal{M} and approximate each component by a subnetwork. Let $\{h_k\}_{k=1}^{C_{\mathcal{M}}}$ be a partition of unity of \mathcal{M} such that h_k is supported on U_k . We decompose f as

$$f = \sum_{k=1}^{C_{\mathcal{M}}} h_k f.$$

Note that for each k , $h_k f$ is a function defined on \mathcal{M} supported on U_k , and $(h_k f) \circ \phi_k^{-1}$ is a function defined in \mathbb{R}^{d_0} and supported on $[-2B, 2B]^{d_0}$. The following lemma shows that $h_k f$ is in the same space as f :

Lemma 20. Suppose Assumption 8 holds. Let $\{U_k, \phi_k\}_{k=1}^{C_{\mathcal{M}}}$ be defined in **Step 1**. For each k , we have $h_k f \in \mathcal{C}^{0,1}(\mathcal{M})$ and $\|h_k f\|_{\mathcal{C}^{0,1}(\mathcal{M})}$ is bounded by a constant depending on d_0, h_k, f and ϕ_k .

Lemma 20 can be proved by following the proof of [13, Lemma 2]. The proof is omitted here. According to Lemma 20 and since ϕ_k is a linear projection, we have $(h_k f) \circ \phi_k^{-1} \in \mathcal{C}^{0,1}([-2B, 2B]^{d_0})$. Lemma 8 implies that there exists a neural network $\tilde{f}_k \in \mathcal{F}_{\text{NN}}(1, L_2, p_2, M)$ with

$$L_2 = O(\tilde{L}_2 \log \tilde{L}_2), p_2 = O(\tilde{p}_2 \log \tilde{p}_2), M = R.$$

for any $\tilde{L}_2, \tilde{p}_2 > 0$ such that

$$\|\tilde{f}_k - (h_k f) \circ \phi_k^{-1}\|_\infty \leq C_1 \tilde{L}_2^{-\frac{2}{d_0}} \tilde{p}_2^{-\frac{2}{d_0}}$$

for some constant C_1 depending on d_0, B, R .

Step 4. We then assemble all subnetworks constructed in the previous steps and approximate f by

$$\tilde{f} = \sum_{k=1}^{C_M} \tilde{\times}(\tilde{f}_k \circ \phi_k, \tilde{\mathbb{I}}_k \circ \tilde{d}_k^2). \quad (114)$$

In (114), according to Lemma 19, we set $\tilde{\times} \in \mathcal{F}_{\text{NN}}(1, 4L_3, 9p_3 + 1, M)$ as an approximation of \times with $M = R$ and error $24R^2 p_3^{-4L_3}$. The following lemma gives an upper bound of the approximation error of \tilde{f} (see a proof in Appendix I):

Lemma 21. The error of \tilde{f} can be decomposed as

$$\|\tilde{f} - f\|_\infty \leq \sum_{k=1}^{C_M} A_{k,1} + A_{k,2} + A_{k,3}$$

with

$$\begin{aligned} A_{k,1} &= \left\| \tilde{\times}(\tilde{f} \circ \phi_k^{-1}, \tilde{\mathbb{I}}_\Delta \circ \tilde{d}_k^2) - (\tilde{f} \circ \phi_k^{-1}) \times (\tilde{\mathbb{I}}_\Delta \circ \tilde{d}_k^2) \right\|_\infty \leq 24R^2 p_3^{-2} L_3^{-2}, \\ A_{k,2} &= \left\| (\tilde{f} \circ \phi_k^{-1}) \times (\tilde{\mathbb{I}}_\Delta \circ \tilde{d}_k^2) - [(h_k f) \circ \phi_k^{-1}] \times (\tilde{\mathbb{I}}_\Delta \circ \tilde{d}_k^2) \right\|_\infty \leq C_{12} \tilde{L}_2^{-\frac{2}{d_0}} \tilde{p}_2^{-\frac{2}{d_0}}, \\ A_{k,3} &= \left\| [(h_k f) \circ \phi_k^{-1}] \times (\tilde{\mathbb{I}}_\Delta \circ \tilde{d}_k^2) - [(h_k f) \circ \phi_k^{-1}] \times \mathbb{1}_{[0, r^2]} \right\|_\infty \leq \frac{C_{13}(\pi + 1)}{r(1 - r/\tau)} \Delta, \end{aligned}$$

for some constant C_{12} depending on d_0, τ, B, R , and C_{13} depending on R .

According to Lemma 21, for any $\tilde{L}, \tilde{p} > 0$, we set

- $\tilde{f}_k \in \mathcal{F}_{\text{NN}}(1, L_2, p_2, M)$ with $L_2 = O(\tilde{L} \log \tilde{L}), p_2 = O(\tilde{p} \log \tilde{p})$.
- $\tilde{\times} \in \mathcal{F}_{\text{NN}}(1, 4L_3, 9p_3 + 1, M)$ with $L_3 = O(\tilde{L}), p_3 = O(\tilde{p})$,
- $\tilde{d}_k^2 \in \mathcal{F}_{\text{NN}}(1, 4L_1, d_\chi(9p_1 + 1), M)$ with $\Delta = \tilde{L}^{-\frac{2}{d_0}} \tilde{p}^{-\frac{2}{d_0}}, L_1 = \tilde{L} + \log(12d_\chi B^2), p_1 = \tilde{p}$ such that

$$\begin{aligned} 24d_\chi B^2 p_1^{-4L_1} &= 24d_\chi B^2 \tilde{p}^{-4\tilde{L} - \log(48d_\chi B^2)} = 24d_\chi B^2 \tilde{p}^{-\log(48d_\chi B^2)} \tilde{p}^{-4\tilde{L}} \\ &= 24d_\chi B^2 (48d_\chi B^2)^{-\log \tilde{p}} \tilde{p}^{-4\tilde{L}} \leq 24d_\chi B^2 (24d_\chi B^2)^{-1} \tilde{p}^{-4\tilde{L}} \\ &\leq \tilde{p}^{-4\tilde{L}} \leq \tilde{p}^{-2(\tilde{L}+1)} \leq \tilde{p}^{-2} 2^{-2\tilde{L}} \leq \tilde{p}^{-2} \tilde{L}^{-2} < \Delta, \end{aligned} \quad (115)$$

- $\tilde{\mathbb{1}}_\Delta \in \mathcal{F}_{\text{NN}}(1, 2, 1, 1)$.

The total approximation error is bounded by $C_3 \tilde{L}^{-\frac{2}{d_0}} \tilde{p}^{-\frac{2}{d_0}}$ for some C_3 depending on d_0, R, B, τ and the surface area of \mathcal{M} . The constant hidden in $O(\cdot)$ depends on d_0, R, B, τ and the surface area of \mathcal{M} . The resulting network is in $\mathcal{F}_{\text{NN}}(1, L, p, M)$ with L, p, M defined in (85). □

I Proof of Lemma 21

Proof. For $A_{k,1}$, since $\tilde{\varphi} \in \mathcal{F}_{\text{NN}}(1, 4L_3, p_3, R)$, by Lemma 19, we have

$$A_{k,1} \leq 24R^2 p_3^{-4L_3} \leq 24R^2 p_3^{-2(L_3+1)} \leq 24R^2 p_3^{-2} 2^{-2L_3} \leq 24R^2 p_3^{-2} L_3^{-2}.$$

For $A_{k,2}$, since $\tilde{\mathbb{1}}_k \circ \tilde{a}_k^2 \in [0, 1]$, we have

$$A_{k,2} \leq \left\| \tilde{f} \circ \phi_k^{-1} - (h_k f) \circ \phi_k^{-1} \right\|_\infty \leq \left\| \tilde{f} - (h_k f) \right\|_\infty \leq C_{12} \tilde{L}_2^{-\frac{2}{d_0}} \tilde{p}_2^{-\frac{2}{d_0}}.$$

The upper bound of $A_{k,3}$ is proved in [13, Proof of Lemma 3]. □