

Operator learning with PCA-Net: upper and lower complexity bounds

Samuel Lanthaler

SLANTH@CALTECH.EDU

*Computing and Mathematical Sciences
California Institute of Technology
Pasadena, CA 91125, USA*

Editor: Aarti Singh

Abstract

PCA-Net is a recently proposed neural operator architecture which combines principal component analysis (PCA) with neural networks to approximate operators between infinite-dimensional function spaces. The present work develops approximation theory for this approach, improving and significantly extending previous work in this direction: First, a novel universal approximation result is derived, under minimal assumptions on the underlying operator and the data-generating distribution. Then, two potential obstacles to efficient operator learning with PCA-Net are identified, and made precise through lower complexity bounds; the first relates to the complexity of the output distribution, measured by a slow decay of the PCA eigenvalues. The other obstacle relates to the inherent complexity of the space of operators between infinite-dimensional input and output spaces, resulting in a rigorous and quantifiable statement of a “curse of parametric complexity”, an infinite-dimensional analogue of the well-known curse of dimensionality encountered in high-dimensional approximation problems. In addition to these lower bounds, upper complexity bounds are finally derived. A suitable smoothness criterion is shown to ensure an algebraic decay of the PCA eigenvalues. Furthermore, it is shown that PCA-Net can overcome the general curse for specific operators of interest, arising from the Darcy flow and the Navier-Stokes equations.

Keywords: neural operator, operator learning, curse of dimensionality, principal component analysis

1. Introduction

The application of neural networks (Goodfellow et al., 2016) to computational science and engineering is receiving growing interest. At their core, many problems of scientific interest involve the approximation of an underlying operator, which defines a mapping between two infinite-dimensional spaces of functions. *Neural operators* (Anandkumar et al., 2020; Bhattacharya et al., 2021; Lu et al., 2021; Kovachki et al., 2023) are a generalization of neural networks to such an infinite-dimensional setting. They aim to approximate, or “learn”, operators from data given in the form of input and output pairs. Neural operators hold promise as surrogate models to accelerate and complement traditional numerical methods in many-query problems, and they can be used for data-driven discovery of the underlying input-output map, even when no mathematical model is available.

Recent years have seen the emergence of several neural operator architectures. This includes deep operator networks (DeepONet) (Lu et al., 2021), building on early work on operator learning in (Chen and Chen, 1995). We also mention subsequent extensions of DeepONets, for example (Jin et al., 2022; Prasthofer et al., 2022; Seidman et al., 2022; Lanthaler et al., 2023b; Patel et al., 2022). DeepONets have been deployed with success in a variety of applications (Di Leoni et al., 2021; Mao et al., 2021; Cai et al., 2021). Another popular approach is based on a class of neural operators introduced in (Anandkumar et al., 2020; Kovachki et al., 2023). Here, neural operators are defined in close analogy with conventional neural networks, where the weight matrices in the hidden layers are generalized to integral operators. Special cases of this framework include the graph neural operator (Anandkumar et al., 2020; Li et al., 2020) and the Fourier neural operator (FNO) (Li et al., 2021). In this context, we also mention related frameworks in (Raonic et al., 2023; Tripura and Chakraborty, 2023; Gupta et al., 2021). Another notable, and somewhat distinct, approach to operator learning is the operator-valued random feature model proposed in (Nelsen and Stuart, 2021).

Universal approximation results for many of these frameworks are known in a variety of settings; the universality of DeepONets is established in (Chen and Chen, 1995; Lu et al., 2021; Lanthaler et al., 2022), FNOs are shown to be universal in (Kovachki et al., 2021), another universal approximation result for neural operators is derived in (Kovachki et al., 2023). We also mention recent work (Lanthaler et al., 2023a), which proves a general universal approximation result for the so-called “averaging neural operator” (ANO), a minimal architecture that is at the core of many other frameworks, thereby allowing to unify much of the analysis of this emerging zoo of neural operator architectures.

Going beyond universality, it is crucial to improve our understanding of the required computational complexity of neural operators in order to assess when the methods will be effective. Pertinent numerical experiments may be found in (De Hoop et al., 2022). Relevant analysis of linear problems from this point of view has been given in (Boullé and Townsend, 2022; de Hoop et al., 2023). The required complexity of neural network-based methods for specific PDE operators of interest are studied from a approximation theoretic point of view, in e.g. (Kovachki et al., 2021; Schwab and Zech, 2019; Kutyniok et al., 2022; Lanthaler et al., 2022; Ryck and Mishra, 2022). A focus of these papers is on beating the “*curse of dimensionality*”.

Since the input and output spaces are infinite-dimensional in these problems, clarification may be needed as to the meaning of beating the curse of dimensionality: it is interpreted as identifying conditions under which the required size (number of tunable parameters) of the operator approximation grows only algebraically with the inverse of the desired error. This notion of the “curse of dimensionality” is specific to operator learning and somewhat distinct from the conventional meaning of this term in the context of high-dimensional problems. To disambiguate the well-known curse of dimensionality in high-dimensional approximation problems from the analogue of this curse in infinite-dimensional approximation (i.e. operator learning), we will refer to the latter as the “curse of parametric complexity”. The present work will provide further clarification and motivation for the use of this term in the context of operator learning, and will connect the curse of dimensionality with this curse of parametric complexity.

The focus of the present work is the so-called *PCA-Net*, a methodology which combines ideas from principal component analysis with neural networks (Hesthaven and Ubbiali, 2018; Bhattacharya et al., 2021). Principal component analysis (PCA) is a standard tool for dimension reduction in high-dimensional statistics and unsupervised learning (Jolliffe, 2002). In (Bhattacharya et al., 2021), a combination of PCA with neural networks has been proposed as a data-driven operator learning framework. As indicated above, the goal of operator learning is to approximate an unknown operator $\Psi^\dagger : \mathcal{X} \rightarrow \mathcal{Y}$, mapping between two infinite-dimensional spaces \mathcal{X} and \mathcal{Y} . Given data in the form of pairs of inputs and outputs, we seek to determine an accurate, data-driven approximation of Ψ^\dagger . The PCA-Net operator learning architecture achieves this goal by (i) using PCA to reduce the dimensions of the input and output spaces and (ii) approximating a map between the resulting finite-dimensional latent spaces (Bhattacharya et al., 2021). First analysis, including a universal approximation result, have been derived in (Bhattacharya et al., 2021). Furthermore, in the same work, the efficacy of the proposed architecture has been demonstrated empirically for prototypical problems, including the solution operator of the viscous Burgers equation and the Darcy flow equation. However, so far, a detailed mathematical analysis providing a theoretical underpinning for this empirically observed efficiency of PCA-Net, has been outstanding.

The present work fills this gap by developing relevant approximation theory for PCA-Net. The main contributions of this paper are the following:

- **Universal approximation:** We prove a novel universal approximation theorem for PCA-Net, Theorem 6, under significantly relaxed conditions on the distribution of the data-generating measure and the underlying operator Ψ^\dagger compared to previous work; the universality of PCA-Net is here shown under natural minimal conditions, which are in fact necessary for PCA to be well-defined on the input and output spaces.
- **Curse of parametric complexity:** A rigorous result is proven which demonstrates that the curse of parametric complexity, an infinite-dimensional scaling limit of the conventional curse of dimensionality, cannot be overcome by PCA-Net in general (cp. Theorem 9); more precisely, this result shows that it is impossible to derive algebraic complexity bounds when considering general classes of operators, such as the class of all Lipschitz- or even \mathcal{C}^k -continuous operators. Hence, we conclude that at this level of generality, the curse is unavoidable.
- **Overcoming the curse of parametric complexity:** Given the negative result on the general curse, we argue that a central challenge in operator learning is to identify the relevant class of operators which *do allow* for efficient approximation by a given operator learning framework. To gain further insight into the relevant mathematical structure that can be leveraged by PCA-Net, we restrict attention to two prototypical PDE operators of interest arising from the Darcy flow and Navier-Stokes equations. In both cases, we show that PCA-Net can overcome the general curse of parametric complexity; algebraic error and complexity estimates are established in Theorems 18 and 25, demonstrating that these operators belong to a restricted class which is efficiently approximated by PCA-Net.

1.1 Overview

In section 2, relevant background on PCA and the PCA-Net methodology is provided; First, PCA and empirical PCA are reviewed in section 2.1, and two error estimates for the PCA projection error are stated in section 2.2. Next, it is explained how PCA-Net combines PCA with a neural network, resulting in an operator learning architecture.

In Section 3, we develop approximation theory for PCA-Net. A new universal approximation result for PCA-Net is derived in section 3.1. In section 3.2, two potential obstacles to the efficacy of PCA-Net are identified through rigorous lower complexity bounds. Upper complexity bounds are the subject of the remaining sections: In section 3.3, a smoothness criterion is shown to rule out the first potential obstacle to efficient operator learning. We then finally show, in section 3.4, how PCA-Net can overcome the curse of parametric complexity for two prototypical PDE operators arising in the context of the Darcy flow and Navier-Stokes equations, respectively. Conclusions and perspectives for future work are summarized in Section 4.

1.2 Notation

Throughout the following discussion, $\mathcal{H}, \mathcal{X}, \mathcal{Y}$ denote separable Hilbert spaces. We will use $\|\cdot\|_{\mathcal{H}}$ and $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ to denote the norm and inner product on \mathcal{H} ; if it is clear from the context, we may occasionally omit the subscript to aid readability. On finite-dimensional Euclidean spaces, $|\cdot|$ is used for the Euclidean norm, and $|\cdot|_{\infty}$ denotes the maximum-norm. The space of probability measures on \mathcal{H} is denoted $\mathcal{P}(\mathcal{H})$. We denote by $u \sim \mu$ a random variable distributed according to probability measure μ . $\mathbb{E}_{u \sim \mu}[F(u)]$ denotes the expectation of F with respect to μ . We consistently use Ψ^{\dagger} to denote the underlying (truth) operator, and Ψ will denote a (PCA-Net) approximation of Ψ^{\dagger} . For two numbers a, b , we will write $a \sim b$ for equivalence up to constants, i.e. there exists $C > 0$ such that $C^{-1}a \leq b \leq Ca$. Similarly $a \lesssim b$, $a \gtrsim b$ denotes inequality up to a constant, i.e. $a \leq Cb$ and $Ca \geq b$, respectively. On occasion, we write a subscript $a \lesssim_k b$ to emphasize the dependence of the implied constant $C = C(k)$ on a given parameter k . We follow the convention that constants in estimates can change their value from line to line; their dependence on the relevant parameters will always be indicated. Other notation is introduced as needed.

2. PCA-Net methodology

PCA-Net combines principal component analysis (PCA) for dimension reduction of the input and output spaces, with a neural network mapping between the resulting finite-dimensional latent spaces. Before summarizing PCA-Net, we first review PCA in subsection 2.1, and derive two high-probability estimates for the PCA projection error in subsection 2.2. Next, we summarize how PCA-Net combines PCA with a neural network, resulting in an operator learning architecture in subsection 2.3.

2.1 Principal Component Analysis

In the present section, we provide necessary background material on PCA, and we prove a high-probability estimate for the PCA projection error, building on previous results (Reiß and Wahl, 2020; Milbradt and Wahl, 2020).

PCA. Given a Hilbert space \mathcal{H} , a probability measure μ on \mathcal{X} , and a projection dimension d , PCA aims to minimize the average reconstruction error $\mathbb{E}_{u \sim \mu}[\|u - Pu\|^2]$ over the set Π_d of orthogonal projections $P : \mathcal{H} \rightarrow \mathcal{H}$ of rank d . It is well-known (e.g. (Jolliffe, 2002; Bhattacharya et al., 2021; Reiß and Wahl, 2020)) that this can be achieved by considering the covariance operator $\Sigma = \mathbb{E}_{u \sim \mu}[u \otimes u]$, which is diagonalizable; i.e., there exists a sequence $\lambda_1 \geq \lambda_2 \geq \dots \geq 0$ of eigenvalues and corresponding orthonormal basis of eigenvectors $\phi_1, \phi_2, \dots \in \mathcal{H}$, such that $\Sigma \phi_j = \lambda_j \phi_j$ for all j . The optimal PCA projection of dimension d , $P_{\leq d} : \mathcal{H} \rightarrow \mathcal{H}$, can then be written as a composition $P_{\leq d} = \mathcal{D}_{\mathcal{H}}^{\text{opt}} \circ \mathcal{E}_{\mathcal{H}}^{\text{opt}}$, where the optimal PCA encoder $\mathcal{E}_{\mathcal{H}}^{\text{opt}}$ is given by,

$$\mathcal{E}_{\mathcal{H}}^{\text{opt}} : \mathcal{H} \rightarrow \mathbb{R}^d, \quad \mathcal{E}_{\mathcal{H}}^{\text{opt}}(u) := (\langle u, \phi_1 \rangle, \dots, \langle u, \phi_d \rangle), \quad (1)$$

and the corresponding PCA decoder $\mathcal{D}_{\mathcal{H}}^{\text{opt}}$ is the mapping,

$$\mathcal{D}_{\mathcal{H}}^{\text{opt}} : \mathbb{R}^d \rightarrow \mathcal{H}, \quad \mathcal{D}_{\mathcal{H}}^{\text{opt}}(\eta) = \sum_{j=1}^d \eta_j \phi_j. \quad (2)$$

It can be shown that $P_{\leq d}$ defines an orthogonal projection on \mathcal{H} , such that

$$\mathbb{E}_{u \sim \mu}[\|u - P_{\leq d} u\|_{\mathcal{H}}^2] = \min_{P \in \Pi_d} \mathbb{E}_{u \sim \mu}[\|u - Pu\|_{\mathcal{H}}^2].$$

In the following we denote by

$$\mathcal{R}_d^{\text{opt}}(\mu) := \min_{P \in \Pi_d} \mathbb{E}_{u \sim \mu}[\|u - Pu\|_{\mathcal{H}}^2], \quad (3)$$

this optimal PCA projection error. One can show, e.g. (Lanthaler et al., 2022, Thm. 3.8), that $\mathcal{R}_d^{\text{opt}}(\mu)$ is related to the PCA eigenvalues by

$$\mathcal{R}_d^{\text{opt}}(\mu) = \sum_{j>d} \lambda_j. \quad (4)$$

Empirical PCA. Empirical PCA applies the above procedure to the empirical distribution $\mu_N = \frac{1}{N} \sum_{k=1}^N \delta_{u_k}$, obtained by sampling from μ : Given a finite number of independent and identically distributed (i.i.d.) samples $u_1, \dots, u_N \stackrel{iid}{\sim} \mu$, define the covariance operator Σ_N , by

$$\Sigma_N = \frac{1}{N} \sum_{k=1}^N u_k \otimes u_k. \quad (5)$$

Letting $\hat{\lambda}_1 \geq \hat{\lambda}_2 \geq \dots \geq 0$ and $\hat{\phi}_1, \hat{\phi}_2, \dots \in \mathcal{H}$ denote the eigenvalues and corresponding orthonormal eigenbasis. The empirical PCA projection $\hat{P}_{\leq d}$ of dimension d is given by $\hat{P}_{\leq d} = \mathcal{D}_{\mathcal{H}} \circ \mathcal{E}_{\mathcal{H}}$, where $\mathcal{E}_{\mathcal{H}}$, $\mathcal{D}_{\mathcal{H}}$ are defined as in (1), (2), but replacing the eigenvectors ϕ_j by their empirical counterparts $\hat{\phi}_j$.

2.2 Projection error of empirical PCA

We first note that the empirical PCA projection error approximates the optimal projection error, provided that a sufficient amount of data is available. We state the following result in high probability.

Proposition 1 *Let \mathcal{H} be a separable Hilbert space. Let $\mu \in \mathcal{P}(\mathcal{H})$ be a probability measure with finite second moments, $\mathbb{E}_{u \sim \mu}[\|u\|_{\mathcal{H}}^2] < \infty$. Then for any $\delta, \epsilon > 0$, there exists a requisite amount of data $N_0 = N_0(\mu, d, \delta, \epsilon)$, such that the encoding error for empirical PCA with dimension d , and based on $N \geq N_0$ samples $u_1, \dots, u_N \stackrel{iid}{\sim} \mu$, satisfies*

$$\mathbb{E}_{u \sim \mu} [\|u - \mathcal{D}_{\mathcal{H}} \circ \mathcal{E}_{\mathcal{H}}(u)\|_{\mathcal{H}}^2] \leq \mathcal{R}_d^{\text{opt}}(\mu) + \epsilon, \quad (6)$$

with probability at least $1 - \delta$.

The proof of Proposition 1 relies on a well-known bound on the excess risk for empirical PCA in terms of the Hilbert-Schmidt distance $\|\Sigma - \widehat{\Sigma}\|_2$ of the true and empirical covariance operators (Reiß and Wahl, 2020, e.g. Sect. 2.2). This is combined with a general Monte-Carlo estimate to prove Proposition 1. We present the details in Appendix A.

The result of Proposition 1 is purely qualitative, as it doesn't give us any estimate on the required amount of data N . Our next goal is to establish a *quantitative* bound, under additional assumptions on μ . We will call a probability measure $\mu \in \mathcal{P}(\mathcal{H})$ **sub-Gaussian**, if there exists $K_{\mu} \geq 0$, such that

$$\mathbb{E}_{u \sim \mu} [\|u\|_{\mathcal{H}}^p]^{1/p} \leq K_{\mu} \sqrt{p}, \quad \forall p \geq 1. \quad (7)$$

According to this definition (7), μ is sub-Gaussian if the real-valued random variable $\|u\|_{\mathcal{H}}$, with $u \sim \mu$, is sub-Gaussian in the conventional sense. The moment bound (7) is one of many equivalent characterizations of real-valued sub-Gaussian random variables (see e.g. (Vershynin, 2018, Sect. 2.5.1)). We then have:

Proposition 2 *Let \mathcal{H} be a separable Hilbert space, and let μ be a sub-Gaussian probability measure on \mathcal{H} . Fix $\delta \in (0, 1/2)$. The encoding error for empirical PCA with dimension d , and based on $N \geq \log(2/\delta)$ samples $u_1, \dots, u_N \stackrel{iid}{\sim} \mu$, satisfies the following upper bound,*

$$\mathbb{E}_{u \sim \mu} [\|u - \mathcal{D}_{\mathcal{H}} \circ \mathcal{E}_{\mathcal{H}}(u)\|_{\mathcal{H}}^2] \leq \mathcal{R}_d^{\text{opt}}(\mu) + \sqrt{\frac{Qd \log(2/\delta)}{N}}. \quad (8)$$

with probability at least $1 - \delta$. Here, $Q = Q(K_{\mu})$ depends only on the constant K_{μ} in (7).

Proposition 2 above is a natural high-probability analogue of a previous result on empirical PCA of (Bhattacharya et al., 2021), derived in expectation in that work. Proposition 2 uses the same bound on the PCA excess risk as Proposition 1, but combines it with a general Bernstein concentration bound for \mathcal{H} -valued random variables to derive quantitative rates. We include the details in Appendix A.

Remark 3 *We note that under more fine-grained information on the underlying measure μ , considerable improvements to the upper bound of Proposition 2 are possible; this has*

e.g. been achieved in (Milbradt and Wahl, 2020) for a different notion of a “sub-Gaussian” distribution, requiring that

$$\sup_{p \in \mathbb{N}} \frac{\mathbb{E}_{u \sim \mu} [|\langle u, v \rangle_{\mathcal{H}}|^p]^{1/p}}{\sqrt{p}} \leq K'_\mu \mathbb{E}_{u \sim \mu} [|\langle u, v \rangle_{\mathcal{H}}|^2]^{1/2}, \quad \forall v \in \mathcal{H}, \quad (9)$$

for a constant K'_μ depending only on μ . The condition (9) is stronger than (7); Assuming (9), it can in fact be shown that the empirical PCA projection error is of order $\mathcal{O}(\mathcal{R}_d^{\text{opt}}(\mu))$ whenever $N \gtrsim d \log(1/\delta)$, thereby achieving essentially optimal PCA convergence rates. However, in the present context, PCA will be applied for dimension reduction on both the input and output spaces under a non-linear mapping Ψ^\dagger . Unfortunately, even if μ satisfies (9), it is unclear whether this property is preserved under the push-forward by Ψ^\dagger , i.e. whether a bound of the form (9) continues to hold for $\Psi^\dagger_{\#}\mu$. In contrast, the bound (7) is robust under such a push-forward. Therefore, we contend ourselves with the more pessimistic bound of Proposition 2, and leave potential improvements, possibly building on (Milbradt and Wahl, 2020), as a challenge for future work.

2.3 PCA-Net architecture

We next recall the PCA-Net architecture proposed in (Bhattacharya et al., 2021), which combines empirical PCA with a neural network mapping, to approximate an underlying operator $\Psi^\dagger : \mathcal{X} \rightarrow \mathcal{Y}$. In the following, let \mathcal{X} and \mathcal{Y} be separable Hilbert spaces and let $\Psi^\dagger : \mathcal{X} \rightarrow \mathcal{Y}$ be a non-linear operator. The goal of the PCA-Net methodology is to approximate Ψ^\dagger from a finite number of input-/output-samples $\{u_k, \Psi^\dagger(u_k)\}_{k=1}^N$. To this end, PCA-Net combines an encoding $\mathcal{E}_{\mathcal{X}} : \mathcal{X} \rightarrow \mathbb{R}^{d_x}$, a neural network $\psi : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_y}$ and a decoding $\mathcal{D}_{\mathcal{Y}} : \mathbb{R}^{d_y} \rightarrow \mathcal{Y}$ (cp. Figure 1).




Figure 1: Diagrammatic illustration of PCA-Net based on a PCA encoder $\mathcal{E}_{\mathcal{X}}$, a neural network ψ , and a PCA decoder $\mathcal{D}_{\mathcal{Y}}$.

Here, $\mathcal{E}_{\mathcal{X}}$ and $\mathcal{D}_{\mathcal{Y}}$ are chosen as an empirical PCA encoder and decoder, respectively. A precise definition is given below, see equations (12), (13). Given these ingredients, the resulting PCA-Net is defined as the mapping

$$\Psi : \mathcal{X} \rightarrow \mathcal{Y}, \quad \Psi(u) := \mathcal{D}_{\mathcal{Y}} \circ \psi \circ \mathcal{E}_{\mathcal{X}}(u). \quad (10)$$

The encoder $\mathcal{E}_{\mathcal{X}}$ and decoder $\mathcal{D}_{\mathcal{Y}}$ perform a dimension reduction on the input and output spaces. The neural network ψ approximates a mapping on the resulting finite-dimensional latent spaces.

Associated with the encoder $\mathcal{E}_{\mathcal{X}}$, we will define a natural PCA decoder $\mathcal{D}_{\mathcal{X}}$ on \mathcal{X} . Similarly $\mathcal{D}_{\mathcal{Y}}$ is accompanied by an encoder $\mathcal{E}_{\mathcal{Y}}$ on \mathcal{Y} . The intuition is that these underlying

encoder/decoder pairs, and the neural network ψ , should satisfy the following approximate identities (Bhattacharya et al., 2021):

$$\mathcal{D}_{\mathcal{X}} \circ \mathcal{E}_{\mathcal{X}} \approx \text{id}_{\mathcal{X}}, \quad \mathcal{D}_{\mathcal{Y}} \circ \mathcal{E}_{\mathcal{Y}} \approx \text{id}_{\mathcal{Y}}, \quad \mathcal{D}_{\mathcal{Y}} \circ \psi \circ \mathcal{E}_{\mathcal{X}} \approx \Psi^\dagger.$$

Here, $\text{id}_{\mathcal{X}}$ and $\text{id}_{\mathcal{Y}}$ denote the identity mappings on \mathcal{X} and \mathcal{Y} , respectively. Our analysis of the PCA-methodology aims to quantify the accuracy of these approximations when the encoders/decoders are defined by PCA and ψ is chosen to be a neural network.

PCA encoding and decoding. We will now specify the particular choice of $\mathcal{E}_{\mathcal{X}}$ and $\mathcal{D}_{\mathcal{Y}}$. In the following, we assume that the input samples $\{u_k\}_{k=1}^N$ are i.i.d. samples from a probability measure μ on \mathcal{X} . Note that the output-samples $\{\Psi^\dagger(u_k)\}_{k=1}^N$ are then i.i.d. with respect to the corresponding push-forward measure $\Psi^\dagger_{\#}\mu$ on \mathcal{Y} . For a given choice of latent dimensions $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$, we apply empirical PCA to the samples $\{u_k\}_{k=1}^N \subset \mathcal{X}$ and $\{\Psi^\dagger(u_k)\}_{k=1}^N \subset \mathcal{Y}$. In the following, we denote by

$$\phi_1^{\mathcal{X}}, \dots, \phi_{d_{\mathcal{X}}}^{\mathcal{X}} \in \mathcal{X}, \quad \phi_1^{\mathcal{Y}}, \dots, \phi_{d_{\mathcal{Y}}}^{\mathcal{Y}} \in \mathcal{Y}, \quad (11)$$

the empirical PCA bases on \mathcal{X} and \mathcal{Y} , respectively. We emphasize that the empirical PCA bases are themselves random variables, as they depend on the random input-/output-samples $\{u_k, \Psi^\dagger(u_k)\}_{k=1}^N$. The first basis, $\{\phi_j^{\mathcal{X}}\}_{j=1}^{d_{\mathcal{X}}}$, defines an encoder on \mathcal{X} ,

$$\mathcal{E}_{\mathcal{X}} : \mathcal{X} \rightarrow \mathbb{R}^{d_{\mathcal{X}}}, \quad \mathcal{E}_{\mathcal{X}}(u) := (\langle u, \phi_1^{\mathcal{X}} \rangle, \dots, \langle u, \phi_{d_{\mathcal{X}}}^{\mathcal{X}} \rangle), \quad (12)$$

with corresponding decoder, $\mathcal{D}_{\mathcal{X}} : \mathbb{R}^{d_{\mathcal{X}}} \rightarrow \mathcal{X}$, $\mathcal{D}_{\mathcal{X}}(\xi) := \sum_{j=1}^{d_{\mathcal{X}}} \xi_j \phi_j^{\mathcal{X}}$. Similarly, $\phi_1^{\mathcal{Y}}, \dots, \phi_{d_{\mathcal{Y}}}^{\mathcal{Y}}$ defines an encoder on \mathcal{Y} , $\mathcal{E}_{\mathcal{Y}} : \mathcal{Y} \rightarrow \mathbb{R}^{d_{\mathcal{Y}}}$, $\mathcal{E}_{\mathcal{Y}}(v) := (\langle v, \phi_j^{\mathcal{Y}} \rangle)_{j=1}^{d_{\mathcal{Y}}}$, with corresponding decoder,

$$\mathcal{D}_{\mathcal{Y}} : \mathbb{R}^{d_{\mathcal{Y}}} \rightarrow \mathcal{Y}, \quad \mathcal{D}_{\mathcal{Y}}(\eta) := \sum_{j=1}^{d_{\mathcal{Y}}} \eta_j \phi_j^{\mathcal{Y}}. \quad (13)$$

Neural networks. Given a depth $L \in \mathbb{N}$, layer widths d_k ($k = 0, \dots, L+1$), and weights and biases $A_k \in \mathbb{R}^{d_{k+1} \times d_k}$, $b_k \in \mathbb{R}^{d_{k+1}}$, $k = 0, \dots, L$, a (*deep*) *neural network* (DNN) ψ is a mapping $\xi \mapsto \psi(\xi)$, defined as a composition of non-linear layers,

$$\begin{cases} \xi_0 := \xi, \\ \xi_{k+1} := \sigma(A_k \xi_k + b_k), & \text{for } k = 0, \dots, L-1, \\ \psi(\xi) := A_L \xi_L + b_L. \end{cases} \quad (14)$$

The non-linear layers are expressed in terms of an activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ which is applied componentwise. In the following, we will restrict attention to the ReLU activation function $\sigma(\xi) := \max(\xi, 0)$. Extension of our results to more general choices of σ is possible.

Given a DNN ψ , we define $\text{size}(\psi) := \sum_{k=0}^L (\|A_k\|_0 + \|b_k\|_0)$ as the total number of non-zero weights and biases in the architecture, and we define $\text{depth}(\psi) := L$ as the number of layers. With these definitions, $\text{size}(\psi)$ and $\text{depth}(\psi)$ provide a measure of the complexity of the DNN.

Remark 4 *In the present work, we employ the above notion of size as a measure of the complexity of a neural network. This is sufficient for our present purposes, but it should be pointed out that this notion of complexity does not take into account the number of bits which would be required to encode each parameter in a practical implementation. Recent work such as Yarotsky (2018); Siegel (2022), and references therein, has shown that it is possible for very carefully constructed ReLU neural networks to achieve unexpectedly high convergence rates for the approximation of Sobolev functions. This comes at the expense of requiring a large number of bits to encode the individual weights, e.g. (Siegel, 2022, discussion on page 5).*

Neural network training. Ideally, the neural network would be chosen as a minimizer of the expected loss

$$\mathcal{L}(\psi) := \mathbb{E}_{(\xi, \eta)} [|\psi(\xi) - \eta|^2],$$

where the expectation is over pairs $(\xi, \eta) \in \mathbb{R}^{d_x} \times \mathbb{R}^{d_y}$ of encoded input-/output-pairs of the form $\xi = \mathcal{E}_{\mathcal{X}}(u)$, $\eta = \mathcal{E}_{\mathcal{Y}}(\Psi^\dagger(u))$, with $u \sim \mu$. In practice, the neural network ψ in (10) is usually trained to minimize the following empirical loss,

$$\widehat{\mathcal{L}}(\psi) := \frac{1}{N} \sum_{k=1}^N |\psi(\xi_k) - \eta_k|^2,$$

where $\xi_k := \mathcal{E}_{\mathcal{X}}(u_k) \in \mathbb{R}^{d_x}$ and $\eta_k := \mathcal{E}_{\mathcal{Y}}(\Psi^\dagger(u_k)) \in \mathbb{R}^{d_y}$ are the encoded input and output samples.

Remark 5 *In the present work, we will not address the practical training of the neural network ψ , analysis of which is a notoriously difficult problem. Instead, we will contend ourselves with an “approximation theoretic” approach; it will be shown that certain ψ exists, but no a priori guarantee is given that this ψ can be found by the numerical optimization of the empirical loss $\widehat{\mathcal{L}}$. The samples u_1, \dots, u_N will, however, enter our analysis of the empirical PCA encoder $\mathcal{E}_{\mathcal{X}}$ and decoder $\mathcal{D}_{\mathcal{Y}}$.*

3. Approximation theory

This section develops approximation theory for PCA-Net, and is divided into four subsections. First, in Section 3.1, we discuss a new universality theorem for PCA-Net. Next, we point out two potential obstacles to efficient operator learning with PCA-Net in Section 3.2. As explained there, one such obstacle relates to the complexity of the output distribution on \mathcal{Y} , which could in principle lead to an arbitrarily slow decay of the PCA-Net error with the PCA-dimension $d_{\mathcal{Y}}$ (cp. Proposition 8). The second potential obstacle to efficient operator learning relates to the inherent complexity of the space of operators between infinite-dimensional spaces (cp. Theorem 9); this makes rigorous a notion of “curse of dimensionality” for operator learning intuited in earlier work (Kovachki et al., 2021; Lanthaler et al., 2022). The lower bounds of Section 3.2 are complemented by upper bounds in the subsequent sections; In Section 3.3, we show that a suitable smoothness condition rules out the first obstacle, ensuring an algebraic decay of the PCA encoding error on \mathcal{Y} .

Finally, in Section 3.4, we demonstrate that PCA-Net can leverage additional structure for two operators of interest, allowing it to overcome the general curse of dimensionality when approximating the solution operator of the Darcy flow and Navier-Stokes equations.

3.1 Universal approximation

We first state the following universal approximation theorem in high probability:

Theorem 6 (Universal approximation) *Let \mathcal{X}, \mathcal{Y} be separable Hilbert spaces and let $\mu \in \mathcal{P}(\mathcal{X})$ be a probability measure on \mathcal{X} . Let $\Psi^\dagger : \mathcal{X} \rightarrow \mathcal{Y}$ be a μ -measurable mapping. Assume the following moment conditions,*

$$\mathbb{E}_{u \sim \mu}[\|u\|_{\mathcal{X}}^2], \quad \mathbb{E}_{u \sim \mu}[\|\Psi^\dagger(u)\|_{\mathcal{Y}}^2] < \infty.$$

Then for any $\delta, \epsilon > 0$, there are dimensions $d_{\mathcal{X}} = d_{\mathcal{X}}(\epsilon, \delta)$, $d_{\mathcal{Y}} = d_{\mathcal{Y}}(\epsilon, \delta)$, a requisite amount of data $N = N(d_{\mathcal{X}}, d_{\mathcal{Y}}, \mu, \Psi^\dagger)$, a neural network $\psi = \psi(\cdot; \{u_k\})$ depending on this data, such that the PCA-Net, $\Psi = \mathcal{D}_{\mathcal{Y}} \circ \psi \circ \mathcal{E}_{\mathcal{X}}$, satisfies

$$\mathbb{E}_{u \sim \mu} \left[\|\Psi^\dagger(u) - \Psi(u; \{u_k\})\|_{\mathcal{Y}}^2 \right] \leq \epsilon,$$

with probability at least $1 - \delta$ in the input data $u_1, \dots, u_N \sim \mu$.

We have written $\Psi(u) = \Psi(u; \{u_k\})$ to emphasize the dependency of the PCA-Net encoder $\mathcal{E}_{\mathcal{X}}$, decoder $\mathcal{D}_{\mathcal{Y}}$ and the neural network $\psi = \psi(\cdot; \{u_k\})$ on the given data u_1, \dots, u_N . The detailed proof, included in Appendix B, combines well-known universal approximation results for neural networks with the high-probability bounds on the PCA projection error in Proposition 1 in Section 2.2, above; more precisely, it is shown in Lemma 32 that the PCA error can be decomposed into an error due to the encoding on \mathcal{X} , the decoding error on \mathcal{Y} and a neural network approximation error. The encoding and decoding errors are expressed in terms of the PCA projection error, which can be bounded by invoking Proposition 2. The neural network approximation error can then be made arbitrarily small by a suitable choice of the neural network ψ .

Remark 7 *As already pointed out in Remark 5, the practical training of the neural network ψ is not discussed in the present work. Hence the above theorem should be viewed as a pure existence result, in the sense that the proof does not provide an explicit algorithmic recipe for determining the neural network ψ , given the data u_1, \dots, u_N .*

Theorem 6 shows that PCA-Net is able to approximate almost arbitrary operators to any desired accuracy, provided a sufficient number of data points are available for empirical PCA, and provided that the underlying neural network is sufficiently large. A previous universal approximation result in (Bhattacharya et al., 2021, Theorem 3.1) was stated only for Lipschitz continuous Ψ^\dagger and under an assumption of finite fourth moments. In contrast, Theorem 6 requires no regularity condition on the underlying operator Ψ^\dagger , and shows that a bound on the second moments of the input measure μ and the push-forward measure $\Psi^\dagger_{\#}$ suffices for universal approximation. We have formulated Theorem 6 in high-probability, whereas (Bhattacharya et al., 2021, Theorem 3.1) is derived in expectation. This last

difference is mostly for consistency with the high probability results in later sections; indeed, Theorem 6 is in fact derived from a corresponding result in expectation (cp. Proposition 31).

The main drawback of universal approximation results is that they are purely qualitative, and do not provide any information about the required size of the data or the neural network; hence, universal approximation cannot provide information on the efficiency of operator learning with PCA-Net. Deriving more quantitative bounds is particularly relevant in view of the two potential obstacles to efficient operator learning, elaborated upon in the next section.

3.2 Obstacles to effective operator learning

Theorem 6 does not provide any quantitative information on the required complexity to achieve a given accuracy. For the practical success of PCA-Net in operator learning, it is crucial that the PCA-Net approximation is not only possible in principle, but also efficient in practice; we interpret “efficiency” as the statement that the PCA dimensions $d_{\mathcal{X}}$, $d_{\mathcal{Y}}$, the requisite amount of data N and the size of the neural network ψ , that is required to achieve a desired accuracy $\epsilon > 0$, should grow at most at an algebraic rate $\epsilon^{-\gamma}$, with quantifiable exponent $\gamma > 0$. As explained in this section, there are at least two potential obstacles to the efficiency of PCA-Net.

Complexity of the output distribution. The first potential reason for the inefficiency of PCA-Net is a consequence of the following lower bound on the approximation error.

Proposition 8 *Let $\mu \in \mathcal{P}(\mathcal{X})$ be a probability measure, and let $\Psi^\dagger \in L_\mu^2(\mathcal{X}; \mathcal{Y})$ be an operator. Let $\lambda_1 \geq \lambda_2 \geq \dots \geq 0$ be the PCA eigenvalues of the push-forward measure $\Psi^\dagger_\# \mu$ on \mathcal{Y} . Then we have the following lower bound,*

$$\mathbb{E}_{u \sim \mu} \left[\|\Psi^\dagger(u) - \Psi(u)\|_{\mathcal{Y}}^2 \right] \geq \sum_{j > d_{\mathcal{Y}}} \lambda_j, \quad (15)$$

for any PCA-Net Ψ with PCA dimension $d_{\mathcal{Y}}$ on \mathcal{Y} .

The proof of Proposition 8 is an almost verbatim repetition of the argument in (Lanthaler et al., 2022, Theorem 3.6); we won’t repeat the proof here. As a consequence of (15), the approximation error that can be achieved with a PCA dimension $d_{\mathcal{Y}}$ is lower bounded by the decay of the PCA eigenvalues of the push-forward measure $\Psi^\dagger_\# \mu$, i.e. by the optimal PCA projection error $\mathcal{R}_d^{\text{opt}}(\Psi^\dagger_\# \mu) = \sum_{j > d_{\mathcal{Y}}} \lambda_j$. In particular, if this decay is very slow, e.g. $\mathcal{R}_d^{\text{opt}}(\Psi^\dagger_\# \mu) \gtrsim \log(d_{\mathcal{Y}})^{-1}$, then an *exponentially* large PCA dimension $d_{\mathcal{Y}}(\epsilon) \sim \exp(\epsilon^{-1})$ is required, which entails that also an exponential number of samples $N \gtrsim \exp(\epsilon^{-1})$ and an exponential neural network size $\text{size}(\psi) \gtrsim \exp(\epsilon^{-1})$ are required: The bound on N follows from the fact that the empirical PCA covariance operator (5) is of rank at most N , and hence the determination of $d_{\mathcal{Y}}$ non-trivial PCA eigendirections requires at least $N \geq d_{\mathcal{Y}}$ samples. Furthermore, since the neural network ψ must have $d_{\mathcal{Y}}$ non-trivial output components, at least $d_{\mathcal{Y}}$ non-zero weights are required. The first potential obstacle thus relates to the *complexity of the output space*, encoded in the PCA eigenvalue decay of the measure $\Psi^\dagger_\# \mu$. We note in passing that the problem of a slow eigenvalue decay can sometimes

be ameliorated by replacing the linear decoder $\mathcal{D}_{\mathcal{Y}}$ by a non-linear mapping, leading to improved results both theoretically and empirically (Seidman et al., 2022; Lanthaler et al., 2023b).

Curse of dimensionality and parametric complexity. The last section shows that the complexity, or “size”, of the output space \mathcal{Y} can be one obstacle to operator learning with PCA-Net. A second potential obstacle to efficient operator learning is that the space of operators $\Psi^\dagger : \mathcal{X} \rightarrow \mathcal{Y}$ itself is very large.

It is well-known that the task of approximating a high-dimensional function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ by ordinary methods, such as polynomial approximation, suffers from a curse of dimensionality, where the number of degrees of freedom needed to achieve a desired accuracy scales exponentially in d . Indeed, optimal error bounds for interpolation are typically of the form,

$$\sup_{\xi \in D} |f(\xi) - p(\xi)| \lesssim \|f\|_{C^k} N^{-k/d}, \quad (16)$$

where p denotes the (e.g. polynomial) interpolant and N represents the number of degrees of freedom of the interpolation space. Such upper bounds deteriorate for large d , with convergence rates becoming arbitrarily slow. Similar results have also been established for neural networks, see e.g. (Yarotsky, 2017; Achour et al., 2022; Siegel, 2022), including upper and lower bounds on the required number of weights. In the context of operator learning, the underlying input and output spaces are infinite-dimensional. Given the unfavorable scaling of (16) with the dimension d , it is thus far from obvious that operator learning should be practicable at all.

In analogy with the above, we next introduce the notion of an **algebraic convergence rate**: Given a set $\mathcal{C} \subset L_\mu^2(\mathcal{X}; \mathcal{Y})$ of operators, we say that $\gamma_{\mathcal{C}} > 0$ is an algebraic convergence rate for \mathcal{C} , if for any $\Psi^\dagger \in \mathcal{C}$, there exist a constant $C = C(\Psi^\dagger) > 0$, depending only on properties of the operator Ψ^\dagger , and a PCA-Net mapping $\Psi := \mathcal{D}_{\mathcal{Y}} \circ \psi \circ \mathcal{E}_{\mathcal{X}} : \mathcal{X} \rightarrow \mathcal{Y}$, such that

$$\mathbb{E}_{u \sim \mu} \left[\|\Psi^\dagger(u) - \Psi(u)\|_{\mathcal{Y}}^2 \right] \leq C \text{size}(\psi)^{-\gamma_{\mathcal{C}}}, \quad (17)$$

The central point is that the convergence rate $\gamma_{\mathcal{C}}$ in (17) should be uniform over \mathcal{C} , and that only the multiplicative constant C depends on Ψ^\dagger , thus allowing the approximation of any operator $\Psi^\dagger \in \mathcal{C}$ to a desired accuracy, *at a rate specified by $\gamma_{\mathcal{C}}$* . This is a natural analogue of (16). In (17), the linear encoder $\mathcal{E}_{\mathcal{X}} : \mathcal{X} \rightarrow \mathbb{R}^{d_{\mathcal{X}}}$ and decoder $\mathcal{D}_{\mathcal{Y}} : \mathbb{R}^{d_{\mathcal{Y}}} \rightarrow \mathcal{Y}$ are allowed to be arbitrary; In particular, we do not restrict the dimensions $d_{\mathcal{X}}$, $d_{\mathcal{Y}}$. Even though the additional dependence on $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$ is also of practical relevance, we postulate that any operator $\Psi^\dagger \in L_\mu^2(\mathcal{X}; \mathcal{Y})$, or indeed class of operators $\mathcal{C} \subset L_\mu^2(\mathcal{X}; \mathcal{Y})$, which is “efficiently” approximated by PCA-Net, must (at the very least) possess a finite algebraic convergence rate $\gamma_{\mathcal{C}}$ in the above sense (17).

The next result shows that for often considered classes of operators $\Psi^\dagger : \mathcal{X} \rightarrow \mathcal{Y}$, such as the set of all Lipschitz continuous or r -times Fréchet differentiable operators, efficient operator learning by PCA-Net is in fact impossible. In analogy with the finite-dimensional curse of dimensionality, we refer to this as the “curse of parametric complexity”:

Theorem 9 (Curse of parametric complexity) *Let \mathcal{X}, \mathcal{Y} be separable Hilbert spaces, with $\dim(\mathcal{X}) = \infty$, $\dim(\mathcal{Y}) \geq 1$. Let $\mu \in \mathcal{P}(\mathcal{X})$ be a non-degenerate Gaussian measure. Fix*

$k \in \mathbb{N}$ and let \mathcal{C}^k denote the set of all k -times Fréchet differentiable operators $\Psi^\dagger : \mathcal{X} \rightarrow \mathcal{Y}$ whose k -th total derivative is uniformly bounded on \mathcal{X} . For any $\gamma > 0$, there exists $\Psi_\gamma^\dagger \in \mathcal{C}^k$ and a constant $c_\gamma > 0$, such that

$$\mathbb{E}_{u \sim \mu} \left[\|\Psi_\gamma^\dagger(u) - \Psi(u)\|_{\mathcal{Y}}^2 \right] \geq c_\gamma \text{size}(\psi)^{-\gamma},$$

for any PCA-Net $\Psi = \mathcal{D}_{\mathcal{Y}} \circ \psi \circ \mathcal{E}_{\mathcal{X}}$. In particular, there cannot exist a finite algebraic convergence rate $\gamma_{\mathcal{C}}$ for \mathcal{C}^k , as in (17).

Remark 10 *Theorem 9 makes precise a notion of “curse of dimensionality” in operator learning with PCA-Net, which was anticipated in earlier work by Kovachki et al. (2021); Lanthaler et al. (2022). As will be explained, this result can indeed be interpreted as a scaling limit of the conventional curse of dimensionality. In contrast to the approximation of a high-dimensional function in $d \gg 1$ dimensions, the goal of operator learning is to approximate a mapping defined on an infinite-dimensional space, where $d = \infty$. The “curse of parametric complexity” can be viewed as a scaling limit of the finite-dimensional curse of dimensionality, when $d \rightarrow \infty$.*

Remark 11 *The underlying input/output spaces \mathcal{X}, \mathcal{Y} are fixed in our setting. In particular, if \mathcal{X} is a function space e.g. consisting of functions $u : D \rightarrow \mathbb{R}$ for some domain $D \subset \mathbb{R}^n$, then the dimension n of this domain D is fixed. The curse of parametric complexity identified in Theorem 9 is independent of the presence (or absence) of an additional curse of dimensionality that may arise due to the (potentially large) dimension n of the domain.*

Remark 12 *To simplify the statement of Theorem 9, it is assumed that μ is a non-degenerate Gaussian. Inspection of the proof shows that the non-degeneracy requirement could be relaxed to only assuming that the covariance operator $\Sigma = \mathbb{E}_{u \sim \mu}[u \otimes u]$ have infinitely many non-trivial eigendirections (with eigenvalue $\lambda_j \neq 0$), thus allowing for some degeneracy. The crucial assumption is that the input distribution μ be “truly infinite-dimensional”.*

Proof [Sketch of proof] The proof of Theorem 9 is based on the following lower bound for ReLU neural network approximation of functions in the unit cube of the Sobolev space $W^{k,\infty}([0,1]^d)$,

$$F_{k,d} := \left\{ f \in W^{k,\infty}([0,1]^d) \mid \|f\|_{W^{k,\infty}} \leq 1 \right\},$$

which is derived in the present work and may be of independent interest:

Proposition 13 *Fix $k, d \in \mathbb{N}$. There exists $c_{k,d} > 0$ depending only on k and d , $f \in F_{k,d}$ and an absolute constant $\lambda > 0$ independent of both d and k , such that for any neural network ψ , we have the lower bound*

$$\|f - \psi\|_{L^2([0,1]^d)} \geq c_{k,d} \text{size}(\psi)^{-\lambda k/d}. \quad (18)$$

This finite-dimensional lower bound (18) builds on the recent work (Achour et al., 2022). Given the lower bound of Proposition 13, the proof of Theorem 9 is then based on the

intuition that d can be chosen arbitrarily large if the underlying input space \mathcal{X} is infinite-dimensional, and hence the exponent in (18) can be arbitrarily small; additional work is needed to make this intuition rigorous. Detailed proofs of Theorem 9 and Proposition 13 are given in Appendix C. \blacksquare

In particular, Theorem 9 shows that it is *impossible* to derive algebraic error and complexity estimates for PCA-Net, when e.g. assuming only Lipschitz regularity of Ψ^\dagger . It should be emphasized that this result holds even when the relevant space of input and output functions can be efficiently approximated; indeed, no restriction on the decay of the PCA eigenvalues is assumed in Theorem 9, allowing them to decay at an arbitrarily fast rate on both \mathcal{X} and \mathcal{Y} . In fact, Theorem 9 even allows for $\mathcal{Y} = \mathbb{R}$, in which case reconstruction on \mathcal{Y} is *trivial* (in contrast, the assumptions do imply that infinitely many PCA eigenvalues on \mathcal{X} are non-zero). As alluded to above, the reason for the obstacle to efficient operator learning expressed by Theorem 9 is the intrinsic complexity of the space of all \mathcal{C}^k -regular operators (or functionals), defined on an infinite-dimensional input space \mathcal{X} .

Remark 14 *Under slightly stronger assumptions than Theorem 9 (in particular, assuming an algebraic decay of the PCA eigenvalues λ_j), one can likely show that there in fact exists $\Psi^\dagger \in \mathcal{C}^k$ and constants $c, \gamma > 0$, depending only on μ and k , such that*

$$\mathbb{E}_{u \sim \mu} \left[\|\Psi^\dagger(u) - \Psi(u)\|_{\mathcal{Y}}^2 \right] \geq c \log(\text{size}(\psi))^{-\gamma}.$$

for any PCA-Net $\Psi = \mathcal{D}_{\mathcal{Y}} \circ \psi \circ \mathcal{E}_{\mathcal{X}}$. Thus, achieving accuracy ϵ requires an exponential complexity of the underlying neural network, $\text{size}(\psi) \gtrsim \exp(c\epsilon^{-1/\gamma})$. Similar exponential lower bounds will be the subject of forthcoming work (Lanthaler and Stuart, 2023).

Given the negative result of Theorem 9, we posit that a central challenge in operator learning is to *identify* the relevant class of operators $\mathcal{C} \subset L_{\mu}^2(\mathcal{X}; \mathcal{Y})$, which allow for efficient approximation by a given operator learning framework, and which possess a prescribed (finite) algebraic convergence rate $\gamma_{\mathcal{C}}$ in (17).

3.3 Quantitative encoding error bounds

As pointed out above, one potential obstacle to the efficacy of PCA-Nets is a slow decay of the PCA eigenvalues on the output function space. In this section, we provide a general estimate on the PCA projection error, $\mathcal{R}_d^{\text{opt}}(\nu) = \sum_{\ell > d} \lambda_{\ell}$, based on smoothness properties of the underlying functions. Given a domain $D \subset \mathbb{R}^n$, we recall that the Sobolev space $H^s(D; \mathbb{R}^n)$ is defined as the space of all functions $u : D \rightarrow \mathbb{R}^n$ possessing L^2 -integral weak derivatives up to order s . We note that $H^s(D; \mathbb{R}^n)$ is a Hilbert space. We then have:

Proposition 15 *Let n, n' be integers. Let $\mathcal{Y} = H^s(D; \mathbb{R}^{n'})$, $s \geq 0$, be a Sobolev space defined on either a Lipschitz domain $D \subset \mathbb{R}^n$ or the periodic torus $D = \mathbb{T}^n$. Assume that $\nu \in \mathcal{P}(\mathcal{Y})$ is a probability measure and that there exists $\zeta > 0$, such that $\mathbb{E}_{u \sim \nu} [\|u\|_{H^{s+\zeta}}^2] < \infty$. Then, there exists a constant $C = C(n, n') > 0$, such that*

$$\mathcal{R}_d^{\text{opt}}(\nu) \leq Cd^{-2\zeta/n} \mathbb{E}_{u \sim \nu} [\|u\|_{H^{s+\zeta}}^2], \quad \forall d \in \mathbb{N}.$$

We recall that $\mathcal{R}_d^{\text{opt}}(\nu)$ is the minimal projection error over all projections P of rank d (cp. 3).

Thus, smoothness of the output functions ensures an algebraic decay of the PCA eigenvalues and thus rules out the first potential obstacle to efficient operator learning by PCA-Net, pointed out after Proposition 8. A proof of Proposition 15 is provided in Appendix D. Proposition 15 extends a result in (Lanthaler et al., 2022, Prop. 3.14), which was restricted to the periodic case. The main novel ingredient in our proof here, and which allows the generalization to arbitrary Lipschitz domains D , is based on a general Sobolev extension result by Stein (Kovachki et al., 2021, e.g. Appendix B).

3.4 Overcoming the curse of dimensionality

The last section provides a general criterion to tame the lower bound of Proposition 8, which could limit the efficiency of PCA-Net when the relevant distribution of outputs in \mathcal{Y} is very complex (cf. Section 3.2). As proved in Theorem 9, a second obstacle to efficient operator learning is the *curse of dimensionality*. It would be very desirable to find a useful mathematical characterization of the entire class of operators $\mathcal{C} \subset L^2_\mu(\mathcal{X}; \mathcal{Y})$ for which an algebraic convergence rate, as in (17), can be established. At present, this appears a very distant goal. Instead, in this section, we aim to develop additional intuition of the basic mechanisms that PCA-Net can exploit to achieve algebraic convergence rates for specific examples. To this end, we consider two prototypical operators arising in the context of PDEs; the Darcy flow and Navier-Stokes equations.

3.4.1 DARCY FLOW

Let $D \subset \mathbb{R}^n$ be a bounded domain. We consider the operator $\Psi^\dagger : a \mapsto w$, mapping the coefficient field a to the solution w of the following elliptic problem:

$$\begin{cases} -\nabla \cdot (a(x)\nabla w(x)) = f(x), & (x \in D), \\ w(x) = 0, & (x \in \partial D). \end{cases} \quad (19)$$

Here, $D \subset \mathbb{R}^n$ is a given smooth domain, the right-hand side f is fixed, and we assume Dirichlet boundary conditions. Let $H_0^1(D)$ be the Sobolev space consisting of weakly differentiable functions $w : D \rightarrow \mathbb{R}$ which vanish on the boundary ∂D , and whose gradient is square-integrable. It is well-known, e.g. (Evans, 2010, Chapt. 6), that if the coefficient field $a \in L^\infty(D)$, satisfies two-sided (coercivity and upper) bounds

$$0 < \lambda \leq a(x) \leq \Lambda < \infty, \quad \forall x \in D, \quad (20)$$

for constant λ, Λ , and if $f \in H^{-1}(D)$ belongs to the dual space of $H_0^1(D)$, then there exists a unique solution $w \in H_0^1(D)$ to (19). Furthermore, there exists a constant $C = C(\lambda, \Lambda, D) > 0$, such that

$$\|w\|_{H_0^1} \leq C\|f\|_{H^{-1}}. \quad (21)$$

It is thus natural to consider the output space $\mathcal{Y} = H_0^1(D)$. We will follow the celebrated work by Cohen, Devore and Schwab (Cohen et al., 2010), and subsequent extensions in (Cohen et al., 2011; Schwab and Zech, 2019; Opschoor et al., 2022, 2020), and consider the

following setting: We assume that the underlying measure μ can be written as the law of random coefficient fields $a(x) = a(x; \mathbf{z})$, of the parametrized form

$$a(x; \mathbf{z}) = \bar{a}(x) + \sum_{\ell=1}^{\infty} \gamma_{\ell} z_{\ell} \rho_{\ell}(x), \quad (22)$$

where $\mathbf{z} = (z_1, z_2, \dots)$ is a sequence of random variables (not necessarily independent), such that $|z_{\ell}| \leq 1$. In our analysis, we will assume that the functions $\rho_1, \rho_2, \dots \in \mathcal{X}$ are orthonormal in a Hilbert space \mathcal{X} , where $\mathcal{X} \hookrightarrow L^{\infty}$ possesses a continuous embedding. Note that the series (22) converges, if

$$0 \leq \gamma_{\ell} \leq M \ell^{-1-\alpha}, \quad (23)$$

decays at an algebraic rate with constants $M, \alpha > 0$. We will assume this bound (23). To ensure coercivity (cp. Remark 16 below) we assume that there exists $\kappa > 0$, such that

$$\sum_{\ell=1}^{\infty} \gamma_{\ell} \|\rho_{\ell}\|_{L^{\infty}} \leq \frac{\kappa}{1 + \kappa} \bar{a}_{\min}, \quad (24)$$

with $\bar{a}_{\min} := \operatorname{ess\,inf}_{x \in D} \bar{a}(x) > 0$.

Remark 16 *The upper bound (24) ensures uniform coercivity, since*

$$a(x; \mathbf{z}) \geq \bar{a}_{\min} - \sum_{\ell=1}^{\infty} \gamma_{\ell} \|\rho_{\ell}\|_{L^{\infty}} \geq \frac{1}{1 + \kappa} \bar{a}_{\min} =: \lambda > 0,$$

for any $\mathbf{z} \in U := [-1, 1]^{\mathbb{N}}$. On the other hand, the assumed embedding $\mathcal{X} \hookrightarrow L^{\infty}(D)$ implies that there exists a constant $C > 0$, such that $\|\cdot\|_{L^{\infty}} \leq C \|\cdot\|_{\mathcal{X}}$, and hence (23) ensures a uniform upper bound,

$$a(x; \mathbf{z}) \leq \|\bar{a}\|_{L^{\infty}} + CM \sum_{\ell=1}^{\infty} \ell^{-1-\alpha} =: \Lambda < \infty.$$

In particular, all a in the support of the probability measure μ satisfy the two-sided bounds (20), and hence the elliptic PDE (19) is well-posed.

Remark 17 *We do not assume any (explicit) knowledge of the functions \bar{a} , ρ_{ℓ} , the parameters γ_{ℓ} , α , λ , Λ or indeed any information on the law of the joint random variable $\mathbf{z} = (z_1, z_2, \dots) \in [-1, 1]^{\mathbb{N}}$. In particular, the random variables z_1, z_2, \dots need not be independent. For the following arguments, it is sufficient that an expansion of the form (22) exists.*

Given this setting, we prove the following theorem:

Theorem 18 *Under the setting and the prevailing assumptions of this section and with $\mu \in \mathcal{P}(\mathcal{X})$ the law of $a(\cdot; \mathbf{z})$ given by (22). For any $\delta, \eta > 0$ and $\epsilon > 0$, there exists a PCA-Net $\Psi = \mathcal{D}_{\mathcal{Y}} \circ \psi \circ \mathcal{E}_{\mathcal{X}}$ satisfying the error bound,*

$$\mathbb{E}_{a \sim \mu} \left[\|\Psi^{\dagger}(a) - \Psi(a)\|_{H_0^1}^2 \right] \leq C\epsilon,$$

with probability at least $1 - \delta$, and with constant $C = C(\mu, \eta) > 0$ depending only on μ and η . With the same implied constant, the required PCA dimensions are at most $d_X = d_Y = d \sim \epsilon^{-\frac{1}{2\alpha} - \eta}$ and the required number of samples for PCA is at most $N \sim d^{1+4\alpha} \log(1/\delta)$. Furthermore, the following complexity bounds hold for the ReLU network ψ ,

$$\text{size}(\psi) \leq C\epsilon^{-\frac{1}{\alpha} - \eta}, \quad \text{depth}(\psi) \leq C \log(\epsilon^{-1})^2,$$

independently of the data $a_1, \dots, a_N \sim \mu$.

We provide a sketch of the proof of Theorem 18 at the end of this subsection.

Remark 19 Regarding the fudge factor $\eta > 0$, we note that in our estimates leading to Theorem 18, the constant $C = C(\mu, \eta)$ blows up as $\eta \rightarrow 0$, i.e. $C(\mu, \eta) \rightarrow \infty$. Therefore, while $\eta > 0$ can be chosen arbitrarily small, our estimates break down when $\eta = 0$.

Remark 20 Theorem 18 shows that approximation of the Darcy flow operator Ψ^\dagger is possible with algebraic bounds on the PCA dimensions d_X, d_Y , the number of required PCA samples N , and the number of neural network parameters, $\text{size}(\psi)$. In fact, even under a mild decay rate $\alpha > 1$, the required size of ψ is at most linear in the desired accuracy.

We conjecture that the scaling of $N \sim d^{1+4\alpha}$ in Theorem 18 is highly pessimistic; indeed, under a potential improvement of the empirical PCA estimate of Proposition 2, as discussed in Remark 3, the much more favorable scaling $N \sim d$ appears natural.

Sketch of proof of Theorem 18. To prove Theorem 18, we first define a parametric mapping $\mathcal{F} : U \rightarrow \mathcal{Y}$, where $U = [-1, 1]^{\mathbb{N}}$, $\mathcal{Y} = H_0^1(D)$, by

$$\mathcal{F}(\mathbf{z}) := \Psi^\dagger(\bar{a} + \sum_{\ell=1}^{\infty} \gamma_\ell z_\ell \rho_\ell). \quad (25)$$

This mapping has been studied in a series of papers (Cohen et al., 2010, 2011; Schwab and Zech, 2019; Opschoor et al., 2022, 2020), and is known to allow for a convergent Taylor series expansion (in the variables $\mathbf{z} \in U$). To state the next Lemma, which implies a suitable convergence of Taylor series, we recall that $\boldsymbol{\nu}$ is called a multi-index in this infinite-dimensional context, if $\boldsymbol{\nu} = (\nu_1, \nu_2, \dots)$ is a sequence of non-negative integers, such that $\nu_\ell = 0$ for almost all ℓ . For $\mathbf{z} \in U$ and a multi-index $\boldsymbol{\nu}$, we define the monomial $\mathbf{z}^\boldsymbol{\nu} = \prod_{\nu_\ell \neq 0} z_\ell^{\nu_\ell}$. The following lemma follows immediately from (Cohen et al., 2011, Thm. 1.3).

Lemma 21 Assume the setting and prevailing assumptions of this section. Then there exists a set of coefficients $t_\boldsymbol{\nu} \in \mathcal{Y}$ (the Taylor coefficients), indexed by multi-indices $\boldsymbol{\nu}$, such that for any $m \in \mathbb{N}$, there is a set Λ_m of multi-indices $\boldsymbol{\nu}$, with cardinality $|\Lambda_m| = m$, such that

$$\sup_{\mathbf{z} \in U} \|\mathcal{F}(\mathbf{z}) - \sum_{\boldsymbol{\nu} \in \Lambda_m} t_\boldsymbol{\nu} \mathbf{z}^\boldsymbol{\nu}\|_{\mathcal{Y}} \leq C m^{-\alpha + \eta}, \quad (26)$$

for any fixed constant $\eta > 0$. Here $C = C(\mathcal{F}, \mu, \eta) > 0$ is a constant depending on η , but is independent of m .

For completeness, we provide the details in Appendix E.1. As a consequence of this lemma, we can estimate the optimal PCA encoding errors not only on \mathcal{X} but also on \mathcal{Y} . Indeed, we will derive the following result.

Proposition 22 *Under the setting and prevailing assumptions of this section. For any $\eta > 0$, there exists a constant $C = C(\mathcal{F}, \mu, \eta) > 0$, such that*

$$\mathcal{R}_{d_{\mathcal{X}}}^{\text{opt}}(\mu) \leq C d_{\mathcal{X}}^{-2\alpha-1}, \quad \mathcal{R}_{d_{\mathcal{Y}}}^{\text{opt}}(\Psi_{\#}^{\dagger} \mu) \leq C d_{\mathcal{Y}}^{-2\alpha+\eta}.$$

The \mathcal{X} estimate of Proposition 22 is a straightforward consequence of the assumed expansion (22). The \mathcal{Y} estimate follows from the observation that (26) provides a bound of the PCA projection error by comparing it to the projection onto $\text{span}\{t_{\nu} \mid \nu \in \Lambda_{d_{\mathcal{Y}}}\} \subset \mathcal{Y}$. Details of the required argument are given in Appendix E.2. Since the input measure μ and its push-forward $\Psi_{\#}^{\dagger} \mu$ are concentrated on a bounded set of functions in the respective norms on \mathcal{X} and \mathcal{Y} , Proposition 2 immediately implies that empirical PCA with $d_{\mathcal{X}} = d_{\mathcal{Y}} = d$ and with a sufficient number of $N \gtrsim d^{1+4\alpha} \log(\delta/2)$ samples achieves, up to a constant and with high probability, the same asymptotic error as the optimal PCA projection on \mathcal{X} and \mathcal{Y} .

The main remaining challenge is then to construct a neural network ψ , such that the composition $\Psi = \mathcal{D}_{\mathcal{Y}} \circ \psi \circ \mathcal{E}_{\mathcal{X}}$ approximates Ψ^{\dagger} to within a prescribed tolerance. The construction of such ψ relies on a neural network approximation result for the parametric mapping $\mathcal{F}(\mathbf{z})$. The following result follows from (Opschoor et al., 2020, Theorem 4.11) (the present statement is closer in formulation to (Schwab and Zech, 2019, Theorem 3.9), and has appeared with slightly sharper bounds in (Lanthaler et al., 2022)):

Lemma 23 *Under the prevailing assumption and with \mathcal{F} defined by (25). Let $\eta > 0$ be a small (fudge) constant. There exists a constant $C = C(\mathcal{F}, \mu, \eta) > 0$, depending only on \mathcal{F} , on the decay rate α and on η , such that for any PCA encoder $\mathcal{E}_{\mathcal{Y}} : \mathcal{Y} \rightarrow \mathbb{R}^{d_{\mathcal{Y}}}$, and for every $m \in \mathbb{N}$, there exists a ReLU network $\psi^* : \mathbb{R}^m \rightarrow \mathbb{R}^{d_{\mathcal{Y}}}$, $z \mapsto \psi^*(z_1, \dots, z_m)$ with*

$$\sup_{z \in [-1, 1]^m} \|\mathcal{E}_{\mathcal{Y}} \circ \mathcal{F}(\mathbf{z}) - \psi^*(z_1, \dots, z_m)\|_{\ell^2} \leq C m^{-\alpha+\eta}, \quad (27)$$

and such that $\text{size}(\psi^*) \leq C m (\log(m)^2 + d_{\mathcal{Y}})$ and $\text{depth}(\psi^*) \leq C \log(m)^2$.

The main additional problem in the present context of PCA-Net is that the neural network ψ in the definition of a PCA-Net, $\Psi = \mathcal{D}_{\mathcal{Y}} \circ \psi \circ \mathcal{E}_{\mathcal{X}}$, acts on the *encoded input* $\mathcal{E}_{\mathcal{X}}(a)$, with encoder $\mathcal{E}_{\mathcal{X}}$ obtained from empirical PCA. Therefore, we cannot directly access the coefficients z_1, z_2, \dots in the parametric expansion (22) of a , and indeed there is no way to exactly recover these coefficients from the PCA encoding of a . A careful discussion of this ‘‘compatibility issue’’ between the PCA encoding and the a priori expansion (22) is therefore necessary. This is the main issue addressed in Appendix E.3, where we show that such ψ can indeed be constructed and the additional error due to the incompatibility can be controlled. This then leads to the statement of Theorem 18.

3.4.2 NAVIER-STOKES EQUATIONS

In the previous section, we showed that the solution operator Ψ^{\dagger} of the Darcy flow equations can be efficiently approximated by PCA-Net. With some additional effort, this result

could likely be extended to a more general class of so-called (\mathbf{b}, ϵ) -holomorphic operators, e.g. (Schwab and Zech, 2019, Section 2.1). In particular, the relevant underlying class of operators is here characterized by *analytic regularity*; this assumption goes well beyond \mathcal{C}^k -regularity, and thereby, PCA-Net can overcome the general curse of dimensionality of Theorem 9, for this restricted class of operators.

However, many operators of interest, in particular in the context of advection dominated problems such as hyperbolic PDEs, are not holomorphic in this sense. In the present section, we therefore discuss another mechanism by which polynomial complexity estimates can be obtained, even in the absence of holomorphy: namely, PCA-Net can efficiently emulate numerical methods. Following ideas developed in (Lanthaler et al., 2022; Kovachki et al., 2021), and starting from a known (and convergent) numerical method, such an emulation result provides an upper bound on the required complexity of Ψ , by showing that a specific choice of the neural network weights can emulate the given numerical method. The derivation of explicit estimates requires us to fix a particular numerical method for the analysis, but the intuition behind these emulation results is that Ψ can, in principle, efficiently emulate a very rich class of numerical methods. This includes methods with high convergence rates, which the neural network can explore during optimization; this expressive power of neural networks can thus provide a theoretical rationale for their efficiency within the PCA-Net methodology.

In the following, we will focus on an emulation result for the Navier-Stokes equations, based on spectral methods. The underlying idea is similar to a recent emulation result for Fourier neural operators (Kovachki et al., 2021); however, while Fourier neural operators very naturally (and by design) provide the necessary ingredients to build a spectral method, the PCA-Net methodology requires an extension of the results of (Kovachki et al., 2021), including a detailed discussion of compatibility of the PCA-projection with such emulation results. In addition, at a more technical level, it is here shown that the smoothness assumption on the activation function, which was *essential* in all proofs of (Kovachki et al., 2021), can be substantially relaxed. Indeed, the results of the present work are based on the popular ReLU activation function $\sigma(x) = \max(x, 0)$, leading to comparable complexity estimates as in (Kovachki et al., 2021) differing only by log-factors.

To illustrate the general approach, we consider the periodic Navier-Stokes equations in spatial dimension $n = 2$, over a fixed time interval $[0, T]$:

$$\begin{cases} \partial_t u + u \cdot \nabla u + \nabla p = \nu \Delta u, \\ \operatorname{div}(u) = 0, u(t=0) = \bar{u}. \end{cases} \quad (28)$$

Here $u : \mathbb{T}^n \times [0, T] \rightarrow \mathbb{R}^n$ is the flow vector field, $p : \mathbb{T}^n \times [0, T] \rightarrow \mathbb{R}$ is the scalar pressure and $\nu \geq 0$ is the viscosity (we allow $\nu = 0$, corresponding to the incompressible Euler equations). We have denoted by $\bar{u} : \mathbb{T}^n \rightarrow \mathbb{R}^n$ the (divergence-free) initial data. Since the solution operator Ψ^\dagger associated with (28) is only known to be well-defined in two spatial dimensions, we focus on this case. We however point out that all results readily extend to the three-dimensional case, under additional (unproven) smoothness assumptions.

Remark 24 *Classical well-posedness results for the Navier-Stokes ($\nu > 0$) and Euler ($\nu = 0$) equations, e.g. (Majda and Bertozzi, 2001) and references therein, imply in the two-dimensional case $n = 2$, that if the random initial data $\bar{u} \sim \mu$, is uniformly bounded $\|\bar{u}\|_{H^r} \leq$*

\bar{M} for $r > n/2 + 1$, then the corresponding solution $u(t)$ at a fixed later time $t \in [0, T]$ satisfies a similar bound $\|u(t)\|_{H^r} \leq M$, for some $M = M(T, \bar{M})$.

The main result of the present section is Theorem 25, below:

Theorem 25 *Consider the two-dimensional, periodic Navier-Stokes equations. Fix parameters $M, T > 0$, and integer $r > n/2 + 1$. Assume that $\mu \in \mathcal{P}(L^2(\mathbb{T}^2; \mathbb{R}^2))$ is a probability measure on initial data \bar{u} of (28), such that $\|\bar{u}\|_{H^r} \leq M$ μ -almost surely. Let Ψ^\dagger be the forward solution operator of the Navier-Stokes equations, mapping initial data to the solution at the final time T , $\Psi^\dagger : u(0) \mapsto u(T)$. For any $\epsilon, \delta > 0$, there exists a PCA-Net $\Psi = \mathcal{D}_y \circ \psi \circ \mathcal{E}_x$, such that*

$$\mathbb{E}_{u \sim \mu} \left[\|\Psi^\dagger(u) - \Psi(u)\|_{L_x^2}^2 \right] \leq C\epsilon,$$

with probability at least $1 - \delta$, and with a constant $C = C(M, r, T) > 0$. The PCA dimensions d_x and d_y are bounded by $d_x, d_y \leq C\epsilon^{-1/r}$, the requisite amount of data $N \leq Cd_x^{1+2r} \log(1/\delta)$ and the neural network ψ satisfies the complexity bounds

$$\begin{cases} \text{size}(\psi) \leq C\epsilon^{-1/r} \left(\epsilon^{-1/2} \log(\epsilon^{-1})^2 + \epsilon^{-1/r} \right), \\ \text{depth}(\psi) \leq C\epsilon^{-1/2} \log(\epsilon^{-1})^2. \end{cases}$$

A sketch of the proof of Theorem 25 is provided below.

Remark 26 *In fact, ψ in Theorem 25 can be written as an n_T -fold composition*

$$\psi = Q \circ \underbrace{\psi_* \circ \dots \circ \psi_*}_{n_T\text{-fold}} \circ R,$$

where the mappings $R : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_H}$, $Q : \mathbb{R}^{d_H} \rightarrow \mathbb{R}^{d_y}$ are linear (input and output layers), $n_T \leq C\epsilon^{-1/2}$, and $\psi_* : \mathbb{R}^{d_H} \rightarrow \mathbb{R}^{d_H}$ is a ReLU neural network with “hidden layer” dimension $d_H \leq C\epsilon^{-1/r}$, such that

$$\begin{cases} \text{size}(\psi_*) \leq C\epsilon^{-1/r} \log(\epsilon^{-1})^2, \\ \text{depth}(\psi_*) \leq C \log(\epsilon^{-1})^2. \end{cases}$$

Sketch of proof of Theorem 25. The derivation of the quantitative error and complexity bounds of Theorem 25 is based on an *emulation result*; the idea is to show that for any choice of PCA bases, there exists a ReLU neural network ψ which can efficiently emulate a numerical method which is known to converge at a precisely quantifiable rate. The complete details of the required argument, including all proofs, will be provided in Appendix F. Here, we instead give a general overview of the main ideas, to aid intuition.

As a first step towards this emulation result, we review a convergent spectral scheme in section F.1. Then, we construct a ReLU neural network emulation of this spectral scheme in section F.2, leading to Algorithm 1 and the neural network size estimates of Lemma 41. The constructed neural network emulator defines a mapping from the truncated Fourier coefficients of the initial data $u(0)$, to (an approximation of) the truncated Fourier

coefficients of the solution $u(T)$. The relevant set of truncated Fourier modes with cut-off parameter $K \in \mathbb{N}$ is given by

$$\mathcal{K} = \mathcal{K}_K := \{k = (k_1, k_2) \in \mathbb{Z}^2 \mid |k|_\infty := \max(|k_1|, |k_2|) \leq K\}.$$

The mapping on these truncated Fourier coefficients defines a mapping between two-finite dimensional Euclidean spaces, upon identifying $\mathbb{C}^{\mathcal{K}} \simeq \mathbb{R}^{2\mathcal{K}}$. This mapping can be represented by an ordinary neural network. Given this construction, we then proceed to analyze the approximation error of the neural network emulation in section F.3, leading to the following proposition. This is our core emulation result, essentially stating the fact that ReLU neural networks can indeed efficiently emulate the underlying spectral scheme.

Proposition 27 *Let $M, r, T > 0$ be given. For any $\epsilon > 0$, there exists $K \in \mathbb{N}$, $K \sim \epsilon^{-1/r}$, and a ReLU neural network $\widehat{\psi} : \mathbb{C}^{\mathcal{K}} \rightarrow \mathbb{C}^{\mathcal{K}}$, $\mathcal{K} = \mathcal{K}_K$, such that*

$$\|\widehat{\psi}(\widehat{u}(0)) - \widehat{u}(T)\|_{\ell^2} \leq \epsilon, \quad \text{whenever } \|u(0)\|_{H^r} \leq M,$$

where $\widehat{u}(t) = \{\widehat{u}_k(t)\}_{|k|_\infty \leq K}$ denote the Fourier coefficients of the solution $u(t)$ of (28), with initial data $u(0)$. Furthermore, we have the following complexity estimates:

$$\begin{cases} \text{size}(\widehat{\psi}) \leq C\epsilon^{-2/r-1} \log(\epsilon^{-1})^2, \\ \text{depth}(\widehat{\psi}) \leq C\epsilon^{-1} \log(\epsilon^{-1})^2. \end{cases}$$

The constant $C = C(M, T, r) > 0$ depends only M, T, r , but is independent of ϵ . Furthermore, $\widehat{\psi}$ can be written as a n_T -fold composition $\widehat{\psi} = \widehat{\psi}_* \circ \dots \circ \widehat{\psi}_*$, where $\widehat{\psi}_* : \mathbb{C}^{\mathcal{K}} \rightarrow \mathbb{C}^{\mathcal{K}}$ is a ReLU neural network with

$$\begin{cases} \text{size}(\widehat{\psi}_*) \leq C\epsilon^{-2/r} \log(\epsilon^{-1})^2, \\ \text{depth}(\widehat{\psi}_*) \leq C \log(\epsilon^{-1})^2, \end{cases}$$

and $n_T \leq C\epsilon^{-1}$, corresponding to the number of time-steps of the underlying scheme.

In fact, Proposition 27 is the two-dimensional case of a general d -dimensional result derived in the appendix (cf. Proposition 45). Given this neural network emulation result, the remaining issue is that the empirical PCA encoder $\mathcal{E}_{\mathcal{X}}$ and decoder $\mathcal{D}_{\mathcal{Y}}$ do *not* act on Fourier coefficients. And hence, additional work is necessary to suitably adapt the construction of the neural network in Proposition 27, ultimately resulting in an efficient PCA-Net approximation $\Psi = \mathcal{D}_{\mathcal{Y}} \circ \psi \circ \mathcal{E}_{\mathcal{X}}$ in Section F.4. We summarize the result in the following lemma:

Lemma 28 *For any $\epsilon > 0$, there exists a PCA-Net $\Psi = \mathcal{D}_{\mathcal{Y}} \circ \psi \circ \mathcal{E}_{\mathcal{X}}$, such that*

$$\begin{aligned} \mathbb{E}_{u \sim \mu} \left[\|\Psi(u) - \Psi^\dagger(u)\|_{L^2(\mu)}^2 \right]^{1/2} &\leq C\epsilon + C\mathbb{E}_{u \sim \mu} \left[\|u - \mathcal{D}_{\mathcal{X}} \circ \mathcal{E}_{\mathcal{X}}(u)\|_{L^2}^2 \right]^{1/2} \\ &\quad + \mathbb{E}_{v \sim \Psi^\dagger_{\#}\mu} \left[\|v - \mathcal{D}_{\mathcal{Y}} \circ \mathcal{E}_{\mathcal{Y}}(v)\|_{L^2}^2 \right]^{1/2}, \end{aligned}$$

where ψ is a neural network of size

$$\text{size}(\psi) \leq C\epsilon^{-2/r} (\epsilon^{-1} \log(\epsilon^{-1})^2 + (d_{\mathcal{X}} + d_{\mathcal{Y}})), \quad \text{depth}(\psi) \leq C\epsilon^{-1} \log(\epsilon^{-1}),$$

and $C = C(M, r, T) > 0$ is a constant independent of $\epsilon, d_{\mathcal{X}}, d_{\mathcal{Y}}$.

Finally, the smoothness bound on the PCA eigenvalues of Proposition 15 can be used to estimate the PCA projection errors with high probability, see Lemma 47 for details. Combining the above lemma with the estimates on the PCA-projection errors, results in Theorem 25.

4. Conclusion

PCA-Net is a data-driven operator learning methodology introduced in (Bhattacharya et al., 2021; Hesthaven and Ubbiali, 2018), which (i) uses PCA to reduce the dimensions of the input and output spaces and (ii) uses neural networks to approximate a map between the resulting finite-dimensional latent spaces. The main aim of the present work is to develop relevant approximation theory for PCA-Net.

Our first main result is a novel universal approximation theorem for PCA-Net, Theorem 6. Compared to previous work (Bhattacharya et al., 2021), this theorem establishes universality under significantly relaxed conditions on the distribution of the data-generating measure and the underlying operator Ψ^\dagger . The present assumptions are in fact minimal conditions to ensure that PCA is well-defined on the input and output spaces.

The next main contribution of the present work is a detailed discussion of two potential obstacles to efficient operator learning with PCA-Net in Section 3.2; the first obstacle relates to the complexity of the output distribution. The second obstacle relates to the inherent complexity of the space of operators between infinite-dimensional input and output spaces, and gives rigorous meaning to the notion of a curse of dimensionality; Theorem 9 shows that it is impossible to derive algebraic complexity bounds, when considering general classes of operators, such as the class of all Lipschitz- or even C^k -continuous operators. Hence, we conclude that at this level of generality, the curse of dimensionality is inevitable.

Given this negative result demonstrating the curse of dimensionality over general classes of operators, we posit that a central challenge in the approximation theory of PCA-Net (and other operator learning methodologies) is to identify and characterize operators, and classes of such operators, which allow for efficient approximation by the methodology. To obtain first insight into this problem for PCA-Net, we focus our attention on two prototypical PDE operators of interest, arising from the Darcy flow and Navier-Stokes equations. In both cases, we show that PCA-Net can overcome the general curse of dimensionality, establishing algebraic error and complexity estimates in Theorems 18 and 25. This demonstrates that these operators belong to a restricted class which is efficiently approximated by PCA-Net. In the case of Darcy flow, our proof relies on the analytic regularity (holomorphy) of the underlying operator. For the Navier-Stokes equations, we rely on an emulation result, showing that PCA-Net can emulate a known spectral scheme to efficiently approximate the underlying operator.

It is an open challenge for future work to improve our understanding of the relevant class of operators for which operator learning is feasible, and to derive a useful mathematical characterization of relevant features that enable efficient approximation by PCA-Net and other operator learning architectures. Future work could also aim to derive more precise lower bounds which characterize the curse of dimensionality. In this context, we mention the forthcoming article (Lanthaler and Stuart, 2023), where similar ideas are refined and considerably generalized, and *exponential* lower bounds are derived for a more general class

of neural operators, albeit with respect to the supremum norm; those results do not translate to any lower bounds in the L^2_μ -norm considered here. In a different research direction, we point out that the present work has focused only on an approximation theoretic point of view, leaving out important questions related to optimization and generalization errors, given a finite amount of data. A significant challenge for future work is to address the practical training of the underlying neural network, and in particular, to determine bounds on the amount of training data that is necessary to achieve a desired accuracy.

Finally, it could be of interest to link the results of the present work with recent work on overcoming the curse of dimensionality in the solution of high-dimensional PDEs, e.g. (Chow et al., 2017, 2019; Darbon and Osher, 2016; Darbon et al., 2020; Marwah et al., 2021; Chen et al., 2021; Marwah et al., 2023). In the context of these high-dimensional PDEs, the focus has usually been on approximating an individual solution rather than learning the underlying solution operator. It would be of interest to further explore the possibility of operator learning for such high-dimensional PDEs, with the aim of deriving complexity estimates that exhibit good scaling with respect to the infinite-dimensional input and output function spaces, and in addition, scale favorably with respect to the spatial dimension $n \gg 1$ of the underlying spatial domain on which the PDE is defined. We leave these general research directions as interesting avenues for future work.

Acknowledgments

The author would like to thank Siddhartha Mishra for helpful discussions and guidance when developing many of the ideas that have gone into this work. This work has been supported by Postdoc.Mobility grant P500PT-206737 from the Swiss National Science Foundation.

References

- E. M. Achour, A. Foucault, S. Gerchinovitz, and F. Malgouyres. A general approximation lower bound in l^p norm, with applications to feed-forward neural networks. *Advances in Neural Information Processing Systems*, 35:22396–22408, 2022.
- 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. URL <https://openreview.net/forum?id=fg2ZFmXF03>.
- K. Bhattacharya, B. Hosseini, N. B. Kovachki, and A. M. Stuart. Model Reduction And Neural Networks For Parametric PDEs. *The SMAI journal of computational mathematics*, 7:121–157, 2021. doi: 10.5802/smai-jcm.74. URL <https://smai-jcm.centre-mersenne.org/articles/10.5802/smai-jcm.74/>.
- N. Boullé and A. Townsend. Learning elliptic partial differential equations with randomized linear algebra. *Foundations of Computational Mathematics*, pages 1–31, 2022.

- 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. Publisher: Elsevier.
- A. Caponnetto and E. De Vito. Optimal rates for the regularized least-squares algorithm. *Foundations of Computational Mathematics*, 7(3):331–368, 2007.
- 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.
- Z. Chen, J. Lu, and Y. Lu. On the representation of solutions to elliptic pdes in barron spaces. *Advances in neural information processing systems*, 34:6454–6465, 2021.
- Y. T. Chow, J. Darbon, S. Osher, and W. Yin. Algorithm for overcoming the curse of dimensionality for time-dependent non-convex Hamilton–Jacobi equations arising from optimal control and differential games problems. *Journal of Scientific Computing*, 73(2):617–643, 2017.
- Y. T. Chow, J. Darbon, S. Osher, and W. Yin. Algorithm for overcoming the curse of dimensionality for state-dependent Hamilton–Jacobi equations. *Journal of Computational Physics*, 387:376–409, 2019.
- A. Cohen, R. DeVore, and C. Schwab. Convergence rates of best n-term galerkin approximations for a class of elliptic sPDEs. *Foundations of Computational Mathematics*, 10(6):615–646, 2010.
- A. Cohen, R. Devore, and C. Schwab. Analytic regularity and polynomial approximation of parametric and stochastic elliptic PDE’s. *Analysis and Applications*, 9(01):11–47, 2011.
- J. Darbon and S. Osher. Algorithms for overcoming the curse of dimensionality for certain Hamilton–Jacobi equations arising in control theory and elsewhere. *Research in the Mathematical Sciences*, 3(1):1–26, 2016.
- J. Darbon, G. P. Langlois, and T. Meng. Overcoming the curse of dimensionality for some Hamilton–Jacobi partial differential equations via neural network architectures. *Research in the Mathematical Sciences*, 7(3):1–50, 2020.
- M. De Hoop, D. Z. Huang, E. Qian, and A. M. Stuart. The cost-accuracy trade-off in operator learning with neural networks. *To appear, Journal of Machine Learning, arXiv:2203.13181*, 2022.
- M. V. de Hoop, N. B. Kovachki, N. H. Nelsen, and A. M. Stuart. Convergence rates for learning linear operators from noisy data. *SIAM/ASA Journal on Uncertainty Quantification*, 11(2):480–513, 2023.
- P. C. Di Leoni, L. Lu, C. Meneveau, G. Karniadakis, and T. A. Zaki. DeepONet prediction of linear instability waves in high-speed boundary layers. *arXiv preprint arXiv:2105.08697*, 2021. URL <https://arxiv.org/abs/2105.08697>.

- L. C. Evans. *Partial differential equations*, volume 19. American Mathematical Society, Providence, Rhode Island, 2nd edition, 2010.
- I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT press, 2016.
- G. Gupta, X. Xiao, and P. Bogdan. Multiwavelet-based operator learning for differential equations. *Advances in Neural Information Processing Systems*, 34:24048–24062, 2021.
- J. S. Hesthaven and S. Ubbiali. Non-intrusive reduced order modeling of nonlinear problems using neural networks. *Journal of Computational Physics*, 363:55–78, 2018. Publisher: Elsevier.
- P. Jin, S. Meng, and L. Lu. Mionet: learning multiple-input operators via tensor product. *SIAM Journal on Scientific Computing*, 44(6):A3490–A3514, 2022. doi: 10.1137/22M1477751. URL <https://doi.org/10.1137/22M1477751>.
- I. T. Jolliffe. *Principal component analysis*. Springer Series in Statistics. Springer New York, NY, 2002.
- 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. URL <http://jmlr.org/papers/v22/21-0806.html>.
- N. B. Kovachki, Z. Li, B. Liu, K. Azizzadenesheli, K. Bhattacharya, A. M. Stuart, and A. Anandkumar. Neural operator: Learning maps between function spaces with applications to pdes. *J. Mach. Learn. Res.*, 24(89):1–97, 2023.
- G. Kutyniok, P. Petersen, M. Raslan, and R. Schneider. A theoretical analysis of deep neural networks and parametric PDEs. *Constructive Approximation*, 55(1):73–125, Feb 2022. ISSN 1432-0940. doi: 10.1007/s00365-021-09551-4. URL <https://doi.org/10.1007/s00365-021-09551-4>.
- S. Lanthaler and A. M. Stuart. The curse of dimensionality in operator learning. *arXiv preprint arXiv:2306.15924*, 2023.
- S. Lanthaler, S. Mishra, and G. E. Karniadakis. Error estimates for DeepONets: a deep learning framework in infinite dimensions. *Transactions of Mathematics and Its Applications*, 6(1), 03 2022. ISSN 2398-4945. doi: 10.1093/imatrm/tnac001. URL <https://doi.org/10.1093/imatrm/tnac001>. tnac001.
- S. Lanthaler, Z. Li, and A. M. Stuart. The nonlocal neural operator: Universal approximation. (*in preparation*), 2023a.
- S. Lanthaler, R. Molinaro, P. Hadorn, and S. Mishra. Nonlinear reconstruction for operator learning of PDEs with discontinuities. In *11th International Conference on Learning Representations (ICLR)*. OpenReview.net, 2023b.
- Z. Li, N. B. Kovachki, K. Azizzadenesheli, B. Liu, A. M. Stuart, K. Bhattacharya, and A. Anandkumar. Multipole graph neural operator for parametric partial differential equations. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors,

- Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 6755–6766. Curran Associates, Inc., 2020.
- Z. Li, N. B. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, and A. Anandkumar. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=c8P9NQVtmn0>.
- L. Lu, P. Jin, G. Pang, Z. Zhang, and G. E. Karniadakis. Learning nonlinear operators via deepnet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3):218–229, 2021.
- A. J. Majda and A. L. Bertozzi. Vorticity and incompressible flow. *Vorticity and Incompressible Flow*, page 558, 2001.
- Z. Mao, L. Lu, O. Marxen, T. A. Zaki, and G. E. Karniadakis. DeepMandMnet for hypersonics: Predicting the coupled flow and finite-rate chemistry behind a normal shock using neural-network approximation of operators. *Journal of Computational Physics*, 447:110698, 2021. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2021.110698>. URL <https://www.sciencedirect.com/science/article/pii/S0021999121005933>.
- T. Marwah, Z. Lipton, and A. Risteski. Parametric complexity bounds for approximating pdes with neural networks. *Advances in Neural Information Processing Systems*, 34:15044–15055, 2021.
- T. Marwah, Z. C. Lipton, J. Lu, and A. Risteski. Neural network approximations of pdes beyond linearity: A representational perspective. In *International Conference on Machine Learning*, pages 24139–24172. PMLR, 2023.
- A. Maurer and M. Pontil. Concentration inequalities under sub-Gaussian and sub-exponential conditions. *Advances in Neural Information Processing Systems*, 34:7588–7597, 2021.
- C. Milbradt and M. Wahl. High-probability bounds for the reconstruction error of PCA. *Statistics & Probability Letters*, 161:108741, 2020.
- N. H. Nelsen and A. M. Stuart. The random feature model for input-output maps between banach spaces. *SIAM Journal on Scientific Computing*, 43(5):A3212–A3243, 2021.
- J. A. Opschoor, C. Schwab, and J. Zech. Exponential relu dnn expression of holomorphic maps in high dimension. *Constructive Approximation*, 55(1):537–582, 2022.
- J. A. A. Opschoor, C. Schwab, and J. Zech. Deep learning in high dimension: ReLU network expression rates for bayesian PDE inversion. Technical Report 2020-47, Seminar for Applied Mathematics, ETH Zürich, Switzerland, 2020.
- D. Patel, D. Ray, M. R. A. Abdelmalik, T. J. R. Hughes, and A. A. Oberai. Variationally mimetic operator networks. *arXiv preprint arXiv:2209.12871*, 2022.

- I. F. Pinelis and A. I. Sakhanenko. Remarks on inequalities for large deviation probabilities. *Theory of Probability & Its Applications*, 30(1):143–148, 1986.
- M. Prasthofer, T. De Ryck, and S. Mishra. Variable-input deep operator networks. *arXiv preprint arXiv:2205.11404*, 2022.
- B. Raonic, R. Molinaro, T. Rohner, S. Mishra, and E. de Bezenac. Convolutional neural operators. In *ICLR 2023 Workshop on Physics for Machine Learning*, 2023.
- M. Reiß and M. Wahl. Nonasymptotic upper bounds for the reconstruction error of PCA. *The Annals of Statistics*, 48(2):1098 – 1123, 2020. doi: 10.1214/19-AOS1839. URL <https://doi.org/10.1214/19-AOS1839>.
- A. Rudi and L. Rosasco. Generalization properties of learning with random features. *Advances in neural information processing systems*, 30, 2017.
- T. D. Ryck and S. Mishra. Error analysis for deep neural network approximations of parametric hyperbolic conservation laws. *arXiv preprint arXiv:2207.07362*, 2021.
- T. D. Ryck and S. Mishra. Generic bounds on the approximation error for physics-informed (and) operator learning. In A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, editors, *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=bF4eYy3LTR9>.
- C. Schwab and J. Zech. Deep learning in high dimension: Neural network expression rates for generalized polynomial chaos expansions in uq. *Analysis and Applications*, 17(01): 19–55, 2019.
- J. H. Seidman, G. Kissas, P. Perdikaris, and G. J. Pappas. NOMAD: Nonlinear manifold decoders for operator learning. In A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, editors, *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=50WV-sZvM1>.
- J. W. Siegel. Optimal approximation rates for deep relu neural networks on sobolev spaces. *arXiv preprint arXiv:2211.14400*, 2022.
- T. Tripura and S. Chakraborty. Wavelet neural operator for solving parametric partial differential equations in computational mechanics problems. *Computer Methods in Applied Mechanics and Engineering*, 404:115783, 2023. ISSN 0045-7825. doi: <https://doi.org/10.1016/j.cma.2022.115783>. URL <https://www.sciencedirect.com/science/article/pii/S0045782522007393>.
- R. Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- D. Yarotsky. Error bounds for approximations with deep ReLU networks. *Neural Networks*, 94:103–114, 2017. Publisher: Elsevier.
- D. Yarotsky. Optimal approximation of continuous functions by very deep relu networks. In *Conference on learning theory*, pages 639–649. PMLR, 2018.

Appendix A. PCA projection error bound

Let $\widehat{P}_{\leq d} = \mathcal{D}_{\mathcal{H}} \circ \mathcal{E}_{\mathcal{H}}$ denote the empirical PCA projection, given samples $u_1, \dots, u_N \stackrel{iid}{\sim} \mu$. We recall that the excess risk is defined by

$$\mathcal{E}_{\leq d}^{\text{PCA}} = \mathbb{E}_{u \sim \mu} \left[\|u - \widehat{P}_{\leq d} u\|_{\mathcal{H}}^2 \right] - \min_{P \in \Pi_d} \mathbb{E}_{u \sim \mu} \left[\|u - Pu\|_{\mathcal{H}}^2 \right]. \quad (29)$$

The following general bound on the excess risk is well-known (Reiß and Wahl, 2020, Sect. 2.2):

$$\mathcal{E}_{\leq d}^{\text{PCA}} \leq \sqrt{2d} \|\Sigma - \widehat{\Sigma}\|_{HS}, \quad (30)$$

where $\|\cdot\|_{HS}$ denotes the Hilbert-Schmidt norm, and where $\widehat{\Sigma} = \frac{1}{N} \sum_{k=1}^N u_k \otimes u_k$ the empirical covariance operator, and $\Sigma = \mathbb{E}[\widehat{\Sigma}]$ is its expectation. We recall that the set of Hilbert-Schmidt operators $HS(\mathcal{H})$ is itself a Hilbert space with norm $\|\cdot\|_{HS}$, and that $\|A\|_{HS}^2 = \text{tr}(A^T A)$ can be expressed in terms of the trace.

A.1 Proof of Proposition 1

Let $\mathcal{E}_{\leq d}^{\text{PCA}}$ be the excess risk (29). Note that the claimed bound (6) in Proposition 1 can be written in terms of the excess risk, as $\mathcal{E}_{\leq d}^{\text{PCA}} \leq \epsilon$. Rather than proving the high-probability result of Proposition 1 directly, we will deduce it from the following expectation result:

Lemma 29 *Let \mathcal{H} be a separable Hilbert space. Let $\mu \in \mathcal{P}(\mathcal{H})$ be a probability measure with finite second moments, $\mathbb{E}_{u \sim \mu} [\|u\|_{\mathcal{H}}^2] < \infty$. Then for any $\epsilon > 0$ and integer $d \in \mathbb{N}$, there exists a requisite amount of data $N = N(\mu, d, \epsilon)$, such that the encoding error for empirical PCA with dimension d , and based on N samples $u_1, \dots, u_N \stackrel{iid}{\sim} \mu$, satisfies*

$$\mathbb{E}_{\{u_j\} \sim \mu^{\otimes N}} [\mathcal{E}_{\leq d}^{\text{PCA}}] \leq \sqrt{2d} \mathbb{E} [\|\Sigma - \widehat{\Sigma}\|_{HS}] \leq \epsilon.$$

The corresponding high-probability result, claimed in Proposition 1, then follows readily from Lemma 29, as shown next.

Proof [Proof of Proposition 1 from Lemma 29] Let $\epsilon, \delta > 0$ be given. To deduce the claimed high-probability bound of Proposition 1, we apply Lemma 29 with ϵ replaced by $\epsilon' := \epsilon\delta$, to deduce that there exists a required amount of data $N = N(\epsilon, \delta)$, such that

$$\mathbb{E} [\mathcal{E}_{\leq d}^{\text{PCA}}] \leq \epsilon' = \epsilon\delta.$$

Using the non-negativity of the random variable $\mathcal{E}_{\leq d}^{\text{PCA}}$ together with Markov's inequality, this now immediately implies,

$$\text{Prob} [\mathcal{E}_{\leq d}^{\text{PCA}} > \epsilon] \leq \frac{\mathbb{E} [\mathcal{E}_{\leq d}^{\text{PCA}}]}{\epsilon} \leq \delta,$$

i.e. with probability at least $1 - \delta$, we have $\mathcal{E}_{\leq d}^{\text{PCA}} \leq \epsilon$. ■

We finally come to the proof of the result in expectation, Lemma 29.

Proof [Proof of Lemma 29] We define $Z := u \otimes u$, with $u \sim \mu$, as an $(HS(\mathcal{H}), \|\cdot\|_{HS})$ -valued random variable; note that Z takes values in the space of Hilbert-Schmidt operators $HS(\mathcal{H})$. Note that $\widehat{\Sigma} = \frac{1}{N} \sum_{k=1}^N Z_k$ is the sum of i.i.d. copies of Z . Given a cut-off $M > 0$, we define random variables $Z_{k,\leq M} := (u_k \otimes u_k)1_{[\|u_k\| \leq M]}$, $Z_{k,>M} := (u_k \otimes u_k)1_{[\|u_k\| > M]}$. Then, by the triangle inequality:

$$\|\Sigma - \widehat{\Sigma}\|_{HS} \leq \left\| \frac{1}{N} \sum_{k=1}^N Z_{k,\leq M} - \mathbb{E}[Z_{k,\leq M}] \right\|_{HS} + \frac{1}{N} \sum_{k=1}^N \|Z_{k,>M} - \mathbb{E}[Z_{k,>M}]\|_{HS}$$

We next observe that, upon taking expectations, the second contribution can be estimated by

$$\mathbb{E} [\|Z_{k,>M} - \mathbb{E}[Z_{k,>M}]\|_{HS}] \leq 2\mathbb{E} [\|Z_{k,>M}\|_{HS}] = 2\mathbb{E}_{u \sim \mu} [\|u\|^2 1_{[\|u\| > M]}].$$

On the other hand, since $\|Z_{k,\leq M}\|_{HS} = \|u_k\|^2 1_{\|u_k\| \leq M} \leq M^2$ is uniformly bounded, we can apply Jensen's inequality and the standard Monte-Carlo estimate on the expectation of the first term to obtain

$$\begin{aligned} \mathbb{E} \left[\left\| \frac{1}{N} \sum_{k=1}^N Z_{k,\leq M} - \mathbb{E}[Z_{k,\leq M}] \right\|_{HS} \right] &\leq \mathbb{E} \left[\left\| \frac{1}{N} \sum_{k=1}^N Z_{k,\leq M} - \mathbb{E}[Z_{k,\leq M}] \right\|_{HS}^2 \right]^{1/2} \\ &\leq \frac{M^2}{\sqrt{N}}. \end{aligned}$$

Combining these estimates, it follows that

$$\mathbb{E} [\|\Sigma - \widehat{\Sigma}\|_{HS}] \leq \frac{M^2}{\sqrt{N}} + 2\mathbb{E}_{u \sim \mu} [\|u\|^2 1_{[\|u\| > M]}]. \quad (31)$$

By assumption, the second moment $\mathbb{E}_{u \sim \mu} [\|u\|^2]$ is finite. Therefore, for any $\epsilon > 0$, there exists $M = M(\epsilon, d, \mu) > 0$, such that

$$2\mathbb{E}_{u \sim \mu} [\|u\|^2 1_{[\|u\| > M]}] \leq \frac{\epsilon}{2\sqrt{2d}}.$$

Choosing next $N = N(\epsilon, d, \mu)$ sufficiently large, we ensure that $M^2/\sqrt{N} \leq \epsilon/2\sqrt{2d}$. With this choice of N , it follows from (31), that

$$\mathbb{E} [\|\Sigma - \widehat{\Sigma}\|_{HS}] \leq \epsilon/\sqrt{2d},$$

and hence

$$\mathbb{E}_{\{u_j\} \sim \mu^{\otimes N}} [\mathcal{E}_{\leq d}^{\text{PCA}}] \leq \sqrt{2d} \mathbb{E} [\|\Sigma - \widehat{\Sigma}\|_{HS}] \leq \epsilon. \quad \blacksquare$$

A.2 Proof of Proposition 2

We first recall a vector-valued Bernstein inequality that is applicable on infinite-dimensional Hilbert spaces (Caponnetto and De Vito, 2007; Maurer and Pontil, 2021; Rudi and Rosasco, 2017). The following result is an immediate consequence of (Pinelis and Sakhnenko, 1986, Cor. 1, p. 144):

Theorem 30 (Vector-valued Bernstein inequality) *Let Z be an H -valued random variable, where H is a separable Hilbert space. Suppose there exist positive numbers $b > 0$ and $\sigma > 0$ such that*

$$\mathbb{E}\|Z - \mathbb{E}Z\|_H^p \leq \frac{1}{2}p!\sigma^2b^{p-2} \quad \text{for all } p \geq 2. \quad (32)$$

Then for any $\delta \in (0, 1)$ and $N \in \mathbb{N}$, denoting by $\{Z_k\}_{k=1}^N$ a sequence of N iid copies of Z , it holds that

$$\text{Prob} \left\{ \left\| \frac{1}{N} \sum_{k=1}^N Z_k - \mathbb{E}Z \right\|_H \leq \frac{2b \log(2/\delta)}{N} + \sqrt{\frac{2\sigma^2 \log(2/\delta)}{N}} \right\} \geq 1 - \delta. \quad (33)$$

A random variable Z satisfying the *Bernstein moment condition* (32) is subexponential in the sense that $\|Z - \mathbb{E}Z\|_H$ is subexponential on \mathbb{R} . The parameters σ^2 and b represent the variance and variance proxy/subexponential norm, respectively.

Proof [Proof of Proposition 2] Let $\hat{P}_{\leq d} = \mathcal{D}_{\mathcal{H}} \circ \mathcal{E}_{\mathcal{H}}$ denote the empirical PCA projection, given samples $u_1, \dots, u_N \stackrel{iid}{\sim} \mu$. We recall that the excess risk is defined by

$$\mathcal{E}_{\leq d}^{\text{PCA}} = \mathbb{E}_{u \sim \mu} \left[\|u - \hat{P}_{\leq d} u\|_{\mathcal{H}}^2 \right] - \min_{P \in \Pi_d} \mathbb{E}_{u \sim \mu} \left[\|u - Pu\|_{\mathcal{H}}^2 \right].$$

Note that the claimed bound (8) in Proposition 2 can be written in the form

$$\mathcal{E}_{\leq d}^{\text{PCA}} \leq \sqrt{\frac{Qd \log(2/\delta)}{N}} \quad (34)$$

We again rely on the following general bound on the excess risk (cp. eq. 30):

$$\mathcal{E}_{\leq d}^{\text{PCA}} \leq \sqrt{2d} \|\Sigma - \hat{\Sigma}\|_{HS}. \quad (35)$$

with $\|\cdot\|_{HS}$ the Hilbert-Schmidt norm, and where $\hat{\Sigma} = \frac{1}{N} \sum_{k=1}^N u_k \otimes u_k$ the empirical covariance operator, and we have $\Sigma = \mathbb{E}[\hat{\Sigma}]$. We recall that the set of Hilbert-Schmidt operators $HS(\mathcal{H})$ is itself a Hilbert space with norm $\|\cdot\|_{HS}$. Let $Z := u \otimes u$, $u \sim \mu$, be an $HS(\mathcal{H})$ -valued random variable. Then $\hat{\Sigma} = \frac{1}{N} \sum_{k=1}^N Z_k$ is the sum of i.i.d. copies of Z . Since $\|u\|_{\mathcal{H}}$ is sub-Gaussian by assumption, it follows that $\|Z\|_2 = \|u\|_{\mathcal{H}}^2$ is sub-exponential, and as a consequence, $\|Z - \mathbb{E}Z\|_2$ is also sub-exponential. In particular, Theorem 30 applied on the Hilbert space $H := HS(\mathcal{H})$ now implies that there exists a constant $q > 0$, depending only on the law μ of u (which also determines Z), such that

$$\|\hat{\Sigma} - \Sigma\|_{HS} = \left\| \frac{1}{N} \sum_{k=1}^N Z_k - \mathbb{E}[Z] \right\|_2 \leq q \max \left\{ \frac{\log(2/\delta)}{N}, \sqrt{\frac{\log(2/\delta)}{N}} \right\},$$

with probability at least $1 - \delta$. By assumption, we have $N \geq \log(2/\delta)$, and hence upon defining $Q := 2q^2$, we obtain

$$\|\widehat{\Sigma} - \Sigma\|_{HS} \leq q \max \left\{ \frac{\log(2/\delta)}{N}, \sqrt{\frac{\log(2/\delta)}{N}} \right\} \leq \sqrt{\frac{Q \log(2/\delta)}{2N}},$$

with probability at least $1 - \delta$. Substitution in (35) now yields (34). \blacksquare

Appendix B. Universal approximation

B.1 Proof of Theorem 6

The high-probability result of Theorem 6 can easily be derived from the following result in expectation:

Proposition 31 (Universal approximation in expectation) *Let \mathcal{X}, \mathcal{Y} be separable Hilbert spaces and let $\mu \in \mathcal{P}(\mathcal{X})$ be a probability measure on \mathcal{X} . Let $\Psi^\dagger : \mathcal{X} \rightarrow \mathcal{Y}$ be a μ -measurable mapping. Assume the following moment conditions,*

$$\mathbb{E}_{u \sim \mu}[\|u\|_{\mathcal{X}}^2], \quad \mathbb{E}_{u \sim \mu}[\|\Psi^\dagger(u)\|_{\mathcal{Y}}^2] < \infty.$$

Then for any $\epsilon > 0$, there are dimensions $d_{\mathcal{X}} = d_{\mathcal{X}}(\epsilon)$, $d_{\mathcal{Y}} = d_{\mathcal{Y}}(\epsilon)$, a requisite amount of data $N = N(d_{\mathcal{X}}, d_{\mathcal{Y}}, \mu, \Psi^\dagger)$, and a neural network ψ such that the PCA-Net, $\Psi = \mathcal{D}_{\mathcal{Y}} \circ \psi \circ \mathcal{E}_{\mathcal{X}}$, satisfies

$$\mathbb{E}_{\{u_k\} \sim \mu^{\otimes N}} \left[\mathbb{E}_{u \sim \mu} \left[\|\Psi^\dagger(u) - \Psi(u; \{u_k\})\|_{\mathcal{Y}}^2 \right] \right] \leq \epsilon.$$

Here, the expectation is over the given data $u_1, \dots, u_N \sim \mu$, which determine the input/output-pairs $\{u_k, \Psi^\dagger(u_k)\}_{k=1}^N$.

We first show that Proposition 31 implies Theorem 6.

Proof [Proof of Theorem 6 from Proposition 31] Define $Z \in \mathbb{R}$ by

$$Z := \mathbb{E}_{u \sim \mu} \|\Psi^\dagger(u) - \Psi(u; \{u_k\})\|_{\mathcal{Y}}^2.$$

Note that Z is itself a random variable due to its dependence on the random data $\{u_k\} \sim \mu^{\otimes N}$. Applying Proposition 31 with $\epsilon' := \epsilon\delta$ instead of ϵ , we conclude that there exist dimensions $d_{\mathcal{X}}, d_{\mathcal{Y}}$, a requisite amount of data $N \in \mathbb{N}$, and a neural network ψ , such that

$$\mathbb{E}_{\{u_k\} \sim \mu^{\otimes N}} [Z] \leq \epsilon' = \epsilon\delta.$$

We can now estimate the probability of the event $\{Z > \epsilon\}$,

$$\text{Prob}[Z > \epsilon] \leq \frac{\mathbb{E}_{\{u_k\} \sim \mu^{\otimes N}} [Z]}{\epsilon} \leq \delta.$$

In particular with this choice of $d_{\mathcal{X}}, d_{\mathcal{Y}}, N$ and ψ , we have

$$\mathbb{E}_{u \sim \mu} \|\Psi^\dagger(u) - \Psi(u; \{u_k\})\|_{\mathcal{Y}}^2 \leq \epsilon,$$

with probability at least $1 - \delta$. \blacksquare

It thus remains to prove Proposition 31, which is the main aim of this appendix. The detailed proof is the subject of the following section.

B.2 Proof of Proposition 31 (in expectation)

We are interested in the following PCA-Net approximation error,

$$\widehat{\mathcal{E}} = \mathbb{E}_{u \sim \mu} \left[\|\Psi^\dagger(u) - \Psi(u)\|_{\mathcal{Y}}^2 \right]^{1/2}, \quad (36)$$

where we have suppressed the dependence of $\Psi(u) = \Psi(u; \{u_k\})$ on the data, arising from the data-dependence of empirical PCA. We also recall that a PCA-Net is written as a composition of three mappings $\Psi(a) = \mathcal{D}_{\mathcal{Y}} \circ \psi \circ \mathcal{E}_{\mathcal{X}}$. Intuitively, we expect each of these mappings to give a contribution to the overall error. Our analysis of $\widehat{\mathcal{E}}$ will thus be based on a decomposition of the approximation error $\widehat{\mathcal{E}}$ in terms of three contributions: PCA encoding/decoding errors $\widehat{\mathcal{E}}_{\mathcal{X}}$ and $\widehat{\mathcal{E}}_{\mathcal{Y}}$ on \mathcal{X} and \mathcal{Y} , and a neural network error $\widehat{\mathcal{E}}_{\psi}$. At times, it is more convenient not to consider the encoding and neural network error separately, we will denote the combined error (associated with the composition $\psi \circ \mathcal{E}_{\mathcal{X}}$) by $\widehat{\mathcal{E}}_{\psi}^*$.

Lemma 32 (Error decomposition) *Let $\mu \in \mathcal{P}(\mathcal{X})$ be a probability measure on a Hilbert space \mathcal{X} . Let $\Psi^\dagger : \mathcal{X} \rightarrow \mathcal{Y}$ be a μ -measurable, non-linear operator, with $\mathbb{E}_{u \in \mu} [\|\Psi^\dagger(u)\|_{\mathcal{Y}}^2] < \infty$. Let $\Psi(u) := \mathcal{D}_{\mathcal{Y}} \circ \psi \circ \mathcal{E}_{\mathcal{X}}(u)$ be a PCA-Net with PCA dimensions $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$. Then the approximation error satisfies the following estimates,*

$$\widehat{\mathcal{E}} \leq \widehat{\mathcal{E}}_{\mathcal{Y}} + \widehat{\mathcal{E}}_{\psi}^* \leq \widehat{\mathcal{E}}_{\mathcal{Y}} + \text{Lip}(\Psi^\dagger) \widehat{\mathcal{E}}_{\mathcal{X}} + \widehat{\mathcal{E}}_{\psi}. \quad (37)$$

Here we have introduced the following encoding errors on \mathcal{X} and \mathcal{Y} ,

$$\widehat{\mathcal{E}}_{\mathcal{X}} = \mathbb{E}_{u \sim \mu} \left[\|u - \widehat{P}_{\leq d_{\mathcal{X}}}^{\mathcal{X}} u\|_{\mathcal{X}}^2 \right]^{1/2}, \quad (38a)$$

$$\widehat{\mathcal{E}}_{\mathcal{Y}} = \mathbb{E}_{v \sim \Psi_{\#}^\dagger \mu} \left[\|v - \widehat{P}_{\leq d_{\mathcal{Y}}}^{\mathcal{Y}} v\|_{\mathcal{Y}}^2 \right]^{1/2}, \quad (38b)$$

with $\widehat{P}_{\leq d_{\mathcal{X}}}^{\mathcal{X}} u := \mathcal{D}_{\mathcal{X}} \circ \mathcal{E}_{\mathcal{X}}(u)$ and $\widehat{P}_{\leq d_{\mathcal{Y}}}^{\mathcal{Y}} u := \mathcal{D}_{\mathcal{Y}} \circ \mathcal{E}_{\mathcal{Y}}(u)$ the empirical PCA projections, and we introduce the following neural network errors,

$$\widehat{\mathcal{E}}_{\mathcal{X}}^* = \mathbb{E}_{u \sim \mu} \left[\|\mathcal{E}_{\mathcal{Y}} \circ \Psi^\dagger(u) - \psi \circ \mathcal{E}_{\mathcal{X}}(u)\|_{\ell^2(\mathbb{R}^{d_{\mathcal{Y}}})}^2 \right]^{1/2}, \quad (39a)$$

$$\widehat{\mathcal{E}}_{\psi} = \mathbb{E}_{\alpha \sim (\mathcal{E}_{\mathcal{X}})_{\#} \mu} \left[\|\mathcal{E}_{\mathcal{Y}} \circ \Psi^\dagger \circ \mathcal{D}_{\mathcal{X}}(\alpha) - \psi(\alpha)\|_{\ell^2(\mathbb{R}^{d_{\mathcal{Y}}})}^2 \right]^{1/2}. \quad (39b)$$

The derivation of Lemma 32 is straight-forward, and has been provided in detail for the conceptually similar DeepONet architecture in (Lanthaler et al., 2022, Theorem 3.3). We will not repeat the derivation here.

We next state a general approximation result for operators, which will be a key ingredient in our proof of the universal approximation result, Proposition 6. This next Lemma 33 is formulated with respect to the supremum norm, which prepares the stage for a corresponding approximation result in the L_{μ}^2 -norm in Lemma 34.

Lemma 33 *Let $\Psi^\dagger : K \rightarrow \mathcal{Y}$ be a continuous mapping, where $K \subset \mathcal{X}$ is compact, and \mathcal{X}, \mathcal{Y} are Banach spaces. Assume that $\sup_{u \in K} \|\Psi^\dagger(u)\|_{\mathcal{Y}} \leq M$ for some $M > 0$. Then for any $\epsilon > 0$, there exists a Lipschitz continuous $\Psi^* : \mathcal{X} \rightarrow \mathcal{Y}$, such that*

$$\sup_{u \in K} \|\Psi^\dagger(u) - \Psi^*(u)\|_{\mathcal{Y}} \leq \epsilon,$$

and $\|\Psi^*(u)\|_{\mathcal{Y}} \leq M$, for all $u \in \mathcal{X}$.

Our proof of Lemma 33 uses only classical tools of topology; we suspect that similar results may have been derived before. For completeness, we provide a detailed proof, below. **Proof** To prove the above lemma, we first construct a suitable Lipschitz continuous partition of unity $\{\psi_j : \mathcal{X} \rightarrow [0, 1]\}$ on K , which has a Lipschitz continuous extension to all of \mathcal{X} . Then, we show that an interpolation of Ψ^\dagger of the form $\Psi^*(u) := \sum_j \Psi^\dagger(u_j)\psi_j(u)$ satisfies the claimed properties.

Since K is compact, Ψ^\dagger is in fact uniformly continuous. In particular, there exists $\delta > 0$, such that for all $u, u' \in K$, we have

$$\|u - u'\|_{\mathcal{X}} \leq \delta \quad \Rightarrow \quad \|\Psi^\dagger(u) - \Psi^\dagger(u')\|_{\mathcal{Y}} \leq \epsilon. \quad (40)$$

Using the compactness of K , let $\{u_j\}_{j=1}^N$ be a finite collection of elements of K , such that the open balls of radius $\delta > 0$ form a cover of K , i.e.

$$K \subset \bigcup_{j=1}^N B_\delta(u_j).$$

Define a ‘‘hat’’ function $\rho_\delta : [0, \infty) \rightarrow [0, \infty)$, by

$$\rho_\delta(r) := \begin{cases} 1 - r/\delta, & (0 \leq r \leq \delta), \\ 0, & (r > \delta). \end{cases}$$

We note that $\Psi : \mathcal{X} \rightarrow \mathbb{R}$, given by

$$\Psi(u) := \text{dist}(u, K) + \sum_{j=1}^N \rho_\delta(\|u - u_j\|_{\mathcal{X}}), \quad (41)$$

is a Lipschitz continuous function, and that there exists $\delta_0 > 0$, such that $\Psi(u) \geq \delta_0$ for all $u \in \mathcal{X}$. To see the latter, we note that since K is compact, there exists $0 < \delta_0 \leq \delta$, such that

$$K \subset B_{\delta_0}(K) := \{u \in \mathcal{X} \mid \text{dist}(u, K) < \delta_0\} \subset \bigcup_{j=1}^N B_\delta(u_j).$$

After possibly decreasing δ_0 , we may wlog assume that $\delta_0 < 1/2$ and $\delta_0 < \delta/2$. It then follows from (41) that, for any $u \notin B_{\delta_0}(K)$, we have $\Psi(u) \geq \text{dist}(u, K) \geq \delta_0$. On the other hand, if $u \in B_{\delta_0}(K)$, then there exists u_j , such that $\|u - u_j\| < \delta_0 \leq \delta/2$. And hence,

$$\Psi(u) \geq \rho_\delta(\|u - u_j\|_{\mathcal{X}}) \geq 1 - 1/2 = 1/2 \geq \delta_0.$$

Thus, $\Psi(u) \geq \delta_0$ for all $u \in \mathcal{X}$. Next, define

$$\psi_j(u) := \frac{\rho_\delta(\|u - u_j\|_{\mathcal{X}})}{\Psi(u)}, \quad (j = 1, \dots, N).$$

Then ψ_j is globally Lipschitz continuous for all $j = 1, \dots, N$, $0 \leq \psi_j \leq 1$, and $\sum_{j=1}^N \psi_j(u) = 1$ for all $u \in K$, i.e. the ψ_j form a partition of unity on K . Furthermore, it is easy to see that $\sum_{j=1}^N \psi_j(u) \leq 1$ for all $u \in \mathcal{X} \setminus K$. Given this construction, we now define

$$\Psi^*(u) := \sum_{j=1}^N \psi_j(u)\Psi^\dagger(u_j).$$

Since $\|\Psi^\dagger(u_j)\|_{\mathcal{Y}} \leq \sup_{u \in K} \|\Psi^\dagger(u)\|_{\mathcal{Y}} \leq M$ for all $j = 1, \dots, N$, and since $\sum_{j=1}^N \psi_j(u) \leq 1$, it follows that for any $u \in \mathcal{X}$:

$$\|\Psi^*(u)\|_{\mathcal{Y}} \leq \sum_{j=1}^N \psi_j(u) \|\Psi^\dagger(u_j)\|_{\mathcal{Y}} \leq \sum_{j=1}^N \psi_j(u) M \leq M.$$

This shows the upper bound $\sup_{u \in \mathcal{X}} \|\Psi^*(u)\|_{\mathcal{Y}} \leq M$. We finally claim that $\sup_{u \in K} \|\Psi^\dagger(u) - \Psi^*(u)\|_{\mathcal{Y}} \leq \epsilon$. To prove this, we note that for $u \in K$:

$$\|\Psi^\dagger(u) - \Psi^*(u)\|_{\mathcal{Y}} \leq \sum_{j=1}^N \psi_j(u) \|\Psi^\dagger(u) - \Psi^\dagger(u_j)\|_{\mathcal{Y}} \leq \sum_{j=1}^N \psi_j(u) \epsilon = \epsilon,$$

where we have used the uniform continuity (40) of Ψ^\dagger in the second inequality and the fact that $\psi_j(u) \neq 0$, only if $\|u - u_j\|_{\mathcal{X}} \leq \delta$. This completes the proof. \blacksquare

Based on Lemma 33, we can prove the following L_μ^2 -approximation result:

Lemma 34 *Let \mathcal{X}, \mathcal{Y} be separable Banach spaces, let $\mu \in \mathcal{P}(\mathcal{X})$ be a probability measure on \mathcal{X} . Assume that $\Psi^\dagger : \mathcal{X} \rightarrow \mathcal{Y}$ is a Borel measurable mapping, such that $\mathbb{E}_{u \sim \mu} [\|\Psi^\dagger(u)\|_{\mathcal{Y}}^2] < \infty$. Then for any $\epsilon > 0$, there exists a bounded Lipschitz continuous mapping $\Psi^* : \mathcal{X} \rightarrow \mathcal{Y}$, such that*

$$\mathbb{E}_{u \sim \mu} \left[\|\Psi^\dagger(u) - \Psi^*(u)\|_{\mathcal{Y}}^2 \right]^{1/2} < \epsilon,$$

and $\sup_{u \in \mathcal{X}} \|\Psi^*(u)\|_{\mathcal{Y}} < \infty$.

Proof Let $\epsilon > 0$ be given. Define Ψ_M^\dagger by

$$\Psi_M^\dagger(u) := \begin{cases} \Psi^\dagger(u), & (\|\Psi^\dagger(u)\|_{\mathcal{Y}} \leq M), \\ M \frac{\Psi^\dagger(u)}{\|\Psi^\dagger(u)\|_{\mathcal{Y}}}, & (\|\Psi^\dagger(u)\|_{\mathcal{Y}} > M). \end{cases}$$

For sufficiently large $M > 0$, we have $\mathbb{E}_{u \sim \mu} \left[\|\Psi^\dagger(u) - \Psi_M^\dagger(u)\|_{\mathcal{Y}}^2 \right]^{1/2} \leq \epsilon/2$. It will thus suffice to show that there exists $\Psi^* : \mathcal{X} \rightarrow \mathcal{Y}$ which is globally Lipschitz, satisfies $\|\Psi^*(u)\|_{\mathcal{Y}} \leq M$ for all $u \in \mathcal{X}$, and

$$\mathbb{E}_{u \sim \mu} \left[\|\Psi^*(u) - \Psi_M^\dagger(u)\|_{\mathcal{Y}}^2 \right]^{1/2} \leq \epsilon/2.$$

To simplify notation, we will drop the subscript M , and assume directly that $\|\Psi^\dagger(u)\|_{\mathcal{Y}} \leq M$ for all $u \in \mathcal{X}$. We now prove the existence Ψ^* . First, by Lusin's theorem, we note that there exists a compact $K \subset \mathcal{X}$, such that

$$\mu(\mathcal{X} \setminus K) \leq \left(\frac{\epsilon}{8M} \right)^2, \quad \Psi^\dagger|_K : K \rightarrow \mathcal{Y} \text{ is continuous.} \quad (42)$$

By Lemma 33, there exists Lipschitz continuous $\Psi^* : \mathcal{X} \rightarrow \mathcal{Y}$, such that $\sup_{u \in \mathcal{X}} \|\Psi^*(u)\|_{\mathcal{Y}} \leq M$, and

$$\sup_{u \in K} \|\Psi^\dagger(u) - \Psi^*(u)\|_{\mathcal{Y}} \leq \epsilon/4.$$

Thus, it follows that

$$\begin{aligned} \mathbb{E}_{u \sim \mu} \left[\|\Psi^\dagger(u) - \Psi^*(u)\|_{\mathcal{Y}}^2 \right]^{1/2} &\leq \mathbb{E}_{u \sim \mu} \left[1_K(u) \|\Psi^\dagger(u) - \Psi^*(u)\|_{\mathcal{Y}}^2 \right]^{1/2} \\ &\quad + \mathbb{E}_{u \sim \mu} \left[1_{\mathcal{X} \setminus K}(u) \|\Psi^\dagger(u) - \Psi^*(u)\|_{\mathcal{Y}}^2 \right]^{1/2} \\ &\leq \epsilon/4 + 2M [\mu(\mathcal{X} \setminus K)]^{1/2} \leq \epsilon/2. \end{aligned}$$

The last inequality follows from the choice of K in (42). This proves the claim. \blacksquare

After these preparatory results, we can now finally provide the detailed proof of Proposition 31 (i.e. universal approximation in expectation):

Proof [Proof of Proposition 31] By assumption, $\Psi^\dagger : \mathcal{X} \rightarrow \mathcal{Y}$ is Borel measurable and has finite second moment, $\mathbb{E}_{u \sim \mu} [\|\Psi^\dagger(u)\|_{\mathcal{Y}}^2] < \infty$. Our goal is to show that for any $\epsilon > 0$, there exist PCA dimensions $d_{\mathcal{X}}$, $d_{\mathcal{Y}}$, a requisite amount of data N , depending only on ϵ , and a neural network ψ , such that the PCA-Net approximation $\Psi = \mathcal{D}_{\mathcal{Y}} \circ \psi \circ \mathcal{E}_{\mathcal{X}}$ of Ψ^\dagger , with empirical PCA encoder/decoder and based on $u_1, \dots, u_N \stackrel{iid}{\sim} \mu$ i.i.d. samples, satisfies

$$\mathbb{E}_{\{u_k\} \sim \mu^{\otimes N}} \mathbb{E}_{u \sim \mu} \left[\|\Psi^\dagger - \Psi\|_{\mathcal{Y}}^2 \right]^{1/2} \leq \epsilon.$$

To this end, let $\epsilon > 0$ be given. The proof will follow in six steps 1–6, each constructing and estimating a different part of the PCA-Net Ψ . We start from Lemma 32, which gives the following general error decomposition:

$$\mathbb{E}_{u \sim \mu} \left[\|\Psi^\dagger(u) - \Psi(u)\|_{\mathcal{Y}}^2 \right]^{1/2} \leq \widehat{\mathcal{E}}_{\mathcal{Y}} + \widehat{\mathcal{E}}_{\mathcal{X}}^*,$$

where

$$\begin{aligned} \widehat{\mathcal{E}}_{\mathcal{Y}} &= \mathbb{E}_{v \sim \Psi_{\#}^\dagger \mu} \left[\|v - \mathcal{D}_{\mathcal{Y}} \circ \mathcal{E}_{\mathcal{Y}}(v)\|_{\mathcal{Y}}^2 \right]^{1/2}, \\ \widehat{\mathcal{E}}_{\mathcal{X}}^* &= \mathbb{E}_{u \sim \mu} \left[\|\mathcal{E}_{\mathcal{Y}} \circ \Psi^\dagger(u) - \psi \circ \mathcal{E}_{\mathcal{X}}(u)\|_{\ell^2(\mathbb{R}^{d_{\mathcal{X}}})} \right]^{1/2}, \end{aligned}$$

Step 1: (Bounding the decoding error on \mathcal{Y}) The first error, $\widehat{\mathcal{E}}_{\mathcal{Y}}$, is the PCA encoding error on \mathcal{Y} . To control this error, we first choose $d_{\mathcal{Y}} = d_{\mathcal{Y}}(\epsilon)$ sufficiently large, such that the (optimal) PCA projection error of the push-forward measure $\Psi_{\#}^\dagger \mu$ on \mathcal{Y} , is bounded by

$$\mathcal{R}_{d_{\mathcal{Y}}}^{\text{opt}}(\Psi_{\#}^\dagger \mu) \equiv \mathbb{E}_{v \sim \Psi_{\#}^\dagger \mu} \left[\|v - \mathcal{D}_{\mathcal{Y}}^{\text{opt}} \circ \mathcal{E}_{\mathcal{Y}}^{\text{opt}}(v)\|_{\mathcal{Y}}^2 \right] \leq \epsilon/8.$$

Next, invoking Lemma 29 for the probability measure $\Psi_{\#}^\dagger \mu$ and with ϵ replaced by $\epsilon/8$, it follows that there exists a requisite amount of data $N = N(\Psi^\dagger, \mu, \epsilon)$, such that

$$\mathbb{E}_{\{u_j\} \sim \mu^{\otimes N}} \mathbb{E}_{v \sim \Psi_{\#}^\dagger \mu} \left[\|v - \mathcal{D}_{\mathcal{Y}} \circ \mathcal{E}_{\mathcal{Y}}(v)\|_{\mathcal{Y}}^2 \right] \leq \frac{\epsilon}{8} + R_{d_{\mathcal{Y}}}^{\text{opt}}(\Psi_{\#}^\dagger \mu) \leq \frac{\epsilon}{4}. \quad (43)$$

Step 2: (Lipschitz approximation $\Psi^* \approx \Psi$) Next, we aim to estimate the approximation error $\widehat{\mathcal{E}}_{\mathcal{X}}^*$. To this end, we note that for $\Psi^* \approx \Psi^\dagger$ a Lipschitz continuous approximation of Ψ^\dagger , we have

$$\begin{aligned} \widehat{\mathcal{E}}_{\mathcal{X}}^* &\leq \mathbb{E}_{u \sim \mu} \left[\|\mathcal{E}_{\mathcal{Y}} \circ \Psi^*(u) - \psi \circ \mathcal{E}_{\mathcal{X}}(u)\|_{\ell^2(\mathbb{R}^{d_{\mathcal{X}}})}^2 \right]^{1/2} \\ &\quad + \mathbb{E}_{u \sim \mu} \left[\|\mathcal{E}_{\mathcal{Y}} \circ \Psi^\dagger(u) - \mathcal{E}_{\mathcal{Y}} \circ \Psi^*(u)\|_{\ell^2(\mathbb{R}^{d_{\mathcal{X}}})}^2 \right]^{1/2} \\ &\leq \mathbb{E}_{u \sim \mu} \left[\|\mathcal{E}_{\mathcal{Y}} \circ \Psi^*(u) - \psi \circ \mathcal{E}_{\mathcal{X}}(u)\|_{\ell^2(\mathbb{R}^{d_{\mathcal{X}}})}^2 \right]^{1/2} \\ &\quad + \mathbb{E}_{u \sim \mu} \left[\|\Psi^\dagger(u) - \Psi^*(u)\|_{\mathcal{Y}}^2 \right]^{1/2}, \end{aligned}$$

where we used that the PCA encoding $\mathcal{E}_{\mathcal{Y}} : \mathcal{Y} \rightarrow \mathbb{R}^{d_{\mathcal{Y}}}$ is a contraction in the last step. Note that the last term is independent of the PCA encoding, and only depends on the approximation $\Psi^* \approx \Psi^\dagger$. Using Lemma 34, we can find a Lipschitz continuous bounded mapping $\Psi^* : \mathcal{X} \rightarrow \mathcal{Y}$, such that

$$\mathbb{E}_{u \sim \mu} \left[\|\Psi^\dagger(u) - \Psi^*(u)\|_{\mathcal{Y}}^2 \right]^{1/2} \leq \frac{\epsilon}{4}, \quad \sup_{u \in \mathcal{X}} \|\Psi^*(u)\|_{\mathcal{Y}} \leq M < \infty, \quad (44)$$

for some constant $M > 0$, depending only on ϵ .

Step 3: (Decomposition in \mathcal{X} -encoding and neural network approximation errors) Next, we continue to bound the first term in the estimate on $\widehat{\mathcal{E}}_{\mathcal{X}}^*$ as follows:

$$\begin{aligned} &\mathbb{E}_{u \sim \mu} \left[\|\mathcal{E}_{\mathcal{Y}} \circ \Psi^*(u) - \psi \circ \mathcal{E}_{\mathcal{X}}(u)\|_{\ell^2(\mathbb{R}^{d_{\mathcal{X}}})}^2 \right]^{1/2} \\ &\leq \mathbb{E}_{u \sim \mu} \left[\|\mathcal{E}_{\mathcal{Y}} \circ \Psi^*(u) - \mathcal{E}_{\mathcal{Y}} \circ \Psi^* \circ \mathcal{D}_{\mathcal{X}} \circ \mathcal{E}_{\mathcal{X}}(u)\|_{\ell^2(\mathbb{R}^{d_{\mathcal{X}}})}^2 \right]^{1/2} \\ &\quad + \mathbb{E}_{u \sim \mu} \left[\|\mathcal{E}_{\mathcal{Y}} \circ \Psi^* \circ \mathcal{D}_{\mathcal{X}} \circ \mathcal{E}_{\mathcal{X}}(u) - \psi \circ \mathcal{E}_{\mathcal{X}}(u)\|_{\ell^2(\mathbb{R}^{d_{\mathcal{X}}})}^2 \right]^{1/2} \\ &\leq \text{Lip}(\Psi^*) \mathbb{E}_{u \sim \mu} \left[\|u - \mathcal{D}_{\mathcal{X}} \circ \mathcal{E}_{\mathcal{X}}(u)\|_{\mathcal{X}}^2 \right]^{1/2} \\ &\quad + \mathbb{E}_{u \sim \mu} \left[\|\mathcal{E}_{\mathcal{Y}} \circ \Psi^* \circ \mathcal{D}_{\mathcal{X}} \circ \mathcal{E}_{\mathcal{X}}(u) - \psi \circ \mathcal{E}_{\mathcal{X}}(u)\|_{\ell^2(\mathbb{R}^{d_{\mathcal{X}}})}^2 \right]^{1/2} \\ &=: \text{Lip}(\Psi^*) \widehat{\mathcal{E}}_{\mathcal{X}} + \widehat{\mathcal{E}}_{\Psi^*}, \end{aligned}$$

where we denote

$$\begin{aligned} \widehat{\mathcal{E}}_{\Psi^*} &:= \mathbb{E}_{u \sim \mu} \left[\|\mathcal{E}_{\mathcal{Y}} \circ \Psi^* \circ \mathcal{D}_{\mathcal{X}} \circ \mathcal{E}_{\mathcal{X}}(u) - \psi \circ \mathcal{E}_{\mathcal{X}}(u)\|_{\ell^2(\mathbb{R}^{d_{\mathcal{X}}})}^2 \right]^{1/2} \\ &= \mathbb{E}_{w \sim (\mathcal{E}_{\mathcal{X}})_{\#} \mu} \left[\|\mathcal{E}_{\mathcal{Y}} \circ \Psi^* \circ \mathcal{D}_{\mathcal{X}}(w) - \psi(w)\|_{\ell^2(\mathbb{R}^{d_{\mathcal{X}}})}^2 \right]^{1/2}. \end{aligned}$$

Step 4: (Bounding the encoding error on \mathcal{X}) Again, employing the estimate from Lemma 29 similar to the argument leading to (43), we note that there exists a PCA dimension $d_{\mathcal{X}} = d_{\mathcal{X}}(\mu, \text{Lip}(\Psi^*), \epsilon)$ and a requisite amount of data $N = N(\mu, \text{Lip}(\Psi^*), \epsilon)$, such that we can bound the encoding error $\mathcal{E}_{\mathcal{X}}$ by

$$\mathbb{E}_{\{u_j\} \sim \mu^{\otimes N}} \mathbb{E}_{u \sim \mu} \left[\|u - \mathcal{D}_{\mathcal{X}} \circ \mathcal{E}_{\mathcal{X}}(u)\|_{\mathcal{X}}^2 \right]^{1/2} \leq \frac{\epsilon}{4 \text{Lip}(\Psi^*)}.$$

This ensures that

$$\text{Lip}(\Psi^*) \mathbb{E}_{\{u_j\} \sim \mu^{\otimes N}}[\widehat{\mathcal{E}}_{\mathcal{X}}] \leq \frac{\epsilon}{4}. \quad (45)$$

Step 5: (Bounding the neural network approximation error) Finally, we define the Lipschitz continuous function $G : \mathbb{R}^{d_{\mathcal{X}}} \rightarrow \mathbb{R}^{d_{\mathcal{Y}}}$ by $G(w) := \mathcal{E}_{\mathcal{Y}} \circ \Psi^* \circ \mathcal{D}_{\mathcal{X}}(w)$, and denote $\nu := (\mathcal{E}_{\mathcal{X}})_{\#} \mu$. We note that (44) implies that

$$\sup_{w \in \mathbb{R}^{d_{\mathcal{X}}}} |G(w)| = \sup_{w \in \mathbb{R}^{d_{\mathcal{X}}}} |\mathcal{E}_{\mathcal{Y}}(\Psi^*(\mathcal{D}_{\mathcal{X}}(w)))| \leq \sup_{u \in \mathcal{X}} \|\mathcal{E}_{\mathcal{Y}}\|_{\mathcal{X} \rightarrow \ell^2(\mathbb{R}^{d_{\mathcal{X}}})} \|\Psi^*(u)\|_{\mathcal{Y}} \leq M,$$

where we made use of the fact that the PCA-encoder $\mathcal{E}_{\mathcal{Y}}$ is contractive, i.e.

$$|\mathcal{E}_{\mathcal{Y}}(v)| \equiv \|\mathcal{E}_{\mathcal{Y}}(v)\|_{\ell^2(\mathbb{R}^{d_{\mathcal{X}}})} \leq \|v\|_{\mathcal{Y}}, \quad \forall v \in \mathcal{Y}.$$

For any $R > 0$, we then obtain the following estimate on $\widehat{\mathcal{E}}_{\Psi^*}$:

$$\begin{aligned} \widehat{\mathcal{E}}_{\Psi^*} &\leq \left\{ M + \sup_{w \in \mathbb{R}^{d_{\mathcal{X}}}} |\psi(w)| \right\} \nu \left(\mathbb{R}^{d_{\mathcal{X}}} \setminus [-R, R]^{d_{\mathcal{X}}} \right)^{1/2} \\ &\quad + \sup_{w \in [-R, R]^{d_{\mathcal{X}}}} |G(w) - \psi(w)|. \end{aligned}$$

Since ν is a probability measure on $\mathbb{R}^{d_{\mathcal{X}}}$, by choosing $R > 0$ sufficiently large, we can ensure that

$$3M\nu \left(\mathbb{R}^{d_{\mathcal{X}}} \setminus [-R, R]^{d_{\mathcal{X}}} \right)^{1/2} \leq \epsilon/8.$$

Next, by the universal approximation theorem of ReLU neural networks, there exists a (sufficiently large) DNN $\tilde{\psi} : \mathbb{R}^{d_{\mathcal{X}}} \rightarrow \mathbb{R}^{d_{\mathcal{Y}}}$, such that

$$\sup_{w \in [-R, R]^{d_{\mathcal{X}}}} |G(w) - \tilde{\psi}(w)| \leq \epsilon/8.$$

Finally, adding the two-layer ReLU network $\gamma_R : \mathbb{R}^{d_{\mathcal{X}}} \rightarrow \mathbb{R}^{d_{\mathcal{Y}}}$, given by $\gamma_R(x) := \min(\max(x, -R), R)$ ensures that $\psi(w) := \tilde{\psi}(\gamma_R(w))$ satisfies

$$\psi(w) = \tilde{\psi}(w), \quad \forall w \in [-R, R]^{d_{\mathcal{X}}},$$

and (wlog assuming $\epsilon \leq M$):

$$\sup_{w \in \mathbb{R}^{d_{\mathcal{X}}}} |\psi(w)| = \sup_{w \in [-R, R]^{d_{\mathcal{X}}}} |\tilde{\psi}(w)| \leq M + \epsilon \leq 2M.$$

Thus, with this construction of ψ , we have

$$\begin{aligned} \widehat{\mathcal{E}}_{\Psi^*} &\leq \left\{ M + \sup_{w \in \mathbb{R}^{d_{\mathcal{X}}}} |\psi(w)| \right\} \nu \left(\mathbb{R}^{d_{\mathcal{X}}} \setminus [-R, R]^{d_{\mathcal{X}}} \right)^{1/2} \\ &\quad + \sup_{w \in [-R, R]^{d_{\mathcal{X}}}} |G(w) - \psi(w)| \\ &= 3M\nu \left(\mathbb{R}^{d_{\mathcal{X}}} \setminus [-R, R]^{d_{\mathcal{X}}} \right)^{1/2} + \sup_{w \in [-R, R]^{d_{\mathcal{X}}}} |G(w) - \tilde{\psi}(w)| \\ &\leq \epsilon/8 + \epsilon/8 = \epsilon/4. \end{aligned}$$

Thus, for this neural network ψ , we have

$$\widehat{\mathcal{E}}_{\Psi^*} \leq \frac{\epsilon}{4}. \quad (46)$$

Step 6: (Conclusion) To conclude, we have shown that for any $\epsilon > 0$, we can choose sufficiently large $d_{\mathcal{X}} = d_{\mathcal{X}}(\mu, \Psi^\dagger, \epsilon)$, $d_{\mathcal{Y}} = d_{\mathcal{Y}}(\mu, \Psi^\dagger, \epsilon)$, a sufficient requisite amount of data $N = N(\mu, \Psi^\dagger, \epsilon)$, and we can construct a ReLU neural network $\psi : \mathbb{R}^{d_{\mathcal{X}}} \rightarrow \mathbb{R}^{d_{\mathcal{Y}}}$, such that the resulting PCA-Net $\Psi = \mathcal{D}_{\mathcal{Y}} \circ \psi \circ \mathcal{E}_{\mathcal{X}}$, satisfies

$$\begin{aligned} \mathbb{E}_{u \sim \mu} \left[\|\Psi^\dagger - \Psi\|_{\mathcal{Y}}^2 \right]^{1/2} &\leq \widehat{\mathcal{E}}_{\mathcal{Y}} + \widehat{\mathcal{E}}_{\mathcal{X}}^* \\ &\leq \widehat{\mathcal{E}}_{\mathcal{Y}} + \text{Lip}(\Psi^*) \widehat{\mathcal{E}}_{\mathcal{X}} + \widehat{\mathcal{E}}_{\Psi^*} + \mathbb{E}_{u \sim \mu} \left[\|\Psi^\dagger(u) - \Psi^*(u)\|_{\mathcal{Y}}^2 \right]^{1/2}, \end{aligned}$$

where the expectation of the terms on the right is bounded by (43),(44), (45) and (46). Taking expectation over the data $u_1, \dots, u_N \sim \mu^{\otimes N}$, we conclude that the constructed PCA-Net satisfies

$$\mathbb{E}_{\{u_j\} \sim \mu^{\otimes N}} \mathbb{E}_{u \sim \mu} \left[\|\Psi^\dagger - \Psi\|_{\mathcal{Y}}^2 \right]^{1/2} \leq \epsilon.$$

This concludes our proof of the universal approximation property of PCA-Nets. \blacksquare

Appendix C. Curse of Dimensionality

Our goal in this appendix is to prove Theorem 9. To this end, it suffices to prove the following

Claim 1 *Let \mathcal{X}, \mathcal{Y} be Hilbert spaces, with $\dim(\mathcal{X}) = \infty$, $\dim(\mathcal{Y}) \geq 1$. Let $\mu \in \mathcal{P}(\mathcal{X})$ be a non-degenerate Gaussian measure. For any $\gamma > 0$, there exists a k -times Fréchet differentiable operator $\Psi_\gamma^\dagger : \mathcal{X} \rightarrow \mathcal{Y}$ and a constant $c_\gamma > 0$, such that for any PCA-Net $\Psi = \mathcal{D}_{\mathcal{Y}} \circ \psi \circ \mathcal{E}_{\mathcal{X}}$, we have*

$$\mathbb{E}_{u \sim \mu} \left[\|\Psi_\gamma^\dagger(u) - \Psi(u)\|_{\mathcal{Y}}^2 \right] \geq c_\gamma \text{size}(\psi)^\gamma. \quad (47)$$

To prove this claim, it is sufficient to prove that such $\Psi_\gamma^\dagger : \mathcal{X} \rightarrow \mathcal{Y}$ exists with the *additional constraint* that $\text{Im}(\Psi_\gamma^\dagger)$ belongs to a one-dimensional subspace of \mathcal{Y} . In this case, we can clearly identify Ψ_γ^\dagger with a mapping $\mathcal{X} \rightarrow \mathbb{R}$, and thus, our goal is to prove the proposition in the (most constrained) case, where the output space $\mathcal{Y} = \mathbb{R}$ is one-dimensional. This will be the goal of the remainder of this section.

C.1 A lower bound in finite dimensions

We let $F_{k,d} := \{f : \mathbb{R}^d \rightarrow \mathbb{R} \mid \|f\|_{W^{k,\infty}} \leq 1\}$ denote the unit ball in the space of k -times weakly differentiable functions on \mathbb{R}^d , and where

$$\|f\|_{W^{k,\infty}} = \max_{|\alpha| \leq k} \|D^\alpha f\|_{L^\infty(\mathbb{R}^d)} < \infty.$$

In the following, we denote by ψ_θ an arbitrary, but fixed, feedforward neural network architecture (as defined in Section 2.3), with a total of W non-zero parameters collected in a vector $\theta \in \mathbb{R}^W$ (concatenating all weights and biases). We first recall the following result, which is an immediate consequence of (Achour et al., 2022, Corollary 1 and Lemma 2):

Lemma 35 *Fix $k \in \mathbb{N}$. There exists $c_{k,d} > 0$ and $W_{\min} \in \mathbb{N}$ depending only on k and d , such that for any $W \geq W_{\min}$, we have the following lower bound:*

$$\sup_{f \in F_{k,d}} \inf_{\theta \in \mathbb{R}^W} \|f - \psi_\theta\|_{L^2([0,1]^d)} \geq c_{k,d} W^{-3k/d}. \quad (48)$$

We note the close connection of the above lower bound (48) and the claimed bound (47). Indeed, choosing d sufficiently large, we can ensure that $3k/d < \gamma$ and hence (48) implies a lower bound of the form $c_\gamma W^{-\gamma}$. The main difference here is that (a) the architecture of ψ_θ , including its potentially sparse structure, is *fixed* in (48) and (b) the estimate (48) involves a supremum over all $f \in F_{k,d}$, whereas (47) asserts a similar bound for *one specific*, fixed operator and an arbitrary size of the network ψ . The main challenge will be to prove Proposition 13, repeated below, which gives a relevant lower bound in finite dimensions:

Proposition 36 *Fix $k, d \in \mathbb{N}$. There exists $c_{k,d} > 0$ depending only on k and d , $f \in F_{k,d}$ and an absolute constant $\lambda > 0$ independent of both d and k , such that for any neural network ψ , we have the lower bound*

$$\|f - \psi\|_{L^2([0,1]^d)} \geq c_{k,d} \text{size}(\psi)^{-\lambda k/d}. \quad (18)$$

Let us first show that Proposition 13 implies Claim 1.

Proof [Proof of Claim 1] Let a rate $\gamma > 0$ and a smoothness parameter $k \in \mathbb{N}$ be given. Choose a dimension d such that $\gamma > \lambda k/d$, where λ is the universal constant of Proposition 13. Let $f \in F_{k,d}$ be a function satisfying the lower bound in (18).

We recall that the underlying measure $\mu \in \mathcal{P}(\mathcal{X})$ is a non-degenerate Gaussian measure. Let $\lambda_1 \geq \lambda_2 \geq \dots \geq 0$ and $\phi_1, \phi_2, \dots \in \mathcal{X}$ denote the ordered eigenvalues and eigenfunctions of the covariance operator $\Sigma = \mathbb{E}_{u \sim \mu}[u \otimes u]$. In particular, since μ is non-degenerate, all PCA eigenvalues $\lambda_j > 0$ are positive. With the dimension d chosen above, let $\mathcal{E}_{\mathcal{X}}^{\text{opt}} : \mathcal{X} \rightarrow \mathbb{R}^d$ denote the optimal PCA encoding, $\mathcal{E}_{\mathcal{X}}^{\text{opt}}(u) = (\langle u, \phi_1 \rangle, \dots, \langle u, \phi_d \rangle)$. We denote by $\mathcal{D}_{\mathcal{X}}^{\text{opt}} : \mathbb{R}^d \rightarrow \mathcal{X}$ the PCA decoder, and by $P_{\leq d} = \mathcal{D}_{\mathcal{X}}^{\text{opt}} \circ \mathcal{E}_{\mathcal{X}}^{\text{opt}}$ the PCA projection onto the first d PCA eigenfunctions. We now define

$$\Psi_\gamma^\dagger(u) := f(\mathcal{E}_{\mathcal{X}}^{\text{opt}}(u)). \quad (49)$$

We clearly have $\Psi_\gamma^\dagger \in \mathcal{C}^k$. We claim that Ψ_γ^\dagger also satisfies the asserted lower bound of Claim 1. In particular, we aim to show that there exists a constant $c_\gamma > 0$, such that for any PCA-Net $\Psi = \mathcal{D}_y \circ \psi \circ \mathcal{E}_x$, we have

$$\mathbb{E}_{u \sim \mu} \left[\|\Psi_\gamma^\dagger(u) - \Psi(u)\|_y^2 \right] \geq c_\gamma \text{size}(\psi)^\gamma. \quad (50)$$

Let $u \sim \mu$ be a random variable. Since μ is Gaussian, we can expand u in its Karhunen-Loeve expansion $u = \bar{u} + \sum_{j=1}^{\infty} \sqrt{\lambda_j} Z_j \phi_j$, where the Z_j are iid standard Gaussian random

variables and $\bar{u} = \mathbb{E}_{u \sim \mu}[u]$ is deterministic. In particular, this allows us to express u in the form $u = u_1 + u_2$, where $u_1 \sim \mu_{\leq d}$ and $u_2 \sim \mu_{> d}$ are independent Gaussian random variables with $\mu_{\leq d} := P_{\leq d, \#} \mu$ and $\mu_{> d} := P_{> d, \#} \mu$ denote the push-forward by projecting onto the relevant components.

Let now $\Psi = \psi \circ \mathcal{E}_{\mathcal{X}}$ be a PCA-Net with values in \mathbb{R} , and $\psi : \mathbb{R}^{d_{\mathcal{X}}} \rightarrow \mathbb{R}$ a neural network. Note that the decoder $\mathcal{D}_{\mathcal{Y}}$ is trivial when $\mathcal{Y} = \mathbb{R}$, and may wlog be assumed to be given by multiplication by 1. By definition of $\Psi_{\gamma}^{\dagger}(u)$, we have $\Psi_{\gamma}^{\dagger}(u) = \Psi_{\gamma}^{\dagger}(u_1)$ under the decomposition $u = u_1 + u_2$ from above. Hence, we obtain

$$\begin{aligned} \mathbb{E}_{u \sim \mu} \left[|\Psi_{\gamma}^{\dagger}(u) - \Psi(u)|^2 \right] &= \mathbb{E}_{u_2 \sim \mu_{> d}} \mathbb{E}_{u_1 \sim \mu_{\leq d}} \left[|\Psi_{\gamma}^{\dagger}(u_1) - \Psi(u_1 + u_2)|^2 \right] \\ &\geq \inf_{u_2} \mathbb{E}_{u_1 \sim \mu_{\leq d}} \left[|\Psi_{\gamma}^{\dagger}(u_1) - \Psi(u_1 + u_2)|^2 \right], \end{aligned} \quad (51)$$

where the infimum is over all $u_2 \in \text{Im}(P_{> d}) = \text{span}\{\phi_j \mid j > d\}$. Let us fix such u_2 for the moment. Identifying u_1 with its unique encoding $\xi = \mathcal{E}_{\mathcal{X}}^{\text{opt}}(u_1) \in \mathbb{R}^d$, we note that, by definition of Ψ_{γ}^{\dagger} , we have $\Psi_{\gamma}^{\dagger}(u_1) = f(\xi)$. Furthermore, we can write

$$\Psi(u_1 + u_2) = \psi(A\xi + b) =: \psi_{A,b}(\xi),$$

where A is a matrix representing the linear map $\mathcal{E}_{\mathcal{X}} \circ \mathcal{D}_{\mathcal{X}}^{\text{opt}} : \mathbb{R}^d \rightarrow \mathbb{R}^d$, and $b := \mathcal{E}_{\mathcal{X}}(u_2) \in \mathbb{R}^d$ is a fixed vector depending on the fixed choice of u_2 . Substitution now implies that

$$\mathbb{E}_{u_1 \sim \mu_{\leq d}} \left[|\Psi_{\gamma}^{\dagger}(u_1) - \Psi(u_1 + u_2)|^2 \right] = \mathbb{E}_{\xi \sim \mathcal{N}(m_{\leq d}, \Sigma_{\leq d})} \left[|f(\xi) - \psi_{A,b}(\xi)|^2 \right],$$

where $\mathcal{N}(m_{\leq d}, \Sigma_{\leq d}) = \mathcal{E}_{\mathcal{X}, \#}^{\text{opt}} \mu_{\leq d}$ is a Gaussian distribution with (encoded) mean $m_{\leq d} := \mathcal{E}_{\mathcal{X}}(\bar{u})$ and diagonal covariance matrix $\Sigma_{\leq d} = \text{diag}(\lambda_1, \dots, \lambda_d)$. Since by assumption, μ is a non-degenerate Gaussian measure, it follows that all PCA eigenvalues $\lambda_j > 0$ for all j . In particular, this implies that $\mathcal{N}(m_{\leq d}, \Sigma_{\leq d})$ has a strictly positive density over the unit cube $[0, 1]^d$. Thus, there exists a constant c'_d , such that

$$\mathbb{E}_{\xi \sim \mathcal{N}(m_{\leq d}, \Sigma_{\leq d})} \left[|f(\xi) - \psi_{A,b}(\xi)|^2 \right] \geq c'_d \|f - \psi_{A,b}\|_{L^2([0,1]^d)}^2.$$

By construction of f (cp. Proposition 13) and the fact that $\psi_{A,b}$ is a neural network, the latter quantity can be bounded from below by $c_{k,d} \text{size}(\psi_{A,b})^{-\lambda k/d}$. Furthermore, it is easy to see that $\psi_{A,b}$ can be represented by a neural network with size bound

$$\text{size}(\psi_{A,b}) \leq \text{size}(\psi) + \underbrace{\text{size}(\psi)d}_{\cong \|A\|_0} + \underbrace{d}_{\cong \|b\|_0} \leq 2(d+1)\text{size}(\psi).$$

In particular, combining the above estimates and recalling that $\lambda k/d < \gamma$, it follows that there exists a constant $c_{\gamma} > 0$, depending only on k , d and μ , such that

$$\mathbb{E}_{u_1 \sim \mu_{\leq d}} \left[|\Psi_{\gamma}^{\dagger}(u_1) - \Psi(u_1 + u_2)|^2 \right] \geq c_{\gamma} \text{size}(\psi)^{-\gamma},$$

for any choice of ψ . The right-hand side is independent of u_2 . Hence, taking the infimum over all u_2 on the left and recalling (51), we obtain

$$\mathbb{E}_{u \sim \mu} \left[|\Psi_{\gamma}^{\dagger}(u) - \Psi(u)|^2 \right] \geq c_{\gamma} \text{size}(\psi)^{-\gamma},$$

as claimed. ■

It finally remains to prove Proposition 13. This is the subject of the next section.

C.2 Proof of Proposition 13

The goal of this section is to prove the existence of $f \in F_{k,d}$ and a universal constant $\lambda > 0$, for which the lower bound (18) holds for any neural network ψ . To construct f , we adapt an idea of Yarotsky (Yarotsky, 2017, Section 4). In fact, the present analysis is very closely related to a refinement of Yarotsky’s result (Yarotsky, 2017, Theorem 5), which is derived in the forthcoming work (Lanthaler and Stuart, 2023). The main difference here is that the present work develops a lower bound for the L^2 -norm, while (Lanthaler and Stuart, 2023) focuses on the supremum norm.

We now follow (Yarotsky, 2017), and define $\mathcal{N}(f, \epsilon)$ as the minimal number of hidden computation units in a ReLU network ψ which achieves approximation accuracy $\|f - \psi\|_{L^2([0,1]^d)} \leq \epsilon$. We note that Lemma 35 implies in particular that for a fixed architecture ψ_θ with W weights, the upper bound

$$\sup_{f \in F_{k,d}} \inf_{\theta \in \mathbb{R}^W} \|f - \psi_\theta\|_{L^2([0,1]^d)} \leq \epsilon, \quad (52)$$

requires $W \geq c\epsilon^{-d/3k}$ for some constant $c = c(d, k)$. Repeating the proof of (Yarotsky, 2017, Lemma 3), the following is then immediate:

Lemma 37 *Fix d, k . For any $\epsilon > 0$, there exists $f_\epsilon \in F_{k,d}$ such that $\mathcal{N}(f_\epsilon, \epsilon) \geq c_1\epsilon^{-d/12k}$, with some constant $c_1 = c_1(d, k) > 0$.*

The idea is that any (sparse) neural network ψ with at most N computation units can be embedded in a maximally connected, enveloping architecture ψ_θ , $\theta \in \mathbb{R}^W$, with $W = O(N^4)$ weights. Thus, if ψ achieves error ϵ , then (52) applied to the enveloping network ψ_θ implies that $N^4 \gtrsim W \gtrsim \epsilon^{-d/3k}$, and hence $N \gtrsim \epsilon^{-d/12k}$. Fixing $\epsilon > 0$ and choosing f_ϵ to be an optimizer of (52), the claimed lower bound on $\mathcal{N}(f_\epsilon, \epsilon)$ follows. Further details can be found in (Yarotsky, 2017, Pf. of Lemma 3)

Lemma 37 will be our main technical tool for the proof of Proposition 13. We also recall that $\mathcal{N}(f, \epsilon)$ satisfies the following properties, (Yarotsky, 2017, eq. (42)–(44)):

$$\mathcal{N}(af, |a|\epsilon) = \mathcal{N}(f, \epsilon), \quad (53a)$$

$$\mathcal{N}(f \pm g, \epsilon + \|g\|_{L^\infty}) \leq \mathcal{N}(f, \epsilon), \quad (53b)$$

$$\mathcal{N}(f_1 \pm f_2, \epsilon_1 + \epsilon_2) \leq \mathcal{N}(f_1, \epsilon_1) + \mathcal{N}(f_2, \epsilon_2). \quad (53c)$$

We are now ready to prove Proposition 13.

Proof (Proposition 13) We recall that the goal is to show that there exists $f \in F_{k,d}$ such that

$$\|f - \psi\|_{L^2([0,1]^d)} \geq c_{k,d} \text{size}(\psi)^{-\lambda k/d}.$$

for any neural network ψ . We will achieve this by constructing $f \in F_{k,d}$ of the form

$$f = \sum_{s=1}^{\infty} a_s f_s, \quad \text{with } f_s \in F_{k,d} \forall s \in \mathbb{N}. \quad (54)$$

To this end, we define a rapidly converging sequence $\epsilon_s \rightarrow 0$, by $\epsilon_1 = 1/4$ and recursively set $\epsilon_{s+1} = \epsilon_s^2$. Note by iteration, that $\epsilon_s = 2^{-2^s}$. We also define $a_s := \frac{1}{2}\epsilon_s^{1/2} = \frac{1}{2}\epsilon_{s-1}$. For later reference, since $\epsilon_s \leq 1/2$ for all s , we have

$$\sum_{r=s+1}^{\infty} a_r = \frac{1}{2} \left(\epsilon_s + \epsilon_s^2 + \epsilon_s^{2^2} + \dots \right) \leq \frac{1}{2} \epsilon_s \sum_{r=0}^{\infty} 2^{-r} = \epsilon_s. \quad (55)$$

We also note that $a_s \leq 2^{-s}$. If the function f is of the above form (54) then,

$$\|f\|_{W^{k,\infty}} \leq \sum_{s=1}^{\infty} a_s \|f_s\|_{W^{k,\infty}} \leq 1,$$

so that $f \in F_{k,d}$. This is irrespective of the specific choice of $f_s \in F_{k,d}$. We are going to determine the f_s recursively, formally starting from the empty sum, i.e. $f_0 = 0$. We fix $2\lambda := 1/25$ throughout the following construction. In the recursive step, given f_1, \dots, f_{s-1} , we distinguish two cases:

Case 1: Assume that

$$\mathcal{N} \left(\sum_{r=1}^{s-1} a_r f_r, 2\epsilon_s \right) \geq \epsilon_s^{-2\lambda d/k}.$$

In this case, we set $f_s := 0$, and trivially obtain

$$\mathcal{N} \left(\sum_{r=1}^s a_r f_r, 2\epsilon_s \right) \geq \epsilon_s^{-2\lambda d/k}. \quad (56)$$

Case 2: In the other case, we have

$$\mathcal{N} \left(\sum_{r=1}^{s-1} a_r f_r, 2\epsilon_s \right) < \epsilon_s^{-2\lambda d/k}.$$

By (53c), and the upper bound above:

$$\begin{aligned} \mathcal{N} \left(\sum_{r=1}^s a_r f_r, 2\epsilon_s \right) &\geq \mathcal{N}(a_s f_s, 4\epsilon_s) - \mathcal{N} \left(\sum_{r=1}^{s-1} a_r f_r, 2\epsilon_s \right) \\ &\geq \mathcal{N}(a_s f_s, 4\epsilon_s) - \epsilon_s^{-2\lambda d/k}. \end{aligned} \quad (57)$$

Recalling Lemma 37, there exists $f_s \in F_{k,d}$, such that

$$\stackrel{(53a)}{\mathcal{N}(a_s f_s, 4\epsilon_s)} \stackrel{\downarrow}{=} \mathcal{N}(f_s, 4\epsilon_s/a_s) = \mathcal{N}(f_s, 4\epsilon_s^{1/2}) \geq c_{d,k} \epsilon_s^{-d/24k}. \quad (58)$$

This defines our recursive choice of f_s in Case 2.

Our next goal is to show that with this recursive choice of f_s , (56) holds for all sufficiently small $\epsilon_s \leq \bar{\epsilon}$. This is trivial in Case 1. To prove it for Case 2, we define $\bar{\epsilon} = \bar{\epsilon}(d, k) > 0$, as the unique solution of

$$c_{d,k} \bar{\epsilon}^{-d/24k} = 2\bar{\epsilon}^{-2\lambda d/k}.$$

Then, by (57) and (58), and with our choice of $2\lambda := 1/25 < 1/24$, it follows that for all $\epsilon_s \leq \bar{\epsilon}$, we have

$$c_{d,k}\epsilon_s^{-d/24k} \geq 2\epsilon_s^{-2\lambda d/k},$$

since the exponent on the left is strictly more negative than the exponent on the right. Hence, (57) and (58) imply that

$$\mathcal{N}\left(\sum_{r=1}^s a_r f_r, 2\epsilon_s\right) \geq c_{d,k}\epsilon_s^{-d/24k} - \epsilon_s^{2\lambda d/k} \geq \epsilon_s^{2\lambda d/k}, \quad (59)$$

for all $\epsilon_s \leq \bar{\epsilon}$. This demonstrates inequality (56) for $\epsilon_s \leq \bar{\epsilon}$ also in Case 2. The above recursive construction thus yields a sequence $f_1, f_2, \dots \in F_{k,d}$, such that (59) holds for any $s \in \mathbb{N}$ such that $\epsilon_s \leq \bar{\epsilon}$. Define $f := \sum_{s=1}^{\infty} a_s f_s$. We claim that for *any* $\epsilon \leq \bar{\epsilon}$ we have

$$\mathcal{N}(f, \epsilon) \geq \epsilon^{-\lambda d/k}.$$

To this end, we fix $\epsilon \leq \bar{\epsilon}$. Choose $s \in \mathbb{N}$, such that $\epsilon_s \leq \bar{\epsilon}$ and $\epsilon_s^2 \leq \epsilon \leq \epsilon_s$. Then,

$$\begin{aligned} \mathcal{N}(f, \epsilon) &\geq \mathcal{N}(f, \epsilon_s) = \mathcal{N}\left(\sum_{r=1}^{\infty} a_r f_r, \epsilon_s\right) \\ &\stackrel{(53b)}{\geq} \mathcal{N}\left(\sum_{r=1}^s a_r f_r, \epsilon_s + \left\|\sum_{r=s+1}^{\infty} a_r f_r\right\|_{L^\infty}\right) \\ &\geq \mathcal{N}\left(\sum_{r=1}^s a_r f_r, \epsilon_s + \sum_{r=s+1}^{\infty} a_r\right) \\ &\stackrel{(55)}{\geq} \mathcal{N}\left(\sum_{r=1}^s a_r f_r, 2\epsilon_s\right) \\ &\stackrel{(59)}{\geq} \epsilon_s^{-2\lambda d/k} \geq \epsilon^{-\lambda d/k}, \end{aligned}$$

where the last inequality follows from $\epsilon_s \leq \epsilon^{1/2}$. The claim of Proposition 13 thus follows for all $\epsilon \leq \bar{\epsilon}$, and with universal constant $\lambda = 1/50$. \blacksquare

Appendix D. Proof of Proposition 15

Proof We first assume that $D = \mathbb{T}^n$ is the periodic, n -dimensional torus. To prove the theorem in this case, let $d \in \mathbb{N}$ be given, and let $\phi_1, \dots, \phi_d, \dots \in L^2(\mathbb{T}^n; \mathbb{R})$ an enumeration of the standard Fourier basis; to be more precise, we require that for all $k \in \mathbb{N}$, $\Delta\phi_k = \gamma_k\phi_k$ is an eigenfunction of the periodic Laplacian, and that the eigenvalues $\gamma_1 \leq \gamma_2 \leq \dots$ are monotonically ordered. To give an example for $n = 1$, a suitable choice is $\phi_1(x) \equiv 1/\sqrt{2\pi}$, $\phi_2(x) = \pi^{-1/2} \cos(x)$, $\phi_3(x) = \pi^{-1/2} \sin(x)$, $\phi_4(x) = \pi^{-1/2} \cos(2x)$, etc. We note that there exist constants $A = A(n)$, $B = B(n) > 0$, such that

$$Ak^{1/n} \leq \sqrt{\gamma_k} \leq Bk^{1/n}, \quad \forall k > 1.$$

Let $V_d := \text{span}\{\phi_1, \dots, \phi_d\}$. We now use the fact that

$$u - \Pi_{V_d} u = \sum_{k>d} \langle u, \phi_k \rangle \phi_k,$$

and combine this with the following two-sided bounds,

$$\left\| \sum_{k>d} c_k (-\Delta)^{s/2} \phi_k \right\|_{L^2}^2 \lesssim \left\| \sum_{k>d} c_k \phi_k \right\|_{H^s}^2 \lesssim \left\| \sum_{k>d} c_k (-\Delta)^{s/2} \phi_k \right\|_{L^2}^2,$$

and the identity,

$$\left\| \sum_{k>d} c_k (-\Delta)^{s/2} \phi_k \right\|_{L^2}^2 = \sum_{k>d} \gamma_k^s |c_k|^2,$$

which are valid for any sequence (c_k) for which the last expression is finite. Using these bounds with $c_k = \langle u, \phi_k \rangle$, we can estimate (with C changing from line to line):

$$\begin{aligned} \|u - \Pi_{V_d} u\|_{H^s}^2 &= \left\| \sum_{k>d} c_k \phi_k \right\|_{H^s}^2 \leq C \left\| \sum_{k>d} c_k (-\Delta)^{s/2} \phi_k \right\|_{L^2}^2 \\ &= C \sum_{k>d}^{\infty} \gamma_k^s |\langle u, \phi_k \rangle|^2 = C \sum_{k>d}^{\infty} \frac{\gamma_k^{s+\zeta}}{\gamma_k^\zeta} |\langle u, \phi_k \rangle|^2 \\ &\leq C \frac{1}{\gamma_d^\zeta} \sum_{k>d} \gamma_k^{s+\zeta} |\langle u, \phi_k \rangle|^2 \leq \frac{C}{d^{2\zeta/n}} \|u\|_{H^{s+\zeta}}^2, \end{aligned}$$

and where the constant $C = C(n, s, \zeta)$ only depends on n , s and ζ , but is independent of d . In particular, this implies that

$$R_d(\mu) = \min_{V_d} \mathbb{E}_{u \sim \mu} [\|u - \Pi_{V_d} u\|_{H^s}^2] \leq C d^{-2\zeta/n} \mathbb{E}_{u \sim \mu} [\|u\|_{H^{s+\zeta}}^2].$$

To prove the general case, we note that for any Lipschitz domain $D \subset \mathbb{R}^n$, there exists a bounded, linear extension operator $L : H^s(D; \mathbb{R}^{n'}) \hookrightarrow H^s(\mathbb{T}^n; \mathbb{R}^{n'})$, which also defines a bounded mapping, $L : H^{s+\zeta}(D; \mathbb{R}^{n'}) \hookrightarrow H^{s+\zeta}(\mathbb{T}^n; \mathbb{R}^{n'})$ (see (Kovachki et al., 2021, Lemma B.3)), with bounded left-inverse $L^+ : H^s(\mathbb{T}^n; \mathbb{R}^{n'}) \rightarrow H^s(D; \mathbb{R}^{n'})$, and such that $L^+ L = \text{id}$. Here, L can be thought of as an extension operator, which extends a function $u : D \rightarrow \mathbb{R}^{n'}$ to a periodic function with domain a bounding box $B \supset D$. And L^+ can be thought of as a restriction operator, restricting $u : B \rightarrow \mathbb{R}^{n'}$ to $u|_D : D \rightarrow \mathbb{R}^{n'}$.

The above argument in the periodic case shows that the optimal d -dimensional PCA subspace $\bar{V}_d \subset H^s(\mathbb{T}^n; \mathbb{R}^{n'})$ of the push-forward measure $L\#\mu$ (which is defined on periodic functions) satisfies

$$\mathbb{E}_{u \sim L\#\mu} \left[\|u - \Pi_{\bar{V}_d} u\|_{H^s}^2 \right] \leq C d^{-2\zeta/n} \mathbb{E}_{u \sim L\#\mu} [\|u\|_{H^{s+\zeta}}^2].$$

Since the support of $\text{supp}(L\#\mu) \subset \text{Im}(L)$ is contained in the image of L , it follows that $\bar{V}_d \subset \text{Im}(L)$. Since L is an isomorphism onto its image, there exists a d -dimensional

subspace $V_d \subset H^s(D; \mathbb{R}^{n'})$, such that $L(V_d) = \bar{V}_d$. Notice that for any $\bar{u} \in H^s(\mathbb{T}^n; \mathbb{R}^{n'})$, we have

$$L^+ \Pi_{\bar{V}_d} \bar{u} \in L^+ \bar{V}_d = L^+ L V_d = V_d.$$

Using the fact that $\Pi_{V_d} : H^s(D; \mathbb{R}^{n'}) \rightarrow H^s(D; \mathbb{R}^{d'})$ is the orthogonal projection onto V_d , we can thus estimate, for any $\bar{u} \in H^s(\mathbb{T}^n; \mathbb{R}^{n'})$,

$$\|L^+ \bar{u} - \Pi_{V_d} L^+ \bar{u}\|_{H^s} = \inf_{v \in V_d} \|L^+ \bar{u} - v\|_{H^s} \leq \|L^+ \bar{u} - L^+ \Pi_{\bar{V}_d} \bar{u}\|_{H^s}.$$

We thus conclude from $\mu = \text{id}_{\#} \mu = L_{\#}^+ L_{\#} \mu$, that

$$\begin{aligned} \mathbb{E}_{u \sim \mu} [\|u - \Pi_{V_d} u\|_{H^s}^2] &= \mathbb{E}_{u \sim L_{\#}^+ L_{\#} \mu} [\|u - \Pi_{V_d} u\|_{H^s}^2] \\ &= \mathbb{E}_{\bar{u} \sim L_{\#} \mu} [\|L^+ \bar{u} - \Pi_{V_d} L^+ \bar{u}\|_{H^s}^2] \\ &\leq \mathbb{E}_{\bar{u} \sim L_{\#} \mu} [\|L^+ \bar{u} - L^+ \Pi_{\bar{V}_d} \bar{u}\|_{H^s}^2] \\ &\leq \|L^+\|^2 \mathbb{E}_{\bar{u} \sim L_{\#} \mu} [\|\bar{u} - \Pi_{\bar{V}_d} \bar{u}\|_{H^s}^2] \\ &\leq C \|L^+\|^2 d^{-2\zeta/n} \mathbb{E}_{\bar{u} \sim L_{\#} \mu} [\|\bar{u}\|_{H^{s+\zeta}}^2] \\ &= C \|L^+\|^2 d^{-2\zeta/n} \mathbb{E}_{u \sim \mu} [\|Lu\|_{H^{s+\zeta}}^2] \\ &\leq C \|L^+\|^2 \|L\|^2 d^{-2\zeta/n} \mathbb{E}_{u \sim \mu} [\|u\|_{H^{s+\zeta}}^2] \end{aligned}$$

Absorbing the additional factor $\|L^+\|^2 \|L\|^2$ in the constant C , this implies the claim for general Lipschitz domains $D \subset \mathbb{R}^n$. \blacksquare

Appendix E. Proofs for Darcy flow

The following subsections will discuss the derivation of the error and complexity bounds of Theorem 18.

E.1 Proof of Lemma 21

Thm. 1.3 of (Cohen et al., 2011) shows that if $a(x; \mathbf{z})$ is of the parametrized form (22), and if the sequence $(\gamma_\ell)_{\ell \in \mathbb{N}}$ belongs to $\ell^p(\mathbb{N})$ for some $0 < p < 1$, then the family of Taylor coefficients $(t_\nu)_\nu$ is also ℓ^p -summable. In addition, if Λ_m denotes the set of the largest m Taylor coefficients, then

$$\sup_{\mathbf{z} \in U} \|\mathcal{F}(\mathbf{z}) - \sum_{\nu \in \Lambda_m} t_\nu \mathbf{z}^\nu\|_y \leq \|(\|t_\nu\|_y)_\nu\|_{\ell^p} m^{-s}, \quad s := \frac{1}{p} - 1.$$

By assumption, we have $|\gamma_\ell| \lesssim \ell^{-1-\alpha}$, and hence, γ_ℓ belongs to ℓ^p whenever $p(1+\alpha) > 1$. Hence, for any $s = p^{-1} - 1 < \alpha$, there exists $C > 0$, possibly depending on s , such that

$$\sup_{\mathbf{z} \in U} \|\mathcal{F}(\mathbf{z}) - \sum_{\nu \in \Lambda_m} t_\nu \mathbf{z}^\nu\|_y \leq C m^{-s}.$$

Writing $s = \alpha - \eta$ for $\eta > 0$ yields the claim of Lemma 21.

E.2 Proof of Proposition 22

The goal of the present section is to prove Proposition 22, which we recall here:

Proposition 38 *Under the setting and prevailing assumptions of this section. For any $\eta > 0$, there exists a constant $C = C(\mathcal{F}, \mu, \eta) > 0$, such that*

$$\mathcal{R}_{d\mathcal{X}}^{\text{opt}}(\mu) \leq Cd\mathcal{X}^{-2\alpha-1}, \quad \mathcal{R}_{d\mathcal{Y}}^{\text{opt}}(\Psi_{\#}^{\dagger}\mu) \leq Cd\mathcal{Y}^{-2\alpha+\eta}.$$

Proof We recall that by assumption, the input measure is the law of random functions a of the form

$$a = \bar{a} + \sum_{\ell=1}^{\infty} \gamma_{\ell} z_{\ell} \rho_{\ell},$$

where $\rho_{\ell} \in \mathcal{X}$ are orthonormal, $\gamma_{\ell} \leq C\ell^{-1-\alpha}$, and $|z_{\ell}| \leq 1$ for all ℓ . We estimate the optimal PCA projection error $\mathcal{R}_d^{\text{opt}}(\mu)$ by comparing it against the (at most) d -dimensional projection P' onto $\text{span}\{\bar{a}, \rho_1, \dots, \rho_{d-1}\}$. This yields,

$$\begin{aligned} \mathcal{R}_d^{\text{opt}}(\mu) &= \min_{P \in \Pi_d} \mathbb{E}_{a \sim \mu} [\|a - Pa\|_{\mathcal{X}}^2] \leq \mathbb{E}_{a \sim \mu} [\|a - P'a\|_{\mathcal{X}}^2] \\ &= \mathbb{E}_z \sum_{\ell=d}^{\infty} \gamma_{\ell}^2 z_{\ell}^2 \leq \sum_{\ell=d}^{\infty} \gamma_{\ell}^2 \leq Cd^{-2\alpha-1}. \end{aligned}$$

To estimate the PCA projection error on \mathcal{Y} , we note that by Lemma 21, for any $\eta > 0$, there exists a constant $C > 0$, depending only on \mathcal{F} , μ and η , and there exist coefficients $t_{\nu} \in \mathcal{Y}$ (the Taylor coefficients of \mathcal{F}) and a family of multi-index sets Λ_m , with $|\Lambda_m| = m$ for all $m \in \mathbb{N}$, such that

$$\sup_{z \in U} \|\mathcal{F}(z) - \sum_{\nu \in \Lambda_m} t_{\nu} z^{\nu}\|_{\mathcal{Y}} \leq Cm^{-\alpha+\eta}.$$

Choosing $m = d$ to be the PCA dimension, and letting $P'' : \mathcal{Y} \rightarrow \mathcal{Y}$ denote the orthogonal projection onto $\text{span}\{t_{\nu} \mid \nu \in \Lambda_d\} \subset \mathcal{Y}$, it follows that

$$\begin{aligned} \mathcal{R}_d^{\text{opt}}(\Psi_{\#}^{\dagger}\mu) &= \min_{P \in \Pi_d} \mathbb{E}_{a \sim \mu} [\|\Psi^{\dagger}(a) - P\Psi^{\dagger}(a)\|_{\mathcal{Y}}^2] \leq \mathbb{E}_{a \sim \mu} [\|\Psi^{\dagger}(a) - P''\Psi^{\dagger}(a)\|_{\mathcal{Y}}^2] \\ &= \mathbb{E}_z [\|\mathcal{F}(z) - P''\mathcal{F}(z)\|_{\mathcal{Y}}^2] \leq \mathbb{E}_z [\|\mathcal{F}(z) - \sum_{\nu \in \Lambda_m} t_{\nu} z^{\nu}\|_{\mathcal{Y}}^2] \\ &\leq \sup_{z \in U} \|\mathcal{F}(z) - \sum_{\nu \in \Lambda_m} t_{\nu} z^{\nu}\|_{\mathcal{Y}}^2 \leq Cd^{-2\alpha+2\eta}. \end{aligned}$$

The claimed result now easily follows by noticing that $\eta > 0$ in the last estimate is arbitrary; we can thus repeat the argument with $\eta/2$ in place of η to arrive at the claimed result. \blacksquare

E.3 Proof of Theorem 18

Proof Our goal is to estimate the following PCA-approximation error,

$$\widehat{\mathcal{E}} = \mathbb{E}_{a \sim \mu} \left[\|\Psi(a) - \Psi^\dagger(a)\|_{\mathcal{Y}}^2 \right]^{1/2},$$

where $\mathcal{Y} = H_0^1(D)$ is the output function space. Let $\phi_1^{\mathcal{X}}, \dots, \phi_{d_{\mathcal{X}}}^{\mathcal{X}} \in \mathcal{X}$ denote the empirical PCA-basis on \mathcal{X} . Let $\mathcal{E}_{\mathcal{X}} : \mathcal{X} \rightarrow \mathbb{R}^{d_{\mathcal{X}}}$ denote the PCA-encoder $a \mapsto \langle a, \phi_j^{\mathcal{X}} \rangle_{j=1}^{d_{\mathcal{X}}}$ on \mathcal{X} . According to Lemma 32, the PCA-Net approximation error can be bounded by $\widehat{\mathcal{E}} \leq \widehat{\mathcal{E}}_{\mathcal{X}}^* + \widehat{\mathcal{E}}_{\mathcal{Y}}$, in terms of a reconstruction error on \mathcal{Y} ,

$$\widehat{\mathcal{E}}_{\mathcal{Y}} = \mathbb{E}_{w \sim \Psi_{\#}^\dagger \mu} \left[\|w - \mathcal{D}_{\mathcal{Y}} \circ \mathcal{E}_{\mathcal{Y}}(w)\|_{\mathcal{Y}}^2 \right]^{1/2},$$

and a neural network approximation error

$$\widehat{\mathcal{E}}_{\mathcal{X}}^* = \mathbb{E}_{a \sim \mu} \left[\|\psi \circ \mathcal{E}_{\mathcal{X}}(a) - \mathcal{E}_{\mathcal{Y}} \circ \Psi^\dagger(a)\|_{\ell^2}^2 \right]^{1/2}.$$

By the assumptions of Section 3.4.1, any $a \in \text{supp}(\mu)$ can be written in the form $a(\mathbf{z}) = a(\cdot; \mathbf{z}) = \bar{a} + \sum_{\ell=1}^{\infty} \gamma_{\ell} z_{\ell} \rho_{\ell}$, for an orthonormal set of functions $\rho_{\ell} \in \mathcal{X}$, and for coefficients $\mathbf{z} = (z_{\ell})_{\ell \in \mathbb{N}} \in U := [-1, 1]^{\mathbb{N}}$. Thus, $\widehat{\mathcal{E}}_{\mathcal{X}}^*$ can be estimated trivially by

$$\widehat{\mathcal{E}}_{\mathcal{X}}^* \leq \sup_{\mathbf{z} \in U} \left\| \psi(\langle a(\mathbf{z}), \phi_j^{\mathcal{X}} \rangle) - \mathcal{E}_{\mathcal{Y}} \circ \Psi^\dagger(a(\mathbf{z})) \right\|_{\ell^2}.$$

Let us define a mapping $\mathcal{F} : U \rightarrow \mathcal{Y}$, where $U = [-1, 1]^{\mathbb{N}}$, by

$$\mathcal{F}(\mathbf{z}) := \Psi^\dagger(\bar{a} + \sum_{\ell=1}^{\infty} \gamma_{\ell} z_{\ell} \rho_{\ell}).$$

We note that by Lemma 23, for any $\eta > 0$, and any $m \in \mathbb{N}$, there exists a neural network $\psi^* : [-1, 1]^m \rightarrow \mathbb{R}^{d_{\mathcal{Y}}}$, such that

$$\sup_{\mathbf{z} \in U} \|\mathcal{E}_{\mathcal{Y}} \circ \mathcal{F}(\mathbf{z}) - \psi^*(z_1, \dots, z_m)\|_{\ell^2(\mathbb{R}^{d_{\mathcal{Y}}})} \leq C m^{-\alpha + \eta/2}, \quad (60)$$

and with a constant $C = C(\mathcal{F}, \alpha, s) > 0$ (we choose $\eta/2$ rather than η for later convenience). Furthermore, we have the following complexity estimate:

$$\text{size}(\psi^*) \leq C m (\log(m)^2 + d_{\mathcal{Y}}), \quad \text{depth}(\psi^*) \leq C \log(m)^2.$$

Our next goal will be to find a map $\psi^{**} : \mathbb{R}^{d_{\mathcal{X}}} \rightarrow [-1, 1]^m$, such that for any $a = a(\mathbf{z})$, the composition

$$\psi : \mathbb{R}^{d_{\mathcal{X}}} \rightarrow \mathbb{R}^{d_{\mathcal{Y}}}, \quad \psi(w) := \psi^* \circ \psi^{**}(w),$$

provides a good approximation

$$\begin{aligned} \psi(\mathcal{E}_{\mathcal{X}}(a(\mathbf{z}))) &= \psi(\langle a(\mathbf{z}), \phi_1^{\mathcal{X}} \rangle, \dots, \langle a(\mathbf{z}), \phi_m^{\mathcal{X}} \rangle) \\ &\approx \psi^*(z_1, \dots, z_m) \approx \mathcal{E}_{\mathcal{Y}} \circ \mathcal{F}(\mathbf{z}). \end{aligned}$$

We note that by the orthonormality of the ρ_ℓ in \mathcal{X} , we have

$$z_\ell = \gamma_\ell^{-1} \langle a - \bar{a}, \rho_\ell \rangle,$$

so that in principle, the coefficients z_ℓ could be readily obtained from knowledge of $a = a(\mathbf{z})$. The main difficulty here is that the neural network ψ only has access to the PCA projections $\langle a, \phi_j^{\mathcal{X}} \rangle$, which allows us to reconstruct the orthogonal projection $\Pi_{\mathcal{X}} a := \mathcal{D}_{\mathcal{X}} \circ \mathcal{E}_{\mathcal{X}}(a)$ onto the empirical PCA subspace in \mathcal{X} , but it doesn't generally allow us to reconstruct a itself. Instead, we define

$$\tilde{z}_\ell := \gamma_\ell^{-1} \langle \Pi_{\mathcal{X}}(a - \bar{a}), \rho_\ell \rangle = \gamma_\ell^{-1} \left\{ \sum_{j=1}^{d_{\mathcal{X}}} \langle a, \phi_j^{\mathcal{X}} \rangle \langle \phi_j^{\mathcal{X}}, \rho_\ell \rangle \right\} - \gamma_\ell^{-1} \langle \Pi_{\mathcal{X}} \bar{a}, \rho_\ell \rangle.$$

We note that the map $\mathbb{R}^{d_{\mathcal{X}}} \rightarrow \mathbb{R}^m$, taking $\alpha = (\alpha_1, \dots, \alpha_{d_{\mathcal{X}}})$ with $\alpha_j := \langle a, \phi_j^{\mathcal{X}} \rangle$ and mapping it to $(\tilde{z}_1, \dots, \tilde{z}_m)$, with $\tilde{z}_\ell = \tilde{z}_\ell(\alpha)$ as above, defines an affine mapping. Based on these $\tilde{z}_\ell(\alpha)$, we now define $\psi^{**} : \mathbb{R}^{d_{\mathcal{X}}} \rightarrow [-1, 1]^m$, $\alpha \mapsto (\psi_1^{**}(\alpha), \dots, \psi_m^{**}(\alpha))$ by

$$\psi_\ell^{**}(\alpha) := \text{shrink}(\tilde{z}_\ell(\alpha)), \quad \text{where } \text{shrink}(z) := \begin{cases} 1 & z > 1, \\ z & z \in [-1, 1], \\ -1 & z < -1. \end{cases}$$

We point out that ψ^{**} can be represented by a ReLU neural network, with

$$\text{size}(\psi^{**}) \leq C m d_{\mathcal{X}}, \quad \text{depth}(\psi^{**}) \leq 2.$$

We now consider the composition $\psi(\alpha) := \psi^* \circ \psi^{**}(\alpha)$: Let $\mathbf{z} \in U$ be given. For the following estimate, we denote $\hat{z}_\ell := \text{shrink}(\tilde{z}_\ell(\alpha))$, with $\alpha \in \mathbb{R}^{d_{\mathcal{X}}}$ having components $\alpha_j = \langle a(\mathbf{z}), \phi_j^{\mathcal{X}} \rangle$. We note that, by definition, $\psi^{**}(\alpha) = (\hat{z}_1, \dots, \hat{z}_m)$. Let us denote $\hat{\mathbf{z}} = (\hat{z}_1, \dots, \hat{z}_m, 0, 0, \dots) \in U$. By construction, we then have

$$\begin{aligned} \|\psi(\mathcal{E}_{\mathcal{X}}(a(\mathbf{z}))) - \mathcal{E}_{\mathcal{Y}} \circ \mathcal{F}(\mathbf{z})\|_{\ell^2(\mathbb{R}^{d_{\mathcal{Y}}})} &= \|\psi^*(\hat{z}_1, \dots, \hat{z}_m) - \mathcal{E}_{\mathcal{Y}} \circ \mathcal{F}(\mathbf{z})\|_{\ell^2(\mathbb{R}^{d_{\mathcal{Y}}})} \\ &\leq \|\psi^*(\hat{z}_1, \dots, \hat{z}_m) - \mathcal{E}_{\mathcal{Y}} \circ \mathcal{F}(\hat{\mathbf{z}})\|_{\ell^2(\mathbb{R}^{d_{\mathcal{Y}}})} \\ &\quad + \|\mathcal{E}_{\mathcal{Y}} \circ \mathcal{F}(\hat{\mathbf{z}}) - \mathcal{E}_{\mathcal{Y}} \circ \mathcal{F}(\mathbf{z})\|_{\ell^2(\mathbb{R}^{d_{\mathcal{Y}}})} \\ &\leq C m^{-\alpha + \eta/2} + \|\mathcal{F}(\hat{\mathbf{z}}) - \mathcal{F}(\mathbf{z})\|_{\mathcal{Y}}, \end{aligned} \quad (61)$$

where we have made use of (60) in the last step, and the fact that the PCA encoding $\mathcal{E}_{\mathcal{Y}} : \mathcal{Y} \rightarrow \mathbb{R}^{d_{\mathcal{Y}}}$ is a contractive mapping. We note that

$$\begin{aligned} \|\mathcal{F}(\hat{\mathbf{z}}) - \mathcal{F}(\mathbf{z})\|_{\mathcal{Y}} &= \left\| \Psi^\dagger \left(\bar{a} + \sum_{\ell=1}^{\infty} \gamma_\ell \hat{z}_\ell \rho_\ell \right) - \Psi^\dagger \left(\bar{a} + \sum_{\ell=1}^{\infty} \gamma_\ell z_\ell \rho_\ell \right) \right\|_{\mathcal{Y}} \\ &\leq \text{Lip}(\Psi^\dagger) \left\| \sum_{\ell=1}^{\infty} \gamma_\ell [\hat{z}_\ell - z_\ell] \rho_\ell \right\|_{L^\infty} \\ &\leq C^* \text{Lip}(\Psi^\dagger) \left\| \sum_{\ell=1}^{\infty} \gamma_\ell [\hat{z}_\ell - z_\ell] \rho_\ell \right\|_{\mathcal{X}}. \end{aligned} \quad (62)$$

Where the last estimate follows from the assumed embedding $\mathcal{X} \hookrightarrow L^\infty$. To estimate the last expression, we note that by the properties of the shrink operator, since $\hat{z}_\ell = \text{shrink}(\tilde{z}_\ell)$, and the fact that $z_\ell \in [-1, 1]$, we have

$$|\hat{z}_\ell - z_\ell| = |\text{shrink}(\tilde{z}_\ell) - \text{shrink}(z_\ell)| \leq |\tilde{z}_\ell - z_\ell|.$$

Furthermore, we note that

$$\sum_{\ell=1}^m \gamma_\ell z_\ell \rho_\ell = P_m(a - \bar{a}),$$

with P_m the orthogonal projection onto $\text{span}\{\rho_1, \dots, \rho_m\}$, and similarly, by definition of \tilde{z}_ℓ , we have

$$\sum_{\ell=1}^m \gamma_\ell \tilde{z}_\ell \rho_\ell = P_m \Pi_{\mathcal{X}}(a - \bar{a}).$$

Together with the orthonormality of the ρ_ℓ , and the fact that $\tilde{z}_\ell = 0$ for $\ell > m$, this then implies that

$$\begin{aligned} \left\| \sum_{\ell=1}^{\infty} \gamma_\ell [\widehat{z}_\ell - z_\ell] \rho_\ell \right\|_{\mathcal{X}}^2 &= \sum_{\ell=1}^{\infty} \gamma_\ell^2 |\widehat{z}_\ell - z_\ell|^2 \leq \sum_{\ell=1}^{\infty} \gamma_\ell^2 |\tilde{z}_\ell - z_\ell|^2 \\ &= \left\| \sum_{\ell=1}^m \gamma_\ell [\tilde{z}_\ell - z_\ell] \rho_\ell \right\|_{\mathcal{X}}^2 + \sum_{\ell>m} \gamma_\ell^2 z_\ell^2 \\ &\leq \|\Pi_{\mathcal{X}}(a - \bar{a}) - (a - \bar{a})\|_{\mathcal{X}}^2 + \sum_{\ell>m} \gamma_\ell^2, \end{aligned} \quad (63)$$

where we made use of the fact that P_m is an orthogonal projection, and hence

$$\left\| \sum_{\ell=1}^m \gamma_\ell [\tilde{z}_\ell - z_\ell] \rho_\ell \right\|_{\mathcal{X}} = \|P_m \Pi_{\mathcal{X}}(a - \bar{a}) - P_m(a - \bar{a})\|_{\mathcal{X}} \leq \|\Pi_{\mathcal{X}}(a - \bar{a}) - (a - \bar{a})\|_{\mathcal{X}},$$

to pass to the last line.

By assumption on $\gamma_\ell \leq C\ell^{-1-\alpha}$, we can estimate

$$\sum_{\ell>m} \gamma_\ell^2 \leq C \sum_{\ell>m} \ell^{-2-2\alpha} \leq C \int_m^{\infty} \ell^{-2-2\alpha} d\ell = \frac{C}{1+2\alpha} m^{-1-2\alpha}.$$

We also note that the Lipschitz constant of $\Psi^\dagger : \text{supp}(\mu) \subset \mathcal{X} \rightarrow \mathcal{Y}$ depends only on λ, Λ and the domain $D \subset \mathbb{R}^n$. Thus, combining (61), (62) and (63) with the last estimate, it follows that

$$\begin{aligned} \|\psi \circ \mathcal{E}_{\mathcal{X}}(a) - \mathcal{E}_{\mathcal{Y}} \circ \Psi^\dagger(a)\|_{\ell^2(\mathbb{R}^{d_{\mathcal{Y}}})}^2 &= \|\psi \circ \mathcal{E}_{\mathcal{X}}(a(z)) - \mathcal{E}_{\mathcal{Y}} \circ \mathcal{F}(z)\|_{\ell^2(\mathbb{R}^{d_{\mathcal{Y}}})}^2 \\ &\leq C \left(m^{-2\alpha+\eta} + \|\Pi_{\mathcal{X}}(a - \bar{a}) - (a - \bar{a})\|_{\mathcal{X}}^2 \right), \end{aligned}$$

for a constant $C = C(\lambda, \Lambda, d, \Psi^\dagger, \alpha, \eta) > 0$. In the last estimate, we have absorbed the additional term proportional to $m^{-2\alpha-1}$ in the first (larger) term $m^{-2\alpha+\eta}$. Taking the expectation over $a \sim \mu$, and noting that $\Pi_{\mathcal{X}}\bar{a} - \bar{a} = \mathbb{E}_{a \sim \mu}[\Pi_{\mathcal{X}}a - a]$ implies,

$$\mathbb{E}_{a \sim \mu} [\|\Pi_{\mathcal{X}}(a - \bar{a}) - (a - \bar{a})\|_{\mathcal{X}}^2] \leq \mathbb{E}_{a \sim \mu} [\|\Pi_{\mathcal{X}}a - a\|_{\mathcal{X}}^2] = \widehat{\mathcal{E}}_{\mathcal{X}}^2,$$

we conclude that

$$\widehat{\mathcal{E}}_{\mathcal{X}}^{*2} \leq C \left(m^{-2\alpha+\eta} + \widehat{\mathcal{E}}_{\mathcal{X}}^2 \right),$$

where we are free to choose $\eta > 0$, and $C = C(\lambda, d, \Psi^\dagger, \alpha, \eta) > 0$ is independent of m , $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$. Combining the above estimates, we obtain

$$\mathbb{E}_{a \sim \mu}[\|\Psi^\dagger(a) - \Psi(a)\|_{\mathcal{Y}}^2] = \widehat{\mathcal{E}}^2 \leq 2\widehat{\mathcal{E}}_{\mathcal{X}}^{*2} + 2\widehat{\mathcal{E}}_{\mathcal{Y}}^2 \leq C \left(m^{-2\alpha+\eta} + \widehat{\mathcal{E}}_{\mathcal{X}}^2 + \widehat{\mathcal{E}}_{\mathcal{Y}}^2 \right).$$

Next, we recall that μ and $\Psi_{\#}^\dagger \mu$ are concentrated on bounded sets; indeed, this is the case by assumption for μ , and is a consequence of the a priori bound (21) in the case of $\Psi_{\#}^\dagger \mu$. Invoking Proposition 2, we can therefore estimate

$$\begin{aligned} \widehat{\mathcal{E}}_{\mathcal{X}}^2 &\leq \mathcal{R}_{d_{\mathcal{X}}}^{\text{opt}}(\mu) + \sqrt{\frac{Q_{\mathcal{X}} d_{\mathcal{X}} \log(1/\delta)}{N}}, \\ \widehat{\mathcal{E}}_{\mathcal{Y}}^2 &\leq \mathcal{R}_{d_{\mathcal{Y}}}^{\text{opt}}(\mu) + \sqrt{\frac{Q_{\mathcal{Y}} d_{\mathcal{Y}} \log(1/\delta)}{N}}, \end{aligned}$$

with probability at least $1 - \delta$, and with constants $Q_{\mathcal{X}}, Q_{\mathcal{Y}}$ depending only on μ and $\Psi_{\#}^\dagger \mu$, respectively. The optimal PCA projection errors have been bounded in Proposition 22, showing that $\mathcal{R}_{d_{\mathcal{X}}}^{\text{opt}}(\mu) \leq C d_{\mathcal{X}}^{-2\alpha-1}$ and $\mathcal{R}_{d_{\mathcal{Y}}}^{\text{opt}}(\Psi_{\#}^\dagger \mu) \leq C d_{\mathcal{Y}}^{-2\alpha+\eta}$. Therefore, with the assumed choice of $d_{\mathcal{X}} = d_{\mathcal{Y}} = d$ and number of PCA samples $N \geq d^{1+4\alpha} \log(1/\delta)$, we can now estimate

$$\widehat{\mathcal{E}}_{\mathcal{X}}^2 + \widehat{\mathcal{E}}_{\mathcal{Y}}^2 \leq C d^{-2\alpha+\eta},$$

with probability at least $1 - \delta$.

To summarize: Fix $\eta > 0$. We have shown that there exists a constant $C = C(\lambda, \Lambda, d, \Psi^\dagger, \alpha, \eta) > 0$, such that for any $m \in \mathbb{N}$, there exists a neural network $\psi : \mathbb{R}^{d_{\mathcal{X}}} \rightarrow \mathbb{R}^{d_{\mathcal{Y}}}$, satisfying the complexity estimate

$$\text{size}(\psi) \leq C m (\log(m)^2 + d), \quad \text{depth}(\psi) \leq C \log(m)^2,$$

and such that the (squared) PCA-Net error $\widehat{\mathcal{E}}^2$ for $\Psi = \mathcal{D}_{\mathcal{Y}} \circ \psi \circ \mathcal{E}_{\mathcal{X}}$ can be bounded by

$$\begin{aligned} \mathbb{E}_{a \sim \mu}[\|\Psi^\dagger(a) - \Psi(a)\|_{\mathcal{Y}}^2] &\leq C \left(\widehat{\mathcal{E}}_{\mathcal{X}}^2 + \widehat{\mathcal{E}}_{\mathcal{Y}}^2 + m^{-2\alpha+\eta} \right) \\ &\leq C (d^{-2\alpha+\eta} + m^{-2\alpha+\eta}), \end{aligned}$$

with probability at least $1 - \delta$ in the $N \sim d^{1+4\alpha} \log(1/\delta)$ PCA samples. Setting $m \sim d \sim \epsilon^{-1/(2\alpha-\eta)} = \epsilon^{-\frac{1}{2\alpha}-\eta'}$, for some $\eta' = \eta'(\eta)$, now yields the claimed error and complexity bounds. \blacksquare

Appendix F. Proofs for Navier-Stokes

This appendix contains the background and proof of the PCA-Net emulation result for the Navier-Stokes equations, Theorem 25. For simplicity, the Theorem in the main text was only stated for $n = 2$. Here, we allow for arbitrary n under a suitable smoothness assumption on the underlying solutions. As indicated in the main text, the relevant smoothness assumption is only known to hold for spatial dimension $n = 2$.

F.1 A convergent numerical scheme

For $K \in \mathbb{N}$, let $\mathbb{P}_K : L^2(\mathbb{T}^n; \mathbb{R}^n) \rightarrow L^2(\mathbb{T}^n; \mathbb{R}^n)$ denote the L^2 -orthogonal projection onto divergence-free real-valued vector fields of the form,

$$u_K(x) := \sum_{|k|_\infty \leq K} \widehat{u}_k e^{i\langle k, x \rangle}.$$

We will also denote $L_K^2 := \mathbb{P}_K(L^2(\mathbb{T}^n; \mathbb{R}^n))$ as the image of L^2 under \mathbb{P}_K .

We now consider the numerical discretization of (28), obtained by the following recursion: Let $\Delta t > 0$ be a (small) time-step. We assume we are given (discrete) initial data $u_K^0 \in L_K^2$, $u_K^0 \approx \mathbb{P}_K \bar{u}$. For $m \geq 0$, determine $u_K^{m+1} \in L_K^2$ by solving

$$\frac{u_K^{m+1} - u_K^m}{\Delta t} = -\mathbb{P}_K \left(u_K^m \cdot \nabla u_K^{m+1/2} \right) + \nu \Delta u_K^{m+1/2}, \quad (64)$$

where $u_K^{m+1/2} := \frac{1}{2} (u_K^{m+1} + u_K^m)$. The recursion (64) is equivalent to a recursion for the corresponding Fourier coefficients $\widehat{u}_k^m \in \mathbb{C}^n$, $k \in \mathbb{Z}^n$, $|k|_\infty \leq K$, where

$$u_K^m(x) = \sum_{|k|_\infty \leq K} \widehat{u}_k^m e^{i\langle k, x \rangle},$$

which are uniquely defined by (and equivalent to) the scheme (64): The Fourier transformed version of (64) is given by

$$\frac{\widehat{u}_k^{m+1} - \widehat{u}_k^m}{\Delta t} = -\widehat{\mathbb{P}}_K \left(u_K^m \cdot \nabla u_K^{m+1/2} \right) - \nu |k|^2 \widehat{u}_k^{m+1/2}, \quad (65)$$

where we introduce the shorthand notation

$$\widehat{\mathbb{P}}_K \left(u_K^m \cdot \nabla u_K^{m+1/2} \right) := \left(1 - \frac{k \otimes k}{|k|^2} \right) \sum_{|\ell|_\infty, |k-\ell|_\infty \leq K} i(\ell \cdot \widehat{u}_{k-\ell}^{m+1/2}) \widehat{u}_\ell^m, \quad (66)$$

which is the Fourier transform of the non-linear term $\mathbb{P}_K \left(u_K^m \cdot \nabla u_K^{m+1/2} \right)$.

In (Kovachki et al., 2021, Lem. 53), it has been shown that this numerical scheme is well-defined for time-steps satisfying $\|\bar{u}\|_{L^2} K^{n/2+1} \Delta t \leq \frac{1}{2}$, and that, if $t \mapsto u(t) \in L^2(\mathbb{T}^n; \mathbb{R}^n)$ denotes a smooth solution of (28), and $t^m := m\Delta t$, then we have

$$\begin{aligned} \frac{u(t^{m+1}) - u(t^m)}{\Delta t} &= -\mathbb{P}_K \left(u(t^m) \cdot \nabla \frac{1}{2} [u(t^{m+1}) + u(t^m)] \right) \\ &\quad + \nu \Delta \frac{1}{2} [u(t^{m+1}) + u(t^m)] + R_K^m(u), \end{aligned} \quad (67)$$

where the remainder $R_K^m(u)$ can be estimated by

$$\|R_K^m(u)\|_{L^2}^2 \leq C (\Delta t^2 + M^{-2r}), \quad (68)$$

for any $r \geq n/2 + 1$, and where the constant $C > 0$ depends only on $\|u\|_{C_t H_x^r}$, $\|u\|_{C_t^1 H_x^{r-1}}$. Here, we have introduced $C_t H_x^r := C([0, T]; H_x^r)$, $C_t^1 H_x^{r-1} := C^1([0, T]; H_x^{r-1})$, given as the space of vector fields $u : [0, T] \rightarrow H_x^r(\mathbb{T}^n; \mathbb{R}^n)$ with continuous dependence on time, and the space of vector fields $u : [0, T] \rightarrow H_x^{r-1}(\mathbb{T}^n; \mathbb{R}^n)$ with continuously differentiable dependence on time, respectively.

F.2 Neural network emulation of scheme (64)

Our first goal is to show that the numerical scheme (65) can be efficiently emulated by a neural network. To this end, we will first construct a neural network $\mathcal{NL}_{M,\epsilon}(\hat{u}, \hat{v})$, acting on the Fourier coefficients with $|k|_\infty \leq K$, such that

$$\|\mathcal{NL}_{M,\epsilon}(\hat{u}, \hat{v}) - \widehat{\mathbb{P}}_K(u_K \cdot \nabla v_K)\|_{L^2} \leq \epsilon,$$

for all \hat{u}, \hat{v} with $\|\hat{u}\|_{\ell^2}, \|\hat{v}\|_{\ell^2} \leq M$. In the following, we will denote for a given $K \in \mathbb{N}$, by $\mathcal{K} = \mathcal{K}_K$ the following set of multi-indices:

$$\mathcal{K}_K = \{k \in \mathbb{Z}^n \mid |k|_\infty \leq K\}.$$

We will denote by \mathcal{J}_K (\mathcal{J}_{2K}) an equidistant grid with $2K+1$ ($4K+1$) grid points $\{x_j\}_{j \in \mathcal{J}_K}$ in each direction. We will identify $\mathbb{C}^{\mathcal{K}} \sim \mathbb{R}^{2|\mathcal{K}|}$, i.e. the real and imaginary parts of any vector $\hat{w} \in \mathbb{C}^{\mathcal{K}}$ are considered as separate components of the corresponding element of $\mathbb{R}^{2|\mathcal{K}|}$ under this identification, and all neural networks are understood to act componentwise. We can then state the following approximation result:

Lemma 39 *There exists a constant $C > 0$, such that for any $\epsilon, M > 0$ and $K \in \mathbb{N}$, there exists a ReLU neural network $\mathcal{NL}_{M,\epsilon} : \mathbb{C}^{\mathcal{K}} \times \mathbb{C}^{\mathcal{K}} \rightarrow \mathbb{C}^{\mathcal{K}}$, $\mathcal{K} = \mathcal{K}_K$, satisfying the error bound*

$$\|\mathcal{NL}_{M,\epsilon}(\hat{u}, \hat{v}) - \widehat{\mathbb{P}}_K(u_K \cdot \nabla v_K)\|_{\ell^2} \leq \epsilon, \quad (69)$$

and the Lipschitz bound

$$\|\mathcal{NL}_{M,\epsilon}(\hat{u}, \hat{v}) - \mathcal{NL}_{M,\epsilon}(\hat{u}, \hat{w})\|_{\ell^2} \leq CK^{n/2+1}M\|\hat{v} - \hat{w}\|_{\ell^2}, \quad (70)$$

for all $\hat{u}, \hat{v}, \hat{w} \in \mathbb{C}^{\mathcal{K}}$, such that $\|\hat{u}\|_{\ell^2}, \|\hat{v}\|_{\ell^2}, \|\hat{w}\|_{\ell^2} \leq M$. Furthermore, the size of the neural network $\mathcal{NL}_{M,\epsilon}$ can be estimated by

$$\text{size}(\mathcal{NL}_{M,\epsilon}) \leq C|\mathcal{K}_K| \log(MK/\epsilon), \quad \text{depth}(\mathcal{NL}_{M,\epsilon}) \leq C \log(MK/\epsilon).$$

where $C > 0$ is independent of K, M, ϵ .

The proof of Lemma 39 is provided in Appendix F.5. Using the above lemma, we can then show that the approximate equation

$$\frac{\hat{u}_k^{m+1} - \hat{u}_k^m}{\Delta t} = -\mathcal{NL}_{M,\epsilon}(\hat{u}^m, \hat{u}^{m+1/2}) - \nu|k|^2\hat{u}_k^{m+1/2}, \quad (71)$$

has a unique solution at each time-step, for suitably chosen $\Delta t, M, \epsilon > 0$.

Lemma 40 *If $\|\hat{u}^m\|_{L^2} \leq M$ for some $M > 0$, and if the non-linearity $\mathcal{NL}_{\overline{M},\epsilon}(\hat{u}, \hat{v}) \mapsto \mathcal{NL}_{\overline{M},\epsilon}(\hat{u}, \hat{v})$ satisfies the error estimate (69) for a given $0 < \epsilon \leq 1$ with the Lipschitz estimate (70) for $\overline{M} = 2(M+2)$, and if the time-step $\Delta t \in (0, 1]$ satisfies the CFL condition*

$$C\Delta t K^{n/2+1}\overline{M} \leq 1, \quad (72)$$

then the update rule (71) is well-posed, in the sense that a unique solution \widehat{u}^{m+1} exists with $\|\widehat{u}^{m+1}\|_{L^2} \leq \overline{M}$. Furthermore, \widehat{u}^{m+1} can be approximated via the following fixed-point iteration: Setting $\widehat{w}_k^0 = 0$, and

$$\widehat{w}_k^{\ell+1} := \frac{1}{1 + \frac{1}{2}\Delta t \nu |k|^2} \left[\widehat{u}^m + \Delta t \mathcal{NL}_{\overline{M}, \epsilon} \left(\widehat{u}^m, \frac{\widehat{u}^m + \widehat{w}^\ell}{2} \right)_k - \frac{1}{2} \Delta t \nu |k|^2 \widehat{u}^m \right],$$

for $\ell = 0, 1, \dots$. Then we have the following a priori estimates, valid for any $\ell = 0, 1, \dots$:

$$\|\widehat{w}^\ell - \widehat{u}^{m+1}\|_{\ell^2} \leq \frac{\exp(2\Delta t \epsilon)}{2^\ell} \|\widehat{u}^m\|_{\ell^2},$$

and

$$\|\widehat{w}^\ell\|_{\ell^2} \leq \left(1 + 2^{-\ell}\right) \exp(2\Delta t \epsilon) \|\widehat{u}^m\|_{\ell^2}.$$

Given the well-posedness of the update rule in the emulation scheme (71), and the approximation by a simple fixed-point iteration provided in Lemma 40, it is natural to define the following neural network-based emulation algorithm:

Algorithm 1: Navier-Stokes NN-emulation

Input: parameters $M \geq 1$, $r > 0$, $r > n/2 + 1$, final time $T > 0$,

\widehat{u}^0 initial data, with $\|\widehat{u}^0\|_{\ell^2} \leq M$.

Result: \widehat{u}^{m_T} , where $n_T \Delta t = T$.

1

2 Let C be the constant of estimate (70) for $\mathcal{NL}_{\overline{M}, \epsilon}$.

▷ **Note:** C is independent of $\overline{M}, \epsilon!$

3 Set $\overline{M} \leftarrow 2(eM + 1)$.

▷ $e = \exp(1)$ is Euler's number.

4 Set $\overline{T} := \max(T, 1)$

5 Set $\Delta t \in (0, 1]$ maximal, s.t.:

$$CK^{n/2+1}\overline{M}\Delta t \leq 1, \quad \Delta t \leq K^{-r}, \quad n_T = T/\Delta t \in \mathbb{N}.$$

Set $\epsilon \leftarrow \Delta t / (3M\overline{T})$.

6 Set $L \leftarrow \lceil \log_2(1/\Delta t \epsilon) \rceil = \lceil \log_2(3M\overline{T}/\Delta t^2) \rceil$.

7

8 Construct a neural network $\mathcal{NL}_{\overline{M}, \epsilon}$ as in Lemma 39.

9 **for** $m = 0, 1, \dots, n_T - 1$ **do**

10 | Set $\widehat{w}^{m,0} \leftarrow 0$.

11 | Define $F(\widehat{w}) = F(\widehat{u}^m; \widehat{w})$ by (78) with $\mathcal{NL}_{\overline{M}, \epsilon}$.

12 | **for** $\ell = 0, 1, \dots, L - 1$ **do**

13 | | Set $\widehat{w}^{m, \ell+1} \leftarrow F(\widehat{u}^m; \widehat{w}^{m, \ell})$.

14 | **end**

15 | Set $\widehat{u}^{m+1} \leftarrow \widehat{w}^{m, L}$.

16 **end**

We first note that the recursive algorithm 1 can be represented by either a deep neural network, or a recursive neural network:

Lemma 41 *There exists $C > 0$, such that for any $K \in \mathbb{N}$ and $r > n/2 + 1$, $M \geq 1$, $T > 0$, there exists a neural network $\psi : \mathbb{C}^{\mathcal{K}} \rightarrow \mathbb{C}^{\mathcal{K}}$ of size*

$$\begin{cases} \text{size}(\psi) \leq CM\bar{T}K^{n+r} \log(M\bar{T}K)^2, \\ \text{depth}(\psi) \leq CM\bar{T}K^r \log(M\bar{T}K), \end{cases}$$

where $\bar{T} := \max(T, 1)$, which maps the input of algorithm 1 to its output. In fact, ψ can be written in the form of a recurrent neural network, or more precisely, as a composition

$$\psi = \underbrace{\psi^* \circ \dots \circ \psi^*}_{n_T\text{-fold}} : \mathbb{C}^{\mathcal{K}} \rightarrow \mathbb{C}^{\mathcal{K}}$$

of a neural network $\psi^* : \mathbb{C}^{\mathcal{K}} \rightarrow \mathbb{C}^{\mathcal{K}}$, where $n_T \leq CMTK^r$ and

$$\text{size}(\psi^*) \leq CK^n \log(M\bar{T}K)^2, \quad \text{depth}(\psi^*) \leq C \log(M\bar{T}K).$$

For completeness, a detailed proof of Lemma 41 is provided in Appendix F.7. Before providing an error estimate on the sequence $\hat{u}^m \approx \hat{u}(t^m)$ generated in Algorithm 1, we derive an improved upper bound on the norms $\|\hat{u}^m\|_{L^2}$, which is required in order to ensure that the emulating neural network nonlinearity $\mathcal{NL}_{\bar{M}, \epsilon}$ in (71) provides an accurate approximation of the non-linear term $\mathbb{P}_K(u_K^m \cdot u_K^{m+1})$, uniformly over all time-steps.

Lemma 42 *If $\hat{u}^{m+1} = \hat{w}^{m,L}$, for $m = 0, \dots, n_T - 1$, is determined as in Algorithm 1, then we have*

$$\|\hat{u}^m\|_{\ell^2} \leq \exp(3m\Delta t \epsilon) M \leq eM, \tag{73}$$

for all $m = 0, \dots, n_T$.

The simple proof of this lemma is provided in Appendix F.8. Based on the last lemma, we then obtain the following:

Lemma 43 *If $\hat{u}^{m+1} = \hat{w}^{m,L}$, for $m = 0, \dots, n_T - 1$, is determined as in Algorithm 1, then the corresponding function $u_K^{m+1} := \sum_{|k|_\infty \leq K} \hat{u}_k^{m+1} e^{i\langle k, x \rangle}$ satisfies*

$$\frac{u_K^{m+1} - u_K^m}{\Delta t} = -\mathbb{P}_K \left(u_K^m \cdot \nabla u_K^{m+1/2} \right) + \nu \Delta u_K^{m+1/2} + Q_K^m, \tag{74}$$

where the remainder Q_K^m satisfies a bound of the form

$$\sup_m \|Q_K^m\|_{L^2} \leq C\Delta t, \tag{75}$$

where $C > 0$ is independent of K , M and T .

F.3 Error estimate for neural network emulation

With Lemma 43 in hand, we can now prove the following error estimate for the approximation of the underlying solution, $u_K^m \approx u(t^m)$:

Lemma 44 *Let $u_K^m = \sum_{|k|_\infty \leq K} \widehat{u}_k^m e^{i\langle k, x \rangle}$ be obtained by the emulation algorithm 1 with parameters $M, r, T > 0$, $r > n/2 + 1$, and initial data $\|\widehat{u}^0\|_{\ell^2} \leq M$. Let $u(t)$ be the exact solution of (28) with the initial data $\bar{u} \in H^r(\mathbb{T}^n; \mathbb{R}^n)$. Assume that $u \in C([0, T]; H_x^r) \cap C^1([0, T]; H_x^{r-1})$. Then u_K^m satisfies*

$$\max_{m=0, \dots, n_T} \|u_K^m - u(t^m)\|_{L^2} \leq C \left(K^{-r} + \|\widehat{u}^0 - \widehat{\mathbb{P}}_K u(0)\|_{\ell^2} \right),$$

where $C = C(\|u\|_{C_t H_x^r}, \|u\|_{C_t^1 H_x^{r-1}}, T, M, n, r) > 0$, is independent of K , and where $\widehat{\mathbb{P}}_K u(0) = \{\widehat{u}_k(t=0)\}_{|k|_\infty \leq K}$ denote the Fourier coefficients of the initial data $\bar{u} = u(0)$ with wavenumbers $|k|_\infty \leq K$.

To simplify notation in the following, we introduce $C_t H_x^r := C([0, T]; H_x^r)$, $C_t^1 H_x^{r-1} := C^1([0, T]; H_x^{r-1})$, and define

$$\mathcal{H}_{M,T}^r = \{u \in C_t H_x^r \cap C_t^1 H_x^{r-1} \mid u \text{ solves (28) and } \|u\|_{\mathcal{H}^r} \leq M\}, \quad (76)$$

where

$$\|u\|_{\mathcal{H}^r} := \|u\|_{C_t H_x^r} + \|u\|_{C_t^1 H_x^{r-1}}.$$

Finally, combining Lemma 44 with Lemma 41, we obtain the following estimate for the emulation of the spectral method (64) by ReLU neural networks:

Proposition 45 *Let $M, r, T > 0$ be given. For any $\epsilon > 0$, there exists $K \in \mathbb{N}$, $K \sim \epsilon^{-1/r}$, and a ReLU neural network $\widehat{\psi} : \mathbb{C}^{\mathcal{K}} \rightarrow \mathbb{C}^{\mathcal{K}}$, $\mathcal{K} = \mathcal{K}_K$, such that*

$$\|\widehat{\psi}(\widehat{u}(0)) - \widehat{u}(T)\|_{\ell^2} \leq \epsilon, \quad \forall u \in \mathcal{H}_{M,T}^r,$$

where $\widehat{u}(t) = \{\widehat{u}_k(t)\}_{|k|_\infty \leq K}$ denote the Fourier coefficients of $u(t)$, and such that

$$\begin{cases} \text{size}(\widehat{\psi}) \leq C \epsilon^{-n/r-1} \log(\epsilon^{-1})^2, \\ \text{depth}(\widehat{\psi}) \leq C \epsilon^{-1} \log(\epsilon^{-1})^2. \end{cases}$$

The constant $C > 0$, depends on M, T, r, n , but is independent of ϵ . Furthermore, $\widehat{\psi}$ can be written as a n_T -fold composition $\widehat{\psi} = \widehat{\psi}_* \circ \dots \circ \widehat{\psi}_*$, where $\widehat{\psi}_* : \mathbb{C}^{\mathcal{K}} \rightarrow \mathbb{C}^{\mathcal{K}}$ is a ReLU neural network with

$$\begin{cases} \text{size}(\widehat{\psi}_*) \leq C \epsilon^{-n/r} \log(\epsilon^{-1})^2, \\ \text{depth}(\widehat{\psi}_*) \leq C \log(\epsilon^{-1})^2, \end{cases}$$

and $n_T \leq C \epsilon^{-1}$.

Proposition 27 in the two-dimensional case, $n = 2$, is an immediate consequence of the above result.

F.4 An error estimate for Navier-Stokes

In the last section, we have obtained quantitative estimates on the ReLU DNN emulation of the spectral scheme (64). This provides an approximate mapping from the Fourier coefficients of the input function, to the Fourier coefficients of the output function. In the present section, we combine this idea with the PCA projection, to construct a PCA-Net approximation of the Navier-Stokes solution operator.

To obtain an estimate on the approximation error $\widehat{\mathcal{E}}(\Psi)$ for the Navier-Stokes equations, we fix parameters $M, r > 0$, $r > n/2 + 1$, and a final time $T > 0$, and consider the operator $\Psi^\dagger : L^2(\mathbb{T}^n; \mathbb{R}^n) \rightarrow L^2(\mathbb{T}^n; \mathbb{R}^n)$, given by the forward solution operator from time $t = 0$ to $t = T$ of (28). To ensure that $\Psi^\dagger(u)$ is well-defined for μ -almost every u , we assume that

$$\mu(\{u(0) \mid u \in \mathcal{H}_{M,T}^r\}) = 1, \quad (77)$$

i.e. that μ is concentrated on initial data for which the corresponding solution u belongs to $\mathcal{H}_{M,T}^r$ (cp. (76) for the definition of $\mathcal{H}_{M,T}^r$).

Lemma 46 *Given the above setting, for any $K \in \mathbb{N}$ (corresponding to Fourier cut-off for wavenumbers $|k|_\infty \leq K$), there exists a PCA-Net $\Psi = \mathcal{D}_y \circ \psi \circ \mathcal{E}_x$, such that*

$$\begin{aligned} \mathbb{E}_{u \sim \mu} \left[\|\Psi^\dagger(u) - \Psi(u)\|_{L^2}^2 \right]^{1/2} &\leq CK^{-r} + C\mathbb{E}_{u \sim \mu} \left[\|u - P_x u\|_{L^2}^2 \right]^{1/2} \\ &\quad + \mathbb{E}_{v \sim \mathcal{G}_{\# \mu}} \left[\|v - P_y v\|_{L^2}^2 \right]^{1/2}, \end{aligned}$$

where ψ is a neural network of size

$$\text{size}(\psi) \leq CK^n (K^r \log(K)^2 + (d_x + d_y)), \quad \text{depth}(\psi) \leq CK^r \log(K),$$

and $C = C(M, r, T, n) > 0$ is a constant independent of K, d_x, d_y .

The proof of Lemma 46 is given in Appendix F.11. Setting $K \sim \epsilon^{-1/r}$ and specializing to $n = 2$ immediately implies Lemma 28 in the main text. We finally provide an estimate on the PCA projection errors, which is the general n -dimensional version of Lemma 28 in the main text:

Lemma 47 *Assume that*

$$\mu(\{u(0) \mid u \in \mathcal{H}_{M,T}^r\}) = 1,$$

and the operator Ψ^\dagger is given by

$$\Psi^\dagger(u(0)) = u(T), \quad \forall u \in \mathcal{H}_{M,T}^r,$$

where $u(t)$ solves the Navier-Stokes equations (28) with initial data $u(0)$. Then there exists a constant $C = C(M, r, T, n) > 0$, such that for PCA dimensions $d_x = d_y$, arbitrary $\delta > 0$ and $N \gtrsim d_x^{1+4r/n} \log(1/\delta)$ PCA samples, the empirical PCA projection errors based on N samples $u_1, \dots, u_N \stackrel{iid}{\sim} \mu$ satisfy,

$$\mathbb{E}_{u \sim \mu} \left[\|u - P_x u\|_{L^2}^2 \right] + \mathbb{E}_{v \sim \Psi_{\#}^\dagger \mu} \left[\|v - P_y v\|_{L^2}^2 \right] \leq C d_x^{-2r/n},$$

with probability at least $1 - \delta$.

The proof of Lemma 47 is immediate by combining the general estimate on empirical PCA, Proposition 2, with the upper bound on the PCA projection error $R_p(\mu)$ based on the smoothness of the underlying functions, Proposition 15. We only need to observe that μ and $\Psi_{\#}^\dagger \mu$ are concentrated on a bounded set $\{\|u\|_{H^r} \leq M\}$, by assumption.

Proof [Proof of Theorem 25] Setting $K \sim \epsilon^{-1/r}$ in Lemma 28, with $\Psi = \mathcal{D}_y \circ \psi \circ \mathcal{E}_X$, and combining it with the encoding error estimate of Lemma 47 for $d_X = d_y$, we arrive at

$$\mathbb{E}_{u \sim \mu} \left[\|\Psi^\dagger(u) - \Psi(u)\|_{L^2(\mu)}^2 \right] \leq \epsilon^2 + Cd^{-2r/n},$$

where ψ is a neural network of size

$$\text{size}(\psi) \leq C\epsilon^{-n/r} (\epsilon^{-1} \log(\epsilon^{-1})^2 + d_X), \quad \text{depth}(\psi) \leq C\epsilon^{-1} \log(\epsilon^{-1}),$$

and $C = C(M, r, T, n) > 0$ is a constant independent of ϵ, d_X . For spatial dimension $n = 2$, and upon replacing ϵ by $\epsilon^{1/2}$ throughout, and choosing $d_X \sim \epsilon^{-1/r}$, this yields the estimate of Theorem 25. \blacksquare

The remaining sections contain detailed proofs of all Lemmas above.

F.5 Proof of Lemma 39

Proof We note that the mapping $\mathbb{C}^{\mathcal{K}} \times \mathbb{C}^{\mathcal{K}} \mapsto \mathbb{C}^{\mathcal{K}}, (\hat{u}, \hat{v}) \mapsto \widehat{\mathbb{P}}_K(u_K \cdot \nabla v_K)$ can be written as a composition,

$$\begin{aligned} (\hat{u}_k, \hat{v}_k)_{k \in \mathcal{K}} &\mapsto (\hat{u}_k, k \otimes \hat{v}_k)_{k \in \mathcal{K}} \\ &\mapsto (u_K(x_j), \nabla v_K(x_j))_{j \in \mathcal{J}_{2K}} \\ &\mapsto ([u_K \cdot \nabla v_K](x_j))_{j \in \mathcal{J}_{2K}} \\ &\mapsto \widehat{\mathbb{P}}_K(u_K \cdot \nabla v_K). \end{aligned}$$

The first mapping can trivially be represented by a matrix multiplication with $O(|\mathcal{K}|)$ non-zero entries, the second step can be efficiently computed via the fast Fourier-transform, which can be represented exactly by a ReLU neural network with $O(|\mathcal{K}| \log(|\mathcal{K}|))$ non-zero entries (and depth $O(\log(|\mathcal{K}|))$). Thus, there exists a (absolute) constant $C = C(n) > 0$, and a ReLU neural network $\psi^* : \mathbb{C}^{\mathcal{K}} \times \mathbb{C}^{\mathcal{K}} \rightarrow [\mathbb{R}^n]^{\mathcal{J}} \times [\mathbb{R}^{n \times n}]^{\mathcal{J}}$, such that

$$\psi^*(\hat{u}_k, \hat{v}_k) = (u_K(x_j), \nabla v_K(x_j))_{j \in \mathcal{J}},$$

and $\text{size}(\psi^*) \leq C|\mathcal{K}| \log(|\mathcal{K}|)$, $\text{depth}(\psi^*) \leq C \log(|\mathcal{K}|)$. Similarly, for the last mapping, there exists a ReLU neural network $\psi^{**} : [\mathbb{R}^n]^{\mathcal{J}} \rightarrow \mathbb{C}^{\mathcal{K}}$, such that

$$\psi^{**}(w(x_j)) = \widehat{\mathbb{P}}_K(w), \quad \forall w \in L_{2K}^2,$$

and $\text{size}(\psi^{**}) \leq C|\mathcal{K}| \log(|\mathcal{K}|)$, $\text{depth}(\psi^{**}) \leq C \log(|\mathcal{K}|)$. The main task is thus to show that the point-wise multiplication

$$(u_K(x_j), \nabla v_K(x_j)) \mapsto u_K(x_j) \cdot \nabla v_K(x_j),$$

can be suitably approximated by a neural network. To this end, we will rely on the following lemma:

Lemma 48 ((Ryck and Mishra, 2021, Lemma 2.6)) *Let $M, L > 0$ and denote by $\times : \mathbb{R}^2 \rightarrow \mathbb{R}$, $(x, y) \mapsto xy$ the multiplication operator. For any $m \in \mathbb{N}$, there exists a realization of a ReLU neural network $\widehat{\times}_m : [-M, M] \times [-L, L] \rightarrow \mathbb{R}$, such that $\widehat{\times}_m$ satisfies for all $y \in [-L, L]$ the error bound*

$$\|\times(\cdot, y) - \widehat{\times}_m(\cdot, y)\|_{\text{Lip}([-M, M]; \mathbb{R})} \leq \frac{M + L}{2^{m+1}}.$$

Furthermore, it holds that $\text{size}(\widehat{\times}_m) \lesssim m$, $\text{depth}(\widehat{\times}_m) = m + 1$.

We now note that for any $w \in L^2_{2K}$, $w(x) = \sum_{|k|_\infty \leq 2K} \widehat{w}_k e^{i\langle k, x \rangle}$, we have

$$\|w\|_{L^2}^2 = (2\pi)^n \sum_{|k|_\infty \leq 2K} |\widehat{w}_k|^2 = \frac{(2\pi)^n}{|\mathcal{J}_{2K}|} \sum_{j \in \mathcal{J}_{2K}} |w(x_j)|^2.$$

and we have

$$\|u_K\|_{L^\infty} \leq CK^{n/2} \|u_K\|_{L^2} \leq C_n K^{n/2} \|\widehat{u}_k\|_{\ell^2} \leq C_n K^{n/2} M,$$

and

$$\|\nabla v_K\|_{L^\infty} \leq CK^{n/2} \|\nabla v_K\|_{L^2} \leq C_n K^{n/2+1} \|\widehat{v}_k\|_{\ell^2} \leq C_n K^{n/2+1} M.$$

Applying Lemma 48, it follows that there exists a neural network $\psi^* : \mathbb{R}^n \times \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^n$, such that

$$\|\psi^*(u_K(x_j), \nabla v_K(x_j)) - u_K(x_j) \cdot \nabla v_K(x_j)\|_{\text{Lip}} \leq \epsilon,$$

and $\text{size}(\psi^*) \leq C \log(MK/\epsilon)$, $\text{depth}(\psi^*) \leq C \log(MK/\epsilon)$, where $C = C(n) > 0$ is independent of M, K, ϵ . Clearly, the parallelized mapping $\psi^\epsilon : [\mathbb{R}^n]^\mathcal{J} \times [\mathbb{R}^{n \times n}]^\mathcal{J} \rightarrow [\mathbb{R}^n]^\mathcal{J}$, defined by the point-wise application $(u_K(x_j), \nabla v_K(x_j)) \mapsto \psi^*(u_K(x_j), \nabla v_K(x_j))$, $j \in \mathcal{J}_{2K}$, can be represented by a neural network with $\text{size}(\psi^\epsilon) \leq C |\mathcal{J}_{2K}| \log(MK/\epsilon)$. Furthermore, let $\widehat{\psi}_\epsilon$ denote the coefficients $|k|_\infty \leq K$ of the (fast) Fourier transform of ψ^ϵ based on the values on the grid x_j , $j \in \mathcal{J}_{2K}$, i.e. $\widehat{\psi}_\epsilon = \psi^{**} \circ \psi^\epsilon$. Then, we have

$$\begin{aligned} & \|\widehat{\psi}_\epsilon(u_K(x_j), \nabla v_K(x_j)) - \widehat{\mathbb{P}}_K(u_K(x_j) \cdot \nabla v_K(x_j))\|_{\ell^2(\mathcal{K})} \\ & \leq \frac{1}{|\mathcal{J}_{2K}|^{1/2}} \|\psi^\epsilon(u_K(x_j), \nabla v_K(x_j)) - u_K(x_j) \cdot \nabla v_K(x_j)\|_{\ell^2(\mathcal{J}_{2K})} \\ & \leq |\mathcal{J}_{2K}|^{1/2} \|\psi^\epsilon(u_K(x_j), \nabla v_K(x_j)) - u_K(x_j) \cdot \nabla v_K(x_j)\|_{\ell^\infty(\mathcal{J}_{2K})} \\ & \leq |\mathcal{J}_{2K}|^{1/2} \epsilon. \end{aligned}$$

In particular, replacing $\epsilon \rightarrow \epsilon/|\mathcal{J}_{2K}|^{1/2}$, we conclude that there exists a neural network $\psi : [\mathbb{R}^n]^\mathcal{J} \times [\mathbb{R}^{n \times n}]^\mathcal{J} \rightarrow [\mathbb{R}^n]^\mathcal{J}$, such that

$$\text{size}(\psi) \leq C |\mathcal{K}| \log(MK/\epsilon), \quad \text{depth}(\psi) \leq C \log(MK/\epsilon),$$

where $C = C(n) > 0$ is independent of M, K, ϵ , and

$$\|\widehat{\psi}(u_K(x_j), \nabla v_K(x_j)) - \widehat{\mathbb{P}}_K(u_K(x_j) \cdot \nabla v_K(x_j))\|_{\ell^2(\mathcal{K})} \leq \epsilon.$$

Furthermore, we note that for any v_K, w_K , we have

$$\begin{aligned}
 & \|\widehat{\psi}_\epsilon(u_K(x_j), \nabla v_K(x_j)) - \widehat{\psi}_\epsilon(u_K(x_j), \nabla w_K(x_j))\|_{\ell^2(\mathcal{K})}^2 \\
 &= \frac{1}{|\mathcal{J}_{2K}|} \sum_{j \in \mathcal{J}_{2K}} |\widehat{\psi}_*(u_K(x_j), \nabla v_K(x_j)) - \widehat{\psi}_*(u_K(x_j), \nabla w_K(x_j))|^2 \\
 &\leq \frac{1}{|\mathcal{J}_{2K}|} \sum_{j \in \mathcal{J}_{2K}} \text{Lip}(\psi^*(u_K(x_j), \cdot)) |\nabla v_K(x_j) - \nabla w_K(x_j)|^2.
 \end{aligned}$$

Since

$$\begin{aligned}
 \text{Lip}(\psi^*(u_K(x_j), \cdot)) &\leq \text{Lip}(\psi^*(u_K(x_j), \cdot) - u_K(x_j) \cdot (\cdot)) + \text{Lip}(u_K(x_j) \cdot (\cdot)) \\
 &\leq \epsilon + \|u_K\|_{L^\infty} \\
 &\leq \epsilon + CK^{n/2}M \leq 2CK^{n/2}M,
 \end{aligned}$$

where we have wlog assumed that $\epsilon \leq CK^{n/2}M$, it follows that there exists a constant $C = C(n) > 0$, such that

$$\begin{aligned}
 & \|\widehat{\psi}_\epsilon(u_K(x_j), \nabla v_K(x_j)) - \widehat{\psi}_\epsilon(u_K(x_j), \nabla w_K(x_j))\|_{\ell^2(\mathcal{K})}^2 \\
 &\leq \frac{CK^n M^2}{|\mathcal{J}_{2K}|} \sum_{j \in \mathcal{J}_{2K}} |\nabla v_K(x_j) - \nabla w_K(x_j)|^2 \\
 &= CK^n M^2 \|\nabla v_K - \nabla w_K\|_{L^2}^2 \\
 &\leq CK^{n+2} M^2 \|\widehat{v} - \widehat{w}\|_{\ell^2}^2.
 \end{aligned}$$

We conclude that

$$\begin{aligned}
 & \|\widehat{\psi}(u_K(x_j), \nabla v_K(x_j)) - \widehat{\psi}(u_K(x_j), \nabla w_K(x_j))\|_{\ell^2(\mathcal{K})} \\
 &\leq CK^{n/2+1} M \|\widehat{v} - \widehat{w}\|_{L^2},
 \end{aligned}$$

for all $\|\widehat{u}\|_{\ell^2}, \|\widehat{v}\|_{\ell^2}, \|\widehat{w}\|_{\ell^2} \leq M$. Thus, the composition $\mathcal{NL}_{M,\epsilon} := \widehat{\psi} \circ \psi^* = \psi^{**} \circ \psi \circ \psi^*$ defines a neural network satisfying the estimates of the present lemma. \blacksquare

F.6 Proof of Lemma 40

Proof Fix \widehat{u}^m as in the statement of this lemma. We will denote $\mathcal{NL} := \mathcal{NL}_{\overline{M},\epsilon}$, in the following. Consider the mapping $F : B_{\overline{M}}(0) \subset \mathbb{C}^{\mathcal{K}} \rightarrow \mathbb{C}^{\mathcal{K}}$, given by

$$F(\widehat{w})_k := \frac{1}{1 + \frac{1}{2}\Delta t \nu |k|^2} \left[\widehat{u}^m - \Delta t \mathcal{NL} \left(\widehat{u}^m, \frac{\widehat{u}^m + \widehat{w}}{2} \right)_k - \frac{1}{2} \Delta t \nu |k|^2 \widehat{u}^m \right]. \quad (78)$$

It is straight-forward to check that \widehat{u}^{m+1} solves the update rule (71) if, and only if, $\widehat{u}^{m+1} = F(\widehat{u}^{m+1})$. Furthermore, we note that

$$\begin{aligned} & \|F(\widehat{w}) - F(\widehat{v})\|_{\ell^2} \\ &= \frac{\Delta t}{1 + \frac{1}{2}\Delta t \nu |k|^2} \left\| \mathcal{NL} \left(\widehat{u}^m, \frac{\widehat{u}^m + \widehat{w}}{2} \right) - \mathcal{NL} \left(\widehat{u}^m, \frac{\widehat{u}^m + \widehat{v}}{2} \right) \right\|_{\ell^2} \\ &\leq \Delta t \text{Lip}(\mathcal{NL}(\widehat{u}^m, \cdot)) \frac{\|\widehat{w} - \widehat{v}\|_{\ell^2}}{2} \\ &\leq \frac{1}{2} CK^{n/2+1} \overline{M} \Delta t \|\widehat{w} - \widehat{v}\|_{\ell^2}. \end{aligned}$$

Thus, under the CFL condition $CK^{n/2+1}\overline{M}\Delta t \leq 1$, it follows that $\text{Lip}(F) \leq 1/2$ is a contraction, with constant $\gamma = 1/2 < 1$. We also note that

$$\begin{aligned} \|F(0)\|_{\ell^2} &\leq \left\| \left(\frac{1 - \frac{1}{2}\Delta t \nu |k|^2}{1 + \frac{1}{2}\Delta t \nu |k|^2} \right) \widehat{u}^m \right\|_{\ell^2} + \Delta t \|\mathcal{NL}(\widehat{u}^m, \widehat{u}^m/2)\|_{\ell^2} \\ &\leq \|\widehat{u}^m\|_{\ell^2} + \Delta t \left(\epsilon + \frac{1}{2} \|\widehat{P}_K(\widehat{u}^m \cdot \nabla \widehat{u}^m)\|_{\ell^2} \right) \\ &\leq M + \Delta t \epsilon + \frac{1}{2} CK^{n/2+1} M \Delta t \\ &\leq M + 2. \end{aligned}$$

In particular, this implies that

$$R := \frac{\|F(0)\|}{1 - \gamma} \leq 2(M + 2) \leq \overline{M}.$$

We can thus apply Lemma 49 to conclude that there exists a unique fixed point $\widehat{u}^{m+1} \in B_{\overline{M}}(0)$ ($= \widehat{w}^\infty$) of F , i.e. such that $F(\widehat{u}^{m+1}) = \widehat{u}^{m+1}$, and satisfying the estimate

$$\|\widehat{w}^\ell - \widehat{u}^{m+1}\|_{\ell^2} \leq \gamma^\ell \|\widehat{u}^{m+1}\|_{\ell^2} = 2^{-\ell} \|\widehat{u}^{m+1}\|_{\ell^2}, \quad (79)$$

and hence

$$\|\widehat{w}^\ell\|_{\ell^2} \leq \|\widehat{u}^{m+1}\|_{\ell^2} + \|\widehat{w}^\ell - \widehat{u}^{m+1}\|_{\ell^2} \leq (1 + 2^{-\ell}) \|\widehat{u}^{m+1}\|_{\ell^2}. \quad (80)$$

To provide a (improved) estimate on $\|\widehat{u}^{m+1}\|_{\ell^2}$, we multiply (71) by $\widehat{u}^{m+1/2} = \frac{1}{2}(\widehat{u}^m + \widehat{u}^{m+1})$, to find

$$\begin{aligned} \frac{\|\widehat{u}^{m+1}\|_{\ell^2}^2 - \|\widehat{u}^m\|_{\ell^2}^2}{\Delta} &= \left\langle \psi(\widehat{u}^m, \widehat{u}^{m+1/2}), \widehat{u}^{m+1/2} \right\rangle_{\ell^2} - \nu \sum_{|k|_\infty \leq K} |k|^2 \left| \widehat{u}_k^{m+1/2} \right|^2 \\ &\leq \left\langle \psi(\widehat{u}^m, \widehat{u}^{m+1/2}), \widehat{u}^{m+1/2} \right\rangle_{\ell^2}. \end{aligned} \quad (81)$$

Next, we use the fact that

$$\begin{aligned} \left\langle \widehat{\mathbb{P}}_K(u_K^m \cdot \nabla u_K^{m+1/2}), \widehat{u}^{m+1/2} \right\rangle_{\ell^2} &= (2\pi)^{-n} \left\langle \mathbb{P}_K(u_K^m \cdot \nabla u_K^{m+1/2}), u_K^{m+1/2} \right\rangle_{L^2} \\ &= (2\pi)^{-n} \left\langle u_K^m \cdot \nabla u_K^{m+1/2}, u_K^{m+1/2} \right\rangle_{L^2} \\ &= 0, \end{aligned}$$

and the fact that

$$\|\widehat{u}^{m+1/2}\|_{\ell^2} \leq \frac{1}{2} (\|\widehat{u}^m\|_{\ell^2} + \|\widehat{u}^{m+1}\|_{\ell^2}) \leq \frac{1}{2}(M + R) \leq \overline{M},$$

to find

$$\begin{aligned} & \left\langle \psi(\widehat{u}^m, \widehat{u}^{m+1/2}), \widehat{u}^{m+1/2} \right\rangle_{\ell^2} \\ &= \left\langle \psi(\widehat{u}^m, \widehat{u}^{m+1/2}) - \widehat{P}_K(u_K^m \cdot \nabla u_K^{m+1/2}), \widehat{u}^{m+1/2} \right\rangle_{\ell^2} \\ &\leq \left\| \psi(\widehat{u}^m, \widehat{u}^{m+1/2}) - \widehat{P}_K(u_K^m \cdot \nabla u_K^{m+1/2}) \right\|_{\ell^2} \left\| \widehat{u}^{m+1/2} \right\|_{\ell^2} \\ &\leq \frac{\epsilon}{2} (\|\widehat{u}^m\|_{\ell^2} + \|\widehat{u}^{m+1}\|_{\ell^2}). \end{aligned}$$

Upon substitution of this estimate in (81), and simple manipulations, we thus conclude that

$$\|\widehat{u}^{m+1}\|_{\ell^2} \leq \frac{1 + \frac{\Delta t \epsilon}{2}}{1 - \frac{\Delta t \epsilon}{2}} \|\widehat{u}^m\|_{\ell^2}.$$

We finally note that

$$\frac{1 + \frac{\Delta t \epsilon}{2}}{1 - \frac{\Delta t \epsilon}{2}} = 1 + \frac{\Delta t \epsilon}{1 - \frac{\Delta t \epsilon}{2}} \leq 1 + 2\Delta t \epsilon \leq \exp(2\Delta t \epsilon),$$

where we used the assumption $\Delta t, \epsilon \in (0, 1]$ in the first bound. The claimed error bounds on the recursive sequence thus follow from $\|\widehat{u}^{m+1}\|_{\ell^2} \leq \exp(2\Delta t \epsilon) \|\widehat{u}^m\|_{\ell^2}$, and (79), (80). ■

F.7 Proof of Lemma 41

Proof By Lemma 39, the neural network $\mathcal{NL}_{\overline{M}, \epsilon}$ in Algorithm 1 can be chosen to be of size

$$\text{size}(\mathcal{NL}_{\overline{M}, \epsilon}) \leq C|\mathcal{K}_K| \log(\overline{M}K/\epsilon), \quad \text{depth}(\mathcal{NL}_{\overline{M}, \epsilon}) \leq C \log(\overline{M}K/\epsilon).$$

We now recall that $|\mathcal{K}_K| = |\{k \in \mathbb{Z}^n \mid |k|_\infty \leq K\}| \sim_n K^n$, $\overline{M} \sim M$, $\epsilon \gtrsim \Delta t/MT \sim \min(K^{-r}, K^{-n/2-1}/M)/MT$, for $K \rightarrow \infty$, and hence

$$\text{size}(\mathcal{NL}_{\overline{M}, \epsilon}) \lesssim_{r,n} K^n \log(MTK), \quad \text{depth}(\mathcal{NL}_{\overline{M}, \epsilon}) \lesssim_{r,n} \log(MTK).$$

The fixed-point function $F(\widehat{w}) = F(\widehat{u}^m; \widehat{w})$ (see (78)) applied in the innermost loop of Algorithm (1), lines 11-13, can clearly be represented by a neural network of a size $\sim \text{size}(\mathcal{NL}_{\overline{M}, \epsilon})$, as it only differs from $\mathcal{NL}_{\overline{M}, \epsilon}$ by a linear input and a linear output. Composing $L \lesssim_{r,n} \log(MTK)$ of these networks, we obtain a representation $\psi^* : \widehat{u}^m \mapsto \widehat{u}^{m+1}$, corresponding to lines 9-14 of Algorithm 1, and such that

$$\text{size}(\psi^*) \leq CK^n \log(MTK)^2, \quad \text{depth}(\psi^*) \leq C \log(MTK)^2,$$

where $C = C(n, r) > 0$ is independent of M, T, K . Finally, composing

$$\psi = \underbrace{\psi^* \circ \dots \circ \psi^*}_{n_T\text{-fold}}$$

allows us to represent the entire algorithm 1 by a ReLU neural network. The estimate on the size of ψ follows from the fact that

$$n_T \sim T/\Delta t \sim T/\min(K^{-n/2-1}/M, K^{-r}) \lesssim_{r,d} MTK^r.$$

■

F.8 Proof of Lemma 42

Proof By definition of M , the claim is true for $m = 0$, since $\|\widehat{u}^0\|_{\ell^2} \leq M \leq 2M$. We prove the claim for $m > 0$ by induction: To this end, we assume that (73) holds true for some m . We aim to show that (73) also holds for $m + 1$. By the choice of $\overline{M} = 2(eM + 2)$, and the choice of $\Delta t, \epsilon$ in Algorithm 1, the assumptions of Lemma 40 are fulfilled (with M replaced by eM), and hence there exists a unique fixed point $\widehat{u}^{*,m+1} \in B_{eM}(0)$ of (71), i.e. such that

$$\frac{\widehat{u}_k^{*,m+1} - \widehat{u}_k^m}{\Delta t} = -\mathcal{NL}_{\overline{M},\epsilon}(\widehat{u}^m, \widehat{u}^{*,m+1/2})_k - \nu|k|^2 \widehat{u}_k^{*,m+1/2}.$$

Furthermore, the iterates $\widehat{w}^{\ell,m}$ defined by Algorithm 1 satisfy the estimates

$$\|\widehat{w}^{\ell,m} - \widehat{u}^{*,m+1}\|_{\ell^2} \leq \frac{\exp(2\Delta t \epsilon)}{2^\ell} \|\widehat{u}^m\|_{\ell^2}, \quad (82)$$

and

$$\|\widehat{w}^{\ell,m}\|_{\ell^2} \leq (1 + 2^{-\ell}) \exp(2\Delta t \epsilon) \|\widehat{u}^m\|_{\ell^2}. \quad (83)$$

In particular, with the choice of $L \geq -\log(\Delta t \epsilon)$ of Algorithm 1, we find that $\widehat{u}^{m+1} := \widehat{w}^{L,m}$ satisfies

$$\|\widehat{u}^{m+1}\|_{\ell^2} \leq (1 + \Delta t \epsilon) \exp(2\Delta t \epsilon) \|\widehat{u}^m\|_{\ell^2} \leq \exp(3\Delta t \epsilon) \|\widehat{u}^m\|_{\ell^2}.$$

By the induction hypothesis on $\|\widehat{u}^m\|_{\ell^2} \leq \exp(3m\Delta t \epsilon)M$, and the fact that $3(m+1)\Delta t \epsilon \leq 3T\epsilon \leq 1$, we can thus estimate

$$\|\widehat{u}^{m+1}\|_{\ell^2} \leq \exp(3(m+1)\Delta t \epsilon)M \leq eM.$$

This completes the induction step. ■

F.9 Proof of Lemma 43

Proof We note that upon Fourier transformation, the claim is equivalent to the following:

$$\frac{\widehat{u}^{m+1} - \widehat{u}^m}{\Delta t} = -\widehat{\mathbb{P}}_K \left(u_K^m \cdot \nabla u_K^{m+1/2} \right) - \nu|k|^2 \widehat{u}^{m+1/2} + \widehat{Q}^m, \quad (74')$$

where the remainder $\widehat{Q}^m \in \mathbb{C}^{\mathcal{K}}$ satisfies an estimate of the form

$$\sup_m \|\widehat{Q}^m\|_{\ell^2} \leq C\Delta t, \quad (75')$$

and $C > 0$ is independent of K , M and T .

By the choice of $\overline{M}, \epsilon > 0$ and $\Delta t > 0$ of Algorithm 1, the approximation of the non-linearity, $\mathcal{NL}_{\overline{M}, \epsilon}$, satisfies (69) and (70) with

$$CK^{n/2+1}\overline{M} \leq \frac{1}{\Delta t},$$

i.e. we have

$$\begin{cases} \|\mathcal{NL}_{\overline{M}, \epsilon}(\widehat{u}, \widehat{v}) - \widehat{\mathbb{P}}_K(u_K \cdot \nabla v_K)\|_{\ell^2} \leq \epsilon \leq \frac{\Delta t}{3}, \\ \|\mathcal{NL}_{\overline{M}, \epsilon}(\widehat{u}, \widehat{v}) - \mathcal{NL}_{\overline{M}, \epsilon}(\widehat{u}, \widehat{w})\|_{\ell^2} \leq \frac{1}{\Delta t} \|\widehat{v} - \widehat{w}\|_{\ell^2}, \end{cases} \quad (84)$$

for all $\widehat{u}, \widehat{v}, \widehat{w}$ such that $\|\widehat{u}\|_{\ell^2}, \|\widehat{v}\|_{\ell^2}, \|\widehat{w}\|_{\ell^2} \leq \overline{M} \equiv 2(eM + 2)$. Furthermore, by Lemma 42, we have $\|\widehat{u}^m\|_{\ell^2} \leq \overline{M}$ for all $m = 0, \dots, n_T$.

As in the proof of Lemma 42, for $m \in \{0, \dots, n_T\}$, we will in the following denote by $\widehat{u}^{*,m}$ the (unique) solution of the emulation scheme with non-linearity $\mathcal{NL}_{\overline{M}, \epsilon}$:

$$\frac{\widehat{u}_k^{*,m+1} - \widehat{u}_k^m}{\Delta t} = -\mathcal{NL}_{\overline{M}, \epsilon}(\widehat{u}^m, \widehat{u}^{*,m+1/2})_k - \nu|k|^2 \widehat{u}_k^{*,m+1/2}, \quad (85)$$

where $\widehat{u}^{*,m+1/2} := \frac{1}{2}(\widehat{u}^m + \widehat{u}^{*,m+1})$. By the estimate (82) on the fixed-point iteration, and the fact that $\exp(2\Delta t \epsilon)2^{-L} \leq e\Delta t \epsilon$ (by the choice of L in Algorithm 1), we then have

$$\|\widehat{u}^{m+1} - \widehat{u}^{*,m+1}\|_{\ell^2} \equiv \|\widehat{w}^{L,m} - \widehat{u}^{*,m+1}\|_{\ell^2} \leq 2eM\Delta t \epsilon \leq \frac{2e\Delta t^2}{3}. \quad (86)$$

Using the fact that $\widehat{u}^{*,m+1}$ solves (85) exactly, we readily obtain

$$\begin{aligned} \frac{\widehat{u}^{m+1} - \widehat{u}^m}{\Delta t} &= \frac{\widehat{u}^{m+1} - \widehat{u}^{*,m+1}}{\Delta t} + \frac{\widehat{u}^{*,m+1} - \widehat{u}^m}{\Delta t} \\ &= \widehat{Q}^m - \widehat{\mathbb{P}}_K\left(u_K^m \cdot \nabla u_K^{m+1/2}\right) - \nu|k|^2 \widehat{u}^{m+1/2}, \end{aligned}$$

where $\widehat{Q}^m = \widehat{Q}_1^m + \widehat{Q}_2^m + \widehat{Q}_3^m$ is given by

$$\begin{aligned} \widehat{Q}_1^m &:= \frac{\widehat{u}^{m+1} - \widehat{u}^{*,m+1}}{\Delta t}, \\ \widehat{Q}_2^m &:= \mathcal{NL}_{\overline{M}, \epsilon}(\widehat{u}^m, \widehat{u}^{m+1/2}) - \mathcal{NL}_{\overline{M}, \epsilon}(\widehat{u}^m, \widehat{u}^{*,m+1/2}), \\ \widehat{Q}_3^m &:= \widehat{\mathbb{P}}_K\left(u_K^m \cdot \nabla u_K^{m+1/2}\right) - \mathcal{NL}_{\overline{M}, \epsilon}(\widehat{u}^m, \widehat{u}^{m+1/2}). \end{aligned}$$

By (86), we can estimate $\|\widehat{Q}_1^m\|_{\ell^2} \leq \frac{2e}{3}\Delta t$. By the second estimate of (84), we similarly find $\|\widehat{Q}_2^m\|_{\ell^2} \leq \frac{e}{3}\Delta t$. By the first estimate of (84), we find $\|\widehat{Q}_3^m\|_{\ell^2} \leq \frac{1}{3}\Delta t$. Combining these estimates, we conclude that

$$\|\widehat{Q}^m\|_{\ell^2} \leq C\Delta t,$$

where $C = e + \frac{1}{3}$ is an absolute constant independent of M, T, m . ■

F.10 Proof of Lemma 44

Proof We denote $U^m := u(t^m)$, and $U^{m+1/2} := \frac{1}{2}[u(t^m) + u(t^{m+1})]$. Define $w_K^m := u_K^m - U^m$, and $w_K^{m+1/2} := \frac{1}{2}[w_K^m + w_K^{m+1}]$. Then w_K^m satisfies

$$\begin{aligned} \frac{w_K^{m+1} - w_K^m}{\Delta t} &= -\mathbb{P}_K \left(w_K^m \cdot \nabla U^{m+1/2} \right) - \mathbb{P}_K \left(U^m \cdot \nabla w_K^{m+1/2} \right) \\ &\quad + \nu \Delta w_K^{m+1/2} + R_K^m(u) - Q_K^m. \end{aligned}$$

Integrating against $w_K^{m+1/2}$, we find that

$$\begin{aligned} \frac{1}{2} \|w_K^{m+1}\|_{L^2}^2 - \frac{1}{2} \|w_K^m\|_{L^2}^2 &= -\langle w_K^m \cdot \nabla U^{m+1/2}, w_K^{m+1/2} \rangle_{L^2} - \nu \|\nabla w_K^{m+1/2}\|_{L^2}^2 \\ &\quad + \langle R_K^m(u), w_K^{m+1/2} \rangle_{L^2} - \langle Q_K^m, w_K^{m+1/2} \rangle_{L^2}. \end{aligned} \quad (87)$$

We can now estimate (using the Cauchy-Schwarz inequality)

$$\begin{aligned} \left| \langle w_K^m \cdot \nabla U^{m+1/2}, w_K^{m+1/2} \rangle_{L^2} \right| &\leq \frac{1}{2} \|w_K^{m+1}\|_{L^2}^2 + C \|w_K^m\|_{L^2}^2, \\ -\nu \|\nabla w_K^{m+1/2}\|_{L^2}^2 &\leq 0, \\ \left| \langle R_K^m(u), w_K^{m+1/2} \rangle_{L^2} \right| &\leq \|R_K^m(u)\|_{L^2}^2 + \|w_K^{m+1}\|_{L^2}^2 + \|w_K^m\|_{L^2}^2, \\ \left| \langle Q_K^m, w_K^{m+1/2} \rangle_{L^2} \right| &\leq \|Q_K^m\|_{L^2}^2 + \|w_K^{m+1}\|_{L^2}^2 + \|w_K^m\|_{L^2}^2. \end{aligned}$$

Next, we recall that

$$\|R_K^m(u)\|_{L^2}^2 \leq C(\Delta t^2 + K^{-2r}),$$

by (68). Similarly, by (75), we have

$$\|Q_K^m\|_{L^2}^2 \leq C\Delta t^2,$$

where $C > 0$ is independent of K , Δt and m . Combining these estimates with (87), we find that there exists a constant $C > 0$, such that

$$\|w_K^{m+1}\|_{L^2}^2 \leq (1 + C\Delta t) \|w_K^m\|_{L^2}^2 + C\Delta t (\Delta t^2 + K^{-2r}).$$

Applying the discrete Gronwall inequality, we conclude that

$$\max_{m=0, \dots, n_T} \|w_K^m\|_{L^2}^2 \leq e^{CT} \|w_K^0\|_{L^2}^2 + Ce^{CT} (\Delta t^2 + K^{-2r}).$$

Noting that $w_K^0 = u_K^0 - u(0)$, we can use the assumed bound on $\|u\|_{C_t H_x^r}$ to conclude that

$$\begin{aligned} \|w_K^0\|_{L^2} &\leq \|u_K^0 - \mathbb{P}_K u(0)\|_{L^2} + \|(1 - \mathbb{P}_K)u(0)\|_{L^2} \\ &\leq C \|\hat{u}^0 - \hat{\mathbb{P}}_K u(0)\|_{\ell^2} + CK^{-2r}, \end{aligned}$$

where $C = C(n, r, \|u\|_{C_t H_x^r}) > 0$ is independent of K . Using also that $\Delta t \leq K^{-r}$, we finally obtain

$$\max_{m=0, \dots, n_T} \|w_K^m\|_{L^2}^2 \leq Ce^{CT} \left(K^{-2r} + \|\hat{u}^0 - \hat{\mathbb{P}}_K u(0)\|_{\ell^2} \right),$$

for a constant $C > 0$, depending only on $\|u\|_{C_t H_x^r}$, $\|u\|_{C_t^1 H_x^{r-1}}$, n , r . This implies the claimed estimate. \blacksquare

F.11 Proof of Lemma 46

Proof We now note that

$$\begin{aligned} \|\mathcal{D}_Y \circ \psi \circ \mathcal{E}_X - \Psi^\dagger\|_{L^2(\mu)} &\leq \|\mathcal{D}_Y \circ \psi \circ \mathcal{E}_X - \mathcal{D}_Y \circ \mathcal{E}_Y \circ \Psi^\dagger\|_{L^2(\mu)} \\ &\quad + \|\mathcal{D}_Y \circ \mathcal{E}_Y \circ \Psi^\dagger - \Psi^\dagger\|_{L^2(\mu)} \\ &= \|\mathcal{D}_Y \circ \psi \circ \mathcal{E}_X - \mathcal{D}_Y \circ \mathcal{E}_Y \circ \Psi^\dagger\|_{L^2(\mu)} \\ &\quad + \|P_Y - \text{Id}\|_{L^2(\Psi^\dagger_\# \mu)}, \end{aligned}$$

where $P_Y = \mathcal{D}_Y \circ \mathcal{E}_Y : \mathcal{Y} \rightarrow \mathcal{Y}$ is the (orthogonal) PCA projection. The second term is the PCA projection error for the push-forward measure $\Psi^\dagger_\# \mu$. To estimate the first term, we first note that $\mathcal{D}_Y : \mathbb{R}^{d_Y} \rightarrow \mathcal{Y}$ is a linear isometric injection, hence

$$\|\mathcal{D}_Y \circ \psi \circ \mathcal{E}_X - \mathcal{D}_Y \circ \mathcal{E}_Y \circ \Psi^\dagger\|_{L^2(\mu)} = \|\psi \circ \mathcal{E}_X - \mathcal{E}_Y \circ \Psi^\dagger\|_{L^2(\mu)}.$$

Given $K \in \mathbb{N}$, we now let $\psi : \mathbb{R}^{d_X} \rightarrow \mathbb{R}^{d_Y}$ be a ReLU neural network of the form

$$\psi = (\mathcal{E}_Y \circ \mathcal{F}_K^{-1}) \circ \widehat{\psi} \circ (\mathcal{F}_K \circ \mathcal{D}_X),$$

where $\mathcal{F}_K^{-1} : [\mathbb{C}^n]^\mathcal{K} \rightarrow L^2(\mathbb{T}^n; \mathbb{R}^n)$ and $\mathcal{F}_K : L^2(\mathbb{T}^n; \mathbb{R}^n) \rightarrow [\mathbb{C}^n]^\mathcal{K}$ denote the (truncated) inverse and forward Fourier transforms, respectively; we recall that

$$\begin{aligned} \mathcal{F}_K^{-1}(\widehat{u}) &= \sum_{|k|_\infty \leq K} \widehat{u}_k e^{i\langle k, x \rangle}, \\ \mathcal{F}_K(u) &= \left\{ \frac{1}{(2\pi)^n} \int_{\mathbb{T}^n} u(x) e^{-i\langle k, x \rangle} dx \right\}_{|k|_\infty \leq K}. \end{aligned}$$

Since $\mathcal{E}_Y \circ \mathcal{F}_K^{-1} : [\mathbb{C}^n]^\mathcal{K} \rightarrow \mathbb{R}^{d_Y}$ and $\mathcal{F}_K \circ \mathcal{D}_X : \mathbb{R}^{d_X} \rightarrow [\mathbb{C}^n]^\mathcal{K}$ are linear mappings, they can be represented by matrix multiplications, with matrices $Q \in \mathbb{R}^{2n|\mathcal{K}| \times d_Y}$, and $R \in \mathbb{R}^{d_X \times 2n|\mathcal{K}|}$, respectively; i.e. we can write

$$\psi(w) = Q \cdot \widehat{\psi}(R \cdot w),$$

under the canonical identification $[\mathbb{C}^n]^\mathcal{K} \simeq \mathbb{R}^{2n|\mathcal{K}|}$. In particular, if $\widehat{\psi}$ is a ReLU neural network of a given size, then ψ satisfies

$$\text{size}(\psi) \leq \text{size}(\widehat{\psi}) + 2n|\mathcal{K}|(d_Y + d_X), \quad \text{depth}(\psi) = \text{depth}(\widehat{\psi}).$$

Let us assume that $\widehat{\psi}$ is a ReLU neural network as in Proposition 27. Then we can estimate:

$$\begin{aligned} \|\psi \circ \mathcal{E}_X - \mathcal{E}_Y \circ \Psi^\dagger\|_{L^2(\mu)} &= \|(\mathcal{E}_Y \circ \mathcal{F}_K^{-1}) \circ \widehat{\psi} \circ (\mathcal{F}_K \circ \mathcal{D}_X) \circ \mathcal{E}_X - \mathcal{E}_Y \circ \Psi^\dagger\|_{L^2(\mu)} \\ &\leq \|\mathcal{F}_K^{-1} \circ \widehat{\psi} \circ \mathcal{F}_K \circ (\mathcal{D}_X \circ \mathcal{E}_X) - \Psi^\dagger\|_{L^2(\mu)}. \end{aligned}$$

Denote the Fourier projection $P_K = \mathcal{F}_K^{-1} \circ \mathcal{F}_K$, and the PCA projection $P_X = \mathcal{D}_X \circ \mathcal{E}_X$, and estimate the last term by

$$\begin{aligned} &\leq \|\mathcal{F}_K^{-1} \circ \widehat{\psi} \circ \mathcal{F}_K \circ P_X - \mathcal{F}_K^{-1} \circ \mathcal{F}_K \circ \Psi^\dagger\|_{L^2(\mu)} + \|(1 - P_K)\Psi^\dagger\|_{L^2(\mu)} \\ &\leq \|\widehat{\psi} \circ \mathcal{F}_K \circ P_X - \mathcal{F}_K \circ \Psi^\dagger\|_{L^2(\mu)} + \|(1 - P_K)\Psi^\dagger\|_{L^2(\mu)}. \end{aligned} \tag{88}$$

The integrand in the first term is of the form

$$\|\widehat{\psi}(\widehat{u}^0) - \widehat{u}(T)\|_{\ell^2}, \quad u \in \mathcal{H}_{M,T}^r,$$

where $\widehat{u}^0 = [\widehat{P_{\mathcal{X}}u(0)}] \in [\mathbb{C}^n]^{\mathcal{K}}$ collects the Fourier coefficients with $|k|_{\infty} \leq K$ of the PCA projected $P_{\mathcal{X}}u(0)$. By Proposition 27, there exists a constant $C > 0$, depending only on M, r, T , such that

$$\|\widehat{\psi}(\widehat{u}^0) - \widehat{u}(T)\|_{\ell^2} \leq C(K^{-r} + \|\widehat{u}^0 - \widehat{u}(0)\|_{\ell^2}).$$

Using also the definition of $\widehat{u}^0 = [\widehat{P_{\mathcal{X}}u(0)}]$, we can further estimate the right-hand side in the form

$$\|\widehat{\psi}(\widehat{u}^0) - \widehat{u}(T)\|_{\ell^2} \leq C(K^{-r} + \|(1 - P_{\mathcal{X}})u(0)\|_{L^2_x}), \quad \forall u \in \mathcal{H}_{M,T}^r.$$

The other term in (88), $\|(1 - P_K)\Psi^\dagger\|_{L^2(\mu)}$, can be estimated using the uniform bound $\|\Psi^\dagger(u(0))\|_{H_x^r} \leq M$, for $u \in \mathcal{H}_{M,T}^r$, to obtain

$$\|(1 - P_K)\Psi^\dagger\|_{L^2(\mu)} \leq CK^{-r},$$

where $C = C(M, r, n) > 0$ depends only on $r, M, n > 0$. Upon substitution in (88), we thus obtain an estimate of the form

$$\|\psi \circ \mathcal{E}_{\mathcal{X}} - \mathcal{E}_{\mathcal{Y}} \circ \Psi^\dagger\|_{L^2(\mu)} \leq CK^{-r} + C\|\text{Id} - P_{\mathcal{X}}\|_{L^2(\mu)},$$

where $C = C(M, r, n, T) > 0$ is independent of K . We point out in passing that $\|\text{Id} - P_{\mathcal{X}}\|_{L^2(\mu)}$ is precisely the PCA projection error on \mathcal{X} . Combining the above estimates, we obtain the claimed upper bound

$$\|\Psi - \Psi^\dagger\|_{L^2(\mu)} \leq CK^{-r} + C\|\text{Id} - P_{\mathcal{X}}\|_{L^2(\mu)} + \|\text{Id} - P_{\mathcal{Y}}\|_{L^2(\Psi^\dagger_{\#}\mu)},$$

where ψ is a neural network of size

$$\text{size}(\psi) \leq CK^n (K^r \log(K)^2 + (d_{\mathcal{X}} + d_{\mathcal{Y}})), \quad \text{depth}(\psi) \leq CK^r \log(K),$$

and $C = C(M, r, T, n) > 0$ is a constant independent of $K, d_{\mathcal{X}}, d_{\mathcal{Y}}$. ■

Appendix G. A lemma on fixed points

The following lemma is a simple variant of the well-known Banach fixed point theorem. For completeness, we provide a detailed proof (although the result or a variant thereof may likely have appeared elsewhere).

Lemma 49 *Let X be a Banach space and fix $M > 0$. Let $B_M(0) \subset X$ denote the closed ball of radius M . Let $F : B_M(0) \rightarrow X$ be a contractive mapping with $\gamma = \text{Lip}(F) < 1$, and assume that $\|F(0)\| \leq (1 - \gamma)M$. Denote $R := \|F(0)\|/(1 - \gamma) \leq M$. Then there exists a unique fixed point $w^\infty \in B_M(0)$ of F , and the norm of w^∞ is bounded by $\|w^\infty\| \leq R$. Furthermore, the recursively defined sequence $w^0 := 0$, $w^{k+1} := F(w^k)$, is well-defined for all $k = 0, 1, \dots$, and satisfies the a priori estimates:*

$$\|w^k\| \leq R, \quad \|w^{k+1} - w^k\| \leq \gamma^k R, \quad \|w^k - w^\infty\| \leq \gamma^k R.$$

Proof Set $w^0 := 0$. Given w^k , we define $w^{k+1} := F(w^k)$, if $w^k \in B_M(0)$, otherwise we terminate the sequence. Let

$$k_0 = \sup \left\{ k \in \mathbb{N} \mid w^1, \dots, w^k \text{ are well-defined} \right\}.$$

We claim that $k_0 = \infty$, and that in fact $\|w^k\| \leq R$ for all $k \in \mathbb{N}$. Suppose this was not the case. Then we must have $k_0 < \infty$, and $w^{k_0} \notin B_M(0)$. We now note that for any $k < k_0$, we have

$$\begin{aligned} \|w^{k+1} - w^k\| &= \|F(w^k) - F(w^{k-1})\| \leq \gamma \|w^k - w^{k-1}\| \\ &\leq \dots \leq \gamma^k \|w^1 - w^0\| = \gamma^k \|F(0)\|, \end{aligned}$$

where we used the Lipschitz continuity of F , and the fact that $w^0 = 0$ and $w^1 = F(w^0) = F(0)$, in the last step. Combined with the telescoping sum $w^{k_0} = \sum_{\ell=0}^{k_0-1} [w^{\ell+1} - w^\ell]$, this estimate implies that

$$\|w^{k_0}\| \leq \sum_{\ell=0}^{k_0-1} \|w^{\ell+1} - w^\ell\| \leq \sum_{\ell=0}^{k_0-1} \gamma^\ell \|F(0)\| \leq \frac{\|F(0)\|}{1-\gamma} = R,$$

and hence $w^{k_0} \in B_R(0) \subset B_M(0)$. This clearly contradicts the definition of k_0 . Hence, we must have $k_0 = \infty$. Thus, the recursive sequence $w^0 = 0$, $w^{k+1} = F(w^k)$ is well-defined for any $k = 0, 1, \dots$, and the above argument shows furthermore that $\|w^k\| \leq R$, for all k . In fact, since the terms of the series $\sum_{\ell=0}^{\infty} [w^{\ell+1} - w^\ell]$ are summable in norm, we conclude that there exists a limit

$$w^\infty = \sum_{\ell=0}^{\infty} [w^{\ell+1} - w^\ell] = \lim_{k \rightarrow \infty} \sum_{\ell=0}^{k-1} [w^{\ell+1} - w^\ell] = \lim_{k \rightarrow \infty} w^k.$$

Passing to the limit, it is now straight-forward to show that w^∞ is a fixed point of F , that $\|w^\infty\| \leq R$, and that we have

$$\begin{aligned} \|w^k - w^\infty\| &= \|F(w^{k-1}) - F(w^\infty)\| \leq \gamma \|w^{k-1} - w^\infty\| \\ &\leq \dots \leq \gamma^k \|w^0 - w^\infty\| = \gamma^k \|w^\infty\| \leq \gamma^k R, \end{aligned}$$

for any $k \in \mathbb{N}$. To show that w^∞ is unique, we assume that v^∞ is another fixed point of F . Then

$$\|w^\infty - v^\infty\| = \|F(w^\infty) - F(v^\infty)\| \leq \gamma \|w^\infty - v^\infty\|,$$

is only possible if $\|w^\infty - v^\infty\| = 0$. ■