



MOX-Report No. 86/2024

**On the latent dimension of deep autoencoders for reduced order  
modeling of PDEs parametrized by random fields**

Franco, N.R.; Fraulin, D.; Manzoni, A.; Zunino, P.

MOX, Dipartimento di Matematica  
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

[mox-dmat@polimi.it](mailto:mox-dmat@polimi.it)

<https://mox.polimi.it>

# On the latent dimension of deep autoencoders for reduced order modeling of PDEs parametrized by random fields

Nicola Rares Franco<sup>a</sup>, Daniel Fraulin<sup>a</sup>, Andrea Manzoni<sup>a</sup>, Paolo Zunino<sup>a</sup>

<sup>a</sup>*MOX – Department of Mathematics, Politecnico di Milano, Piazza Leonardo da Vinci 32, I-20133, Milano, Italy*

---

## Abstract

Deep Learning is having a remarkable impact on the design of Reduced Order Models (ROMs) for Partial Differential Equations (PDEs), where it is exploited as a powerful tool for tackling complex problems for which classical methods might fail. In this respect, deep autoencoders play a fundamental role, as they provide an extremely flexible tool for reducing the dimensionality of a given problem by leveraging on the nonlinear capabilities of neural networks. Indeed, starting from this paradigm, several successful approaches have already been developed, which are here referred to as Deep Learning-based ROMs (DL-ROMs). Nevertheless, when it comes to stochastic problems parameterized by random fields, the current understanding of DL-ROMs is mostly based on empirical evidence: in fact, their theoretical analysis is currently limited to the case of PDEs depending on a finite number of (deterministic) parameters. The purpose of this work is to extend the existing literature by providing some theoretical insights about the use of DL-ROMs in the presence of stochasticity generated by random fields. In particular, we derive explicit error bounds that can guide domain practitioners when choosing the latent dimension of deep autoencoders. We evaluate the practical usefulness of our theory by means of numerical experiments, showing how our analysis can significantly impact the performance of DL-ROMs.

*Keywords:* Deep Learning, Reduced Order Modeling, parametrized PDEs, random fields, autoencoders

---

## 1. Introduction

Aside from its striking impact on areas such as data science, language processing, and computer vision, Deep Learning is now becoming ubiquitous in all branches of science, with applications ranging from economics [46, 41, 34] to medicine [62, 43, 6], biology [4, 65, 58], chemistry [67, 57, 59], physics [29, 14, 8], mathematics [15, 27, 7] and engineering [38, 56, 50]. In most cases, researchers exploit Deep Learning to make up for our limited understanding and limited computational resources, typically trying to find a suitable compromise between domain knowledge and data-driven approaches. In this contribution, we shall focus our attention on a specific research area, typical of engineering applications, which concerns the development of Reduced Order Models (ROMs).

Simply put, ROMs are model surrogates that aim at replacing expensive numerical simulations with accurate approximations that are obtained at a reduced computational cost. From a practical point of view, ROMs can be remarkably helpful whenever dealing with real-time many-query applications, such as those characterizing digital twins [36], optimal control [54] and uncertainty quantification [30, 12]. As of today, the literature is filled with plenty of different ROM approaches, and whether to choose one or another is typically problem dependent. Here, we shall focus on Deep-Learning based ROMs (DL-ROMs), specifically addressing the framework proposed in [22, 19], which leverages the use of deep autoencoders and has already reported numerous successful applications; see, e.g., [25, 23, 26, 11].

The driving idea behind the DL-ROM approach is to exploit the nonlinear approximation capabilities of deep autoencoders to perform a suitable dimensionality reduction, allowing the representation of complex high-fidelity solutions as small vectors in some latent space of dimension  $n$ . In general, the choice of this latent dimension is problem specific and its value reflects the intrinsic properties of the so-called *solution manifold*. In the case of PDEs that depend on a finite number of deterministic parameters, this fact has

already been thoroughly studied in [19]. There, the authors characterize the latent dimension of the DL-ROM by investigating the behavior of the so-called *manifold  $n$ -width* [17]. Given a compact parameter space  $\Theta \subseteq \mathbb{R}^p$ , a Hilbert state space  $(V, \|\cdot\|)$  and a parameter-to-solution operator  $\Theta \ni \boldsymbol{\mu} \mapsto u_{\boldsymbol{\mu}} \in V$ , the latter can be written as

$$\inf_{\substack{\Psi' \in \mathcal{E}(V, \mathbb{R}^n) \\ \Psi \in \mathcal{D}(\mathbb{R}^n, V)}} \sup_{\boldsymbol{\mu} \in \Theta} \|u_{\boldsymbol{\mu}} - \Psi(\Psi'(u_{\boldsymbol{\mu}}))\| \quad (1)$$

where  $\Psi' : V \rightarrow \mathbb{R}^n$  and  $\Psi : \mathbb{R}^n \rightarrow V$  are the encoder and decoder networks, respectively, each varying in a suitable class of admissible architectures correspondingly named  $\mathcal{E}(V, \mathbb{R}^n)$  and  $\mathcal{D}(\mathbb{R}^n, V)$ .

The purpose of this work is to extend the analysis proposed in [19] to address the case of stochastic PDEs, where the deterministic parameters are replaced by some random field  $\boldsymbol{\mu} \sim \mathbb{P}$ . In general, this extension presents two major challenges. First of all, the realizations of the input field  $\boldsymbol{\mu}$  might be arbitrarily large in norm, which removes any form of compactness and makes the arguments in [19] inapplicable. On top of that, the stochasticity introduced by the random field makes the PDE formally depend on an infinite number of parameters, i.e.  $p = +\infty$ . Then, the bounds provided in [19] become meaningless.

To address these difficulties, we shall first replace the optimization problem in (1) with its probabilistic counterpart, that is

$$\inf_{\substack{\Psi' \in \mathcal{E}(V, \mathbb{R}^n) \\ \Psi \in \mathcal{D}(\mathbb{R}^n, V)}} \mathbb{E}_{\boldsymbol{\mu} \sim \mathbb{P}} \|u_{\boldsymbol{\mu}} - \Psi(\Psi'(u_{\boldsymbol{\mu}}))\| \quad (2)$$

where  $\mathbb{E}$  is the expectation operator. Then, (2) measures the extent to which an autoencoder with latent dimension  $n$  can approximate the high-fidelity solutions arising from the PDE. Our purpose is to characterize the quantitative behavior of (2) and thus provide practical insights that can guide domain practitioners in the complex design of DL-ROM architectures.

More precisely, we shall prove the following facts. If the PDE depends on a finite number of stochastic parameters,  $p \in \mathbb{N}$ , then the autoencoders can achieve arbitrary accuracy in any latent dimension  $n \geq p$ . In general, this result is much stronger than the one obtained in the deterministic setting, since the optimal bound in that case is  $n \geq 2p + 1$ : see Theorem 3 in [19] for further details. On the contrary, if the PDE depends on a general random field (informally,  $p = +\infty$ ), then we can characterize the decay of the reconstruction error (2) in terms of the eigenvalues of the covariance operator of the input random field,  $\boldsymbol{\mu}$ , and of the output field,  $u_{\boldsymbol{\mu}}$ , respectively. In particular, our analysis shows how deep autoencoders can outperform the performance of linear methods by benefiting from the intrinsic regularities available in both fields.

The paper is organized as follows. First, in Section 2, we formally introduce the problem of reduced order modeling for stochastic PDEs, with a brief overview about the DL-ROM approach, while, in Section 3, we present the mathematical tools that are needed for our construction. Our main contribution can be found in Section 4, where we derive several results about the latent dimension of deep autoencoders for PDEs parameterized by random fields. Finally, we devote Section 5 to numerical experiments.

## 2. Problem setup

Let  $\Omega \subset \mathbb{R}^d$  be a bounded domain. We are given a parametrized boundary value problem with random coefficients, e.g.

$$\begin{cases} \mathcal{A}_{\boldsymbol{\mu}} u_{\boldsymbol{\mu}} = f_{\boldsymbol{\mu}} & \text{in } \Omega \\ \mathcal{B}_{\boldsymbol{\mu}} u_{\boldsymbol{\mu}} = g_{\boldsymbol{\mu}} & \text{on } \partial\Omega, \end{cases} \quad (3)$$

with  $\mathcal{A}_{\boldsymbol{\mu}}, \mathcal{B}_{\boldsymbol{\mu}}, f_{\boldsymbol{\mu}}, g_{\boldsymbol{\mu}}$  parameter dependent operators and problem data, respectively. Here,  $\boldsymbol{\mu}$  is a suitable random field that parameterizes the PDE and introduces a corresponding form of stochasticity. For simplicity, we assume that  $\boldsymbol{\mu}$  is defined over  $\Omega$ , although such a restriction is not necessary. Notice also that we limit our attention to steady PDEs: for a deeper discussion about this aspect, we refer the reader to the remark at the end of this Section.

We assume to have at our disposal a trusted high-fidelity numerical solver – possibly expensive – that can approximate the solution of (3) for each fixed realization of the random field  $\boldsymbol{\mu}$ . We further assume that all numerical simulations produced by such solver live in a common high-fidelity state space  $V_h \subset L^2(\Omega)$  of

dimension  $\dim(V_h) = N_h$ . Let  $\{\xi_i\}_{i=1}^{N_h}$  be the basis of  $V_h$ . Then, for each realization of  $\mu$ , the numerical solver provides a way to compute a set of basis coefficients

$$\mathbf{u}_\mu^h := [c_\mu^{(1)}, \dots, c_\mu^{(N_h)}] \quad \text{such that} \quad u_\mu \approx \sum_{i=1}^{N_h} c_\mu^{(i)} \xi_i.$$

In this sense, the numerical solver, also known as the Full Order Model (FOM), defines a solution operator

$$\mathcal{G}_h : \mu \mapsto \mathbf{u}_\mu^h \in \mathbb{R}^{N_h}$$

in a very natural way. The efficient approximation of such an operator can be of great interest in many-query and real-time applications, where multiple calls to the FOM become computationally unbearable. Thus, the main interest becomes finding a cheaper surrogate  $\mu \mapsto \Phi(\mu)$ , called the reduced order model (ROM), such that  $\Phi(\mu) \approx \mathbf{u}_\mu^h$ . More precisely, let  $\mathbb{P}$  denote the probability law of  $\mu$ , and let  $\|\cdot\|_{V_h}$  be the norm over  $\mathbb{R}^{N_h}$  induced by the  $L^2$ -norm over the state space  $V_h$ . Then, giving some tolerance  $\varepsilon > 0$ , one seeks to construct a suitable ROM for which

$$\mathbb{E}_{\mu \sim \mathbb{P}} \|\mathbf{u}_\mu^h - \Phi(\mu)\|_{V_h} < \varepsilon.$$

### 2.1. Reduced order modeling and the DL-ROM approach

The development and construction of accurate ROMs is an extremely active area of research: as of today, the literature features a very broad spectrum of approaches to model order reduction, such as the Reduced Basis method [45, 52, 31, 63] and its non-intrusive data-driven variations, e.g. [32, 28], adaptive ROMs [3, 10, 48, 37], hybrid techniques based on closure modeling [35, 64], deep learning-based ROMs [19, 22, 24, 9, 51], and many others, each with their own benefits and guarantees.

Here, we shall focus primarily on the case of Deep Learning based ROMs (DL-ROMs), following the framework introduced in [22, 19]. These can be of particular interest whenever: i) intrusive approaches are not available, either because they would entail expensive subroutines, such as hyperreduction, or because the FOM is not accessible; ii) linear methods based on, e.g., Principal Orthogonal Decomposition (POD), fail because of intrinsic complexities entailed by the solution operator (see, e.g., the well-known phenomenon of *slow-decay in the Kolmogorov  $n$ -width* [19, 47, 55]).

For simplicity, let us assume that the random field at input,  $\mu$ , has been discretized at the same level as PDE solutions, so that each realization  $\mu \in L^2(\Omega)$  is formally replaced by some  $\boldsymbol{\mu}^h \in \mathbb{R}^{N_h}$ . In their original formulation, DL-ROMs are characterized by three deep neural network architectures,  $\Psi', \Psi, \phi$ , which operate as

$$\Psi' : \mathbb{R}^{N_h} \rightarrow \mathbb{R}^n, \quad \Psi : \mathbb{R}^n \rightarrow \mathbb{R}^{N_h}, \quad \phi : \mathbb{R}^{N_h} \rightarrow \mathbb{R}^n,$$

where  $n \in \mathbb{N}$  is the so-called latent dimension of the model. The idea is that the first two networks are responsible for learning the fundamental features that characterize PDE solutions; conversely, the third map,  $\phi$ , is left to learn the way in which the input  $\boldsymbol{\mu}^h$  affects the output  $\mathbf{u}_{\boldsymbol{\mu}^h}^h$ . This is achieved by constructing the three networks so that

$$\Psi(\Psi'(\mathbf{u}_{\boldsymbol{\mu}^h}^h)) \approx \mathbf{u}_{\boldsymbol{\mu}^h}^h \quad \text{and} \quad \phi(\boldsymbol{\mu}^h) \approx \Psi'(\mathbf{u}_{\boldsymbol{\mu}^h}^h). \quad (4)$$

Then, the parameter-to-solution operator is approximated as

$$\Phi(\boldsymbol{\mu}^h) := \Psi(\phi(\boldsymbol{\mu}^h)) \approx \mathbf{u}_{\boldsymbol{\mu}^h}^h. \quad (5)$$

In this sense, the accuracy of the DL-ROM is ultimately determined by that of the combined network  $\Phi := \Psi \circ \phi$ , while the role of  $\Psi'$  is only auxiliary. However, such a splitting can be beneficial (see, e.g., Figure 22 in [22]) as it makes the two blocks,  $\Psi$  and  $\phi$ , tackle the different complexities that characterize model order reduction: on the one hand, the spatial complexity of the solutions; on the other hand, the intricate dependency of the solutions with respect to the input parameters. In this sense, the presence of  $\Psi'$  is fundamental, as it allows us to decouple the problem.

Following the conventions of the Deep Learning literature, we refer to the composition  $\Psi' \circ \Psi$  as an *autoencoder* architecture of latent dimension  $n$ , while the two maps,  $\Psi'$  and  $\Psi$ , are referred to as encoder and decoder, respectively. The terminology comes from the fact that the two maps ultimately provide a way to represent high-fidelity solutions as small vectors in  $\mathbb{R}^n$ , in fact, in  $\Psi(\Psi'(\mathbf{u}_{\boldsymbol{\mu}^h}^h)) \approx \mathbf{u}_{\boldsymbol{\mu}^h}^h$  and  $\Psi'(\mathbf{u}_{\boldsymbol{\mu}^h}^h) \in \mathbb{R}^n$ .

Then it is clear that the choice of the appropriate value of  $n$  becomes fundamentally important. The purpose of this work is to provide additional information on this aspect, with a particular focus on the case of PDEs parametrized by random fields. In the next subsection, we shall further motivate this fact with a practical example.

For the sake of completeness, before coming to the main objective of this work, let us conclude this overview with a few words about the training and implementation of DL-ROMs. The first step consists of exploiting the FOM to generate a collection of trusted samples

$$\{\boldsymbol{\mu}_i^h, \mathbf{u}_i^h\}_{i=1}^{N_{\text{train}}},$$

where  $\mathbf{u}_i^h = \mathcal{G}_h(\boldsymbol{\mu}^h)$ , which serve as training data for the three networks in the DL-ROM pipeline. More precisely, once the architectures have been designed, the three modules are trained by minimizing a suitable loss function such as the one below,

$$\begin{aligned} \mathcal{L}(\Psi, \Psi', \phi) = \frac{1}{N_{\text{train}}} \sum_{i=1}^{N_{\text{train}}} & \left( \alpha_1 \|\mathbf{u}_i^h - \Psi(\phi(\boldsymbol{\mu}_i^h))\|_{V_h}^2 + \right. \\ & \alpha_2 \|\mathbf{u}_i^h - \Psi'(\mathbf{u}_i^h)\|_{V_h}^2 + \\ & \left. \alpha_3 \|\Psi'(\mathbf{u}_i^h) - \phi(\boldsymbol{\mu}_i^h)\|_{\mathbb{R}^n}^2 \right). \end{aligned} \quad (6)$$

Here,  $\alpha_1, \alpha_2, \alpha_3 \geq 0$  are suitable weights that are used to define the loss function. The idea is that by minimizing (6), one would automatically enforce both (4) and (5). It should be noted that since the term multiplied by  $\alpha_1$  contains information about the actual accuracy of the ROM, the other two can be seen as regularizers. In this sense, a good rule of thumb is to increase  $\alpha_2$  if the solutions to the PDE are characterized by a complex dependency with respect to the space variable  $\mathbf{x} \in \Omega$ ; conversely, one can leverage  $\alpha_3$  whenever the parameter dependency of the PDE becomes highly involved.

In general, training of the three networks can be achieved simultaneously or in multiple steps. For example, in [22], the authors set  $\alpha_2 = 0$  and proceed with a single training; conversely, in [19], the authors propose a two-stage training phase: first with  $\alpha_1 = \alpha_3 = 0$  and then with  $\alpha_1 = \alpha_2 = 0$ . Here, we will keep all the weights active  $\alpha_1, \alpha_2, \alpha_3 > 0$ , and opt for a one-shot training routine.

After the training phase, the DL-ROM can efficiently approximate high-fidelity solutions in milliseconds. The quality of such an approximation is typically assessed by relying on a suitable *test set*, that is, by comparing the outputs of the two models, the DL-ROM and the FOM, for new independent realizations of the input field.

## 2.2. Choosing the latent dimension: a motivating example

To further motivate our analysis and anticipate the practical impact of our results, let us look at a simple problem featuring a PDE with finitely many random parameters. As in [19], we consider the following boundary value problem defined over the unit square  $\Omega = (0, 1)^2$

$$-\nabla \cdot (\sigma_{\boldsymbol{\mu}} \nabla u) = f_{\boldsymbol{\mu}} \quad \text{in } \Omega \quad (7)$$

complemented with a constant Dirichlet boundary condition,  $u \equiv 0.1$  on  $\partial\Omega$ . Here, the PDE depends on three random parameters,  $\boldsymbol{\mu} = [\mu_1, \mu_2, \mu_3]$ , which affect the permeability field  $\sigma = \sigma_{\boldsymbol{\mu}}$  and the right-hand-side  $f = f_{\boldsymbol{\mu}}$ . Those are defined as

$$\sigma_{\boldsymbol{\mu}}(\mathbf{x}) := \frac{1}{2} + \mu_1 \mathbb{1}_{\Omega_0}(\mathbf{x}), \quad f_{\boldsymbol{\mu}}(\mathbf{x}) := \frac{1}{2\pi\epsilon^2} \exp\left(-\frac{1}{2\epsilon^2} |\mathbf{x} - [\mu_2, \mu_3]^T|^2\right),$$

where  $\epsilon := 0.01$  and  $\Omega_0 \subset \Omega$  is as in Figure 1a. The first parameter,  $\mu_1$ , can attain any random value in  $[1, 4]$ , and has the effect of modifying the permeability of the subdomains; conversely,  $\mu_2$  and  $\mu_3$  are responsible for the random location of the concentrated source  $f_{\boldsymbol{\mu}}$ , and they are allowed to range from 0.1 to 0.9. More precisely, we endow the overall parameter space with a uniform probability distribution supported over  $[1, 4] \times [0.1, 0.9]^2 \subset \mathbb{R}^3$ .




Figure 1: Domain of definition (a) and randomly generated PDE solutions (b-d) for Problem (7).

To construct our trusted high-fidelity solver, we rely on a Finite Element discretization of (7) using continuous piecewise linear elements defined over a structured triangular grid of stepsize  $h \approx 0.0236$ , which results in a FOM dimension of  $N_h = 3721$ . We run the FOM multiple times to generate a total of 2000 random solutions (1500 for training and 500 for testing): we refer to Figure 1b-e for a few examples. Our purpose is to exploit these data to explore the behavior of two different dimensionality reduction techniques. POD, which here stands as a representative of linear ROMs, and deep autoencoders, which, instead, are at the core of DL-ROMs.

Precisely, we compute the reconstruction error (estimated via classical Monte Carlo over the test set)

$$\mathbb{E}_{\mu \sim \mathbb{P}} \|\mathbf{u}_\mu^h - \Psi(\Psi'(\mathbf{u}_\mu^h))\|_{V_h},$$

obtained by employing POD (i.e., by letting  $\Psi' \equiv \mathbf{V}^T$  and  $\Psi \equiv \mathbf{V}$  be nothing but the POD projectors) or deep autoencoders (in which case both  $\Psi$  and  $\Psi'$  are DNNs). We repeat this computation for different choices of the latent dimension, namely  $n = 1, 2, \dots, 40$ , to investigate the impact of the ROM dimension on the quality of the reconstruction. We mention that while the POD matrix is uniquely determined by  $n$ , this is not the case for deep autoencoders, as users can play with several components of the architecture (number of layers, intermediate neurons, layer type, etc.). To highlight this fact, we train three different autoencoder architectures for each value of  $n$ . Results are in Figure 2.

In general, autoencoders quickly outperform POD: in fact, their performances at  $n = 3$  remain unmatched even when the POD considers as much as  $n = 40$  latent variables (left panel of Figure 2). However, here we are not really interested in the actual values of the error, but rather in *how* the error decays.

In the linear case, we can observe a very clear and stable trend, where the error decays as some power of the latent dimension  $n$ , namely  $n^{-\gamma}$  for some  $\gamma > 0$ . However, things change quite a bit when moving to deep autoencoders. Here, in fact, although the error always decreases as a function of  $n$ , the rate of such decay is not constant. In fact, we can spot at least two different trends: first, we have a rapid decay from  $n = 1$  to  $n = 3$ , which is then followed by a much slower decay from  $n = 3$  on. In this sense, the latent dimension  $n = 3$  appears to be a turning point: after that, it is not really worth increasing the ROM dimension; rather, the performances could be improved by considering more complex architectures (see, e.g., the discussion in [19], especially Section 4.1 and Figure 7).

Since for the case at hand the PDE depends on  $p = 3$  random parameters, it is natural to ask whether this is a simple coincidence. As we shall prove in Theorem 3, this is not the case. It is shown here that when  $n = p$ , there is an ideal balance between reduction and precision. However, we must mention that several authors had already conjectured this fact and applied it as a rule of thumb; see, e.g., [40, 22]; furthermore, in the case of finitely supported probability measures, some first insights were already provided in [19]. Here, we extend the analysis proposed in [19], showing that this optimality is preserved even if the parameter space is endowed with a probability measure with unbounded support. Furthermore, we shall generalize the idea to the case of random fields, where the PDE formally depends on the parameters  $p = +\infty$ . From a practical point of view, this will allow domain practitioners to choose the latent dimension of deep autoencoders




Figure 2: Left panel: reconstruction error decay for Problem (7) as a function of the ROM dimension: linear reduction through POD vs nonlinear reduction through autoencoders. Right panel: model complexity, measured in terms of free —i.e., trainable— parameters (cf. Remark 1).

beforehand, thus avoiding the tedious procedures based on trial and error. For example, when faced with a problem such as (7), one can safely let  $n = 3$  without having to repeat the analysis in Figure 2.

*Remark 1.* Our focus for this work is to study the role of the latent dimension of deep autoencoders. Seeing this as a first step towards a more comprehensive investigation, our work does not immediately address other relevant matters, such as sample size requirements, optimization strategies, and model complexity. For example, it is reasonable to ask whether, in Problem (7), the improved performance of autoencoders comes at the cost of increased model complexity. Mathematically speaking, the latter can be defined as the number of trainable parameters in the model (also referred to as "free parameters"). For deep autoencoders, this corresponds to the total number of active —that is, nonzero— entries in the weight matrices and the bias vectors; for POD, instead, the latter is simply given by the number of entries in the POD matrix, i.e.  $N_h \times n$ . In other words, we may quantify the complexity of the two approaches by measuring a quantity that is proportional to the amount of physical memory required to store the corresponding model. As shown in the right panel of Figure 2, in this case autoencoders are indeed more complex compared to POD. For example, when  $n = 1$ , the autoencoder module requires 100 times more memory compared to POD. However, with increasing  $n$ , this difference becomes smaller and smaller. This is because, while the number of POD entries grows linearly at a rate proportional to the FOM dimension,  $N_h$ , autoencoders only increase in size near the bottleneck (in fact, only the hidden layers are involved in this process). Furthermore, it is worth mentioning that other architectures can be significantly more memory efficient. For example, as discussed in [19], the use of convolutional layers can result in autoencoders that are in fact cheaper than POD, with memory reduction close to 90%.

### 3. Preliminaries

In the current Section, the reader can find all the key properties and the mathematical concepts that are required for properly following the theoretical analysis proposed in Section 4. Precisely: in Section 3.1, we shall introduce the main ingredient of our recipe, i.e., DNN architectures; then, in Section 3.2, we familiarize ourselves with the notion of *local variation*, a mathematical concept that we later use to characterize locally Lipschitz operators; finally, in Section 3.3, we summarize several properties about Gaussian processes that are relevant to our analysis in Section 4.

In order to keep the paper self-contained, most of the proofs related to this Section have been postponed to the Appendix, Section A.

### 3.1. Expressivity of Deep Neural Networks

DNNs are computational units based on the composition of affine transformations and nonlinear activations, with the latter being applied component-wise on all vector entries. In the context of Deep Learning, in fact, it is very common to encounter the following notation

$$\rho(\mathbf{x}) := [\rho(x_1), \dots, \rho(x_n)], \quad (8)$$

where  $\mathbf{x} \in \mathbb{R}^n$  and  $\rho : \mathbb{R} \rightarrow \mathbb{R}$ . Here, with little abuse of notation, we agree to adopt the same convention.

In particular, DNNs are obtained via the composition of several maps, called *layers*, of the form  $\mathbf{x} \rightarrow \rho(\mathbf{W}\mathbf{x} + \mathbf{b})$ . Each layer is characterized by its own weight matrix  $\mathbf{W}$ , and its bias vector  $\mathbf{b}$ , two learnable parameters that are generally optimized during the training phase. All layers except the terminal layer (which usually comes without an activation function) are called *hidden layers*, and their total number defines the *depth* of the architecture. In mathematical terms, we may synthesize these notions as follows.

**Definition 1.** Let  $\rho : \mathbb{R} \rightarrow \mathbb{R}$  and  $m, n \in \mathbb{N}$ . We define the family of layers from  $\mathbb{R}^m$  to  $\mathbb{R}^n$  with activation function  $\rho$  as

$$\mathcal{L}_\rho(\mathbb{R}^m, \mathbb{R}^n) = \{f : \mathbb{R}^m \rightarrow \mathbb{R}^n \text{ s.t. } f(\mathbf{x}) = \rho(\mathbf{W}\mathbf{x} + \mathbf{b}), \mathbf{W} \in \mathbb{R}^{m \times n}, \mathbf{b} \in \mathbb{R}^n\}. \quad (9)$$

Similarly, we define the family of Deep Neural Networks (DNNs) from  $\mathbb{R}^m \rightarrow \mathbb{R}^n$  with the activation function  $\rho$  as

$$\begin{aligned} \mathcal{N}_\rho(\mathbb{R}^m, \mathbb{R}^n) = \{f_{l+1} \circ \dots \circ f_1, \text{ with } & f_{l+1} \in \mathcal{L}_{\text{id}}(\mathbb{R}^{n_l}, \mathbb{R}^n), l \geq 1, \\ & f_i \in \mathcal{L}_\rho(\mathbb{R}^{n_{i-1}}, \mathbb{R}^{n_i}), i = 1, \dots, l, \\ & n_i \in \mathbb{N}, \text{ and } n_0 = m\}, \end{aligned} \quad (10)$$

where  $\text{id} : \mathbb{R} \rightarrow \mathbb{R}$  is the identity map,  $\text{id}(x) = x$ .

When embedded in classical functional spaces, DNNs can provide remarkable expressivity. In fact, with very few hypotheses on their activation function, DNNs become able to approximate continuous maps over compact sets, as well as integrable maps over finite measure spaces. These kinds of result are known as Universal Approximation Theorems. Here, since we are dealing with stochastic quantities and probability measures, we are interested in the Universal Approximation Theorem proved by Hornik in 1991 [33]. In particular, we report below a slightly different result, which is a direct consequence of Hornik's Theorem.

**Lemma 1.** Let  $\mathbb{P}$  be a probability distribution on  $\mathbb{R}^p$ . Let  $\rho : \mathbb{R} \rightarrow \mathbb{R}$  be a continuous map. Assume that either one of the following holds:

- i)  $\rho$  is bounded and nonconstant;
- ii)  $\rho$  is bounded from below and  $\rho(x) \rightarrow +\infty$  as  $x \rightarrow +\infty$ ;
- iii) there exists some  $a, b \in \mathbb{R}$  such that  $x \rightarrow a\rho(x) + b\rho(-x)$  satisfies (ii).

Then, for every  $\varepsilon > 0$  and every measurable map  $f : \mathbb{R}^p \rightarrow \mathbb{R}^n$  with  $\mathbb{E}_{\mathbf{x} \sim \mathbb{P}} |f(\mathbf{x})| < +\infty$ , there exists  $\Phi \in \mathcal{N}_\rho(\mathbb{R}^p, \mathbb{R}^n)$  such that

$$\mathbb{E}_{\mathbf{x} \sim \mathbb{P}} |f(\mathbf{x}) - \Phi(\mathbf{x})| < \varepsilon.$$

*Proof.* See Appendix A. □

*Remark 2.* The result in Lemma 1 applies to most activation functions used in deep learning applications. For instance, the statement holds for maps such as the sigmoid or the tanh activation, as they satisfy (i),



but it also holds for other popular maps, such as the ReLU, the SELU and the swish activation, which are lower-bounded. Finally, the Lemma can also be applied to maps such as the  $\alpha$ -leaky ReLU activation,

$$\rho(x) = \begin{cases} x & x \geq 0 \\ \alpha x & x < 0 \end{cases}, \quad (11)$$

where  $|\alpha| < 1$ . In fact, the latter satisfies (iii) with  $a = 1$  and  $b = -\alpha$ .

### 3.2. Local variation of (nonlinear) operators

As we mentioned in the Introduction, under deterministic conditions, it is very natural to study reconstruction errors using *worst-case* approaches; for instance, the manifold  $n$ -width, Eq. (1), constitutes a major representative of this paradigm. Due to their global perspective, worst-case metrics can provide uniform bounds that characterize the regularity of the solution manifold as a whole. However, when considering stochastic scenarios, worst-case approaches may be overly restrictive. This is because different events are now associated with different probabilities. As a result, it becomes more natural to shift towards metrics that describe the *average scenario*, thus promoting a higher accuracy for situations that are most likely to occur, while allowing more leeway for errors in rare occurrences.

Notably, this change of perspective also allows us to consider a broader class of parameter-to-solution operators: in fact, average metrics are -in general- more permissive compared to worst-case ones. For instance, previous works has shown that the continuity of the operator is not enough to guarantee a reasonable behavior of the manifold  $n$ -width, while the stronger assumption of Lipschitz continuity, combined with the boundness of the parameter space, suffices (see Theorem 3 and the corresponding remark in [19]). Here, we shall see that these assumptions can be weakened as soon as we move to average-metrics and stochastic scenarios. In particular, the requirement of global Lipschitz continuity can be replaced with that of *local* Lipschitz continuity, that is, a condition in which the Lipschitz constant of the operator  $\mathcal{G}$  is allowed to change from point to point.

Notice that, despite sounding like a minor change, this adaptation is of fundamental importance. In fact, stochastic scenarios typically involve unbounded parameter spaces, where Lipschitz continuity rapidly becomes an excessively stringent property: for instance, even in the finite rank case, an extremely regular operator such as  $\mathcal{G}(\boldsymbol{\mu}) = |\boldsymbol{\mu}|^2$  would not meet the assumptions. Considering that it is fairly common to have either  $\text{supp}(\mu) = L^q(\Omega)$  when  $\mu$  is a random field, or  $\text{supp}(\boldsymbol{\mu}) = \mathbb{R}^p$  when  $\boldsymbol{\mu}$  is a random vector, it is evident that relaxing these assumptions is fundamental.

All these considerations bring us to the following definition, which we use to study the local behavior of nonlinear operators.

**Definition 2. (Local variation)** Let  $\mathcal{G} : (W, \|\cdot\|_W) \rightarrow (V, \|\cdot\|_V)$  be an operator between normed spaces. We define its local variation as the map  $\|\partial\mathcal{G}\|_V : W \rightarrow [0, +\infty]$  given by

$$\begin{aligned} \|\partial\mathcal{G}\|_V(w) &:= \limsup_{h \rightarrow 0} \frac{\|\mathcal{G}(w+h) - \mathcal{G}(w)\|_V}{\|h\|_W} = \\ &= \lim_{r \rightarrow 0^+} \sup_{0 < \|h\|_W \leq r} \frac{\|\mathcal{G}(w+h) - \mathcal{G}(w)\|_V}{\|h\|_W}. \end{aligned} \quad (12)$$

Ultimately, the *local variation* provides a way to bound the Lipschitz constant near every point of the input space, and it can be computed for any (nonlinear) operator. In fact, our definition can be traced back to the notion of *absolute condition number*, which is a concept commonly encountered in error and numerical analysis; see, e.g., [53]. Here, however, we insist on using a different notation to better reflect our context and avoid any sort of ambiguity.

The concept of local variation can be further understood by considering its relationship to other well-established mathematical concepts such as Lipschitz continuity and Frechét differentiability. These facts are briefly summarized by the Proposition below.

**Proposition 1.** *Let  $\mathcal{G} : (W, \|\cdot\|_W) \rightarrow (V, \|\cdot\|_V)$  be an operator between two normed spaces. Then, the following properties hold:*

i)  $\|\partial\mathcal{G}\|_V(w) < +\infty$  for all  $w \in K \subset W \iff \mathcal{G}$  is locally Lipschitz over  $K$ ;

ii) if  $K \subset W$  is compact and  $\|\partial\mathcal{G}\|_V(w) < +\infty$  for all  $w \in K$ , then  $\mathcal{G}$  is Lipschitz over  $K$ ;

iii) if  $C \subseteq W$  is convex, then

$$L_C := \sup_{w \in C} \|\partial\mathcal{G}\|_V(w) < +\infty$$

if and only if  $\mathcal{G}$  is  $L_C$ -Lipschitz over  $C$ .

iv) if  $\mathcal{G}$  is Fréchet differentiable at  $w \in W$ , then  $\|\partial\mathcal{G}\|_V(w)$  coincides with the operator norm of the Fréchet derivative of  $\mathcal{G}$  at  $w$ .

v) given any  $\mathcal{F} : (V, \|\cdot\|_V) \rightarrow (Y, \|\cdot\|_Y)$ , one has the chain-rule inequality

$$\|\partial(\mathcal{F} \circ \mathcal{G})\|_Y(w) \leq \|\partial\mathcal{F}\|_Y(\mathcal{G}(w)) \cdot \|\partial\mathcal{G}\|_V(w),$$

for all  $w \in W$ .

*Proof.* See Appendix A. □

As a straightforward consequence of Proposition 1, we also have the following Corollary, which can be thought of as a form of Taylor-Lagrange inequality.

**Corollary 1.** *Let  $\mathcal{G} : (W, \|\cdot\|_W) \rightarrow (V, \|\cdot\|_V)$  be any operator between two normed spaces. Then, for all  $w, w' \in W$  we have*

$$\|\mathcal{G}(w) - \mathcal{G}(w')\|_V \leq \left( \sup_{0 \leq t \leq 1} \|\partial\mathcal{G}\|_V(tw + (1-t)w') \right) \|w - w'\|_W. \quad (13)$$

*Proof.* Given  $w, w' \in W$ , let  $K$  be the segment between the two points and define  $L := \sup_{v \in K} \|\partial\mathcal{G}\|_V$ . If  $L = +\infty$ , then (13) is obvious; conversely, if  $L < +\infty$ , then the conclusion follows by (iii) in Proposition 1. □

Equation (13) provides a way to control the discrepancy between two different outputs of the operator, but it may also be applied to derive growth conditions. For example, it can be shown that if the local variation grows at most exponentially, then so does the operator. We formalize this consideration below.

**Corollary 2.** *Let  $\mathcal{G} : (W, \|\cdot\|_W) \rightarrow (V, \|\cdot\|_V)$  be an operator between two normed spaces. The following holds true,*

$$\begin{aligned} \exists M, \beta > 0 \text{ s.t. } \|\partial\mathcal{G}\|_V(w) \leq M e^{\beta\|w\|_W} \quad \forall w \in W, \\ \implies \exists \tilde{M}, \gamma > 0 \text{ s.t. } \|\mathcal{G}(w)\|_V \leq \tilde{M} e^{\gamma\|w\|_W} \quad \forall w \in W. \end{aligned}$$

*Proof.* Let  $c := \mathcal{G}(0)$  and fix any  $w \in W$ . By Corollary 1, we have

$$\begin{aligned} \|\mathcal{G}(w)\|_V &\leq c + \|\mathcal{G}(w) - \mathcal{G}(0)\|_V \leq c + \left( \sup_{0 \leq t \leq 1} \|\partial\mathcal{G}\|_V(tw) \right) \|w\|_W \leq \\ &\leq c + \left( \sup_{0 \leq t \leq 1} M e^{\beta\|tw\|_W} \right) \|w\|_W = c + M e^{\beta\|w\|_W} \|w\|_W. \end{aligned}$$

Since  $a < e^a$  for all  $a \in \mathbb{R}$ , and  $c \leq ce^b$  for all  $b \geq 0$ , we have

$$\|\mathcal{G}(w)\|_V \leq \dots \leq (c + M)e^{(\beta+1)\|w\|_W}. \quad \square$$

To conclude, we present a practical example of an operator whose local variation grows at most exponentially. Despite its simplicity, we believe this example to be of high interest, as it describes the case of a stochastic Darcy flow, which, in particular, has direct implications in the study of porous media.

**Proposition 2.** *Let  $\Omega \subset \mathbb{R}^d$  be a bounded domain with Lipschitz boundary, and let  $f \in H^{-1}(\Omega)$  be given. For any  $\sigma \in L^\infty(\Omega)$ , let  $u = u_\sigma$  be the solution to the following boundary value problem,*

$$\begin{cases} -\nabla \cdot (e^\sigma \nabla u) = f & \text{in } \Omega \\ u = 0 & \text{on } \partial\Omega. \end{cases} \quad (14)$$

Let  $\mathcal{G} : L^\infty(\Omega) \rightarrow L^2(\Omega)$  be the operator that maps  $\sigma \mapsto u$ . Then, for all  $\sigma, \sigma' \in L^\infty(\Omega)$

$$\|\mathcal{G}(\sigma) - \mathcal{G}(\sigma')\|_{L^2(\Omega)} \leq C\|f\|_{H^{-1}(\Omega)} e^{3\|\sigma\|_{L^\infty(\Omega)} + 3\|\sigma'\|_{L^\infty(\Omega)}} \|\sigma - \sigma'\|_{L^\infty(\Omega)} \quad (15)$$

and, in particular,

$$\|\partial\mathcal{G}(\sigma)\|_{L^2(\Omega)} \leq C\|f\|_{H^{-1}(\Omega)} e^{6\|\sigma\|_{L^\infty(\Omega)}},$$

where  $C = C(\Omega)$  is some positive constant.

*Proof.* See Appendix A. □

### 3.3. Regularity and energy estimates for Gaussian processes

We now take the opportunity to recall some fundamental facts about stochastic processes and random fields. More specifically, here we shall limit our analysis to Gaussian processes: for further comments about this choice, we refer the reader to the discussion at the end of Section 4.3, see Remark 4. We start with a result that links the regularity of the covariance kernel of a Gaussian process with that of the trajectories of the random field. In doing so, we also include some estimates on the norms of the process, which we shall exploit later on.

For the sake of simplicity, here, expected values will be directly denoted as  $\mathbb{E}[\cdot]$ , without any explicit declaration of the integration variable and its underlying probability distribution – given the context, there should be no ambiguity.

**Lemma 2.** *Let  $\Omega \subset \mathbb{R}^d$  be pre-compact, and let  $Z$  be a mean zero Gaussian random field defined over  $\Omega$ . Assume that, for some  $0 < \alpha \leq 1$ , the covariance kernel of the process,  $\text{Cov} : \Omega \times \Omega \rightarrow \mathbb{R}$ , defined as*

$$\text{Cov}(\mathbf{x}, \mathbf{y}) := \mathbb{E}[Z(\mathbf{x})Z(\mathbf{y})],$$

*is  $\alpha$ -Hölder continuous, with Hölder constant  $L > 0$ . Then,  $Z$  is sample-continuous, that is  $\mathbb{P}(Z \in \mathcal{C}(\Omega)) = 1$ . Furthermore, for  $\sigma^2 := \max_{\mathbf{x} \in \Omega} \text{Cov}(\mathbf{x}, \mathbf{x})$ , one has*

$$\mathbb{E}^{1/2} \|Z\|_{L^\infty(\Omega)}^2 \leq c_1 \sigma \left(1 + \sqrt{\log^+(1/\sigma)}\right) \quad \text{and} \quad \mathbb{E}[e^{\beta\|Z\|_{L^\infty(\Omega)}}] = c_2 < +\infty, \quad (16)$$

*for all  $\beta > 0$ , where  $c_1 = c_1(d, L, \alpha, \Omega)$  and  $c_2 = c_2(d, L, \alpha, \sigma, \beta, \Omega)$  are constants that depend continuously on their parameters (domain excluded). Here,*

$$\log^+(a) := \max\{\log a, 0\}.$$

*Proof.* See Appendix A. □

Another fundamental result that we need for our construction is a Corollary of Mercer's Theorem, known as the KKL series expansion. We report it below, together with some considerations about the covariance kernel and its truncations.

**Lemma 3.** *Let  $\Omega \subset \mathbb{R}^d$  be a compact subset and let  $Z$  be a mean zero Gaussian random field defined over  $\Omega$ . Assume that the covariance kernel of  $Z$ ,  $\text{Cov}$ , is Lipschitz continuous. Then, there exists a nonincreasing summable sequence  $\lambda_1 \geq \lambda_2 \geq \dots \geq 0$  and a sequence of Lipschitz continuous maps,  $\{\varphi_i\}_{i=1}^{+\infty}$ , forming an orthonormal basis of  $L^2(\Omega)$ , such that*

$$\text{Cov}(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^{+\infty} \lambda_i \varphi_i(\mathbf{x}) \varphi_i(\mathbf{y}) \quad (17)$$

for all  $\mathbf{x}, \mathbf{y} \in \Omega$ . Furthermore, there exists a sequence of independent standard normal random variables,  $\{\eta_i\}_{i=1}^{+\infty}$ , such that

$$Z = \sum_{i=1}^{+\infty} \sqrt{\lambda_i} \eta_i \varphi_i \quad (18)$$

almost surely. Finally, the truncated kernels,

$$\text{Cov}_{p,q}(\mathbf{x}, \mathbf{y}) := \sum_{i=p}^q \lambda_i \varphi_i(\mathbf{x}) \varphi_i(\mathbf{y}),$$

defined for varying  $1 \leq p \leq q \leq +\infty$ ,

- i) converge uniformly as  $p, q \rightarrow +\infty$ ;
- ii) are all  $1/2$ -Hölder continuous, with a common Hölder constant.

*Proof.* See Appendix A. □

Both the result in Lemma 2 and that in Lemma 3 require some form of uniform continuity of the covariance kernel, however, they have the advantage of yielding useful estimates to treat the  $L^\infty$ -case. At the same time, these properties may be far too restrictive if one moves to the simpler scenario in which the trajectories of the random field are only assumed to be square-integrable. In light of this, we report below a different result specifically tailored for the  $L^2$ -case, which can be seen as an adaptation of the previous Lemmas.

**Lemma 4.** *Let  $\Omega \subset \mathbb{R}^d$  be a compact subset and let  $Z$  be a mean zero Gaussian random field defined over  $\Omega$ . Assume that the covariance kernel of  $Z$ ,  $\text{Cov}$ , is square-integrable over  $\Omega \times \Omega$ . Then, there exist a non-increasing summable sequence  $\lambda_1 \geq \lambda_2 \geq \dots \geq 0$  and an orthonormal basis of  $L^2(\Omega)$ ,  $\{\varphi_i\}_{i=1}^{+\infty}$ , such that*

$$\text{Cov}(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^{+\infty} \lambda_i \varphi_i(\mathbf{x}) \varphi_i(\mathbf{y})$$

for almost every  $(\mathbf{x}, \mathbf{y}) \in \Omega \times \Omega$ . Furthermore, there exists a sequence of independent standard normal random variables,  $\{\eta_i\}_{i=1}^{+\infty}$ , such that

$$Z = \sum_{i=1}^{+\infty} \sqrt{\lambda_i} \eta_i \varphi_i$$

almost surely. Finally, the  $L^2$ -norm of the process is exponentially integrable, i.e.

$$\mathbb{E} \left[ e^{\beta \|Z\|_{L^2(\Omega)}} \right] < +\infty \quad \forall \beta > 0. \quad (19)$$

*Proof.* See Appendix A. □

The main difference between Lemma 3 and Lemma 4 lies in the regularity that one requires to the covariance kernel. Clearly, stronger assumptions about the latter result in stronger estimates about the random field and its norms.

In conclusion, we also report an abstract version of the KKL expansion for generic Hilbert-valued random variables. In this case, it is convenient to consider an *uncentered* KKL expansion, as the latter retains useful optimality properties: in fact, it is the abstract equivalent of the POD algorithm.

**Lemma 5.** *Let  $(V, \|\cdot\|)$  be a separable Hilbert space and let  $u$  be a squared integrable  $V$ -valued random variable,  $\mathbb{E}\|u\|^2 < +\infty$ . Then, there exists an orthonormal basis  $\{v_i\}_{i=1}^{+\infty} \subset V$ , a sequence of (scalar) random variables  $\{\omega_i\}_{i=1}^{+\infty}$ , with  $\mathbb{E}[\omega_i \omega_j] = \delta_{i,j}$ , and a nonincreasing summable sequence  $\lambda_1 \geq \lambda_2 \geq \dots \geq 0$  such that*

$$u = \sum_{i=1}^{+\infty} \sqrt{\lambda_i} \omega_i v_i,$$

*almost-surely.*

*Proof.* Up to adaptations, this is a standard result; see, e.g. Theorem 3.14 in [39]. However, the interested reader can find a detailed proof in Appendix A.  $\square$

#### 4. Autoencoder-based nonlinear reduction for PDEs parametrized by random fields

We are now ready to put things into action and present the main results of this work. We shall start with a preliminary consideration about the expressivity of deep autoencoders and introduce a suitable notion of *admissibility* that will help us avoid unrealistic/pathological situations. Then, we shall derive error bounds for the reconstruction error of deep autoencoders in the case of: i) PDEs depending on a finite number of random coefficients (Section 4.2); ii) PDEs parameterized by Gaussian random fields (Section 4.3).

For better readability, in what follows, we shall drop the dependency of expected values with respect to their underlying probability distribution. In particular, since all randomness will be encoded in the input variable, which is  $\mu \sim \mathbb{P}$  or  $\mu \sim \mathbb{P}$ , we shall simply write  $\mathbb{E}$  in place of  $\mathbb{E}_{\mu \sim \mathbb{P}}$ .

##### 4.1. Admissible autoencoders and density results

As already mentioned, our main interest is to investigate how the choice of the latent dimension affects the optimization of the reconstruction error, and thus to provide guidelines for the design of deep autoencoders. To this end, we must note that DNN spaces lack many of the properties usually holding for classical functional spaces; furthermore, their topology can easily become rather involved, see, e.g., the discussion in [49]. For this reason, it can be convenient to recast the optimization problem over a broader class of functions, e.g. by relying on suitable density results. For example, in [19], the authors consider a more general framework in which the encoder and decoder are allowed to be any pair continuous maps. Here, we relax these hypotheses even further. In fact, as a direct consequence of Lemma 1, it is easy to see that the only property that is actually needed is measurability. More precisely, we have the following result.

**Theorem 1.** *Let  $\mu$  be a random vector in  $\mathbb{R}^p$ . Let  $\mathcal{G} : \mathbb{R}^p \rightarrow V \cong \mathbb{R}^{N_h}$  be a measurable operator and define  $\mathbf{u}_\mu := \mathcal{G}(\mu)$ . Let  $n \in \mathbb{N}$ . If  $\mathbb{E}\|\mathbf{u}_\mu\| < +\infty$ , then*

$$\inf_{\substack{\Psi' \in \mathcal{M}(V, \mathbb{R}^n) \\ \Psi \in \mathcal{M}(\mathbb{R}^n, V)}} \mathbb{E}\|\mathbf{u}_\mu - \Psi(\Psi'(\mathbf{u}_\mu))\| = \inf_{\substack{\hat{\Psi}' \in \mathcal{N}_\rho(V, \mathbb{R}^n) \\ \hat{\Psi} \in \mathcal{N}_\rho(\mathbb{R}^n, V)}} \mathbb{E}\|\mathbf{u}_\mu - \hat{\Psi}(\hat{\Psi}'(\mathbf{u}_\mu))\|, \quad (20)$$

*for all Lipschitz continuous activations  $\rho$  satisfying the hypothesis of Lemma 1.*

*Proof.* We assume the left-hand side to be finite, as the statement would be trivially true otherwise. Let  $\Psi' \in \mathcal{M}(V, \mathbb{R}^n)$  and  $\Psi \in \mathcal{M}(\mathbb{R}^n, V)$  be such that

$$\mathbb{E}\|\mathbf{u}_\mu - \Psi(\Psi'(\mathbf{u}_\mu))\| < +\infty, \quad (21)$$

and fix any  $\varepsilon > 0$ . Let  $\rho = \tanh$  be the hyperbolic tangent activation. Clearly,  $\Psi$  and  $\Psi'$  have the same reconstruction error as  $\Psi_\rho := \Psi \circ \rho^{-1}$  and  $\Psi'_\rho := \rho \circ \Psi'$ . Therefore, up to replacing  $\Psi$  with  $\Psi_\rho$  and  $\Psi'$  with  $\Psi'_\rho$ , we may assume that  $\Psi'$  is bounded.

Now, let us define the random vector  $\mathbf{v}_\mu := \Psi'(\mathbf{u}_\mu)$ . Since  $\mathbb{E}\|\mathbf{u}_\mu\| < +\infty$  and (21) hold, it follows by linearity that  $\mathbb{E}\|\Psi(\mathbf{v}_\mu)\| < +\infty$ . In particular, we may apply Lemma 1 in order to find some  $\hat{\Psi} \in \mathcal{N}_\rho(\mathbb{R}^n, V)$  such that

$$\mathbb{E}\|\Psi(\mathbf{v}_\mu) - \hat{\Psi}(\mathbf{v}_\mu)\| < \varepsilon/2,$$

Of note,  $\hat{\Psi}$  inherits the Lipschitz continuity of  $\rho$ . Thus, for any  $\hat{\Psi}' \in \mathcal{N}_\rho(V, \mathbb{R}^n)$ ,

$$\begin{aligned} \|\mathbb{E}\|\mathbf{u}_\mu - \Psi(\Psi'(\mathbf{u}_\mu))\| - \mathbb{E}\|\mathbf{u}_\mu - \hat{\Psi}(\hat{\Psi}'(\mathbf{u}_\mu))\|\| &\leq \mathbb{E}\|\Psi(\Psi'(\mathbf{u}_\mu)) - \hat{\Psi}(\hat{\Psi}'(\mathbf{u}_\mu))\| \leq \\ &\leq \mathbb{E}\|\Psi(\Psi'(\mathbf{u}_\mu)) - \hat{\Psi}(\Psi'(\mathbf{u}_\mu))\| + \mathbb{E}\|\hat{\Psi}(\Psi'(\mathbf{u}_\mu)) - \hat{\Psi}(\hat{\Psi}'(\mathbf{u}_\mu))\| \leq \\ &\leq \frac{1}{2}\varepsilon + L\mathbb{E}\|\Psi'(\mathbf{u}_\mu) - \hat{\Psi}'(\mathbf{u}_\mu)\|, \end{aligned}$$

where  $L > 0$  is the Lipschitz constant of  $\hat{\Psi}$ . Then, by lemma 1, we may choose  $\hat{\Psi}'$  so that  $\mathbb{E}\|\Psi'(\mathbf{u}_\mu) - \hat{\Psi}'(\mathbf{u}_\mu)\| < \varepsilon/2L$ , and thus

$$\|\mathbb{E}\|\mathbf{u}_\mu - \Psi(\Psi'(\mathbf{u}_\mu))\| - \mathbb{E}\|\mathbf{u}_\mu - \hat{\Psi}(\hat{\Psi}'(\mathbf{u}_\mu))\|\| \leq \varepsilon. \quad (22)$$

Note that, this time, we were able to apply Lemma 1 due to the boundness of  $\Psi'$  (which in turn ensures its integrability with respect to any probability measure). As  $\varepsilon > 0$  is arbitrary, the inequality in (22) suffices to prove the identity in (20).  $\square$

The result in Theorem 1 is a double-edged sword. On the one hand, it allows us to reframe the optimization problem on a less restrictive class of functions, giving us the possibility, e.g., to study the behavior of the reconstruction error without having to worry about the discretization of the state space: in fact, the spaces  $\mathcal{M}(V, \mathbb{R}^n)$  and  $\mathcal{M}(\mathbb{R}^n, V)$  are well defined even if  $V$  is infinite-dimensional. On the other hand, the autoencoders in Eq. (20) can become extremely irregular and thus more difficult to capture. In fact, one can show that in most cases

$$\inf_{\substack{\Psi' \in \mathcal{M}(V, \mathbb{R}^n) \\ \Psi \in \mathcal{M}(\mathbb{R}^n, V)}} \mathbb{E}\|\mathbf{u}_\mu - \Psi(\Psi'(\mathbf{u}_\mu))\| = 0,$$

for all  $n \geq 1$ , as there always exists a suitable space-filling curve that provides a lossless embedding. However, such a representation would be completely useless, as it would correspond to an architecture that is either impossible to reproduce or train: see, e.g., the discussion at the end of [13] by Cohen et al.

One way to overcome all these issues is to impose certain additional assumptions on the regularity of the autoencoder architecture. Here, we proceed as follows. We define the (enlarged) class of admissible encoders  $V \rightarrow \mathbb{R}^n$  as

$$\mathcal{E}_{B,M}(V, \mathbb{R}^n) := \left\{ \Psi' \in \mathcal{M}(V, \mathbb{R}^n) \text{ s.t. } \mathbb{E}|\Psi'(\mathbf{u}_\mu)| < +\infty \text{ and } \sup_{\mathbf{v} \in B} |\Psi'(\mathbf{v})| \leq M \right\}, \quad (23)$$

where  $B \subset V$  is a *control set* and  $M > 0$  is a suitable upper bound. Both  $B$  and  $M$  are to be considered as hyperparameters: their role is to ensure that, at least in the control set  $B$ , the encoder networks are uniformly well behaved. On the contrary, we define the (enlarged) family of admissible decoders as

$$\mathcal{D}_{M,L}(\mathbb{R}^n, V) := \left\{ \Psi \in \mathcal{M}(\mathbb{R}^n, V) \text{ s.t. } \sup_{\mathbf{c} \in [-M, M]^n} \|\partial \Psi(\mathbf{c})\| \leq L \right\}, \quad (24)$$

where  $L > 0$  is a suitable hyperparameter that controls the regularity of the decoder. In fact, the condition in (24) forces  $\Psi$  to be  $L$ -Lipschitz continuous over the hypercube  $[-M, M]^n$ : cf. (iii) in Proposition 1.

This setup allows us to avoid the phenomenon of space-filling curves and, at the same time, to regain interest in the optimization of the reconstruction error. Furthermore, this formulation comes with a natural adaptation of Theorem 1, which we report below.

**Theorem 2.** Let  $\boldsymbol{\mu}$  be a random vector in  $\mathbb{R}^p$ . Let  $\mathcal{G} : \mathbb{R}^p \rightarrow V \cong \mathbb{R}^{N_h}$  be a measurable map and define  $\mathbf{u}_\mu := \mathcal{G}(\boldsymbol{\mu})$ . Let  $\rho$  be the  $\alpha$ -leaky ReLU activation,  $|\alpha| < 1$ . Let  $B \subset V$  be a bounded set,  $M, L > 0$  and  $n \in \mathbb{N}$ . Assume that the probability law of  $\mathbf{u}_\mu$  is absolutely continuous. If  $\mathbb{E}\|\mathbf{u}_\mu\| < +\infty$ , then

$$\inf_{\substack{\Psi' \in \mathcal{E}_{B,M}(V, \mathbb{R}^n) \\ \Psi \in \mathcal{D}_{M,L}(\mathbb{R}^n, V)}} \mathbb{E}\|\mathbf{u}_\mu - \Psi(\Psi'(\mathbf{u}_\mu))\| = \inf_{\substack{\hat{\Psi}' \in \mathcal{N}_\rho^e(V, \mathbb{R}^n) \\ \hat{\Psi} \in \mathcal{N}_\rho^d(\mathbb{R}^n, V)}} \mathbb{E}\|\mathbf{u}_\mu - \hat{\Psi}(\hat{\Psi}'(\mathbf{u}_\mu))\|, \quad (25)$$

where  $\mathcal{N}_\rho^e(V, \mathbb{R}^n) := \mathcal{N}_\rho(V, \mathbb{R}^n) \cap \mathcal{E}_{B,M}(V, \mathbb{R}^n)$  and  $\mathcal{N}_\rho^d(\mathbb{R}^n, V) := \mathcal{N}_\rho(\mathbb{R}^n, V) \cap \mathcal{D}_{M,L}(\mathbb{R}^n, V)$  are all admissible encoder and decoder networks, respectively.

*Proof.* The proof is roughly the same as the one of Theorem 1, up to replacing the use of Lemma 1 with stronger results that can ensure the preservation of the constraints in Equations (23) and (24). For further details, we refer to Lemmas B2, Lemma B3 and Corollary B1 in Appendix B.  $\square$

The identity in Theorem 2 has two major implications. First, it shows how general autoencoders can be, thus further motivating their usage for nonlinear reduction. Second, it allows us to reframe the optimization problem in an abstract way. This can be very useful, as it allows us to adopt a more general perspective where  $V$  can be either discrete,  $V \cong \mathbb{R}^{N_h}$  or continuous, e.g.  $V = L^2(\Omega)$ . In fact, the families  $\mathcal{E}_{B,M}(V, \mathbb{R}^n)$  and  $\mathcal{D}_{M,L}(\mathbb{R}^n, V)$  can be defined without the need to discretize the state space.

In light of this, for the next few pages, we shall drop our assumption on  $V$  being finite-dimensional. This means that from now on, all the results reported will hold true both in the discrete and in the continuous setting. For better readability, elements of the state space will be indicated as  $u \in V$ , to emphasize the fact that such elements can be vectors (for which the notation  $\mathbf{u}$  would be more fitting) or functions.

*Remark 3.* Although the two definitions in (23) and (24) may sound quite technical, they actually mirror some of the practical strategies that researchers and data scientists commonly use. For example, when training a DNN model, it is very common to use Tychonoff regularizations to avoid excessive growth of the DNN weights. However, this procedure is equivalent to imposing admissibility constraints via a Lagrange multiplier. To see this, let  $\Phi$  be some DNN. For simplicity, assume that  $\Phi$  only has one hidden layer, so that  $\Phi(\mathbf{x}) = \mathbf{W}_2 \rho(\mathbf{W}_1 \mathbf{x} + \mathbf{b}_1) + \mathbf{b}_2$ , where  $\rho$  is some  $\ell$ -Lipschitz activation function. With little abuse of notation, let us denote by  $|\cdot|$  both the Euclidean norm and the Frobenius norm. We have

$$|\Phi(\mathbf{x}) - \Phi(\mathbf{y})| \leq \ell |\mathbf{W}_2| \cdot |\mathbf{W}_1| \cdot |\mathbf{x} - \mathbf{y}|.$$

In particular,  $\Phi$  is Lipschitz continuous and the logarithm of its Lipschitz constant is bounded by  $\log \ell + \log |\mathbf{W}_2| + \log |\mathbf{W}_1|$ . It is then clear that penalizing the mass of weights matrices has a direct impact on the Lipschitz constant of the whole network, thus mimicking our condition on the decoder, Eq. (24). Furthermore, the same argument applies to the constraint for the encoder, Eq. (23). In fact, in our construction, the control set  $B$  is always assumed to be bounded. In particular,

$$|\Phi(\mathbf{x})| \leq |\Phi(0)| + |\Phi(\mathbf{x}) - \Phi(0)| \leq |\mathbf{W}_2| \cdot |\rho(0)| + \ell |\mathbf{W}_2| \cdot |\rho(\mathbf{b}_1)| + |\mathbf{b}_2| + R \ell |\mathbf{W}_2| \cdot |\mathbf{W}_1|,$$

where  $R > 0$  is any radius for which one has  $|\mathbf{x}| \leq R$  for all  $v \in B$ . Thus, the same reasoning can be applied up to including an additional penalty for the biases  $\mathbf{b}_i$ .

#### 4.2. Bounds on the latent dimension: finite rank case

We start by addressing the finite-rank case, in which the source of randomness is given by some random vector  $\boldsymbol{\mu}$ . We report our main result in the following, which, in a way, can be seen as a generalization of Theorem 3 in [19] to the (unbounded) probabilistic setting.

**Theorem 3.** Let  $\mathcal{G} : \mathbb{R}^p \rightarrow V$  be a locally Lipschitz operator, where  $(V, \|\cdot\|)$  is a given Banach space. Let  $\boldsymbol{\mu}$  be a random vector in  $\mathbb{R}^p$  and let  $u_\mu := \mathcal{G}(\boldsymbol{\mu})$  be the  $V$ -valued random variable obtained by mapping  $\boldsymbol{\mu}$

through  $\mathcal{G}$ . Assume that the latter is Bochner integrable, that is,  $\mathbb{E}\|u_\mu\| < +\infty$ . Denote by  $B := \mathcal{G}(\{|c| \leq 1\})$  the image of the unit ball. Then, there exists  $L_0 > 0$  such that for all  $M \geq 1$  and all  $L \geq L_0$  one has

$$\inf_{\substack{\Psi' \in \mathcal{E}_{B,M}(V, \mathbb{R}^n) \\ \Psi \in \mathcal{D}_{M,L}(\mathbb{R}^n, V)}} \mathbb{E}\|u_\mu - \Psi(\Psi'(u_\mu))\| = 0, \quad (26)$$

for all  $n \geq p$ .

*Proof.* It is sufficient to prove the case  $n = p$ . Let  $M \geq 1$ . Since  $\mathcal{G}$  is continuous, the latter admits a measurable right-inverse  $g' : V \rightarrow \mathbb{R}^n$ , that is, a map for which

$$\mathcal{G}(g'(\mathcal{G}(\mu))) = \mathcal{G}(\mu).$$

Furthermore, the latter can be constructed such that  $g'(B) \subseteq \{|c| \leq 1\}$ : for a detailed proof we refer the reader to the Appendix, particularly to Corollary C2. Let now  $\rho : \mathbb{R} \rightarrow \mathbb{R}$  be the following activation function

$$\rho(x) = \begin{cases} x & |x| \leq M \\ M \tanh(x) / \tanh(M) & |x| > M, \end{cases}$$

and define  $\Psi' := \rho \circ g'$ , where the action of  $\rho$  is aimed at components. Then,

$$\mathbb{E}|\Psi'(u_\mu)| \leq M / \tanh(M) < +\infty \quad \text{and} \quad \Psi'(B) = g'(B) \subseteq \{|c| \leq 1\} \subset [-M, M]^n, \quad (27)$$

in fact,  $g'(B) \subset [-M, M]^n$  and  $\rho$  act as the identity over  $[-M, M]^n$ . In particular, it follows from (27) that  $\Psi' \in \mathcal{E}_{B,M}(V, \mathbb{R}^n)$ . Now, noting that  $\rho$  is invertible, we let  $\Psi := \mathcal{G} \circ \rho^{-1}$ , with the convention that  $\rho^{-1} \equiv +\infty$  outside of  $\rho(\mathbb{R})$ . It is straightforward to see that the pair  $(\Psi', \Psi)$  produces a lossless compression: thus, it is sufficient to prove that  $\Psi \in \mathcal{D}_{M,L}(\mathbb{R}^n, V)$  for a suitable choice of  $L$ . To this end, we note that  $\Psi$  is locally Lipschitz over  $[-M, M]^n$ . Then, since the latter is compact, we may simply set

$$L_0 := \sup_{c \in [-M, M]^n} \|\partial \Psi\|(c) < +\infty,$$

see, e.g., (ii) in Proposition 1. □

The takeaway from Theorem 3 is that if the PDE depends on  $p$  (scalar) random variables, then an autoencoder with latent dimension  $n = p$  can compress and reconstruct solutions with arbitrary (average) accuracy. Clearly, since Equation (26) features an infimum, for this to work, the autoencoder must be sufficiently expressive in the remaining parts of the architecture, especially in the decoder. In this respect, domain practitioners can find a valuable help in recent works that address the approximation capabilities of DNNs in high-dimensional spaces; see, for instance [20, 60].

#### 4.3. Bounds on the latent dimension: infinite dimensional case

We are now set to discuss the infinite-dimensional case in which the input variable is given by a random field. In particular, we shall deal with random fields  $\mu$  having trajectories in  $L^q(\Omega)$  and operators of the form  $\mathcal{G} : L^q(\Omega) \rightarrow V$ , for  $V$  a Hilbert state space,  $1 \leq q \leq +\infty$ , and  $\Omega$  a bounded domain. We shall focus on two different scenarios: one in which  $q = 2$  and one in which  $q = +\infty$ . To this end, we recall that for any  $q \leq \tilde{q}$  one has continuous embedding

$$L^{\tilde{q}}(\Omega) \hookrightarrow L^q(\Omega),$$

where each space is considered with its canonical norm. In particular, it follows that  $\mathcal{C}(L^2(\Omega), V) \subset \mathcal{C}(L^\infty(\Omega), V)$ , meaning that case  $q = +\infty$  allows for a broader class of operators and is consequently much harder to handle. However, it is worth addressing both situations, as the two can lead to very different analyses and error bounds. We start with the simpler  $L^2$ -case. For the sake of readability, both proofs have been postponed to Section 4.4.



**Theorem 4.** ( $L^2$ -version) Let  $\Omega$  be a bounded domain and let  $(V, \|\cdot\|)$  be a Hilbert space. Let  $\mu$  be a Gaussian random field defined over  $\Omega$ , with a square integrable mean  $m : \Omega \rightarrow \mathbb{R}$  and a square integrable covariance kernel  $\text{Cov} : \Omega \times \Omega \rightarrow \mathbb{R}$ . Finally, let  $\mathcal{G} : L^2(\Omega) \rightarrow V$  be an operator satisfying the growth condition below,

$$\|\partial \mathcal{G}\|(\nu) \leq A e^{\beta \|\nu\|_{L^2(\Omega)}} \text{ for all } \nu \in L^2(\Omega),$$

for some constants  $A, \beta > 0$ . According to the Lemmae 3 and 5, let

$$\mu = \mathbb{E}[\mu] + \sum_{i=1}^{+\infty} \sqrt{\lambda_i^\mu} \eta_i \varphi_i \quad \text{and} \quad u = \sum_{i=1}^{+\infty} \sqrt{\lambda_i^u} \omega_i v_i, \quad (28)$$

be the KKL expansions of  $\mu$  and  $u$ , respectively. Consider the control set

$$B := \mathcal{G} \left( \left\{ m + \sum_{i=1}^{+\infty} \sqrt{\lambda_i^\mu} \nu_i \varphi_i \text{ s.t. } \sum_{i=1}^{+\infty} |\nu_i|^2 \leq 1 \right\} \right),$$

and let  $\mathbb{P}$  be the probability law of  $\mu$ . Then, there exist positive constants  $C = C(d, \mathcal{G}, \mathbb{P})$  and  $L_0 = L_0(n, \mathcal{G})$ , such that for all  $M \geq 1$  and  $L \geq L_0$  one has

$$\inf_{\substack{\Psi' \in \mathcal{E}_{B, M}(V, \mathbb{R}^n) \\ \Psi \in \mathcal{D}_{M, L}(\mathbb{R}^n, V)}} \mathbb{E} \|u_\mu - \Psi(\Psi'(u_\mu))\| \leq C \min \left\{ \sqrt{\sum_{i>n} \lambda_i^\mu}, \sqrt{\sum_{i>n} \lambda_i^u} \right\}, \quad (29)$$

for all latent dimensions  $n \geq 1$ .

We now report on the  $L^\infty$ -counterpart of Theorem 4. As we mentioned, this is a much more difficult case that requires additional care.

**Theorem 5.** ( $L^\infty$ -version) Let  $\Omega$  be a bounded domain and let  $(V, \|\cdot\|)$  be a Hilbert space. Let  $\mu$  be a Gaussian random field defined over  $\Omega$ , with a bounded mean  $m : \Omega \rightarrow \mathbb{R}$  and a Lipschitz continuous covariance kernel  $\text{Cov} : \Omega \times \Omega \rightarrow \mathbb{R}$ . Finally, let  $\mathcal{G} : L^\infty(\Omega) \rightarrow V$  be an operator satisfying the growth condition below,

$$\|\partial \mathcal{G}\|(\nu) \leq A e^{\gamma \|\nu\|_{L^\infty(\Omega)}} \text{ for all } \nu \in L^\infty(\Omega), \quad (30)$$

for some constants  $A, \gamma > 0$ . According to the Lemmae 3 and 5, let

$$\mu = \mathbb{E}[\mu] + \sum_{i=1}^{+\infty} \sqrt{\lambda_i^\mu} \eta_i \varphi_i \quad \text{and} \quad u = \sum_{i=1}^{+\infty} \sqrt{\lambda_i^u} \omega_i v_i, \quad (31)$$

be the KKL expansions of  $\mu$  and  $u$ , respectively. Consider the control set

$$B := \mathcal{G} \left( \left\{ m + \sum_{i=1}^{+\infty} \sqrt{\lambda_i^\mu} \nu_i \varphi_i \text{ s.t. } \sum_{i=1}^{+\infty} |\nu_i|^2 \leq 1 \right\} \right),$$

and let  $\mathbb{P}$  be the probability law of  $\mu$ . Fix any small  $0 < \epsilon < 1/2$ . Then, there exist two positive constants,  $C = C(d, \mathcal{G}, \mathbb{P})$  and  $L_0 = L_0(n, \mathcal{G})$ , such that for all  $M \geq 1$  and  $L \geq L_0$  one has

$$\inf_{\substack{\Psi' \in \mathcal{E}_{B, M}(V, \mathbb{R}^n) \\ \Psi \in \mathcal{D}_{M, L}(\mathbb{R}^n, V)}} \mathbb{E} \|u_\mu - \Psi(\Psi'(u_\mu))\| \leq C \sqrt{\log(1/\epsilon)} \min \left\{ \left\| \sum_{i>n} \lambda_i^\mu \varphi_i^2 \right\|_{L^\infty(\Omega)}^{1/2}, \sqrt{\sum_{i>n} \lambda_i^u} \right\}, \quad (32)$$

for all latent dimensions  $n \geq 1$  satisfying either  $\sum_{i>n} \lambda_i^\mu \geq \epsilon$  or  $\sum_{i>n} \lambda_i^\mu = 0$ .

The strength of the two Theorems lies in that they show how nonlinear autoencoders can simultaneously benefit from the regularity of both the input and the output fields, something that is clearly not possible with linear methods alone: see, e.g., the discussion by Lanthaler et al. in [39], Section 3.4.1. Furthermore, the results in Theorem 4 and 5 are fairly general, since they are both framed in a purely abstract fashion with mild assumptions on the regularity of the forward operator. For example, following our previous discussion in Section 3.2, we note that the result in Theorem 5 can be easily applied to the case of Darcy flows in porous media.

*Remark 4.* All the results reported within this Section are limited to the case of Gaussian processes. While this is an extremely broad class of stochastic processes, one may wonder whether similar results can be obtained for other probability distributions. In general, the  $L^2$  case, namely Theorem 4, can be readily applied to any random field  $\mu$  that satisfies

$$\mathbb{E} \left[ e^{\beta \|\mu\|_{L^2(\Omega)}} \right] < +\infty \text{ for all } \beta > 0. \quad (33)$$

In fact, with such an exponential integrability, it is straightforward to see that these processes admit a KKL expansion and that the proof of Theorem 4 can be easily adapted. For the  $L^\infty$ -case, instead, stronger assumptions are required. In particular, these should be sufficiently demanding to ensure that the properties analogous to those of Lemmas 2 and 3 hold.

Clearly, one may also go the other way around, i.e. by restricting the analysis to more regular operators, with the advantage of allowing for a larger class of probability distributions. For example, Theorems 4 and 5 impose an exponential bound on the local variation,  $\|\partial\mathcal{G}\|$ . This condition is trivially satisfied by all Lipschitz continuous operators, as Lipschitz continuity is a far more stringent property (cf. Lemma 1). In particular, if one restricts the attention to such operators, the proof of Theorem 5 can be re-adapted with weaker assumptions on the random field. That is, one needs to have  $\mathbb{E}\|\mu\|_{L^\infty(\Omega)} < +\infty$  and

$$\mathbb{E}\|\mu - \mu_n\|_{L^\infty(\Omega)} = O\left(\|\text{Cov}_{\mu-\mu_n}\|_{L^\infty(\Omega)}\right),$$

where  $\text{Cov}_{\mu-\mu_n}$  is the covariance kernel of  $\mu - \mu_n$ . Conversely, the case  $L^2$  becomes trivial, as one can replace the condition in (33) with  $\mathbb{E}\|\mu\|_{L^2(\Omega)} < +\infty$ .

*Remark 5.* Theorems 4 and 5 provide different bounds for the reconstruction error. However, the two become very similar if the eigenfunctions of the input field,  $\varphi_i$ , are uniformly bounded. In fact, if  $\sup_i \|\varphi_i\|_{L^\infty(\Omega)} \leq D$  for some  $D > 0$ , then

$$\left\| \sum_{i>n} \lambda_i^\mu \varphi_i^2 \right\|_{L^\infty(\Omega)}^{1/2} \leq \sqrt{\sum_{i>n} \lambda_i^\mu \|\varphi_i^2\|_{L^\infty(\Omega)}} \leq \sqrt{\sum_{i>n} \lambda_i^\mu D^2} = D \sqrt{\sum_{i>n} \lambda_i^\mu}.$$

In particular, up to fixing the value of  $\epsilon$  in Theorem 5, and adjusting the value of the multiplicative constant  $C > 0$ , one can replace the bounding expression in Theorem 5 with that in Theorem 4. However, whether such a uniform boundness holds or not depends on the problem itself. In fact, although erroneously stated by some authors, this property is not directly implied by the regularity of the covariance kernel: see, e.g., [66] for an instructive counterexample.

#### 4.4. Proofs of Theorems 4 and 5

The interested reader can find below the proofs of the two Theorems, which we have postponed here due to their lengths and technicalities. We start with the proof of Theorem 5, as that is arguably the most difficult and most interesting one; the  $L^2$  case will then follow quite easily.

##### Proof of Theorem 5

To begin, we note that the case  $\sum_{i>n} \lambda_i^\mu = 0$  is already covered by Theorem 3, as it falls into the finite-rank context. Thus, from now on we shall focus on proving the error bound for the remaining case, that is, when  $\sum_{i>n} \lambda_i^\mu \geq \epsilon$ .

Without loss of generality, we assume that  $\mathbb{E}[\mu] \equiv 0$ . Before starting with the proof, we note that the existence of a KKL expansion for  $u$  is guaranteed by the exponential growth condition for  $\mathcal{G}$ . In fact, by (30) and Corollary 2, it follows that

$$\mathbb{E}\|u_\mu\|^2 \leq \mathbb{E} \left[ \left( A' e^{\gamma' \|\mu\|_{L^\infty(\Omega)}} \right)^2 \right] = (A')^2 \mathbb{E} \left[ e^{2\gamma' \|\mu\|_{L^\infty(\Omega)}} \right] < +\infty,$$

for some  $A', \gamma' > 0$ , where the last inequality is a direct consequence of Lemma 2. In particular,  $u_\mu$  is a squared-integrable  $V$ -valued random variable, and thus admits a KKL expansion (cf. Lemma 5).

Let now  $n \in \mathbb{N}$ , with  $n \geq 1$ , and let

$$\delta_{n,M,L} := \inf_{\substack{\Psi' \in \mathcal{E}_{B,M}(V, \mathbb{R}^n) \\ \Psi \in \mathcal{D}_{M,L}(\mathbb{R}^n, V)}} \mathbb{E} \|u_\mu - \Psi(\Psi'(u_\mu))\|.$$

We shall split the proof into several steps. More precisely, we shall prove the following.

**Claim 1.** The definition of the control set,  $B$ , is well-posed.

**Claim 2.**  $\exists \ell_0 = \ell_0(n, \mathcal{G}, \mathbb{P})$  such that  $\delta_{n,M,L} \leq \sqrt{\sum_{i>n} \lambda_i^u}$  for all  $M \geq 1$  and  $L \geq \ell_0$ .

**Claim 3.**  $\exists \ell'_0 = \ell'_0(n, \mathcal{G}, \mathbb{P})$  such that  $\delta_{n,M,L} \leq \mathbb{E} \|u - u_{\mu_n}\|$  for all  $M \geq 1$  and  $L \geq \ell'_0$ . Here,  $u_{\mu_n}$  is the operator image where the input field  $\mu$  has been replaced by its  $n$  th KKL truncation,  $\mu_n$ .

**Claim 4.**  $\mathbb{E} \|u - u_{\mu_n}\| \leq c \mathbb{E}^{1/2} \|\mu - \mu_n\|_{L^\infty(\Omega)}^2$  for some  $c = c(d, \mathcal{G}, \mathbb{P})$ .

**Claim 5.**  $\mathbb{E}^{1/2} \|\mu - \mu_n\|_{L^\infty(\Omega)}^2 \leq c' \sqrt{\log(1/\epsilon)} \|\sum_{i>n} \lambda_i^\mu \varphi_i^2\|_{L^\infty(\Omega)}^{1/2}$  for some  $c' = c'(d, \mathbb{P})$ .

Clearly, once all of the above have been proven, setting  $C = \max\{cc', 1\}$  and  $L_0 := \max\{\ell_0, \ell'_0\}$  quickly yields the conclusion. Thus, we now proceed to prove the five claims one by one.

*Proof of Claim 1.* Let  $B_1 := \{\sum_{i=1}^{+\infty} \sqrt{\lambda_i^\mu} \nu_i \varphi_i \mid \sum_i |\nu_i|^2 \leq 1\}$ . Then, for every element of  $B_1$  and every  $\mathbf{x} \in \Omega$ , we have

$$\left| \sum_{i=1}^{+\infty} \sqrt{\lambda_i^\mu} \nu_i \varphi_i(\mathbf{x}) \right| \leq \sqrt{\sum_{i=1}^{+\infty} |\nu_i|^2} \sqrt{\sum_{i=1}^{+\infty} \lambda_i^\mu \varphi_i(\mathbf{x})^2} \leq \sqrt{\text{Cov}(\mathbf{x}, \mathbf{x})}.$$

Since  $\text{Cov}$  is bounded, this shows that  $B_1$  is a bounded subset of  $L^\infty(\Omega)$ . In particular, since  $B = \mathcal{G}(B_1)$ , the definition of the control set is well-posed. Furthermore, the latter is  $\|\cdot\|$  bounded, as  $\mathcal{G}$  maps bounded sets onto bounded sets (cf. Corollary 2).  $\square$

*Proof of Claim 2.* Let  $P : \mathbb{R}^n \rightarrow V$  and  $P^\dagger : V \rightarrow \mathbb{R}^n$  be the linear operators below

$$P : [c_1, \dots, c_n] \mapsto \sum_{i=1}^n \sqrt{\lambda_i^u} c_i v_i, \quad P^\dagger : v \mapsto \left[ \frac{1}{\sqrt{\lambda_1^u}} \langle v, v_1 \rangle, \dots, \frac{1}{\sqrt{\lambda_n^u}} \langle v, v_n \rangle \right].$$

Since  $P^\dagger$  is both linear and continuous, and  $B$  is bounded, the image  $P^\dagger(B)$  is also bounded. Thus, let  $M_0 := \sup_{c \in P^\dagger(B)} |c| < +\infty$ . Similarly, in light of the linearity and continuity of  $P$ , let  $\text{Lip}(P)$  be the Lipschitz constant of  $P$ . Define the maps

$$\tilde{P}^\dagger := M_0^{-1} P, \quad \tilde{P} := M_0 P.$$

It is straightforward to see that  $\tilde{P}^\dagger \in \mathcal{E}_{B,M}(V, \mathbb{R}^n)$  and  $\tilde{P} \in \mathcal{D}_{M,L}(\mathbb{R}^n, V)$ , for all  $M \geq 1$  and all  $L \geq \ell_0 := M_0 \text{Lip}(P)$ , where  $\ell_0$  ultimately depends on  $n, \mathcal{G}$  and  $\mathbb{P}$ . Since

$$\mathbb{E} \|u - \tilde{P} \tilde{P}^\dagger u\| \leq \sqrt{\mathbb{E}} \|u - P P^\dagger u\| = \mathbb{E}^{1/2} \left\| \sum_{i>n} \sqrt{\lambda_i^u} \omega_i v_i \right\|^2 = \sqrt{\sum_{i>n} \lambda_i^u},$$

this proves Claim 2.  $\square$

*Proof of Claim 3.* As before, it is useful to define the mappings  $Q^\dagger : \mathbb{R}^n \rightarrow L^\infty(\Omega)$  and  $Q : L^\infty(\Omega) \rightarrow \mathbb{R}^n$  as

$$Q : [c_1, \dots, c_n] \mapsto \sum_{i=1}^n \sqrt{\lambda_i^\mu} c_i \varphi_i, \quad Q^\dagger : \nu \mapsto \left[ \frac{1}{\sqrt{\lambda_1^\mu}} \langle \nu, \varphi_1 \rangle, \dots, \frac{1}{\sqrt{\lambda_n^\mu}} \langle \nu, \varphi_n \rangle \right],$$

which are both linear and continuous. Following our previous notation, for any  $R > 0$ , let  $B_R = \{\nu \mid \|\nu\|_{L^\infty(\Omega)} \leq R\} \subseteq L^\infty(\Omega)$  be the closed ball of radius  $R$ . We define the map  $\Psi'_R : V \rightarrow \mathbb{R}^n$  as any measurable selection of

the optimization problem below,

$$\Psi'_R : v \mapsto \operatorname{argmin}_{\nu \in Q^\dagger(B_R)} \|v - \mathcal{G}(Q\nu)\|,$$

The existence of such a map is a straightforward consequence of standard results in set-valued analysis. In fact:

- i) the set of minimizers,  $Q^\dagger(B_R) \subset \mathbb{R}^n$ , is compact. This is because  $B_R$  is both closed and bounded, while the map  $Q^\dagger$  is linear, continuous, and has finite rank;
- ii) the map objective functional,  $(v, \nu) \rightarrow \|v - \mathcal{G}(Q\nu)\|$ , is continuous.

Then, these two properties are enough to guarantee the existence of a measurable map acting as "minimal selection": for the interested reader, we refer to the Appendix, Lemma C3. Let now

$$R_0 := \|\operatorname{Cov}\|_{L^\infty(\Omega \times \Omega)}^{1/2}.$$

In light of our calculations at the beginning of the proof, we note that the control set  $B$  is a subset of  $\mathcal{G}(B_R)$  for all  $R \geq R_0$ . Thus, assuming  $R \geq R_0$ , we set

$$\tilde{\Psi}'_R(v) := \mathbb{1}_B(v)\Psi'_{R_0}(v) + \mathbb{1}_{\mathcal{G}(B_R) \setminus B}(v)\Psi'_R(v).$$

Since both sets  $B$  and  $\mathcal{G}(B_R)$  are measurable (cf. Lemma C1 in the Appendix), and  $\tilde{\Psi}'_R(V) \subseteq Q^\dagger(B_R)$ , the above is both measurable and bounded (thus integrable). Furthermore, the above construction ensures that  $\tilde{\Psi}'_R \equiv 0$  outside of  $\mathcal{G}(B_R)$ , and, most importantly

$$|\tilde{\Psi}'_R(v)| \leq \frac{n|\Omega|^{1/2}R_0}{\sqrt{\lambda_n^\mu}}$$

for all  $v \in B$ . In fact, for any  $\nu$  with  $\|\nu\|_{L^\infty(\Omega)} \leq R_0$  one has

$$\left| \frac{1}{\sqrt{\lambda_i^\mu}} \langle \nu, \varphi_i \rangle \right| \leq \frac{|\Omega|^{1/2}}{\sqrt{\lambda_i^\mu}} R_0,$$

by the Cauchy-Schwarz inequality. In particular,  $\tilde{\Psi}'_R$  is bounded on  $B$ , with a constant that is independent of  $R$ . Furthermore, by very definition,

$$\|\mathcal{G}(\nu) - \mathcal{G}(Q\tilde{\Psi}'_R(\nu))\| \leq \|\mathcal{G}(\nu) - \mathcal{G}(QQ^\dagger\nu)\| \quad (34)$$

for all  $\nu \in B_R$ . Now, exploiting the compactness of  $[-1, 1]^n$ , let  $\operatorname{Lip}(\mathcal{G} \circ Q)$  be the Lipschitz constant of  $\mathcal{G} \circ Q$  over  $[-1, 1]^n$  (recall that the latter has a finite due composition; see, e.g., (v) and (ii) in Proposition 1). Define the maps

$$\hat{\Psi}'_R : v \mapsto \frac{\sqrt{\lambda_n^\mu}}{n|\Omega|^{1/2}R_0} \tilde{\Psi}'_R, \quad \Psi : \mathbf{c} \rightarrow \mathcal{G}\left(\frac{n|\Omega|^{1/2}R_0}{\sqrt{\lambda_n^\mu}} Q\mathbf{c}\right).$$

Then the couple  $(\hat{\Psi}'_R, \Psi)$  forms an admissible encoding-decoding pair for all  $M \geq 1$  and  $L \geq \ell'_0$ , where

$$\ell'_0 := \operatorname{Lip}(\mathcal{G} \circ Q) n|\Omega|^{1/2} R_0 / \sqrt{\lambda_n^\mu}$$

depends only on  $n$ ,  $\mathcal{G}$  and  $\mathbb{P}$ . Consequently, for all  $M \geq 1$  and  $L \geq \ell'_0$ , one has

$$\begin{aligned} \delta_{n,M,L} &\leq \mathbb{E} \|u_\mu - \mathcal{G}_n(\tilde{\Psi}'_R(u_\mu))\| \leq \\ &\leq \mathbb{E} [\mathbb{1}_{B_R}(\mu) \|u_\mu - \mathcal{G}_n(\tilde{\Psi}'_R(u_\mu))\|] + \mathbb{E} [\mathbb{1}_{B_R^c}(\mu) \|u_\mu - \mathcal{G}_n(\tilde{\Psi}'_R(u_\mu))\|]. \end{aligned} \quad (35)$$

For better readability, we now write  $u_{\mu_n} := \mathcal{G}(QQ^\dagger\mu)$ , so that  $u_{\mu_n}$  is the image of the operator that is obtained by replacing the input field  $\mu$  with its  $n$ -truncation. Then, by the very definition of  $\tilde{\Psi}'_R$  and thanks to (34), we can continue (35) as

$$\delta_{n,M,L} \leq \dots \leq \mathbb{E} \|u_\mu - u_{\mu_n}\| + \mathbb{E} [\mathbb{1}_{B_R^c}(\mu) \|u_\mu - u_0\|].$$

Since the above holds for every  $R \geq R_0$ , we can let  $R \rightarrow +\infty$ . In doing so, we note that  $\|u_\mu - u_0\|$  is an integrable random variable. Thus,  $\mathbb{E}[\mathbb{1}_{B_R^c}(\mu)\|u_\mu - u_0\|] \rightarrow 0$  by dominated convergence. Claim 3 follows.  $\square$

*Proof of Claim 4.* As a direct consequence of Corollary 1 and Eq. (30), we have

$$\begin{aligned}\|u_\mu - u_{\mu_n}\| &\leq \left( \sup_{0 \leq t \leq 1} A e^{\gamma \|t\mu + (1-t)\mu_n\|_{L^\infty(\Omega)}} \right) \|\mu - \mu_n\|_{L^\infty(\Omega)} \\ &\leq A e^{\gamma \|\mu\|_{L^\infty(\Omega)} + \gamma \|\mu_n\|_{L^\infty(\Omega)}} \|\mu - \mu_n\|_{L^\infty(\Omega)}.\end{aligned}$$

Then, by the Cauchy-Schwarz inequality, we have

$$\mathbb{E}\|u_\mu - u_{\mu_n}\| \leq A \mathbb{E}^{1/4} [e^{4\gamma \|\mu\|_{L^\infty(\Omega)}}] \mathbb{E}^{1/4} [e^{4\gamma \|\mu_n\|_{L^\infty(\Omega)}}] \mathbb{E}^{1/2} \|\mu - \mu_n\|_{L^\infty(\Omega)}^2.$$

For better readability, let now

$$k_q(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^q \lambda_i^\mu \varphi_i(\mathbf{x}) \varphi_i(\mathbf{y}),$$

so that  $k_\infty = \text{Cov}$  and  $k_n$  is the covariance kernel of  $\mu_n$ . Of note, as a straightforward consequence of the Cauchy-Schwarz inequality, one has

$$\|k_q\|_{L^\infty(\Omega^2)} = \max_{\mathbf{x} \in \Omega} k_q(\mathbf{x}, \mathbf{x}). \quad (36)$$

Now, since both kernels are Lipschitz continuous, by Lemmas 2 and 3, we have

$$\mathbb{E}^{1/4} [e^{4\gamma \|\mu\|_{L^\infty(\Omega)}}] = c_2(d, H(k_\infty), 1/2, \|k_\infty\|_{L^\infty(\Omega^2)}, 4\gamma, \Omega)$$

and

$$\mathbb{E}^{1/4} [e^{4\gamma \|\mu_n\|_{L^\infty(\Omega)}}] = c_2(d, H(k_n), 1/2, \|k_n\|_{L^\infty(\Omega^2)}, 4\gamma, \Omega),$$

where  $H(k_\infty)$  and  $H(k_n)$  are the Hölder constants of the two kernels. We now recall that, as shown in Lemma 3,  $k_n \rightarrow k_\infty$  uniformly and  $\sup_q L(k_q) < +\infty$ . In particular,  $\sup_q \|k_q\|_{L^\infty(\Omega^2)} < +\infty$ , and, since  $c_2$  depends continuously on its parameters, we have

$$\mathbb{E}^{1/4} [e^{4\gamma \|\mu\|_{L^\infty(\Omega)}}] \cdot \mathbb{E}^{1/4} [e^{4\gamma \|\mu_n\|_{L^\infty(\Omega)}}] \leq \tilde{c},$$

for some  $\tilde{c} = \tilde{c}(d, \gamma, \mathbb{P}) = \tilde{c}(d, \mathcal{G}, \mathbb{P})$ . In particular, up to letting  $c := A\tilde{c}$ , we may rewrite our previous bound as,

$$\mathbb{E}\|u_\mu - u_{\mu_n}\| \leq \tilde{C} \mathbb{E}^{1/2} \|\mu - \mu_n\|_{L^\infty(\Omega)}^2.$$

$\square$

*Proof of Claim 5.* Following the same notation as above, we note that  $k_\infty - k_n$  is the covariance kernel of the random field  $\mu - \mu_n$ . Thus, by Lemma 2,

$$\mathbb{E}^{1/2} \|\mu - \mu_n\|_{L^\infty(\Omega)}^2 \leq \tilde{c}' \|k_\infty - k_n\|_{L^\infty(\Omega^2)}^{1/2} \left( 1 + \sqrt{\log^+(1/\|k_\infty - k_n\|_{L^\infty(\Omega^2)})} \right) \quad (37)$$

where, by applying the same arguments as before,  $\tilde{c}'$  can be chosen to depend only on  $d$  and  $\mathbb{P}$ . We now note that

$$\|k_\infty - k_n\|_{L^\infty(\Omega^2)} \geq \|k_\infty - k_n\|_{L^2(\Omega^2)} |\Omega| = \sqrt{\sum_{i>n} \lambda_i} |\Omega| \geq \epsilon^{1/2} |\Omega|.$$

The latter can be then combined with (37) to prove that, up to replacing  $\tilde{c}'$  with a suitable  $c'_0$ , independent of  $\epsilon$ , one has

$$\mathbb{E}^{1/2} \|\mu - \mu_n\|_{L^\infty(\Omega)}^2 \leq c'_0 \|k_\infty - k_n\|_{L^\infty(\Omega^2)}^{1/2} \left( 1 + \sqrt{\log^+(1/\epsilon)} \right),$$

as the map  $a \mapsto \log^+(1/a)$  is monotone nonincreasing. Furthermore,

$$\epsilon < 1/2 \implies \log^+(1/\epsilon) = \log(1/\epsilon) \geq \log(2) > 0.$$

Thus, up to further replacing  $c'_0$  with a proper  $c' = c'(d, \mathbb{P})$ , we may write

$$\mathbb{E}^{1/2} \|\mu - \mu_n\|_{L^\infty(\Omega)}^2 \leq c' \sqrt{\log(1/\epsilon)} \|k_\infty - k_n\|_{L^\infty(\Omega^2)}^{1/2} = c' \sqrt{\log(1/\epsilon)} \left\| \sum_{i>n} \lambda_i^\mu \varphi_i^2 \right\|_{L^\infty(\Omega)}^{1/2},$$

where the last equality follows from (36).  $\square$

Finally, putting together the five Claims yields the inequality in (32), and thus proves the statement in the Theorem.  $\square$

#### Proof of Theorem 4

Let  $n \geq 1$ . We notice that, mutatis mutandis, all the steps in the proof of Theorem 5 can be carried out following the same ideas. The only part that actually changes is the estimate in Claim 5, which now concerns the quantity

$$\mathbb{E}^{1/2} \|\mu - \mu_n\|_{L^2(\Omega)}^2,$$

where  $\mu_n \approx \mu$  is the  $n$ th KKL truncation of the random field  $\mu$ . However, the latter can be estimated trivially as, by orthonormality, one has

$$\mathbb{E} \|\mu - \mu_n\|_{L^2(\Omega)}^2 = \mathbb{E} \left[ \sum_{i>n} \lambda_i^\mu \eta_i^2 \right] = \sum_{i>n} \lambda_i^\mu,$$

so that the conclusion follows.  $\square$

## 5. Numerical experiments

The purpose of this Section is to assess the error estimates in Theorems 4 and 5 through a set of numerical experiments. To do so, we proceed in a schematic way so that we may synthesize as follows. First, we introduce the PDE of interest, together with the corresponding solution operator  $\mathcal{G} : \mu \rightarrow u_\mu$ , and a given probability distribution  $\mathbb{P}$  defined over the input space,  $\mu \sim \mathbb{P}$ . Then, we fix a suitable high-fidelity discretization of the input-output spaces, typically via Finite Elements or Finite Volumes, so that the operator under study becomes  $\mathcal{G}_h : \boldsymbol{\mu}^h \rightarrow \mathbf{u}_{\boldsymbol{\mu}^h}^h$  (here, the superscript  $h$  is used to emphasize the presence of a spatial discretization). The discrete operator is then evaluated relying on a given numerical solver, which we exploit to generate a suitable *training set*, that is, a collection of random independent realizations  $\{\boldsymbol{\mu}_i^h, \mathbf{u}_i^h\}$ , where

$$\mathbf{u}_i^h := \mathcal{G}_h(\boldsymbol{\mu}_i^h) \in \mathbb{R}^{N_h},$$

and the  $\boldsymbol{\mu}_i^h$ 's are sampled according to  $\mathbb{P}$ .

To check whether the error bounds in Theorems 4 and 5 are observed in practice, we consider a sequence of "nested" autoencoder architectures,  $\Psi_n \circ \Psi'_n$ , characterized by an increasing latent dimension  $n$ , so that

$$\Psi'_n : \mathbb{R}^{N_h} \rightarrow \mathbb{R}^n, \quad \Psi_n : \mathbb{R}^n \rightarrow \mathbb{R}^{N_h}.$$

The architectures are nested in that they share the same depth and widths, except for the innermost layers (that is, those mapping to/from the latent space, respectively). All the architectures are then trained over the same training set, and their performances are evaluated in terms of the test error below

$$E_{\text{test}}(n) := \frac{1}{N_{\text{test}}} \sum_{j=1}^N \|\mathbf{u}_{j,\text{test}}^h - \Psi_n(\Psi'_n(\mathbf{u}_{j,\text{test}}^h))\|_{V_h}, \quad (38)$$

which serves as a Monte Carlo estimate of  $\mathbb{E} \|\mathbf{u}_{\boldsymbol{\mu}^h}^h - \Psi_n(\Psi'_n(\mathbf{u}_{\boldsymbol{\mu}^h}^h))\|_{V_h}$ . Here,  $\{\boldsymbol{\mu}_{j,\text{test}}^h, \mathbf{u}_{j,\text{test}}^h\}_{j=1}^{N_{\text{test}}}$  is a suitable test set, generated independently of the training set. On the contrary, the norm  $\|\cdot\|_{V_h}$  corresponds to the discretized  $L^2$ -norm associated to the underlying Finite Element (or Finite Volume) space,  $V_h \subseteq L^2(\Omega)$ , i.e.

$$\|\mathbf{v}\|_{V_h} := \sqrt{\mathbf{v}^T \mathbf{M} \mathbf{v}}$$

for all  $\mathbf{v} \in \mathbb{R}^{N_h}$ , where  $\mathbf{M} \in \mathbb{R}^{N_h \times N_h}$  is the so-called *mass matrix*. The purpose is then to compare the behavior of the three quantities below

$$E_{\text{test}}(n), \quad \sqrt{\sum_{i>n} \lambda_i^\mu}, \quad \sqrt{\sum_{i>n} \lambda_i^u}, \quad (39)$$

for varying  $n$ , as suggested by Theorems 4 and 5. To this end, we exploit the POD algorithm to approximate the first  $n$  eigenvalues of the (uncentered) covariance operators of the two fields. Then, the two tails can be easily approximated by noting that

$$\sum_{i>n} \lambda_i^\mu = \sum_{i=1}^{+\infty} \lambda_i^\mu - \sum_{i=1}^n \lambda_i^\mu \approx \text{Var} \|\boldsymbol{\mu}^h\|_{V_h}^2 - \sum_{i=1}^n \lambda_i^{\mu^h},$$

where  $\lambda_i^{\mu^h} \approx \lambda_i^\mu$  are the eigenvalues computed via POD. The same can be done for  $u$ . For our analysis, we typically let  $n = 1, 2, \dots, 6$ , so that the resulting autoencoder architectures are fairly light. This allows us to keep external sources of error out of the way (such as, e.g., inaccuracies due to an inexact optimization of the loss function or shortage in the training data), and thus provide cleaner results.

Once the theoretical error bounds have been assessed, we also take the opportunity to implement a complete DL-ROM surrogate: while this is not directly related to our analysis in Section 4, it serves the purpose of showing how the whole machinery can be put into action to provide an operative ROM. For simplicity, this step is only repeated once with a fixed latent dimension of choice.

Last but not least, we mention that all the code supporting the forthcoming analysis has been written in Python 3, specifically relying on the FEniCS and Pytorch libraries. The code is available upon request to the authors.

### 5.1. Problems description

We start by introducing the three case studies one by one: all the results are then reported and discussed at the end of this Section. For the sake of our analysis, we have selected two prototypical problems that aim to show different possible behaviors of the solution manifold. The first case study features a diffusive phenomenon, which results in a highly regularizing solution operator: as a consequence, this problem models a scenario in which the eigenvalues of the solution field,  $\lambda_i^\mu$ , decay faster than those of the input field,  $\lambda_i^\mu$ . Conversely, the second case study concerns a nonlinear advection characterized by the presence of shock waves and thus constitutes a remarkable example of the opposite situation (i.e.,  $\lambda_i^\mu \rightarrow 0$  faster than  $\lambda_i^u$ ).

We believe that despite their simplicity, these problems suffice to show the practical counterpart of our theoretical analysis in Section 4. In this respect, we remark that this work does not aim at showcasing the abilities of DL-ROMs in handling complex problems: in fact, the effectiveness of the DL-ROM approach has already been reported elsewhere; see, e.g., [24] for problems concerning fluid dynamics or multiphysics.

#### 5.1.1. Stochastic Darcy flow with random permeability

To start, we consider an elliptic problem that describes diffusion in a porous medium with random permeability, in the same spirit as our previous discussion in Section 3, cf. Proposition 2. More precisely, let  $\Omega := (0, 1)^2$  be the unit square. We consider the boundary value problem below

$$\begin{cases} -\nabla \cdot (e^\mu \nabla u) = 10 & \text{in } \Omega \\ u = 0 & \text{on } \partial\Omega \end{cases}, \quad (40)$$

where  $\mu$  is a centered Gaussian random field with covariance

$$\text{Cov}(\mathbf{x}, \mathbf{y}) := \exp(-|\mathbf{x} - \mathbf{y}|^2).$$

Here, we focus our attention on the solution operator  $\mathcal{G} : \mu \rightarrow u_\mu$  that maps the log-permeability of the medium onto the corresponding solution to (40). We discretize the problem using piecewise linear continuous Finite Elements over a structured triangular mesh of step size  $h = \sqrt{2}/50$ , resulting in a state space  $V_h \subseteq L^2(\Omega)$  with  $N_h = 2601$  degrees of freedom. We generate a total of 5000 snapshots, split between training and testing with a 90:10 ratio. For technical details on the DL-ROM architectures and their training, we refer the reader to the appendix, Section D.

### 5.1.2. Burger's equation with random data

On the segment  $\Omega := (0, L)$ ,  $L = 5$ , we consider the inviscid Burger equation

$$\frac{\partial v}{\partial t} + \frac{1}{2}v \frac{\partial v}{\partial x} = 0, \quad (41)$$

describing the nonlinear transport of a given solute  $v = v(x, t)$  with random initial condition  $v(\cdot, 0) = \mu$ . We complement Eq. (41) with a stationary inflow condition on the left boundary, that is,  $v(0, t) \equiv v(0, 0)$  for all  $t \geq 0$ . Here, we model the trajectories of the random field  $\mu$  as

$$\mu(x) = \frac{1}{2}\rho\left(\phi_0(x) + \sum_{k=1}^{+\infty} \frac{1}{k^2}\eta_k \sin\left(\frac{k\pi x}{L}\right)\right) \quad (42)$$

where  $\{\eta_i\}_{i=1}^{+\infty}$  are i.i.d. standard Gaussians,  $\rho(x) := \min\{\max\{x, 0\}, 1\}$  is a suitable transformation that clamps the data within  $[0, 1]$ , and

$$\phi_0(x) := (x - 1)(2 - x)\mathbb{1}_{[1, 2]}(x).$$

We are interested in approximating the following parameter-to-solution operator

$$\mathcal{G} : \mu \rightarrow u_\mu := v(\cdot, T), \quad (43)$$

which maps any given initial profile onto the state of the system at time  $T = 2$ . We discretize the problem using the Finite Volume method, with a temporal step  $\Delta t = 0.01$  and cells size  $h = 0.01$ , resulting in a high-fidelity state space of dimension  $N_h = 500$ . We exploit the FOM to generate a total of 2000 snapshots, split between training and testing with a 90:10 ratio. Once again, to keep the paper self-contained, we postpone all the technical details about the DL-ROM architectures and their training to the Appendix, Section D.

## 5.2. Results

We start with the first case study, Section 5.1.1. As we mentioned, the parameter-to-solution operator of this problem is highly regularizing because of the diffusive term in the PDE. Indeed, when comparing the tails of the eigenvalues, we see that those of the output field decay 2.5× faster than those of the input field; see Figure 3. The same rate is also achieved by autoencoders, both across the training and the test set. This result is in agreement with our theory since, in this case, Theorem 5 applies (cf. 2).

Since linear methods, such as POD, can directly exploit the decay of the eigenvalues of the output field, and they require fewer data with respect to deep autoencoders, these approaches should be favored when dealing with problems of this type. It should be noted that this conclusion can already be derived from Theorem 5 without having to test multiple autoencoder modules. Nevertheless, if provided with enough data, one may still choose to use autoencoders: at worst, they will match the same accuracy as linear methods (up to optimization errors).

The second case study, on the contrary, shows the opposite situation. This time, the eigenvalues of the output field have a tail that decays 20%-50% slower than that of the input field. Then, our theory suggests that a nonlinear method based on deep autoencoders can exploit this hidden regularity to provide better approximation capabilities. Indeed, this is what we observe in practice; see Figure 4.

Not only do autoencoders surpass the performances of POD by a significant margin, but they also report a much faster decay rate. Notably, when  $n \leq 6$ , the latter is seen to perfectly match the decay rate of the eigenvalues  $\lambda^\mu$ . The situation is similar for  $7 \leq n \leq 12$ , where both quantities are seen to accelerate towards 0. However, as we come to large architectures,  $n = 13, 14, 15, 16$ , we notice an increasingly larger gap between training and test errors. This may be due to the fact that larger models typically require more data in order to generalize properly. This would explain why autoencoders seem to reach the optimal decay rate —according to Theorem 5— over the training set but not over the test set (at least for large  $n$ ). In this sense, although our findings agree with Theorem 5, they also highlight some of its inherent limitations, namely, the lack of a comprehensive analysis encapsulating model accuracy, complexity, and training.

Last but not least, it should be also noted that the results in Figure 4 are even more interesting if we consider that Problem (41) only partially fulfills the hypothesis of Theorem 5. Indeed, while it can be



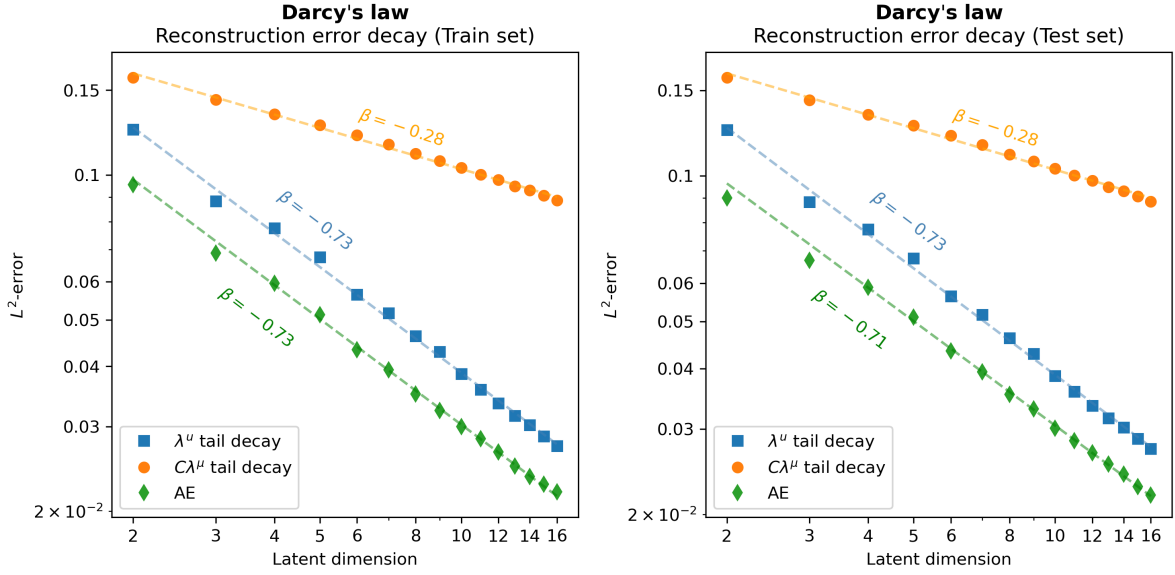


Figure 3: Decay of the reconstruction error for Darcy's law §5.1.1 computed over the training set (left) and the test set (right). In both cases, errors are compared with the tails of the eigenvalues of the input and output fields, respectively: cf. Equations (38) and (39). Eigenvalues of the input field have been scaled by a factor  $C > 0$  to improve readability. Dashed lines are obtained through least-squares in the loglog space.  $\beta$  = rate of decay, computed as the slope of the dashed lines.

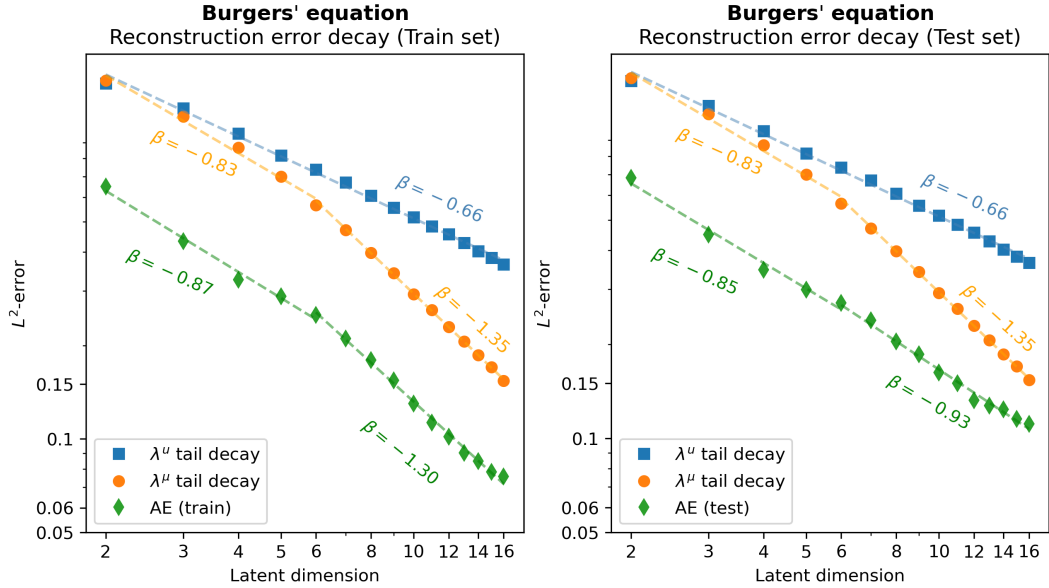


Figure 4: Decay of the reconstruction error for Burger's equation §5.1.2 computed over the training set (left) and the test set (right). In both cases, errors are compared with the tails of the eigenvalues of the input and output fields, respectively: cf. Equations (38) and (39). Dashed lines are obtained through least-squares in the loglog space.  $\beta$  = rate of decay, computed as the slope of the dashed lines.




Figure 5:  $L^\infty$ -norms of the eigenfunctions in the KKL expansion of the input field  $\mu$  for the Burger’s equation §5.1.2: see Equations (31) and (42), respectively.

shown that the parameter-to-solution operator in (43) enjoys a form of  $L^\infty \rightarrow L^2$  Lipschitz continuity, the probability law of the input field is non-Gaussian: see also our previous discussion in Remark 4.

In this concern, we also mention that although Theorem 5 would require the computation of an  $L^\infty$ -tail, that is,  $\|\sum_{i>n} \lambda_i^\mu \varphi_i^2\|_\infty$ , it is sufficient to monitor the behavior of the  $L^2$ -tail,  $\sum_{i>n} \lambda_i^\mu$ . This is because, in this case, the eigenfunctions of the input field are uniformly bounded in the  $L^\infty$ -norm, fact that we can easily appreciate from the plot in Figure 5. Then, the two tails can be shown to decay at the same rate (see also our discussion in Remark 5).

For the sake of completeness, Table 1 shows the overall performances attained by two DL-ROM surrogates for the problems at hand. Although this analysis is not directly related to the theory developed in Section 4, it might still be of interest as it goes back to the global picture (that is, that of reduced order modeling). This time, we train all the architectures in the DL-ROM pipeline simultaneously following our initial discussion in Section 2.1: see, in particular, Equation (6). This means that the autoencoder module,  $\Psi \circ \Psi'$ , is trained together with all the remaining parts of the DL-ROM: consequently, although we are minimizing the reconstruction error made by the autoencoder, there are also other quantities driving the optimization (such as, e.g., the approximation error of the DL-ROM).

In the Darcy flow example, the DL-ROM reports an average  $L^2$ -error of 5.43%. At the same time, its autoencoder module shows a similar accuracy, with a reconstruction error of 4.94%. As expected, this performance is also comparable with the one achieved by a linear approach such as POD: as we mentioned previously, in fact, we do not really need a nonlinear technique to reduce a problem of this type. However, things become quite different when we move to Burger’s equation. In this case, the autoencoder module is almost twice as accurate compared to its linear counterpart, with an average  $L^2$  error of 5.07%. This fact is also reflected in the overall performance of the DL-ROM, which reports an average test error of 5.62%. The interested reader can also find a few examples in Figure 6, where we compare ground-truth simulations and DL-ROM outputs for new random realizations of the input field  $\mu$ .

Inspired by the results in Table 1, we conclude with a short digression on the interplay between reconstruction errors and approximation errors. In general, when it comes to DL-ROMs, the two quantities are only indirectly related. To better explain this fact, let us first consider the opposite case of POD-based

| Problem            | $n$ | POD $L^2$ -error | AE $L^2$ -error | DL-ROM $L^2$ -error |
|--------------------|-----|------------------|-----------------|---------------------|
| Darcy's law §5.1.1 | 16  | 4.46%            | 4.94%           | 5.43%               |
| Burger's §5.1.2    | 16  | 9.74%            | 5.07%           | 5.62%               |

Table 1: Average test errors for the three case studies in Section 5. Here, the latent dimension  $n$  is fixed. POD = projection error, AE = reconstruction error, DL-ROM = model error.




Figure 6: DL-ROM predictions for three unseen initial states of Burger's equation, §5.1.2.

ROMs. Ultimately, the POD projector  $\mathbf{V} \in \mathbb{R}^{N_h \times n}$  operates as a linear autoencoder module, as

$$\mathbf{u}_\mu^h \approx \mathbf{V}\mathbf{V}^T \mathbf{u}_\mu^h,$$

which allows one to represent each  $\mathbf{u}_\mu^h$  with the corresponding set of projection coefficients  $\mathbf{V}^T \mathbf{u}_\mu^h \in \mathbb{R}^n$ . Then, any POD-based ROM, such as, e.g., POD-Galerkin [52], POD-NN [32], POD-DeepONet [42] and POD-GPR [28], can be written abstractly as

$$\mathbf{u}_\mu^h \approx \mathbf{V}\phi_{\text{POD}}(\boldsymbol{\mu}),$$

where  $\phi_{\text{POD}} : \mathbb{R}^{N_h} \rightarrow \mathbb{R}^n$  is some black-box procedure that maps parameters onto reduced coefficients: the latter can be a neural network model, as in POD-NN and POD-DeepONet, a Gaussian process approximator, as in POD-GPR, or a suitable numerical method solving the projected PDE, as in POD-Galerkin. Then, by orthogonality, the error of any such method can be bounded from below as

$$\|\mathbf{u}_\mu^h - \mathbf{V}\phi_{\text{POD}}(\boldsymbol{\mu})\|^2 = \|\mathbf{u}_\mu^h - \mathbf{V}\mathbf{V}^T \mathbf{u}_\mu^h\|^2 + \|\mathbf{V}^T \mathbf{u}_\mu^h - \phi_{\text{POD}}(\boldsymbol{\mu})\|^2 \geq \|\mathbf{u}_\mu^h - \mathbf{V}\mathbf{V}^T \mathbf{u}_\mu^h\|^2.$$

That is: the approximation error of any such ROM is bounded from below by the projection error of the POD. Furthermore, this fact is not an intrinsic property of the POD basis; instead, it is a feature common to all projection-based methods.

In general, the same is not true for DL-ROMs. In fact, since all architectures in the DL-ROM pipeline are optimized simultaneously, cf. Equation (6), the approximation error of the DL-ROM might as well be smaller than the reconstruction error of the autoencoder, and vice versa. This can happen, for example, if it turns out to be simpler to describe latent variables using the input parameters (thus, through  $\phi$ ) rather than using the output features (that is, through the encoder  $\Psi'$ ).

Still, when constructing a DL-ROM, a good rule of thumb is to prefer those architectures for which the two errors behave similarly: this, in fact, ensures a stronger connection between inputs and outputs, and it increases the interpretability of the DL-ROM as a whole.

## 6. Conclusions

This work addresses the practical problem of designing deep autoencoders, specifically in terms of their latent dimension, in the context of reduced order modeling for PDEs parametrized by random fields. This

topic is of particular interest, since deep autoencoders have recently emerged as the pivotal element of Deep Learning based ROMs (DL-ROMs), a novel class of approaches that exploit manifold learning to approximate the solution manifold of a parametrized operator, equipping them with the ability to tackle complex problems for which traditional methods may fall short. The presented research is novel in that it addresses both theoretically and practically the case of stochastic PDEs, a scenario that is particularly relevant to applications involving, e.g., uncertainty quantification or Bayesian inversion, aspects hitherto unexplored in the DL-ROM literature.

Our main contribution is well summarized by Theorems 3-5, in which we provide explicit error bounds that can aid domain practitioners in selecting the appropriate latent dimension for deep autoencoders. Our findings are highly interpretable as they demonstrate the capacity of deep autoencoders to match or surpass the performance achieved by linear techniques. Numerical experiments agree with our theory, effectively highlighting the practical importance of our analysis in the intricate design of DL-ROMs. Future work may include the derivation of complementary results on the complexity of deep autoencoders and their training, possibly exploring a compromise between data-driven and physics-based approaches, in the same spirit of other recent works; see, e.g., [1, 16, 20].

## Acknowledgements

The present research is part of the activities of project Dipartimento di Eccellenza 2023-2027, funded by MUR, and of project FAIR (Future Artificial Intelligence Research) project, funded by the NextGenerationEU program within the PNRR-PE-AI scheme (M4C2, Investment 1.3, Line on Artificial Intelligence).

## Funding

This project has received funding under the ERA-NET ERA PerMed/ FRRB grant agreement No ERAPERMED2018-244, RADprecise - Personalized radiotherapy: incorporating cellular response to irradiation in personalized treatment planning to minimize radiation toxicity and under the Project “Reduced Order Modeling and Deep Learning for the real-time approximation of PDEs (DREAM)” (Starting Grant No. FIS00003154), funded by the Italian Science Fund (FIS) - Ministero dell’Università e della Ricerca.

## Data and code availability

Data and source code are available upon request: all related inquiries should be directed to the authors.

## Appendix A Technical proofs for Section 3

**Lemma 1.** *Let  $\mathbb{P}$  be a probability distribution over  $\mathbb{R}^p$ . Let  $\rho : \mathbb{R} \rightarrow \mathbb{R}$  be a continuous map. Assume that either one of the following holds:*

- i)  $\rho$  is bounded and nonconstant;*
- ii)  $\rho$  is bounded from below and  $\rho(x) \rightarrow +\infty$  as  $x \rightarrow +\infty$ .*
- iii) there exists some  $a, b \in \mathbb{R}$  such that  $x \rightarrow a\rho(x) + b\rho(-x)$  satisfies (ii).*

*Then, for every  $\varepsilon > 0$  and every measurable map  $f : \mathbb{R}^p \rightarrow \mathbb{R}^n$  with  $\mathbb{E}|f(x)| < +\infty$ , there exists  $\Phi \in \mathcal{N}_\rho(\mathbb{R}^p, \mathbb{R}^n)$  such that*

$$\mathbb{E}_{\mathbf{x} \sim \mathbb{P}} |f(\mathbf{x}) - \Phi(\mathbf{x})| < \varepsilon.$$

*Proof.* If (i) holds, then this is just a consequence of Hornik’s Theorem [33]. Then, let (ii) hold. Since  $\rho$  is bounded from below,  $\rho(\mathbb{R}) \subset [A, +\infty)$  for some  $A \in \mathbb{R}$ . Let

$$\sigma(x) := \rho(-\rho(x) + \alpha),$$

where  $\alpha$  is some parameter to be fixed later on. We have,

$$\begin{aligned} \sup_{x \in \mathbb{R}} |\sigma(x)| &= \sup_{x \in \mathbb{R}} |\rho(-\rho(x) + \alpha)| = \sup_{y \in \rho(\mathbb{R})} |\rho(-y + \alpha)| \leq \\ &\leq \sup_{y \geq A} |\rho(-y + \alpha)| = \sup_{z \leq \alpha - A} |\rho(z)| < +\infty, \end{aligned}$$

in fact, by continuity,  $\rho$  is bounded on all intervals of the form  $(-\infty, c]$ , with  $c \in \mathbb{R}$ . In particular,  $\sigma$  is bounded. Furthermore, since  $\rho(x) \rightarrow +\infty$  as  $x \rightarrow +\infty$ ,

$$\limsup_{x \rightarrow \infty} \sigma(x) \leq \limsup_{z \rightarrow -\infty} \rho(z) = c_0 < +\infty.$$

Let now  $\alpha$  be such that

$$\sigma(0) = \rho(\alpha - \rho(0)) > c_0.$$

Then,  $\sigma$  is guaranteed to be both bounded and nonconstant. However, because of the way we defined  $\sigma$ , we also have

$$\mathcal{N}_\sigma(\mathbb{R}^p, \mathbb{R}^n) \subseteq \mathcal{N}_\rho(\mathbb{R}^p, \mathbb{R}^n),$$

and the conclusion now follows by (i). Finally, assume that (iii) holds and let

$$\tilde{\rho}(x) := a\rho(x) + b\rho(-x).$$

As before, we have  $\mathcal{N}_{\tilde{\rho}}(\mathbb{R}^p, \mathbb{R}^n) \subseteq \mathcal{N}_\rho(\mathbb{R}^p, \mathbb{R}^n)$ , and the conclusion follows by (ii).  $\square$

**Proposition 1.** *Let  $\mathcal{G} : (W, \|\cdot\|_W) \rightarrow (V, \|\cdot\|_V)$  be an operator between two normed spaces. We have the following,*

i)  $\|\partial\mathcal{G}\|_V(w) < +\infty$  for all  $w \in K \subset W \iff \mathcal{G}$  is locally Lipschitz over  $K$ ;

ii) if  $K \subset W$  is compact and  $\|\partial\mathcal{G}\|_V(w) < +\infty$  for all  $w \in K$ , then  $\mathcal{G}$  is Lipschitz over  $K$ ;

iii) if  $C \subseteq W$  is convex, then

$$L_C := \sup_{w \in C} \|\partial\mathcal{G}\|_V(w) < +\infty$$

if and only if  $\mathcal{G}$  is  $L_C$ -Lipschitz over  $C$ .

iv) if  $\mathcal{G}$  is Fréchet differentiable at  $w \in W$ , then  $\|\partial\mathcal{G}\|_V(w)$  coincides with the operator norm of the Fréchet derivative of  $\mathcal{G}$  at  $w$ .

v) given any  $\mathcal{F} : (V, \|\cdot\|_V) \rightarrow (Y, \|\cdot\|_Y)$ , one has the chain-rule inequality

$$\|\partial(\mathcal{F} \circ \mathcal{G})\|_Y(w) \leq \|\partial\mathcal{F}\|_Y(\mathcal{G}(w)) \cdot \|\partial\mathcal{G}\|_V(w),$$

for all  $w \in W$ .

*Proof.* Since (i) is trivial, we skip its proof.

ii) Let  $K \subset W$  be a compact subset. Seeking contradiction, let us assume that  $\mathcal{G}$  is not Lipschitz continuous over  $K$ . Then, there exists two sequences  $\{w_n\}_n$  and  $\{v_n\}_n$  such that

$$\frac{\|\mathcal{G}(w_n) - \mathcal{G}(v_n)\|_V}{\|w_n - v_n\|_W} \rightarrow +\infty \quad (44)$$

as  $n \rightarrow +\infty$ . Since  $\mathcal{G}$  is locally Lipschitz, it is also continuous, and thus bounded over  $K$ : therefore, the above implies  $\|w_n - v_n\|_W \rightarrow 0$ . At the same time, by compactness, there exists a subsequence  $\{w_{n_k}\}_k \subseteq \{w_n\}_n$  and an element  $w \in K$  such that  $w_{n_k} \rightarrow w$ . In particular, we have  $w_{n_k}, v_{n_k} \rightarrow w$ . But then (44) would yield  $\|\partial\mathcal{G}\|_V(w) = +\infty$ , absurd.

- iii) Let  $C$  be convex and assume that  $L_C := \sup_{w \in C} \|\partial \mathcal{G}(w)\|_V < +\infty$ . Given any two points  $v, v' \in C$ , let  $K$  be the segment between the two. Since  $K$  is compact, it follows from (ii) that  $\mathcal{G}$  is Lipschitz continuous over  $K$ . Furthermore, since  $K \subseteq C$ ,

$$\sup_{w \in K} \|\partial \mathcal{G}(w)\|_V \leq L_C$$

Exploiting the very definition of limit supremum, for all  $w \in K$ , let  $B(w, r_w)$  be a ball of radius  $r_w > 0$  such that

$$\left| \frac{\|\mathcal{G}(w+h) - \mathcal{G}(w)\|_V}{\|h\|_W} - \|\partial \mathcal{G}\|_V(w) \right| < \varepsilon \quad \forall h \in B(w, r_w).$$

Then, up to rewriting the above, for all  $w \in K$  we have

$$h \in B(w, r_w) \implies \|\mathcal{G}(w+h) - \mathcal{G}(w)\|_V \leq (L_C + \varepsilon)\|h\|_W.$$

We now note that, since  $K \subset \cup_{w \in K} B(w, r_w)$  and  $K$  is compact, there exists a finite sequence of points  $w_1, \dots, w_n \in K$ , and finite sequence of radii  $r_1, \dots, r_n > 0$ , such that  $K \subset \cup_{i=1}^n B(w_i, r_i)$ . Furthermore, upto to removing some of the balls, since  $K$  is actually a segment, we can sort the subcover so that

$$K \cap B(w_i, r_i) \cap B(w_{i+1}, r_{i+1}) \neq \emptyset \quad \forall i = 1, \dots, n-1.$$

For each  $i = 1, \dots, n-1$ , let  $w_i^* \in K \cap B(w_i, r_i) \cap B(w_{i+1}, r_{i+1})$ . We note that, since  $S_{v,v'}$  is a straightline, we have

$$\|v - v'\|_W = \|v - w_1^*\|_W + \dots + \|w_{n-1}^* - v'\|_W.$$

In particular,

$$\begin{aligned} \|\mathcal{G}(v) - \mathcal{G}(v')\|_W &\leq \\ &\leq \|\mathcal{G}(v) - \mathcal{G}(w_1^*)\|_W + \dots + \|\mathcal{G}(w_{n-1}^*) - \mathcal{G}(v')\|_W \leq \\ &\leq (L_C + \varepsilon) (\|v - w_1^*\|_W + \dots + \|w_{n-1}^* - v'\|_W) = \\ &= (L_C + \varepsilon) \|v - v'\|_W. \end{aligned}$$

Since  $v, v' \in C$  and  $\varepsilon > 0$  were arbitrary, this concludes the proof (the other implication, " $\Leftarrow$ ", is trivial and left to the reader).

- iv) Assume that  $\mathcal{G}$  is Fréchet differentiable at  $w \in W$ , and let  $\delta \mathcal{G}[w] : W \rightarrow V$  be the linear operator representing its derivative. Let  $\{h_n\}_n$  be a sequence of unitary increments,  $\|h_n\|_W = 1$ , such that

$$\|\delta \mathcal{G}[w](h_n)\|_V \longrightarrow \|\delta \mathcal{G}[w]\|_{W,V}$$

where  $\|\cdot\|_{W,V}$  is the operator norm for linear maps going from  $W$  to  $V$ . Then,

$$\begin{aligned} \|\partial \mathcal{G}\|_W(w) &\geq \lim_{n \rightarrow +\infty} \frac{\|\mathcal{G}(w + n^{-1}h_n) - \mathcal{G}(w)\|_V}{\|n^{-1}h_n\|_W} = \\ &= \lim_{n \rightarrow +\infty} \frac{\|\delta \mathcal{G}[w](n^{-1}h_n)\|_V}{\|n^{-1}h_n\|_W} = \lim_{n \rightarrow +\infty} \frac{\|\delta \mathcal{G}[w](h_n)\|_V}{\|h_n\|_W} = \\ &= \|\delta \mathcal{G}[w]\|_{W,V}. \end{aligned}$$

Conversely, let  $\{\tilde{h}_n\}_n$  be such that  $\tilde{h}_n \rightarrow 0$  and

$$\|\partial \mathcal{G}\|_W(w) = \lim_{n \rightarrow +\infty} \frac{\|\mathcal{G}(w + \tilde{h}_n) - \mathcal{G}(w)\|_V}{\|\tilde{h}_n\|_W}.$$

Then,

$$\|\partial \mathcal{G}\|_W(w) = \dots = \lim_{n \rightarrow +\infty} \frac{\|\delta \mathcal{G}[w](\tilde{h}_n)\|_V}{\|\tilde{h}_n\|_W} \leq \|\delta \mathcal{G}[w]\|_{W,V},$$

thus  $\|\partial \mathcal{G}\|_W(w) = \|\delta \mathcal{G}[w]\|_{W,V}$  as claimed.

v) Let  $w \in W$ . If either  $\|\partial\mathcal{F}\|_Y(\mathcal{G}(w)) = +\infty$  or  $\|\partial\mathcal{G}\|_V(w) = +\infty$ , then the claim is trivially true. Thus, we assume both quantities to be finite. Fix any  $\varepsilon > 0$ . Then, by definition of limit supremum, there exists  $r_\varepsilon > 0$  such that  $\mathcal{F}$  is  $(\|\partial\mathcal{F}\|_Y(\mathcal{G}(w)) + \varepsilon)$ -Lipschitz over the open ball of radius  $r_\varepsilon$  centered at  $\mathcal{G}(w)$ . It follows that

$$\begin{aligned} \lim_{r \rightarrow 0^+} \sup_{0 < \|h\|_W \leq r} \frac{\|\mathcal{F}(\mathcal{G}(w+h)) - \mathcal{F}(\mathcal{G}(w))\|_V}{\|h\|_W} &\leq \\ &\leq \lim_{r \rightarrow 0^+} \sup_{0 < \|h\|_W \leq r} \frac{(\|\partial\mathcal{F}\|_Y(\mathcal{G}(w)) + \varepsilon) \|\mathcal{G}(w+h) - \mathcal{G}(w)\|_V}{\|h\|_W} = \\ &= (\|\partial\mathcal{F}\|_Y(\mathcal{G}(w)) + \varepsilon) \|\partial\mathcal{G}\|_V(w), \end{aligned}$$

as  $r \leq r_\varepsilon$  definitely. Since  $\varepsilon$  was arbitrary, the conclusion follows.  $\square$

**Proposition 2.** Let  $\Omega \subset \mathbb{R}^d$  be a bounded domain with Lipschitz boundary and let  $f \in H^{-1}(\Omega)$  be given. For any  $\sigma \in L^\infty(\Omega)$ , let  $u = u_\sigma$  be the solution to the following boundary value problem,

$$\begin{cases} -\nabla \cdot (e^\sigma \nabla u) = f & \text{in } \Omega \\ u = 0 & \text{on } \partial\Omega \end{cases}.$$

Let  $\mathcal{G} : L^\infty(\Omega) \rightarrow L^2(\Omega)$  be the operator that maps  $\sigma$  to  $u$ . Then, for all  $\sigma, \sigma' \in L^\infty(\Omega)$

$$\|\mathcal{G}(\sigma) - \mathcal{G}(\sigma')\|_{L^2(\Omega)} \leq C \|f\|_{H^{-1}(\Omega)} e^{3\|\sigma\|_{L^\infty(\Omega)} + 3\|\sigma'\|_{L^\infty(\Omega)}} \|\sigma - \sigma'\|_{L^\infty(\Omega)} \quad (45)$$

and, in particular,

$$\|\partial\mathcal{G}(\sigma)\|_{L^2(\Omega)} \leq C \|f\|_{H^{-1}(\Omega)} e^{6\|\sigma\|_{L^\infty(\Omega)}},$$

where  $C = C(\Omega)$  is some positive constant.

*Proof.* By classical energy estimates on elliptic PDEs, see e.g. Lemma C.1. in [19], we have

$$\begin{aligned} \|\mathcal{G}(\sigma) - \mathcal{G}(\sigma')\|_{L^2(\Omega)} &\leq C \|\mathcal{G}(\sigma) - \mathcal{G}(\sigma')\|_{H_0^1(\Omega)} \leq \\ &\leq C \left[ \min_{\mathbf{x} \in \Omega} e^\sigma(\mathbf{x}) \right]^{-1} \left[ \min_{\mathbf{x} \in \Omega} e^{\sigma'}(\mathbf{x}) \right]^{-1} \|f\|_{H^{-1}(\Omega)} \|e^\sigma - e^{\sigma'}\|_{L^\infty(\Omega)}. \end{aligned}$$

The latter is bounded by

$$C e^{\|\sigma\|_{L^\infty(\Omega)} + \|\sigma'\|_{L^\infty(\Omega)}} \|f\|_{H^{-1}(\Omega)} \|e^{\sigma'}\|_{L^\infty(\Omega)} \|e^{\sigma - \sigma'} - 1\|_{L^\infty(\Omega)}.$$

We now note that, for all  $a \in \mathbb{R}$  one has  $|e^a - 1| \leq |a|e^{|a|}$ . Also,  $\|e^{\sigma'}\|_{L^\infty(\Omega)} = e^{\|\sigma'\|_{L^\infty(\Omega)}}$  by monotonicity of the exponential. It follows that

$$\begin{aligned} \|\mathcal{G}(\sigma) - \mathcal{G}(\sigma')\|_{L^2(\Omega)} &\leq \\ &\leq C e^{\|\sigma\|_{L^\infty(\Omega)} + 2\|\sigma'\|_{L^\infty(\Omega)}} \|f\|_{H^{-1}(\Omega)} \|e^{\sigma - \sigma'}\|_{L^\infty(\Omega)} \|\sigma - \sigma'\|_{L^\infty(\Omega)}. \end{aligned}$$

Since  $\|e^{\sigma - \sigma'}\|_{L^\infty(\Omega)} \leq \|e^\sigma\|_{L^\infty(\Omega)} \|e^{\sigma'}\|_{L^\infty(\Omega)} = e^{\|\sigma\|_{L^\infty(\Omega)} + \|\sigma'\|_{L^\infty(\Omega)}}$ , (45) easily follows.  $\square$

**Lemma 2.** Let  $\Omega \subset \mathbb{R}^d$  be pre-compact, and let  $Z$  be a mean zero Gaussian random field defined over  $\Omega$ . Assume that, for some  $0 < \alpha \leq 1$ , the covariance kernel of the process,

$$\text{Cov} : \Omega \times \Omega \rightarrow \mathbb{R},$$

$$\text{Cov}(\mathbf{x}, \mathbf{y}) := \mathbb{E}[Z(\mathbf{x})Z(\mathbf{y})],$$

is  $\alpha$ -Hölder continuous, with Hölder constant  $L > 0$ . Then,  $Z$  is sample-continuous, that is  $\mathbb{P}(Z \in \mathcal{C}(\Omega)) = 1$ . Furthermore, for  $\sigma^2 := \max_{\mathbf{x} \in \Omega} \text{Cov}(\mathbf{x}, \mathbf{x})$ , one has

$$\mathbb{E}^{1/2} \|Z\|_{L^\infty(\Omega)}^2 \leq c_1 \sigma \left(1 + \sqrt{\log^+(1/\sigma)}\right) \quad \text{and} \quad \mathbb{E}[e^{\beta \|Z\|_{L^\infty(\Omega)}}] = c_2 < +\infty, \quad (46)$$

for all  $\beta > 0$ , where  $c_1 = c_1(d, L, \alpha, \Omega)$  and  $c_2 = c_2(d, L, \alpha, \sigma, \beta, \Omega)$  are constants that depend continuously on their parameters (domain excluded). Here,

$$\log^+(a) := \max\{\log a, 0\}.$$

*Proof.* If  $\sigma = 0$ , the proof is trivial; thus, we let  $\sigma > 0$ . Let

$$d(\mathbf{x}, \mathbf{y}) := \mathbb{E}^{1/2} |Z(\mathbf{x}) - Z(\mathbf{y})|^2 = \sqrt{\text{Cov}(\mathbf{x}, \mathbf{x}) - 2\text{Cov}(\mathbf{x}, \mathbf{y}) + \text{Cov}(\mathbf{y}, \mathbf{y})},$$

be a metric over  $\Omega$  induced by the Gaussian process  $Z$ . We note that the Hölder continuity of the covariance kernel implies

$$d(\mathbf{x}, \mathbf{y}) \leq \sqrt{\text{Cov}(\mathbf{x}, \mathbf{x}) - \text{Cov}(\mathbf{x}, \mathbf{y})} + \sqrt{\text{Cov}(\mathbf{y}, \mathbf{y}) - \text{Cov}(\mathbf{x}, \mathbf{y})} \leq \sqrt{2L|\mathbf{x} - \mathbf{y}|^\alpha}.$$

In particular, the balls  $B_d(\mathbf{x}, \varepsilon) := \{\mathbf{y} \in \Omega : d(\mathbf{x}, \mathbf{y}) < \varepsilon\}$  induced by the metric  $d$  satisfy

$$B_d(\mathbf{x}, \varepsilon) \supseteq B(\mathbf{x}, 2^{-1/\alpha} L^{-1/\alpha} \varepsilon^{2/\alpha}), \quad (47)$$

where  $B(\mathbf{x}, \epsilon)$  is the Euclidean ball of radius  $\epsilon$  centered at  $\mathbf{x}$ . In fact,

$$|\mathbf{x} - \mathbf{y}| \leq 2^{-1/\alpha} L^{-1/\alpha} \varepsilon^{2/\alpha} \implies d(\mathbf{x}, \mathbf{y}) \leq \sqrt{2L(2L)^{-1} \varepsilon^2} = \varepsilon.$$

Now, for any  $\varepsilon > 0$ , let  $N_d(\varepsilon)$  be the minimum number of  $d$ -balls of radius  $\varepsilon$  that are required for covering  $\Omega$ . Similarly, let  $N(\varepsilon)$  be the covering number associated to the Euclidean metric. It is straightforward to see that (47) implies

$$1 \geq N_d(\varepsilon) \leq N(2^{-1/\alpha} L^{-1/\alpha} \varepsilon^{2/\alpha}) \leq \max \left\{ \left( \frac{C\varepsilon^2}{2L} \right)^{-d/\alpha}, 1 \right\}, \quad (48)$$

where  $C = C(d, \alpha, \Omega) > 0$  is an absolute constant. In fact,

$$N(\epsilon) \leq \max\{3^d \text{diam}(\Omega)^d \epsilon^{-d}, 1\},$$

where  $\text{diam}(\Omega)$  is the domain diameter under the Euclidean metric. In particular, we have  $N_d(\varepsilon) < +\infty$  for all  $\varepsilon > 0$ , meaning that  $\Omega$  is  $d$ -compact. Furthermore, (48) also implies that  $N_d(\varepsilon) = 1$  for all  $\varepsilon \geq \sqrt{C/2L}$ .

We now note that for all  $\mathbf{x}, \mathbf{y} \in \Omega$  one has

$$d(\mathbf{x}, \mathbf{y}) \leq \mathbb{E}^{1/2} |Z(\mathbf{x})|^2 + \mathbb{E}^{1/2} |Z(\mathbf{y})|^2 \leq 2\sigma.$$

In particular, the  $d$ -diameter of  $\Omega$  is bounded by  $2\sigma$ . Then, Theorem 1.3.3 in [2] implies

$$\mathbb{E} \left[ \sup_{\mathbf{x} \in \Omega} Z(\mathbf{x}) \right] \leq K \int_0^\sigma \log^{1/2} N_d(\varepsilon) d\varepsilon,$$

where  $K$  is a universal constant. Then, the above together with (48) yields

$$\mathbb{E} \left[ \sup_{\mathbf{x} \in \Omega} Z(\mathbf{x}) \right] \leq K \int_0^{\min\{\sigma, \sqrt{C/2L}\}} \sqrt{-\frac{d}{\alpha} \log \left( \frac{C}{2L} \varepsilon^2 \right)} d\varepsilon.$$



By operating the change of variables  $\epsilon := \sqrt{C/2L}\varepsilon$ , we may rewrite the previous as

$$\mathbb{E} \left[ \sup_{\mathbf{x} \in \Omega} Z(\mathbf{x}) \right] \leq K \sqrt{\frac{4Ld}{C'\alpha}} \int_0^{\min\{\sqrt{C/2L}\sigma, 1\}} \sqrt{\log(1/\epsilon)} d\epsilon.$$

At this point, it is useful to note that for any  $0 < a < 1$  one has

$$\int_0^a \sqrt{\log \epsilon} d\epsilon = a\sqrt{\log(1/a)} + \frac{\sqrt{\pi}}{2} (1 - \operatorname{erf}(\sqrt{\log(1/a)})) \leq a \left( \frac{\sqrt{\pi}}{2} + \sqrt{\log(1/a)} \right)$$

where  $\operatorname{erf}$  is the error function, which is known to satisfy  $1 - \operatorname{erf}(z) \leq e^{-z^2}$ . It follows immediately that

$$\begin{aligned} \mathbb{E} \left[ \sup_{\mathbf{x} \in \Omega} Z(\mathbf{x}) \right] &\leq K \sqrt{\frac{2d}{\alpha}} \sigma \left( \frac{\sqrt{\pi}}{2} + \sqrt{-\log \min \left\{ \sqrt{\frac{C}{2L}} \sigma, 1 \right\}} \right) \leq \\ &\leq C' \sigma \left( 1 + \sqrt{\log^+(1/\sigma)} \right), \end{aligned} \quad (49)$$

where  $C' = C'(d, \alpha, \Omega, L)$  grows less than logarithmically in  $L$ .

The first consequence of this fact, is that the paths of  $Z$  are almost-surely uniformly continuous (c.f. Theorem 1.3.5 in [2]). Additionally, since the Gaussian process  $\tilde{Z} := -Z$  satisfies the same upper bound, it is straightforward to conclude that

$$\mathbb{E} \|Z\|_{L^\infty(\Omega)} \leq \tilde{C} \sigma \left( 1 + \sqrt{\log^+(1/\sigma)} \right),$$

where  $\tilde{C} := 2C'$ . We now recall the celebrated Borell-TIS inequality, see Theorem 2.1.1 in [2],

$$\mathbb{P}(\|Z\|_{L^\infty(\Omega)} - \mathbb{E}\|Z\|_{L^\infty(\Omega)} > z) \leq \exp\left(-\frac{z^2}{2\sigma^2}\right),$$

from which it is straightforward to prove that  $\operatorname{Var} \|Z\|_{L^\infty(\Omega)} \leq 4\sigma^2$ . Then,

$$\mathbb{E}^{1/2} \|Z\|_{L^\infty(\Omega)}^2 \leq \sqrt{\operatorname{Var} \|Z\|_{L^\infty(\Omega)}} + \mathbb{E}\|Z\|_{L^\infty(\Omega)} \leq c_1 \sigma \left( 1 + \sqrt{\log^+(1/\sigma)} \right),$$

as claimed. Finally, the statement concerning  $\mathbb{E}[e^{\beta\|Z\|_{L^\infty(\Omega)}}]$  follows directly from Theorem 2.1.2 in [2]: see also (58) and the reasoning explained thereby.  $\square$

**Lemma 3.** *Let  $\Omega \subset \mathbb{R}^d$  be a compact subset and let  $Z$  be a mean zero Gaussian random field defined over  $\Omega$ . Assume that the covariance kernel of  $Z$ ,  $\operatorname{Cov}$ , is Lipschitz continuous. Then, there exists a nonincreasing summable sequence  $\lambda_1 \geq \lambda_2 \geq \dots \geq 0$  and a sequence of Lipschitz continuous maps,  $\{\varphi_i\}_{i=1}^{+\infty}$ , forming an orthonormal basis of  $L^2(\Omega)$ , such that*

$$\operatorname{Cov}(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^{+\infty} \lambda_i \varphi_i(\mathbf{x}) \varphi_i(\mathbf{y}) \quad (50)$$

for all  $\mathbf{x}, \mathbf{y} \in \Omega$ . Furthermore, there exists a sequence of independent standard normal random variables,  $\{\eta_i\}_{i=1}^{+\infty}$ , such that

$$Z = \sum_{i=1}^{+\infty} \sqrt{\lambda_i} \eta_i \varphi_i \quad (51)$$

almost surely. Finally, the truncated kernels,

$$\operatorname{Cov}_{p,q}(\mathbf{x}, \mathbf{y}) := \sum_{i=p}^q \lambda_i \varphi_i(\mathbf{x}) \varphi_i(\mathbf{y}),$$

defined for varying  $1 \leq p \leq q \leq +\infty$ ,

i) converge uniformly as  $p, q \rightarrow +\infty$ ;

ii) are all  $1/2$ -Hölder continuous, with a common Hölder constant.

*Proof.* Without loss of generality, we shall assume that  $\lambda_i > 0$  for all  $i \in \mathbb{N}$ . The series expansion in (50), and the uniform convergence claimed in (i), are a consequence of Mercer's Theorem [44]. There, the basis functions  $\varphi_i$  are obtained by solving the eigenvalue problem below

$$\lambda_i \varphi_i(\mathbf{x}) = \int_{\Omega} \text{Cov}(\mathbf{x}, \mathbf{y}) \varphi_i(\mathbf{y}) d\mathbf{y}, \quad (52)$$

where the equality holds for almost every  $\mathbf{x} \in \Omega$ , and, similarly, (50) is shown to hold almost everywhere. However, the right-hand-side of (52) is easily shown to be Lipschitz continuous in  $\mathbf{x}$ , since

$$\left| \int_{\Omega} \text{Cov}(\mathbf{x}, \mathbf{y}) \varphi_i(\mathbf{y}) d\mathbf{y} - \int_{\Omega} \text{Cov}(\mathbf{x}', \mathbf{y}) \varphi_i(\mathbf{y}) d\mathbf{y} \right| \leq L |\mathbf{x} - \mathbf{x}'| \int_{\Omega} |\varphi_i(\mathbf{y})| d\mathbf{y} \leq L |\mathbf{x} - \mathbf{x}'| |\Omega|^{1/2}, \quad (53)$$

where  $L$  is the Lipschitz constant of  $\text{Cov}$ . Thus, without loss of generality, we may pick the  $\varphi_i$  to be Lipschitz continuous. Since the series in (50) converges uniformly, both the left-hand-side and the right-hand-side of (50) are continuous: as they coincide a.e. in  $\Omega$ , they must be equal everywhere. To conclude, we shall now prove (ii), as (51) is just the well-known statement of the Kosambi-Karhunen-Loeve Theorem. Pick any  $0 \leq p \leq q \leq +\infty$ . For  $\mathbf{x}, \mathbf{x}', \mathbf{y}, \mathbf{y}' \in \Omega$  we have

$$\begin{aligned} \left| \sum_{i=p}^q \lambda_i \varphi_i(\mathbf{x}) \varphi_i(\mathbf{y}) - \sum_{i=p}^q \lambda_i \varphi_i(\mathbf{x}') \varphi_i(\mathbf{y}') \right| &\leq \sum_{i=p}^q \lambda_i |\varphi_i(\mathbf{x}) \varphi_i(\mathbf{y}) - \varphi_i(\mathbf{x}') \varphi_i(\mathbf{y}')| \leq \\ &\leq \sum_{i=p}^q \lambda_i |\varphi_i(\mathbf{x}) - \varphi_i(\mathbf{x}')| |\varphi_i(\mathbf{y})| + \sum_{i=p}^q \lambda_i |\varphi_i(\mathbf{y}) - \varphi_i(\mathbf{y}')| |\varphi_i(\mathbf{x}')|. \end{aligned} \quad (54)$$

Applying the Cauchy-Schwarz inequality, and using  $\lambda_i = \sqrt{\lambda_i} \sqrt{\lambda_i}$ , allows us to continue inequality (54) as

$$\begin{aligned} \cdots &\leq \sqrt{\sum_{i=p}^q \lambda_i |\varphi_i(\mathbf{x}) - \varphi_i(\mathbf{x}')|^2} \sqrt{\sum_{i=p}^q \lambda_i \varphi_i(\mathbf{y})^2} \\ &\quad + \sqrt{\sum_{i=p}^q \lambda_i |\varphi_i(\mathbf{y}) - \varphi_i(\mathbf{y}')|^2} \sqrt{\sum_{i=p}^q \lambda_i \varphi_i(\mathbf{x}')^2}. \end{aligned} \quad (55)$$

Since all the sums involved concern positive values, we may further bound the above by letting  $p = 0$  and  $q = +\infty$ . Then, thanks to (50), we get

$$\begin{aligned} \left| \sum_{i=p}^q \lambda_i \varphi_i(\mathbf{x}) \varphi_i(\mathbf{y}) - \sum_{i=p}^q \lambda_i \varphi_i(\mathbf{x}') \varphi_i(\mathbf{y}') \right| &\leq \\ &\leq \sqrt{\text{Cov}(\mathbf{x}, \mathbf{x}) - 2\text{Cov}(\mathbf{x}, \mathbf{x}') + \text{Cov}(\mathbf{x}', \mathbf{x}')} \sqrt{\text{Cov}(\mathbf{y}, \mathbf{y})} \\ &\quad + \sqrt{\text{Cov}(\mathbf{y}, \mathbf{y}) - 2\text{Cov}(\mathbf{y}, \mathbf{y}') + \text{Cov}(\mathbf{y}', \mathbf{y}')} \sqrt{\text{Cov}(\mathbf{x}', \mathbf{x}')} \leq \\ &\leq \sqrt{2LM|\mathbf{x} - \mathbf{x}'|} + \sqrt{2LM|\mathbf{y} - \mathbf{y}'|}, \end{aligned} \quad (56)$$

where  $M := \max_{\mathbf{x} \in \Omega} \text{Cov}(\mathbf{x}, \mathbf{x})$ . This shows that the truncated kernel is  $1/2$ -Hölder continuous with Hölder coefficient bounded by  $\sqrt{2LM}$ . As the latter is independent on both  $p$  and  $q$ , this concludes the proof.  $\square$

**Lemma 4.** Let  $\Omega \subset \mathbb{R}^d$  be a compact subset and let  $Z$  be a mean zero Gaussian random field defined over  $\Omega$ . Assume that the covariance kernel of  $Z$ ,  $\text{Cov}$ , is square-integrable over  $\Omega \times \Omega$ . Then, there exists a nonincreasing summable sequence  $\lambda_1 \geq \lambda_2 \geq \cdots \geq 0$  and an orthonormal basis of  $L^2(\Omega)$ ,  $\{\varphi_i\}_{i=1}^{+\infty}$ , such that

$$\text{Cov}(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^{+\infty} \lambda_i \varphi_i(\mathbf{x}) \varphi_i(\mathbf{y})$$

for almost every  $(\mathbf{x}, \mathbf{y}) \in \Omega \times \Omega$ . Furthermore, there exists a sequence of independent standard normal random variables,  $\{\eta_i\}_{i=1}^{+\infty}$ , such that

$$Z = \sum_{i=1}^{+\infty} \sqrt{\lambda_i} \eta_i \varphi_i$$

almost surely. Finally, the  $L^2$ -norm of the process is exponentially integrable, i.e.

$$\mathbb{E} \left[ e^{\beta \|Z\|_{L^2(\Omega)}} \right] < +\infty \quad (57)$$

for all  $\beta > 0$ .

*Proof.* We only need to prove (57), as the rest of the Lemma simply follows from the Kosambi-Karhunen-Loeve Theorem. To this end, we first note that for any random variable  $X$  one has

$$\exists \epsilon > 0 \text{ such that } \mathbb{E} \left[ e^{\epsilon X^2} \right] < +\infty \implies \mathbb{E} \left[ e^{\beta |X|} \right] < +\infty \quad \forall \beta > 0. \quad (58)$$

In fact, given any two positive numbers  $\epsilon$  and  $\beta$ , there exists some  $M > 0$  such that  $M + \epsilon x^2 > \beta |x|$  for all  $x \in \mathbb{R}$ . In light of this, let  $\epsilon > 0$  be a parameter, whose value shall be chosen later on. By orthonormality, we have

$$\|Z\|_{L^2(\Omega)}^2 = \sum_{i=1}^{+\infty} \lambda_i \eta_i^2.$$

As the  $\eta_i$ 's are independent, it follows that

$$\mathbb{E} \left[ e^{\epsilon \|Z\|_{L^2(\Omega)}^2} \right] = \mathbb{E} \left[ e^{\epsilon \sum_{i=1}^{+\infty} \lambda_i \eta_i^2} \right] = \prod_{i=1}^{+\infty} \mathbb{E} \left[ e^{\epsilon \lambda_i \eta_i^2} \right]. \quad (59)$$

For each index  $i$ , if  $\epsilon \lambda_i - 1/2 < 0$ , we have

$$\mathbb{E} \left[ e^{\epsilon \lambda_i \eta_i^2} \right] = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{+\infty} e^{\epsilon \lambda_i z^2} e^{-z^2/2} dz = \sqrt{\frac{1}{1 - 2\epsilon \lambda_i}}.$$

We thus choose  $\epsilon < 1/2\lambda_1$ , so that, by monotonicity, the above holds for all  $i$ . Resuming (59), we get

$$\mathbb{E} \left[ e^{\epsilon \|Z\|_{L^2(\Omega)}^2} \right] = \prod_{i=1}^{+\infty} \sqrt{\frac{1}{1 - 2\epsilon \lambda_i}} = \exp \left( -\frac{1}{2} \sum_{i=1}^{+\infty} \log(1 - 2\epsilon \lambda_i) \right).$$

Since, for  $i \rightarrow +\infty$ ,  $-\frac{1}{2} \log(1 - 2\epsilon \lambda_i)$  is asymptotic to  $\epsilon \lambda_i$ , which is a summable sequence, the conclusion now follows by (58).  $\square$

**Lemma 5.** Let  $(V, \|\cdot\|)$  be a separable Hilbert space and let  $u$  be a squared integrable  $V$ -valued random variable,  $\mathbb{E}\|u\|^2 < +\infty$ . Then, there exists an orthonormal basis  $\{v_i\}_{i=1}^{+\infty} \subset V$ , a sequence of (scalar) random variables  $\{\omega_i\}_{i=1}^{+\infty}$ , with  $\mathbb{E}[\omega_i \omega_j] = \delta_{i,j}$ , and a nonincreasing summable sequence  $\lambda_1 \geq \lambda_2 \geq \dots \geq 0$  such that

$$u = \sum_{i=1}^{+\infty} \sqrt{\lambda_i} \omega_i v_i$$

almost-surely.

*Proof.* Consider the operator  $\mathcal{C} : V \rightarrow V$  defined as  $\mathcal{C}(v) = \mathbb{E}[\langle u, v \rangle u]$ , where the expectation is intended in the Bochner sense. We show that  $\mathcal{C}$  is a symmetric semi-positive definite trace class operator. Indeed, for any  $v, v' \in V$  one has

$$\langle \mathcal{C}(v), v' \rangle = \langle \mathbb{E}[\langle u, v \rangle u], v' \rangle = \mathbb{E}[\langle u, v \rangle \langle u, v' \rangle] = \langle v, \mathcal{C}(v') \rangle,$$

and

$$\langle \mathcal{C}(v), v \rangle = \mathbb{E}[\langle u, v \rangle^2] \geq 0.$$

Furthermore, given any orthonormal basis  $\{e_i\}_{i=1}^{+\infty}$  we have

$$\sum_{i=1}^{+\infty} \langle \mathcal{C}(e_i), e_i \rangle = \sum_{i=1}^{+\infty} \mathbb{E}[\langle u, e_i \rangle^2] = \mathbb{E} \left[ \sum_{i=1}^{+\infty} \langle u, e_i \rangle^2 \right] = \mathbb{E} \|u\|^2 < +\infty,$$

by monotone convergence. Thus, by the well-known Spectral Theorem, there exists an orthonormal basis  $\{v_i\}_{i=1}^{+\infty} \subset V$  and a nonincreasing summable sequence  $\lambda_1 \geq \lambda_2 \geq \dots \geq 0$  such that

$$\mathcal{C}(v) = \sum_{i=1}^{+\infty} \lambda_i \langle v, v_i \rangle v_i \quad \forall v \in V.$$

Furthermore, as directly implied by the above,  $\mathcal{C}(v_i) = \lambda_i v_i$ , meaning that the  $\lambda_i$ 's and the  $v_i$ 's are the eigenvalues and eigenvectors of the (uncentered) covariance operator  $\mathcal{C}$ , respectively. Let now

$$\omega_i := \frac{1}{\sqrt{\lambda_i}} \langle u, v_i \rangle.$$

It is straightforward to see that for all  $i, j \in \mathbb{N}$  we have

$$\mathbb{E}[\omega_i \omega_j] = \frac{1}{\sqrt{\lambda_i \lambda_j}} \mathbb{E}[\langle u, v_i \rangle \langle u, v_j \rangle] = \frac{1}{\sqrt{\lambda_i \lambda_j}} \langle \mathcal{C}(v_i), v_j \rangle = \frac{\lambda_i}{\sqrt{\lambda_i \lambda_j}} \langle v_i, v_j \rangle = \delta_{i,j}.$$

Finally,

$$u = \sum_{i=1}^{+\infty} \sqrt{\lambda_i} \omega_i v_i \tag{60}$$

by definition of the  $\omega_i$ 's. To this end, we also note that, since

$$\sum_{i=1}^{+\infty} \mathbb{E} \left\| \sqrt{\lambda_i} \omega_i v_i \right\|^2 = \mathbb{E} \left[ \sum_{i=1}^{+\infty} \left\| \sqrt{\lambda_i} \omega_i v_i \right\|^2 \right] = \mathbb{E} \|u\|^2 < +\infty,$$

the series in (60) is  $L^2(\mathbb{P}; V)$  convergent (in the Bochner sense [18]), where  $\mathbb{P}$  is the probability law of  $u$ . Thus, (60) holds  $\mathbb{P}$ -almost surely.  $\square$

## Appendix B Complementary results for Section 4

**Lemma B1.** *Let  $(X, \|\cdot\|)$  be a normed space. If  $A \subseteq X$  is dense in  $X$ , then:*

- i)  $A \cap O$  is dense in  $O$  for all open sets  $O \subseteq X$ ;
- ii)  $A \cap C$  is dense in  $C$  for all convex closed sets  $C \subseteq X$ .

*Proof.*

- i) Let  $O \subseteq X$  be open and let  $U$  be an open subset in the subspace topology of  $O$ . Then,  $U = \tilde{U} \cap O$  for some  $\tilde{U} \subseteq X$  open in  $X$ . In particular,  $U$  is also open in the topology of the larger space,  $X$ . Thus,  $U \cap A \neq \emptyset$ , and the conclusion follows.
- ii) Let  $C \subseteq X$  be convex and closed. Fix any  $c_0$  in  $\text{int}(C)$ , the interior of  $C$ . Let  $c \in C$ . It is well known that, under these hypothesis, the segment  $\{(1-t)c_0 + tc\}_{t \in [0,1]}$  can only, at most, intersect  $\partial C$  at  $c$ , as all the remaining points lie in the interior of the set. Let  $c_n := c_0/n + (1-1/n)c$ , so that  $c_n$  is a sequence of interior points converging to  $c$ . Since  $A \cap \text{int}(C)$  is dense in  $\text{int}(C)$ , see (i), for every  $n$  there exists  $c_n^* \in A \cap \text{int}(C) \in A \cap C$  such that  $|c_n - c_n^*| \leq 1/n$ . Then,  $c_n^* \rightarrow c$  as  $n \rightarrow +\infty$ , as wished.

$\square$

**Lemma B2.** Let  $\sigma$  be a finite measure over  $\mathbb{R}^N$ , and let  $B \subset \mathbb{R}^N$  be a bounded set. Let  $\rho: \mathbb{R} \rightarrow \mathbb{R}$  satisfy the assumptions in Lemma 1. Consider the following functional space

$$\mathcal{V} := \left\{ f \in \mathcal{M}(\mathbb{R}^N, \mathbb{R}^n) \text{ s.t. } f|_B \in \mathcal{C}(B), \|f\|_{\mathcal{V}} := \|f\|_{\mathcal{C}(B)} + \int_{\mathbb{R}^N \setminus B} |f(\mathbf{x})| \sigma(d\mathbf{x}) < +\infty \right\}.$$

Then,  $\mathcal{N}_\rho(\mathbb{R}^N, \mathbb{R}^n)$  is  $\|\cdot\|_{\mathcal{V}}$ -dense in  $\mathcal{V}$ .

*Proof.* For the sake of simplicity, we only prove the case  $n = 1$ . We note that  $\mathcal{V} \cong \mathcal{C}(B) \times L_\sigma^1(\mathbb{R}^N \setminus B)$  in the natural way. As a consequence, the dual space of  $\mathcal{V}$  can be given as

$$\mathcal{V}' \cong \mathcal{R}(B) \times L_\sigma^\infty(\mathbb{R}^N \setminus B)$$

where  $\mathcal{R}(B)$  is the set of (signed) Radon measures over  $B$ , considered with the total variation norm. In particular, for every  $F \in \mathcal{V}'$  there exist  $\nu \in \mathcal{R}(B)$  and  $g \in L_\sigma^\infty(\mathbb{R}^N \setminus B)$  such that

$$F(f) = \int_B f d\nu + \int_{\mathbb{R}^N \setminus B} f g d\sigma \quad (61)$$

for all  $f \in \mathcal{V}$ . Assume that  $\mathcal{N}_\rho(\mathbb{R}^N, \mathbb{R}^n)$  is not dense in  $\mathcal{V}$ . Then, there exists some  $F \in \mathcal{V}' \setminus \{0\}$  such that  $F \equiv 0$  over  $\mathcal{N}_\rho(\mathbb{R}^N, \mathbb{R}^n) \subset \mathcal{V}$ . Let  $\nu$  and  $g$  be as in (61), and define  $\hat{\nu} := \nu + g d\sigma$ . Then, we have  $\hat{\nu} \in \mathcal{R}(\mathbb{R}^N)$  and

$$\int_{\mathbb{R}^N} f d\hat{\nu} = 0 \quad \forall f \in \mathcal{N}_\rho(\mathbb{R}^N, \mathbb{R}^n).$$

However, thanks to our assumptions on  $\rho$ , the above implies  $\hat{\nu} \equiv 0 \implies F \equiv 0$  (cf. Theorem 5 in [33]), thus yielding a contraddiction.  $\square$

**Lemma B3.** Let  $\sigma$  be a finite measure over  $\mathbb{R}^n$  and let  $Q := [-M, M]^n$  be a given hypercube,  $M > 0$ . Let  $\rho: \mathbb{R} \rightarrow \mathbb{R}$  be the  $\alpha$ -leaky ReLU activation,  $|\alpha| < 1$ . Consider the following functional space

$$\mathcal{W} := \left\{ f \in \mathcal{M}(\mathbb{R}^n, \mathbb{R}^N) \text{ s.t. } \|f\|_{\mathcal{W}} := \|f\|_{W^{1,\infty}(Q)} + \int_{\mathbb{R}^n \setminus Q} |f(\mathbf{x})| \sigma(d\mathbf{x}) < +\infty \right\}.$$

Then,  $\mathcal{N}_\rho(\mathbb{R}^n, \mathbb{R}^N)$  is  $\|\cdot\|_{\mathcal{W}}$ -dense in  $\mathcal{W}$ .

*Proof.* We recall that, since  $\rho$  is the  $\alpha$ -leaky ReLU activation, the set  $\mathcal{N}_\rho(\mathbb{R}^n, \mathbb{R}^N)$  contains all those functions  $f: \mathbb{R}^n \rightarrow \mathbb{R}^N$  that are piecewise linear over polyhedra.

Following the same idea as in the proof of Lemma B2, let us assume that  $\mathcal{N}_\rho(\mathbb{R}^n, \mathbb{R}^N)$  is not dense in  $\mathcal{W}$ . Then, there exists some nontrivial functional  $F \in \mathcal{W}'$  that vanishes over  $\mathcal{N}_\rho(\mathbb{R}^n, \mathbb{R}^N)$ . Since,  $\mathcal{W} \cong W^{1,+\infty}(Q) \times L_\sigma^1(\mathbb{R}^n \setminus Q)$  in the natural way, we have

$$F(f) = F_1(f|_Q) + F_2(f|_{\mathbb{R}^n \setminus Q})$$

for some  $F_1 \in W^{1,+\infty}(Q)'$  and  $F_2 \in L_\sigma^1(\mathbb{R}^n \setminus Q)'$ . Let now  $g \in W^{1,+\infty}(Q)$ . Since  $Q$  is a polyhedron, it is straightforward to see that for every  $\varepsilon > 0$  there exists  $\phi_\varepsilon \in \mathcal{N}_\rho(\mathbb{R}^n, \mathbb{R}^N)$  such that

$$\|\phi_\varepsilon - g\|_{W^{1,\infty}(Q)} < \varepsilon, \quad \|\phi_\varepsilon\|_{\mathcal{C}(\mathbb{R}^n)} \leq \|g\|_{\mathcal{C}(Q)}, \quad \phi_\varepsilon \equiv 0 \text{ on } \mathbb{R}^n \setminus Q_\varepsilon,$$

where  $Q_\varepsilon := (-M - \varepsilon, M + \varepsilon)^n$ . In fact, such a piecewise linear approximation is easily constructed and it is guaranteed to be a member of  $\mathcal{N}_\rho(\mathbb{R}^n, \mathbb{R}^N)$ . Then,  $\phi_\varepsilon \rightarrow g \cdot \mathbb{1}_Q$  in  $\|\cdot\|_{\mathcal{W}}$ -norm as  $\varepsilon \rightarrow 0$ . Thus, we have

$$F_1(g) = F_1(g) + F_2(0) = \lim_{\varepsilon \rightarrow 0} F(\phi_\varepsilon) = 0,$$

proving that  $F_1 \equiv 0$ . In particular, for  $F$  to vanish over  $\mathcal{N}_\rho(\mathbb{R}^n, \mathbb{R}^N)$  we must have

$$F_2(\phi|_{\mathbb{R}^n \setminus Q}) = 0 \quad \forall \phi \in \mathcal{N}_\rho(\mathbb{R}^n, \mathbb{R}^N).$$

However, since  $\rho$  satisfies the assumptions in Lemma 1, this would imply  $F_2 \equiv 0$ , ultimately yielding a contraddiction.  $\square$

**Corollary B1.** Let  $\sigma$  be a probability measure over  $\mathbb{R}^N$ , with finite moment and absolutely continuous with respect to the Lebesgue measure. Let  $B \subset \mathbb{R}^N$  be a bounded set, and let  $M, L > 0$ ,  $n \in \mathbb{N}$ . Consider the families

$$\mathcal{E}_{B,M}(\mathbb{R}^N, \mathbb{R}^n) := \left\{ \Psi' \in \mathcal{M}(\mathbb{R}^N, \mathbb{R}^n) \text{ s.t. } \int_{\mathbb{R}^N} |\Psi'(\mathbf{v})| \sigma(d\mathbf{v}) < +\infty, \sup_{\mathbf{v} \in B} |\Psi'(\mathbf{v})| \leq M \right\},$$

$$\mathcal{D}_{M,L}(\mathbb{R}^n, V) := \left\{ \Psi \in \mathcal{M}(\mathbb{R}^n, \mathbb{R}^N) \text{ s.t. } \sup_{\mathbf{c} \in [-M,M]^n} \|\partial \Psi\|_V(\mathbf{c}) \leq L \right\}.$$

Let  $\rho$  be the  $\alpha$ -leakyReLU, with  $|\alpha| < 1$ . Then, for every  $\varepsilon > 0$  and every pair  $\Psi' \in \mathcal{E}_{B,M}(\mathbb{R}^N, \mathbb{R}^n)$ ,  $\Psi \in \mathcal{D}_{M,L}(\mathbb{R}^n, \mathbb{R}^N)$  such that

$$\int_{\mathbb{R}^N} |\mathbf{u} - \Psi(\Psi'(\mathbf{u}))| \sigma(d\mathbf{u}) < +\infty, \quad (62)$$

there exists  $\hat{\Psi}' \in \mathcal{E}_{B,M}(\mathbb{R}^N, \mathbb{R}^n) \cap \mathcal{N}_\rho(\mathbb{R}^N, \mathbb{R}^n)$  and  $\hat{\Psi} \in \mathcal{D}_{M,L}(\mathbb{R}^n, \mathbb{R}^N) \cap \mathcal{N}_\rho(\mathbb{R}^n, \mathbb{R}^N)$  such that

$$\left| \int_{\mathbb{R}^N} |\mathbf{u} - \Psi(\Psi'(\mathbf{u}))| \sigma(d\mathbf{u}) - \int_{\mathbb{R}^N} |\mathbf{u} - \hat{\Psi}(\hat{\Psi}'(\mathbf{u}))| \sigma(d\mathbf{u}) \right| < \varepsilon,$$

*Proof.* Let  $\Psi' \in \mathcal{E}_{B,M}(\mathbb{R}^N, \mathbb{R}^n)$  and  $\Psi \in \mathcal{D}_{M,L}(\mathbb{R}^n, \mathbb{R}^N)$  satisfy (62), and let  $\varepsilon > 0$ . Define  $\nu$  as the push-forward measure of  $\sigma$  through the encoder  $\Psi'$ , and let  $(\mathcal{W}, \|\cdot\|_{\mathcal{W}})$  be the normed space in Lemma B3, defined with respect to  $\nu$  and  $M$ . We note that, since  $\sigma$  has finite moment, Eq. (62) implies

$$\int_{\mathbb{R}^N} |\Psi(\mathbf{c})| \nu(d\mathbf{c}) < +\infty \implies \Psi \in L_\nu^1(\mathbb{R}^n).$$

In particular,  $\Psi \in \mathcal{D}_{M,L}(\mathbb{R}^n, \mathbb{R}^N) \cap L_\nu^1(\mathbb{R}^n) \subseteq \mathcal{W}$ . Since the set  $\mathcal{D}_{M,L}(\mathbb{R}^n, \mathbb{R}^N) \cap L_\nu^1(\mathbb{R}^n)$  is both convex and closed in  $(\mathcal{W}, \|\cdot\|_{\mathcal{W}})$ , by Lemmas B1 and B3, there exists some  $\hat{\Psi} \in \mathcal{D}_{M,L}(\mathbb{R}^n, \mathbb{R}^N) \cap \mathcal{N}_\rho(\mathbb{R}^n, \mathbb{R}^N) \cap L_\nu^1(\mathbb{R}^n)$  such that  $\|\Psi - \hat{\Psi}\|_{\mathcal{W}} < \varepsilon/2$ . Let now  $\ell$  be the (global) Lipschitz constant of  $\hat{\Psi}$ . Following the same computations as in the proof of Theorem 1, we get

$$\begin{aligned} \left| \int_{\mathbb{R}^N} |\mathbf{u} - \Psi(\Psi'(\mathbf{u}))| \sigma(d\mathbf{u}) - \int_{\mathbb{R}^N} |\mathbf{u} - \hat{\Psi}(\hat{\Psi}'(\mathbf{u}))| \sigma(d\mathbf{u}) \right| &\leq \\ &\leq \int_{\mathbb{R}^n} |\Psi(\mathbf{c}) - \hat{\Psi}(\mathbf{c})| \nu(d\mathbf{c}) + \ell \int_{\mathbb{R}^N} |\Psi'(\mathbf{v}) - \hat{\Psi}'(\mathbf{v})| \sigma(d\mathbf{v}) \leq \\ &\leq \|\Psi(\mathbf{c}) - \hat{\Psi}(\mathbf{c})\|_{\mathcal{W}} + \ell \int_{\mathbb{R}^N} |\Psi'(\mathbf{v}) - \hat{\Psi}'(\mathbf{v})| \sigma(d\mathbf{v}) \leq \\ &\leq \frac{\varepsilon}{2} + \ell \int_{\mathbb{R}^N} |\Psi'(\mathbf{v}) - \hat{\Psi}'(\mathbf{v})| \sigma(d\mathbf{v}), \quad (63) \end{aligned}$$

for all  $\hat{\Psi}' \in \mathcal{N}_\rho(\mathbb{R}^N, \mathbb{R}^n)$ , where we exploited the fact that  $\nu([-M, M]^n) \leq 1$  due to  $\nu$  being a probability measure.

To conclude the proof, we now wish to bound the second term in (63). This, however, requires some additional care: in fact, we cannot directly repeat the same ideas and apply Lemma B2, as  $\Psi'$  might be discontinuous over  $B$ . To account for this, we shall introduce a proper smoothing step. Let

$$\varphi(\mathbf{x}) := C \exp^{-1/(1-|\mathbf{x}|^2)} \mathbb{1}_{\{|\mathbf{y}| \leq 1\}}(\mathbf{x})$$

be the canonical mollifier in  $\mathbb{R}^N$ , where  $C > 0$  ensures that  $\int_{\mathbb{R}^N} \varphi(\mathbf{x}) d\mathbf{x} = 1$ . For any  $\delta > 0$ , let  $\Psi'_\delta : \mathbb{R}^N \rightarrow \mathbb{R}^n$  be defined as

$$\Psi'_\delta(\mathbf{v}) := \Psi'(\mathbf{v}) \mathbb{1}_{\mathbb{R}^N \setminus B}(\mathbf{v}) + (\Psi' * \varphi_\delta) \mathbb{1}_B(\mathbf{v}),$$

where  $\varphi_\delta(\mathbf{x}) := \delta^{-N} \varphi(\mathbf{x}/\delta)$  and  $*$  denotes the convolution operator. Then, following classical arguments (see, e.g. Theorem 7, Appendix C.2, in [18]), it is straightforward to see that:

- (i)  $\Psi'_\delta$  is continuous over  $B$ ;

(ii)  $\Psi'_\delta \in \mathcal{E}_{B,M}(\mathbb{R}^N, \mathbb{R}^n)$ , as  $|\Psi'_\delta(\mathbf{v})| \leq M \cdot \int_{\mathbb{R}^N} \varphi_\delta(\mathbf{v}) \sigma(d\mathbf{v}) = M$  for all  $\mathbf{v} \in B$ ;

(iii)  $\Psi'_\delta \rightarrow \Psi'$   $\sigma$ -almost everywhere, as  $\sigma$  is absolutely continuous with respect to Lebesgue's measure;

(iv)  $\int_{\mathbb{R}^N} |\Psi'_\delta(\mathbf{v}) - \Psi'(\mathbf{v})| \sigma(d\mathbf{v}) \rightarrow 0$  as  $\delta \rightarrow 0$ , due to dominated convergence and (iii).

In light of (iv), let us fix  $\delta > 0$  such that  $\int_{\mathbb{R}^N} |\Psi'_\delta(\mathbf{v}) - \Psi'(\mathbf{v})| \sigma(d\mathbf{v}) \leq \varepsilon/4\ell$ . As before, let  $(\mathcal{V}, \|\cdot\|_{\mathcal{V}})$  be the normed space in Lemma B2, defined according to the measure  $\sigma$ . Then, due to (i) and (ii), we have  $\Psi'_\delta \in \mathcal{E}_{B,M}(\mathbb{R}^N, \mathbb{R}^n) \cap \mathcal{V}$ , with the latter being a convex closed subset of  $(\mathcal{V}, \|\cdot\|_{\mathcal{V}})$ . In particular, by Lemmas B1 and B2, there exists some

$$\hat{\Psi}'_\delta \in \mathcal{N}_\rho(\mathbb{R}^N, \mathbb{R}^n) \cap \mathcal{E}_{B,M}(\mathbb{R}^N, \mathbb{R}^n) \cap \mathcal{V}$$

such that  $\|\Psi'_\delta - \hat{\Psi}'_\delta\|_{\mathcal{V}} < \varepsilon/4\ell$ . Consequently,

$$\int_{\mathbb{R}^N} |\Psi'(\mathbf{v}) - \hat{\Psi}'_\delta(\mathbf{v})| \sigma(d\mathbf{v}) \leq \int_{\mathbb{R}^N} |\Psi'(\mathbf{v}) - \Psi'_\delta(\mathbf{v})| \sigma(d\mathbf{v}) + \|\Psi'_\delta - \hat{\Psi}'_\delta\|_{\mathcal{V}} < \frac{\varepsilon}{4\ell} \quad (64)$$

where, as before, we exploited the fact that  $\sigma$  is a probability measure, and thus  $\sigma(B) \leq 1$ . Then, setting  $\hat{\Psi} := \hat{\Psi}'_\delta$  and plugging (64) into (63) yields the desired conclusion.  $\square$

## Appendix C Measurable selections and more

In what follows, we use the term *Polish space* to intend a complete separable metric space. We recall, in particular, that all separable Banach spaces are Polish spaces. Finally, given any set  $X$ , we shall write  $2^X$  for its power set, that is, the collection of all subsets of  $X$ ,

$$2^X := \{A \mid A \subseteq X\}.$$

**Definition C1.** Let  $(Y, \mathcal{M})$  be a measurable space and let  $(X, d)$  be a Polish space. Let  $F : Y \rightarrow 2^X$ . We say that  $F$  is a measurable set-valued map if the following conditions hold:

a)  $F(y)$  is closed in  $X$  for all  $y \in Y$ ;

b) for all open sets  $A \subseteq X$  one has  $\mathcal{S}_A \in \mathcal{M}$ , where

$$\mathcal{S}_A := \{y \in Y \mid F(y) \cap A \neq \emptyset\}. \quad (65)$$

**Lemma C1.** Let  $(X, d)$  be a Polish space and let  $(Y, \tau)$  be a topological space equipped with a suitable  $\sigma$ -field  $\mathcal{M}$ . Let  $f : X \rightarrow Y$  be continuous. The following are equivalent:

i)  $f$  maps open sets onto  $\mathcal{M}$ -measurable sets;

ii) the map  $F : Y \rightarrow 2^X$ , defined as

$$F(y) := \begin{cases} \{x \in X \mid f(x) = y\} & y \in f(X) \\ X & \text{otherwise,} \end{cases} \quad (66)$$

is a measurable set-valued map.

*Proof.* Let  $F$  be as in (66). We first note that, since  $f$  is continuous, the preimage of any singleton,  $f^{-1}(\{y\})$ , is closed. In particular, condition (a) in Definition C1 is met. Next, we note that for any  $A \subseteq X$  one has

$$\mathcal{S}_A = f(A) \cup f(X)^c, \quad (67)$$

where  $B^c := Y \setminus B$  denotes the complement of  $B \subseteq Y$ , and  $\mathcal{S}_A$  is as in (65). In fact,

$$y \notin f(X) \implies F(y) = X \implies F(y) \cap A = A \neq \emptyset \implies y \in \mathcal{S}_A,$$

meaning that  $\mathcal{S}_A \cap f(X)^c = f(X)^c$ . On the other hand,

$$y \in f(X) \cap \mathcal{S}_A \iff \exists a \in A \text{ s.t. } a \in \{x \in X \mid f(x) = y\} \iff y \in f(A),$$

implying that  $\mathcal{S}_A \cap f(X) = f(A)$ . Since  $\mathcal{S}_A = (\mathcal{S}_A \cap f(X)^c) \cup (\mathcal{S}_A \cap f(X))$ , the identity in (67) easily follows.

At this point, it is straightforward to see that (i)  $\iff$  (ii). Assume, for instance, that (i) holds. Then, for any open set  $A \subseteq X$ ,  $f(A)$  is  $\mathcal{M}$ -measurable. In particular,  $f(X) \in \mathcal{M} \implies f(X)^c \in \mathcal{M} \implies \mathcal{S}_A \in \mathcal{M}$ , meaning that condition (b) in Definition C1 is met. Conversely, say that  $F$  is measurable, so that  $\mathcal{S}_A \in \mathcal{M}$  for all open sets  $A \subseteq X$ . Let  $A = \emptyset$ . Then,  $f(X)^c \in \mathcal{M}$ . Since  $f(A)$  and  $f(X)^c$  are disjoint,  $f(A) = \mathcal{S}_A \setminus f(X)^c$ , proving that  $f(A) \in \mathcal{M}$ , as claimed.  $\square$

**Lemma C2.** *Let  $f : (X, d_X) \rightarrow (Y, d_Y)$  be a continuous map between two Polish spaces. Let  $\mathcal{M}$  be the  $\mathbb{P}$ -completion of the Borel  $\sigma$ -field defined over  $Y$ , where  $\mathbb{P}$  is a given probability distribution. Then  $f$  maps open sets onto  $\mathcal{M}$ -measurable sets.*

*Proof.* This is a standard result in the theory of Analytic sets, see, e.g., Theorem 4.3.1 in [61]. In fact, any open set  $A \subseteq X$  is Borel measurable in  $X$ . Thus, by very definition,  $f(A)$  is an analytic set in  $Y$ . Since  $\mathcal{M}$  contains all Borel sets of  $Y$  and it is also  $\mathbb{P}$ -complete, it follows that  $f(A) \in \mathcal{M}$ .  $\square$

**Corollary C1.** *Let  $f : (X, d_X) \rightarrow (Y, d_Y)$  be a continuous map between two Polish spaces. Let  $\mathcal{M}$  be the  $\mathbb{P}$ -completion of the Borel  $\sigma$ -field defined over  $Y$ , where  $\mathbb{P}$  is a given probability distribution. Then,  $f$  admits an  $\mathcal{M}$ -measurable right-inverse, that is, a measurable map  $f^{-1} : Y \rightarrow X$  for which*

$$f(f^{-1}(y)) = y \quad \forall y \in f(X).$$

*Proof.* The nontrivial part of the statement lies in the measurability of  $f^{-1}$ , as the existence of a generic right-inverse is already guaranteed by the Axiom of Choice. Let  $F : Y \rightarrow 2^X$  be as in (66), so that, according to Lemmae C1 and C2,  $F$  is a measurable set-valued map. Then, a famous result by Kuratowski and Ryll-Nardzewski states that  $F$  admits a measurable selection (see, e.g., Theorem 8.1.3 and Definition 8.1.1 in [5]). That is, there exists a measurable map  $g : Y \rightarrow X$  such that  $g(y) \in F(y)$  for all  $y \in Y$ . It is straightforward to see that such map retains all the desired properties: in fact, for all  $y \in f(X)$  the condition  $g(y) \in F(y)$  implies  $f(g(y)) = y$ .  $\square$

**Corollary C2.** *Let  $f : (X, d_X) \rightarrow (Y, d_Y)$  be a continuous map between two Polish spaces. Let  $\mathcal{M}$  be the  $\mathbb{P}$ -completion of the Borel  $\sigma$ -field defined over  $Y$ , where  $\mathbb{P}$  is a given probability distribution. Let  $C \subseteq X$  be a closed subset. Then, there exists an  $\mathcal{M}$ -measurable map  $f^{-1} : Y \rightarrow X$  such that*

$$f(f^{-1}(y)) = y \quad \forall y \in f(X), \quad \text{and} \quad f^{-1}(f(C)) \subseteq C.$$

*Proof.* In agreement with Corollary C1, let  $g_0 : Y \rightarrow X$  be an  $\mathcal{M}$ -measurable right-inverse of  $f$ . Consider the metric subspace  $(C, d) \subseteq (X, d)$ . Since  $C$  is closed,  $(C, d)$  is a Polish space. Thus, we may exploit Corollary C1 once again to construct an  $\mathcal{M}$ -measurable map  $g_1 : Y \rightarrow C$  that operates as a right-inverse of  $f|_C$ , the restriction of  $f$  to  $C$ . Define  $g : Y \rightarrow X$  as

$$g(y) := g_1(y) \cdot \mathbb{1}_{f(C)}(y) + g_0(y) \cdot \mathbb{1}_{Y \setminus f(C)}(y).$$

Then  $f^{-1} := g$  fulfills all the requirements.  $\square$



**Lemma C3.** *Let  $(X, d_X)$  be a Polish space and let  $(C, d_C)$  be a compact metric space. Let  $J : X \times C \rightarrow \mathbb{R}$  be continuous. Then, there exists a Borel measurable map  $f : X \rightarrow C$  such that*

$$J(x, f(x)) = \min_{c \in C} J(x, c)$$

for all  $x \in X$ .

*Proof.* First of all, we note that the statement is well-defined as all the minima are attained by compactness of  $C$  and continuity of  $J$ . Let now  $F : X \rightarrow 2^C$  be the following set-valued map

$$F : x \rightarrow \left\{ c \in C \text{ such that } J(x, c) = \min_{c' \in C} J(x, c') \right\},$$

that assigns a (nonempty) subset of  $C$  to each  $x \in X$ . We aim at showing that  $F$  is a measurable set-valued map. To this end, we start by noting that  $F(x)$  is closed in  $C$  for all  $x \in X$ . To see this, fix any  $x \in X$  and let  $j_x := J(x, \cdot)$ , so that  $j_x : C \rightarrow \mathbb{R}$  is continuous. Then,

$$F(x) = j_x^{-1} \left( \left\{ \min_{c' \in C} j_x(c') \right\} \right)$$

is closed as it is the pre-image of a singleton under a continuous transformation.

Conversely, we now claim that, for any compact subset  $K \subseteq C$ , the set

$$\mathcal{S}_K := \{x \in X : F(x) \cap K \neq \emptyset\}$$

is closed. Indeed, let  $\{x_n\}_n \subseteq \mathcal{S}_K$  be a sequence converging to some  $x \in X$ . By definition of  $\mathcal{S}_K$ , for each  $x_n$  there exists a  $c_n \in K$  such that  $c_n \in F(x_n)$ , i.e. for which  $J(x_n, c_n) = \min_{c' \in C} J(x_n, c')$ . Since  $K$  is compact, up to passing to a subsequence, there exists some  $c \in K$  such that  $c_n \rightarrow c$ . Let now  $\tilde{c} \in C$  be a minimizer for  $x$ , i.e. a suitable element for which  $J(x, \tilde{c}) = \min_{c' \in C} J(x, c')$ . By continuity, we have

$$J(x, c) = \lim_{n \rightarrow +\infty} J(x_n, c_n) = \lim_{n \rightarrow +\infty} \min_{c' \in C} J(x_n, c') \leq \lim_{n \rightarrow +\infty} J(x_n, \tilde{c}) = J(x, \tilde{c}),$$

implying that  $c$  is also a minimizer for  $x$ . As a consequence, we have  $c \in K \cap F(x)$  and thus  $x \in \mathcal{S}_K$ . In particular,  $\mathcal{S}_K$  is closed. It is now straightforward to prove that  $\mathcal{S}_A$  is Borel measurable whenever  $A \subseteq C$  is open. In fact, any open set  $A \subseteq C$  can be written as the countable union of compact sets,  $A = \cup_{n \in \mathbb{N}} K_n$ , and clearly  $\mathcal{S}_A = \cup_n \mathcal{S}_{K_n}$ .

We have then proven that  $F$  is a measurable set-valued map. In particular, we may now invoke the measurable selection theorem by Kuratowski–Ryll–Nardzewski [5], which ensures the existence of a measurable map  $f : X \rightarrow C$  such that  $f(x) \in F(x)$ , i.e.  $J(x, f(x)) = \min_{c \in C} J(x, c)$ , as wished.  $\square$

## Appendix D Architectures and training details

We report in this Section the technical details concerning the design of the autoencoder modules (Tables 2-3) and the training of the DL-ROMs (Table 6). As for the latter, Tables 4-5 complete the picture, with a description of the architectures employed for reduced map,  $\phi$ .

We mention that, both for the case of Darcy flow in a porous medium and Burger’s equation, we exploited a combination of classical dense layers together with mesh-informed layers. The latter are a particular class of sparse architectures first introduced in [21] as a way to handle mesh-based functional data. In short, they are obtained through sparsification of dense architectures, which is achieved by means of a mesh-dependent pruning strategy: for further details, we refer the reader to [21]. As for our purposes, it is sufficient to know that mesh-informed layers are characterized by a *support* hyperparameter: the smaller the support, the sparser the architecture; for large supports, the module collapses to a classical dense layer. In all of the following, mesh-informed layers are always constructed by relying on uniform structured grids (with as many dofs as declared in the input-output dimensions).

| Layer # | Type  | Input dim. | Output dim. | Activation |
|---------|-------|------------|-------------|------------|
| 1       | Dense | 2601       | 500         | $\rho$     |
| 2       | Dense | 500        | $n$         | $\rho$     |
| 3       | Dense | $n$        | 500         | $\rho$     |
| 4       | Dense | 500        | 2601        | -          |

Table 2: Autoencoder architecture for Darcy’s law example, §5.1.1. The encoder and decoder modules, here presented together, are divided by a dashed line.  $n$  = latent dimension;  $\rho$  = 0.1-leakyReLU.

| Layer # | Type          | Input dim. | Output dim. | Support | Activation     |
|---------|---------------|------------|-------------|---------|----------------|
| 1       | Dense         | 500        | $n$         | -       | $\rho$         |
| 2       | Dense         | $n$        | 100         | -       | $\rho$         |
| 3       | Mesh-informed | 100        | 500         | 0.4     | $\rho$         |
| 4       | Mesh-informed | 500        | 500         | 0.1     | $\tilde{\rho}$ |

Table 3: Autoencoder architecture for Burger’s example, §5.1.2. Entries read as in Table 2; here,  $\tilde{\rho}(x) := \text{ReLU}(0.5 - \text{ReLU}(0.5 - x))$ .

| Layer # | Type          | Input dim. | Output dim. | Support | Activation |
|---------|---------------|------------|-------------|---------|------------|
| 1       | Mesh-informed | 2601       | 676         | 0.125   | tanh       |
| 2       | Mesh-informed | 676        | 169         | 0.25    | $\rho$     |
| 3       | Dense         | 169        | $n$         | -       | $\rho$     |

Table 4: Architecture of the reduced map network,  $\phi$ , for Darcy’s law example, §5.1.1.  $n$  = latent dimension;  $\rho$  = 0.1-leakyReLU.

| Layer # | Type          | Input dim. | Output dim. | Support | Activation |
|---------|---------------|------------|-------------|---------|------------|
| 1       | Mesh-informed | 500        | 250         | 0.25    | $\rho$     |
| 2       | Mesh-informed | 250        | 125         | 0.5     | $\rho$     |
| 3       | Dense         | 125        | $n$         | -       | $\rho$     |

Table 5: Architecture of the reduced map network,  $\phi$ , for Burger’s example, §5.1.2. Entries read as in Table 4.

## References

- [1] B. Adcock and N. Dexter. The gap between theory and practice in function approximation with deep neural networks. *SIAM Journal on Mathematics of Data Science*, 3(2):624–655, 2021.
- [2] R. J. Adler, J. E. Taylor, et al. *Random fields and geometry*, volume 80. Springer, 2007.
- [3] D. Amsallem and C. Farhat. An online method for interpolating linear parametric reduced-order models. *SIAM Journal on Scientific Computing*, 33(5):2169–2198, 2011.

| Problem            | $\alpha_1$ | $\alpha_2$ | $\alpha_3$ | Training time | Epochs |
|--------------------|------------|------------|------------|---------------|--------|
| Darcy's law §5.1.1 | 1/5        | 1/5        | 1/16       | 7m 20.17s     | 300    |
| Burger's §5.1.2    | rel        | 1          | 1/16       | 4m 01.58s     | 500    |

Table 6: Training of the DL-ROMs: technical details. The  $\alpha_j$ 's define the loss function as in Eq. (6); "rel" means that the term  $\alpha_j \|\mathbf{y}_i - \hat{\mathbf{y}}_i\|^2$  is replaced with a relative error  $\|\mathbf{y}_i - \hat{\mathbf{y}}_i\|/\|\mathbf{y}_i\|$ .

- [4] C. Angermueller, T. Pärnamaa, L. Parts, and O. Stegle. Deep learning for computational biology. *Molecular systems biology*, 12(7):878, 2016.
- [5] J.-P. Aubin and H. Frankowska. *Set-valued analysis*. Springer Science & Business Media, 2009.
- [6] A. Badré, L. Zhang, W. Muchero, J. C. Reynolds, and C. Pan. Deep neural network improves the estimation of polygenic risk scores for breast cancer. *Journal of Human Genetics*, 66(4):359–369, 2021.
- [7] F. Bartolucci, E. De Vito, L. Rosasco, and S. Vigogna. Understanding neural networks with reproducing kernel banach spaces. *Applied and Computational Harmonic Analysis*, 62:194–236, 2023.
- [8] P. G. Breen, C. N. Foley, T. Boekholt, and S. P. Zwart. Newton versus the machine: solving the chaotic three-body problem using deep neural networks. *Monthly Notices of the Royal Astronomical Society*, 494(2):2465–2470, 2020.
- [9] S. Brivio, S. Fresca, N. R. Franco, and A. Manzoni. Error estimates for pod-dl-roms: a deep learning framework for reduced order modeling of nonlinear parametrized pdes enhanced by proper orthogonal decomposition. *arXiv preprint arXiv:2305.04680*, 2023.
- [10] K. Carlberg. Adaptive h-refinement for reduced-order models. *International Journal for Numerical Methods in Engineering*, 102(5):1192–1210, 2015.
- [11] L. Cici, S. Fresca, E. Zappone, S. Pagani, F. Regazzoni, A. Manzoni, A. Quarteroni, et al. Reduced order modeling of the cardiac function across the scales. In *Reduced Order Models for the Biomechanics of Living Organs*, pages 403–433. Elsevier, 2023.
- [12] A. Cohen, W. Dahmen, O. Mula, and J. Nichols. Nonlinear reduced models for state and parameter estimation. *SIAM/ASA Journal on Uncertainty Quantification*, 10(1):227–267, 2022.
- [13] A. Cohen, R. DeVore, G. Petrova, and P. Wojtaszczyk. Optimal stable nonlinear approximation. *Foundations of Computational Mathematics*, 22(3):607–648, 2022.
- [14] A. S. Dalda, J. de la Cruz Rodríguez, B. De Pontieu, and M. Gošić. Recovering thermodynamics from spectral profiles observed by iris: A machine and deep learning approach. *The Astrophysical Journal Letters*, 875(2):L18, 2019.
- [15] I. Daubechies, R. DeVore, S. Foucart, B. Hanin, and G. Petrova. Nonlinear approximation and (deep) relu networks. *Constructive Approximation*, 55(1):127–172, 2022.
- [16] T. De Ryck and S. Mishra. Generic bounds on the approximation error for physics-informed (and) operator learning. *arXiv preprint*, page arXiv:2205.11393, 2022.
- [17] R. A. DeVore, R. Howard, and C. Micchelli. Optimal nonlinear approximation. *Manuscripta mathematica*, 63:469–478, 1989.
- [18] L. C. Evans. *Partial differential equations*, volume 19. American Mathematical Society, 2022.
- [19] N. Franco, A. Manzoni, and P. Zunino. A deep learning approach to reduced order modelling of parameter dependent partial differential equations. *Mathematics of Computation*, 92:483–524, 2023.

- [20] N. R. Franco, S. Fresca, A. Manzoni, and P. Zunino. Approximation bounds for convolutional neural networks in operator learning. *Neural Networks*, 161:129–141, 2023.
- [21] N. R. Franco, A. Manzoni, and P. Zunino. Mesh-informed neural networks for operator learning in finite element spaces. *Journal of Scientific Computing*, 97(35), 2023.
- [22] S. Fresca, L. Dede, and A. Manzoni. A comprehensive deep learning-based approach to reduced order modeling of nonlinear time-dependent parametrized pdes. *Journal of Scientific Computing*, 87(2):1–36, 2021.
- [23] S. Fresca, G. Gobat, P. Fedeli, A. Frangi, and A. Manzoni. Deep learning-based reduced order models for the real-time simulation of the nonlinear dynamics of microstructures. *International Journal for Numerical Methods in Engineering*, 123(20):4749–4777, 2022.
- [24] S. Fresca and A. Manzoni. Pod-dl-rom: enhancing deep learning-based reduced order models for nonlinear parametrized pdes by proper orthogonal decomposition. *Computer Methods in Applied Mechanics and Engineering*, 388:114181, 2022.
- [25] S. Fresca, A. Manzoni, L. Dedè, and A. Quarteroni. Deep learning-based reduced order models in cardiac electrophysiology. *PloS one*, 15(10):e0239416, 2020.
- [26] G. Gobat, S. Fresca, A. Manzoni, and A. Frangi. Reduced order modeling of nonlinear vibrating multiphysics microstructures with deep learning-based approaches. *Sensors*, 23(6):3001, 2023.
- [27] R. Gribonval, G. Kutyniok, M. Nielsen, and F. Voigtlaender. Approximation spaces of deep neural networks. *Constructive approximation*, 55(1):259–367, 2022.
- [28] M. Guo and J. S. Hesthaven. Data-driven reduced order modeling for time-dependent problems. *Computer methods in applied mechanics and engineering*, 345:75–99, 2019.
- [29] K. Hashimoto, S. Sugishita, A. Tanaka, and A. Tomiya. Deep learning and holographic qcd. *Physical Review D*, 98(10):106014, 2018.
- [30] M. Heinkenschloss, B. Kramer, T. Takhtaganov, and K. Willcox. Conditional-value-at-risk estimation via reduced-order models. *SIAM/ASA Journal on Uncertainty Quantification*, 6(4):1395–1423, 2018.
- [31] J. S. Hesthaven, G. Rozza, B. Stamm, et al. *Certified reduced basis methods for parametrized partial differential equations*, volume 590. Springer, 2016.
- [32] 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.
- [33] K. Hornik. Approximation capabilities of multilayer feedforward networks. *Neural networks*, 4(2):251–257, 1991.
- [34] B. Horvath, A. Muguruza, and M. Tomas. Deep learning volatility: a deep neural network perspective on pricing and calibration in (rough) volatility models. *Quantitative Finance*, 21(1):11–27, 2021.
- [35] A. Ivagnes, G. Stabile, A. Mola, T. Iliescu, and G. Rozza. Hybrid data-driven closure strategies for reduced order modeling. *Applied Mathematics and Computation*, 448:127920, 2023.
- [36] M. G. Kapteyn, J. V. Pretorius, and K. E. Willcox. A probabilistic graphical model foundation for enabling predictive digital twins at scale. *Nature Computational Science*, 1(5):337–347, 2021.
- [37] Y. Kazashi, F. Nobile, and F. Zoccolan. Dynamically orthogonal approximation for stochastic differential equations. *arXiv preprint arXiv:2308.11581*, 2023.
- [38] T. Lähivaara, L. Kärkkäinen, J. M. Huttunen, and J. S. Hesthaven. Deep convolutional neural networks for estimating porous material parameters with ultrasound tomography. *The Journal of the Acoustical Society of America*, 143(2):1148–1158, 2018.

- [39] 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):tnac001, 2022.
- [40] K. Lee and K. T. Carlberg. Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders. *Journal of Computational Physics*, 404:108973, 2020.
- [41] W. Long, Z. Lu, and L. Cui. Deep learning-based feature engineering for stock price movement prediction. *Knowledge-Based Systems*, 164:163–173, 2019.
- [42] L. Lu, X. Meng, S. Cai, Z. Mao, S. Goswami, Z. Zhang, and G. E. Karniadakis. A comprehensive and fair comparison of two neural operators (with practical extensions) based on fair data. *Computer Methods in Applied Mechanics and Engineering*, 393:114778, 2022.
- [43] M. C. Massi, F. Gasperoni, F. Ieva, A. M. Paganoni, P. Zunino, A. Manzoni, N. R. Franco, L. Veldeman, P. Ost, V. Fonteyne, et al. A deep learning approach validates genetic risk factors for late toxicity after prostate cancer radiotherapy in a requisite multi-national cohort. *Frontiers in oncology*, 10:541281, 2020.
- [44] J. Mercer. Xvi. functions of positive and negative type, and their connection to the theory of integral equations. *Philosophical transactions of the royal society of London. Series A, containing papers of a mathematical or physical character*, 209(441-458):415–446, 1909.
- [45] F. Negri, G. Rozza, A. Manzoni, and A. Quarteroni. Reduced basis method for parametrized elliptic optimal control problems. *SIAM Journal on Scientific Computing*, 35(5):A2316–A2340, 2013.
- [46] D. M. Nelson, A. C. Pereira, and R. A. De Oliveira. Stock market’s price movement prediction with lstm neural networks. In *2017 International joint conference on neural networks (IJCNN)*, pages 1419–1426. Ieee, 2017.
- [47] M. Ohlberger and S. Rave. Reduced basis methods: Success, limitations and future challenges. *arXiv preprint arXiv:1511.02021*, 2015.
- [48] C. Pagliantini. Dynamical reduced basis methods for hamiltonian systems. *Numerische Mathematik*, 148(2):409–448, 2021.
- [49] P. Petersen, M. Raslan, and F. Voigtlaender. The structure of spaces of neural network functions. In *Wavelets and Sparsity XVIII*, volume 11138, pages 144–151. SPIE, 2019.
- [50] F. Pichi, F. Ballarin, G. Rozza, and J. S. Hesthaven. An artificial neural network approach to bifurcating phenomena in computational fluid dynamics. *Computers & Fluids*, 254:105813, 2023.
- [51] F. Pichi, B. Moya, and J. S. Hesthaven. A graph convolutional autoencoder approach to model order reduction for parametrized pdes. *arXiv preprint arXiv:2305.08573*, 2023.
- [52] A. Quarteroni, A. Manzoni, and F. Negri. *Reduced basis methods for partial differential equations: an introduction*, volume 92. Springer, 2015.
- [53] A. Quarteroni, R. Sacco, and F. Saleri. *Numerical mathematics*, volume 37. Springer Science & Business Media, 2010.
- [54] S. S. Ravindran. A reduced-order approach for optimal control of fluids using proper orthogonal decomposition. *International journal for numerical methods in fluids*, 34(5):425–448, 2000.
- [55] F. Romor, G. Stabile, and G. Rozza. Non-linear manifold reduced-order models with convolutional autoencoders and reduced over-collocation method. *Journal of Scientific Computing*, 94(3):74, 2023.
- [56] L. Rosafalco, M. Torzoni, A. Manzoni, S. Mariani, and A. Corigliano. Online structural health monitoring by model order reduction and deep learning algorithms. *Computers & Structures*, 255:106604, 2021.
- [57] K. Ryan, J. Lengyel, and M. Shatruk. Crystal structure prediction via deep learning. *Journal of the American Chemical Society*, 140(32):10158–10168, 2018.

- [58] K. Sato, M. Akiyama, and Y. Sakakibara. Rna secondary structure prediction using deep learning with thermodynamic integration. *Nature communications*, 12(1):941, 2021.
- [59] K. T. Schütt, H. E. Sauceda, P.-J. Kindermans, A. Tkatchenko, and K.-R. Müller. Schnet—a deep learning architecture for molecules and materials. *The Journal of Chemical Physics*, 148(24), 2018.
- [60] 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.
- [61] S. M. Srivastava. *A course on Borel sets*, volume 180. Springer Science & Business Media, 2008.
- [62] G. S. Tandel, M. Biswas, O. G. Kakde, A. Tiwari, H. S. Suri, M. Turk, J. R. Laird, C. K. Asare, A. A. Ankrah, N. Khanna, et al. A review on a deep learning perspective in brain cancer classification. *Cancers*, 11(1):111, 2019.
- [63] Y. C. Taumhas, G. Dusson, V. Ehrlacher, T. Lelièvre, and F. Madiot. Reduced basis method for non-symmetric eigenvalue problems: application to the multigroup neutron diffusion equations. *arXiv preprint arXiv:2307.05978*, 2023.
- [64] Q. Wang, N. Ripamonti, and J. S. Hesthaven. Recurrent neural network closure of parametric pod-galerkin reduced-order models based on the mori-zwanzig formalism. *Journal of Computational Physics*, 410:109402, 2020.
- [65] J. Wei Tan, S.-W. Chang, S. Abdul-Kareem, H. J. Yap, and K.-T. Yong. Deep learning for plant species classification using leaf vein morphometric. *IEEE/ACM transactions on computational biology and bioinformatics*, 17(1):82–90, 2018.
- [66] D.-X. Zhou. The covering number in learning theory. *Journal of Complexity*, 18(3):739–767, 2002.
- [67] A. Ziletti, D. Kumar, M. Scheffler, and L. M. Ghiringhelli. Insightful classification of crystal structures using deep learning. *Nature communications*, 9(1):2775, 2018.

## MOX Technical Reports, last issues

Dipartimento di Matematica  
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

- 84/2024** Franco, N.R.; Fraulin, D.; Manzoni, A.; Zunino, P.  
*On the latent dimension of deep autoencoders for reduced order modeling of PDEs parametrized by random fields*
- 85/2024** Brivio, S.; Franco, Nicola R.; Fresca, S.; Manzoni, A.  
*Error estimates for POD-DL-ROMs: a deep learning framework for reduced order modeling of nonlinear parametrized PDEs enhanced by proper orthogonal decomposition*
- 83/2024** Conti, P.; Guo, M.; Manzoni, A.; Frangi, A.; Brunton, S. L.; Kutz, J.N.  
*Multi-fidelity reduced-order surrogate modelling*
- 82/2024** Rosafalco, L.; Conti, P.; Manzoni, A.; Mariani, S.; Frangi, A.  
*EKF-SINDy: Empowering the extended Kalman filter with sparse identification of nonlinear dynamics*
- 80/2024** Crippa, B.; Scotti, A.; Villa, A  
*Numerical Solution of linear drift-diffusion and pure drift equations on one-dimensional graphs*
- 79/2024** Baioni, P.J.; Benacchio, T.; Capone, L.; de Falco, C.  
*Portable, Massively Parallel Implementation of a Material Point Method for Compressible Flows*
- 78/2024** Ziarelli, G.; Pagani, S.; Parolini, N.; Regazzoni, F.; Verani, M.  
*A model learning framework for inferring the dynamics of transmission rate depending on exogenous variables for epidemic forecasts*
- 77/2024** Piersanti, R.; Bradley, R.; Ali, S.Y.; Quarteroni A.; Dede', L.; Trayanova, N.A.  
*Defining myocardial fiber bundle architecture in atrial digital twins*
- 75/2024** Cattarossi, L.; Sacco, F.; Giuliani, N.; Parolini, N.; Mola, A.  
*A geometry aware arbitrary order collocation Boundary Element Method solver for the potential flow past three dimensional lifting surfaces*
- 74/2024** Crippa, B., Scotti, A.; Villa, A  
*A mixed-dimensional model for the electrostatic problem on coupled domains*