Departamento de Matemáticas, Facultad de Ciencias
Universidad Autónoma de Madrid

# Physics-Informed Neural Networks

MSc in Mathematics and Applications

*Author:* Daniel López Montero
*Tutor:* Davide Barbieri  Rafael Orive Illera

Course 2024-2025

# Acknowledgement

First and foremost, I want to express my sincere gratitude to the directors of this research work, Davide and Rafael, for their dedication and support throughout this year. Completing my Master's has been a very gratifying experience, allowing me to grow and learn a great deal. I hope to continue collaborating with them in my future endeavors.

I would also like to thank my parents, family, friends, classmates, and coworkers for their support and understanding during this demanding year. Although mathematical research, especially in a niche field, can sometimes feel isolating, the unwavering support of my relatives has been vital and key when facing challenges. They have encouraged me and given me the strength and energy to be the best version of myself every day.

Finally, I want to give a special mention to my mother, who has been a fundamental pillar in my life. She has always been by my side through the good days as well as the bad. She has inspired me and forged the person I am today.

## Resumen

Las redes neuronales se han utilizado ampliamente en muchos campos con resultados impresionantes. La mayoría de las aplicaciones utilizan los datos como única fuente de aprendizaje. Recientemente, los investigadores han propuesto utilizar información adicional sobre los datos latentes que da lugar al aprendizaje automático informado por la información. Una de las aplicaciones recientes ha sido en el campo de las ecuaciones diferenciales, donde queremos que el modelo verifique algunas ecuaciones, condiciones de contorno y, posiblemente, algunos datos. En este manuscrito, nos sumergiremos en la familia de las Redes Neuronales Informadas por la Física, explorando las propiedades matemáticas relacionadas con su convergencia y error.

## Abstract

Neural networks have been used extensively in many fields with impressive results. Most applications use data as the sole learning source. Recently, researchers have proposed to use additional information about the latent data that gives birth to information-informed machine learning. One of the recent applications has been in the field of differential equations where we want the model to verify some equations, boundary conditions, and possibly some data. In this manuscript, we will dive into the family of Physics-Informed Neural Networks, exploring mathematical properties relating to its convergence and error.

# Contents

# Introduction

In our continuous pursuit of enlightenment, knowledge, and understanding of the world around us, we turn to the tool of mathematics. For thousands of years, mathematicians have sought to translate our sensory experiences into a language that captures their essence in the most abstract, universal, and general form.

> “Mathematics is the language in which God has written the universe.”
> — Galileo Galilei

Mathematics has provided a unique framework to express many foundational ideas about the natural world. For example, in 250 BCE, Archimedes proposed a model to describe the upward force exerted by a fluid on a submerged object, i.e., $F = \rho \cdot V \cdot g$ where $F$ is the force, $\rho$ the density, $V$ the volume and $g$ the force of gravity. More than a thousand years later, in the 17th century, Kepler formulated an equation to describe the motion of the planets around the Sun, i.e., $\frac{dA}{dt}$ is constant, where $A$ is the area of the line connecting a planet and the Sun. Shortly afterward, Newton introduced the three Laws of Motion, which became the foundation of classical mechanics and later led to more complex models such as the Naview-Stokes equation.

Most of the equations governing the Universe—though not all—are expressed in terms of differentials. These problems are studied in the field of differential equations. The main objective in this field is to describe the solutions of these equations. Unfortunately, some of them do not have closed-form solutions, requiring us to rely on approximations provided by numerical methods. A key challenge with numerical methods is ensuring that the approximations are reliable. This involves analyzing the consistency, stability, and convergence of the methods. Another significant issue is that numerical computations can be very resource-intensive.

In recent years, alternative approaches to traditional numerical methods have emerged, particularly for handling complex geometries or solving ill-posed problems—such as those with unknown boundary conditions. One such alternative is machine learning, which has gained significant traction in this context. The *leitmotiv* of machine learning is the learn a function from a dataset composed of samples. In our case, our goal may be to learn a *physics phenomenon* from a dataset, which is why this approach is often referred to as *data-driven*. A key drawback of data-driven methods is their heavy reliance on large quantities of data, which may not always be available. Moreover, these methods often disregard prior knowledge about the physical system. This limitation has led to the development of *physics-informed machine learning* Raissi et al. [2017a], Karniadakis et al. [2021]. The motivation behind physics-informed machine learning is twofold: to harness the strengths of machine learning while incorporating prior physical knowledge of the system. Combining these two domains increases the complexity of the problem, and this is the central focus of our study in this manuscript.

The way in which domain knowledge is incorporated—either before or during training—varies from method to method. For example, NeuralODEs learn the vector field rather than the solution itself, using a numerical method to compute the solution. In contrast, Physics-Informed Neural Networks (PINNs) enforce the governing equations by including them in the loss function used during training. This approach introduces two main challenges. First, the loss function must be differentiable in order to apply gradient-based optimization, which means the physical model must also be differentiable and computationally inexpensive for gradient evaluation. Fortunately, this challenge is mitigated by modern automatic differentiation tools. Second, adding multiple terms to the loss function to better align with the physical model can lead to training instabilities, and there is no guarantee that training will converge to a physically meaningful solution. As we will demonstrate throughout this manuscript, minimizing the *empirical loss function* an sometimes result in physically implausible solutions—a phenomenon often referred to as overfitting. However, we will show that incorporating a regularization term into the training process eliminates these abnormal solutions and ensures the *consistency* of the method.

The family of functions use to solve the problem is usually called the *hypothesis set* or model class. The most common family of functions are the neural networks. They provide a good backbone due to the universality principle Cybenko [1989]. Altough we will focus in the neural networks in this manuscript, there are other popular architetures such as Transformers, Convolutional Neural Networks, Kolmogorov Arnold Networks, etc.

This manuscript is divided as follows. In the first chapter, we will introduce some basic concepts in machine learning, some properties of neural networks and we will define the physics-informed framework. The second chapter provides us with the necessary tools in stochastic processes to prove the results that follows. The third chapter includes some of the theoretical results to justify use of the loss function. The forth chapter applies the same results to the NeuralODEs framework. Finally, in the fifth chapter, we motivate the theory with some experiments.

CHAPTER 1

# Physics-Informed Machine Learning

Advances in machine learning and deep learning have led to significant breakthroughs across nearly all areas of science and technology. The most popular building block of machine learning has been the neural networks Rosenblatt [1958]. In this chapter we will briefly introduce the neural networks and some of its properties such us bounds, Lipschitz constants, etc.

Next, we introduce Physics-Informed Neural Networks (PINNs), which leverage neural networks as function approximators while incorporating physical laws directly into the learning process. These physical constraints—such as initial conditions, boundary conditions, and governing differential equations—are embedded into the loss function, guiding the training process through gradient descent. Finally, we define both the theoretical risk and the empirical risk functions, which provide a foundation for understanding the training and generalization behavior of these models.

## 1.1. Neural Networks

Neural Networks are a family of parametrizable functions that are a composition of affine functions combined with a non-linear functions. The goal of the neural network is process an input $x \in \mathcal{X}$ and produce an output $y \in \mathcal{Y}$ that minimizes a *loss function* using a dataset of samples. The process of minimizing the error through the modification of the parameters is called *training* and it is usually used the *gradient descent* method.

Most modern machine learning techniques are based on this very simple concept. For example, *Transformers* uses neural networks in composition with the attention layer, Graph Neural Networks uses neural networks after data aggregation in each node, 1D Convolutional Neural Networks are a special case of neural networks, etc. The prowess of neural networks reside in its simplicity and the ability to leverage modern computing power.

> "The biggest lesson that can be read from 70 years of AI research is that general methods that leverage computation are ultimately the most effective,

and by a large margin."
— Rich Sutton, *The Bitter Lesson*

Despite not being overly complicated, neural networks can produce very complex behaviour. As we will see the class of neural networks are dense in the continuous functions. Let us introduce rigourusly the neural network.

**The class of neural networks** We will use as our building block the class of fully-connected feedforward neural networks with $H \in \mathbb{N}$ hidden layers of sizes $(L_1, \dots, L_H) = (D, \dots, D) \in \mathbb{N}^H$, where we assume, for simplicity, a constant size of the hidden layer. The neural network is a function from $\mathbb{R}^{d_1}$ to $\mathbb{R}^{d_2}$, defined by

$$u_\theta = \mathcal{A}_{H+1} \circ (\phi \circ \mathcal{A}_H) \circ \cdots \circ (\phi \circ \mathcal{A}_1),$$

where $\phi$ is the *activation function*, in our case, and hereafter, we will use the hyperbolic tangent applied point-wise in each dimension, which is preferable over other popular choices, e.g., ReLU, due to being differentiable. Each $\mathcal{A}_k : \mathbb{R}^{L_{k-1}} \to \mathbb{R}^{L_k}$ is an affine function with $W_k \in M_{L_{k-1} \times L_k}(\mathbb{R})$ and $b_k \in \mathbb{R}^{L_k}$. We use for the first and last layer, $L_0 = d_1$, and $L_{H+1} = d_2$, the size of the input and output of the neural network. The parameters learned by the neural network $u_\theta$ are expressed as

$$\theta := (W_1, b_1, \dots, W_{H+1}, b_{H+1}) \in \Theta_{H,D},$$

where $\Theta_{H,D} = \mathbb{R}^{\sum_{i=0}^{H}(L_i L_{i+1} + L_{i+1})}$. We denote the set of all parametrized neural networks as $\mathrm{NN}_H^D := \{u_\theta, \theta \in \Theta_{H,D}\}$.

### 1.1.1. Density of NN in Hölder spaces

A key reason for using neural networks over other parametrizable and differentiable function classes is their *computational efficiency*. The *forward* and the *back-propagation* are computational inexpensive because because neural networks are composition of very simple functions—namely non-linear activation functions, summations, and matrix products—which are highly parallelizable. However, computational efficiency is useful only if the function class is expressive enough to arbitrarily approximate any given function. This is guaranteed by the Universal Approximation Theorem, first proven by Cybenko [1989], which states that a neural network with a single hidden layer can approximate any continuous function to any desired degree of accuracy. This foundational result established neural networks as universal approximators and paved the way for many subsequent theoretical advancements in machine learning.

Before stating the theorem, we will need some technical results related to the derivatives of the individual compontents of the neural network with hyperbolic tangent. Firstly, we will introduce a concept in combinatorics that is useful to calculate the chain-rule of the composition of functions.

**Definition 1.1** (Bell number)**.** The $n$th Bell number $B_n$, corresponds to the number of partitions of the set $\{1, \dots, n\}$. We denote the set of all partitions as $\Pi(n)$. A partition of a set $S$ is a family of nonempty, pairwise disjoint subsets of $S$ whose union is $S$. The first Bell number is defined as $B_0 = 1$ and it is defined recursively as follows

$$B_{n+1} := \sum_{k=0}^{n} \binom{n}{k} B_k. \tag{1.1}$$

For example, the first three bell numbers can be expressed as

$$\begin{aligned} B_1 = 1: &\quad \{\{1\}\} \\ B_2 = 2: &\quad \{\{1\},\{2\}\}, \{\{1,2\}\} \\ B_3 = 5: &\quad \{\{1\},\{2\},\{3\}\}, \{\{1\},\{2,3\}\}, \{\{2\},\{1,3\}\}, \{\{3\},\{1,2\}\}, \{\{1,2,3\}\} \end{aligned}$$

**Lemma 1.2** (Bounding the partial derivatives of a composition function)**.** Let $d_1, d_2 \geq 1, K \geq 0, f \in C^K(\mathbb{R}^{d_1}, \mathbb{R})$, and $g \in C^K(\mathbb{R}, \mathbb{R}^{d_2})$. Then,

$$\|g \circ f\|_{C^K(\mathbb{R}^{d_1})} \leq B_K \|g\|_{C^K(\mathbb{R})} (1 + \|f\|_{C^K(\mathbb{R}^{d_1})})^K$$

where $B_K$ is the $K$-th Bell number.

*Proof.* Let $K_1 \leq K$ and $\Pi(K_1)$ be the set of all partitions of $\{1, \dots, K_1\}$. According to [Hardy, 2006, Proposition 1], for all $h \in C^{K_1}(\mathbb{R}^{K_1+d_1}, \mathbb{R})$,

$$\partial^{K_1}_{1,2,\dots,K_1}(g \circ h) = \sum_{P \in \Pi(K_1)} g^{(|P|)} \circ h \times \prod_{S \in P} \left[ \left( \prod_{j \in S} \partial_j \right) h \right].$$

Let $\alpha = (\alpha_1, \dots, \alpha_{d_1})$ be a multi-index such that $|\alpha| = K_1$. Setting $\alpha_0 = 0, y_j = x_{K_1+j} + (x_{\alpha_1+\dots+\alpha_{j-1}} + \dots + x_{\alpha_1+\dots+\alpha_j-1})$ and letting $h(x_1, \dots, x_{K_1+d_1}) = f(y_1, \dots, y_{d_1})$, we obtain

$$\partial^\alpha (g \circ f) = \sum_{P \in \Pi(K_1)} g^{(|P|)} \circ f \times \prod_{S \in P} \partial^{\alpha(S)} f,$$

where $\alpha(S) = (|\{b \in S, \alpha_1 + \dots + \alpha_{\ell-1} \leq b \leq \alpha_1 + \dots + \alpha_\ell\}|)_{1 \leq \ell \leq d_1}$. Moreover, by the definition of the Bell number, $|\Pi(K_1)| = B_{K_1}$, and by definition of a partition, $|P| \leq K_1$. Hence,

$$\begin{aligned} \|\partial^\alpha (g \circ f)\|_\infty &\leq B_{K_1} \|g\|_{C^{K_1}(\mathbb{R}^{d_1})} \max_{i_1 + 2i_2 + \dots + K_1 i_{K_1} = K_1} \prod_{j=1}^{K_1} \|f\|^{i_j}_{C^j(\mathbb{R}^{d_1})} \\ &\leq B_{K_1} \|g\|_{C^{K_1}(\mathbb{R}^{d_1})} (1 + \|f\|_{C^{K_1}(\mathbb{R}^{d_1})})^{K_1}. \end{aligned} \tag{1.2}$$

This inequality is proved for all $K_1 \leq K$ and for all $|\alpha| = K$. □

**Lemma 1.3** (Bounding derivatives of tanh)**.** Let $K \in \mathbb{N} \cup \{0\}$, Then

$$\| \tanh^{(K)} \|_\infty \leq 2^{K-1}(K+2)!$$

*Proof.* Then tanh is a solution of the equation $y' = 1-y^2$. By induction, we can prove that $\tanh^{(K)} = P_K(\tanh)$ where $P_K$ is a polynomial that can be built by induction. Firstly, $P_0(X) = X$ due to $\tanh^{(0)} = \tanh = P(\tanh)$. Owing to $y' = 1 - y^2$, $P_1(X) = 1 - X^2$. Assuming it is true for $K$, i.e., $\tanh^{(K)}(t) = P_K(\tanh(t))$,

$$\begin{aligned}\tanh^{(K+1)}(t) &= \frac{d}{dt}\tanh^{(K)}|_t = \frac{d}{dt}(P_K \circ \tanh)|_t \\ &= P_K'(\tanh(t))\tanh'(t) = P_K'(\tanh(t))(1-\tanh^2(t)).\end{aligned}$$

Therefore, we obtain the following recursive formula, $P_{K+1}(X) = (1-X^2)P_K'(X)$. The degree of $P_K$ is $2 + (\deg(P_{K-1}) - 1) = P_{K-1} + 1$, hence $\deg(P_K) = K + 1$. Assume that $P_K(X) = a_0^{(K)} + a_1^{(K)}X + \cdots + a^{(K+1)}X^{K+1}$. It verifies that $a_i^{(K+1)} = (i+1)a_{i+1}^{(K)} - (i-1)a_{i-1}^{(K)}$, with $a_{-1}^{(K)} = a_{K+2}^{(K)} = 0$. The largest coefficient $M(P_K) = \max_{0\le i\le K+1}|a_i^{(K)}|$ of $P_K$ satisfies $M(P_{K+1}) \le 2(K+1)M(P_K)$. Since $M(P_1) = 1$, we see that $M(P_K) \le 2^{K-1}K!$. We know that $-1 \le \tanh \le 1$, therefore,

$$\|\tanh^{(K)}\|_\infty = \|P_K(\tanh)\|_\infty \le (K+2)M(P_K) \le 2^{k-1}(K+2)!$$

□

We will not prove the following Lemma as it is mere calculations of derivatives and do not provide any useful insights. The prove can be found in [Doumèche et al., 2024, Lemma C.4 & Corollary C.5]

**Lemma 1.4.** Let $K \in \mathbb{N}$ and $H \in \mathbb{N}$. Then, for all $\epsilon > 0$

$$\lim_{\theta\to\infty} \theta^m \|\tanh_\theta^{\circ H} - \mathrm{sign}\|_{C^K(\mathbb{R}\setminus(-\epsilon,\epsilon))} = 0,$$

where $\tanh_\theta^{\circ H} := \tanh_\theta \circ \overset{H}{\cdots} \circ \tanh_\theta$ where $\tanh_\theta(x) := \tanh(x\theta)$. Moreover, for all $\theta \in \mathbb{R}$, $\|(\tanh_\theta^{\circ H})^{(K)}\|_\infty \le \infty$.

We can now state the density theorem for neural networks in Hölder spaces. This theorem generalizes the universal approximation theorem from [De Ryck et al., 2021, Theorem 5.1], which is limited to the one-dimensional case.

**Proposition 1.5** (Density of NN in Hölder spaces)**.** Let $K \in \mathbb{N}, H \ge 2$, and $\Omega \subset \mathbb{R}^{d_1}$ be a bounded Lipschitz domain. Then, $\cup_{D\in\mathbb{N}}\mathrm{NN}_H^D$ is dense in $(C^\infty(\overline{\Omega}, \mathbb{R}^{d_2}), \|\cdot\|_{C^K(\Omega)})$.

*Proof.* We want to prove that for any function $u \in C^\infty(\overline{\Omega}, \mathbb{R}^{d_2})$, there exists a sequence of neural networks $(u_n)_{n\in\mathbb{N})} \subseteq \mathrm{NN}_H$ such that $\lim_{n\to\infty}\|u - u_n\|_{C^K(\Omega)} = 0$. The Universal Approximation theorem [De Ryck et al., 2021, Theorem 5.1] shows that the class of neural networks with two layers, i.e., $\mathrm{NN}_2$, is dense in $(W^{K+1,\infty}([0,1]^{d_1}, \mathbb{R}), \|\cdot\|_{W^{K+1,\infty}((0,1)^{d_1})})$. Since $C^\infty([0,1]^{d_1}) \subseteq W^{K+1,\infty}([0,1]^{d_1})$ and the norms $\|\cdot\|_{C^K}$ and $\|\cdot\|_{W^{K,\infty}}$ coincide on $C^\infty([0,1]^{d_1})$, then, the class of $\mathrm{NN}_2$ is dense in $(C^\infty([0,1]^{d_1}, \mathbb{R}), \|\cdot\|_{C^K([0,1])^{d_1}})$.

We intend to generalize the previous result to any output dimension $d_2 \in \mathbb{N}$, any bounded Lipschitz domain $\Omega$, and any number of layers $H \ge 2$.

**(Step 1) Generalization to any bounded Lipschitz domain $\Omega$.** We assume that $d_2 = 1$ and we aim to prove that $\mathrm{NN}_2$ is dense in $(C^\infty(\overline{\Omega}, \mathbb{R}), \|\cdot\|_{C^K(\Omega)})$. Let $f \in C^\infty(\overline{\Omega}, \mathbb{R})$. Since $\Omega$ is bounded, there exists an affine transformation $\tau : \mathbb{R}^{d_1} \to \mathbb{R}^{d_1}$ with $x \mapsto A_\tau x + b_\tau$ where $A_\tau \in \mathbb{R}^* \equiv \mathbb{R}^{d_1 \times d_1}$ (dual space) and $b_\tau \in \mathbb{R}^{d_1}$ such that $\tau(\Omega) \subseteq [0,1]^{d_1}$ and that $\tau$ is injective (has null rank). We will compose the inverse of $\tau$ with $f$ and use De Ryck's theorem. Firstly, let $\hat{f} := f \circ \tau^{-1}$ that maps $[0,1]^{d_1} \to \mathbb{R}$. Owing to the extension theorem for Lipschitz domains [Stein, 1971, Theorem 5, Chapter VI.3.3], the function can be extended to a function $\tilde{f} \in W^{K,\infty}([0,1]^{d_1})$ such that $\tilde{f}|_{\tau(\Omega)} = \tilde{f}_{\tau(\Omega)}$. The Universal Approximation Theorem of Neural Networks [De Ryck et al., 2021, Theorem 5.1] assures us that for every $\epsilon > 0$, there exists $u_\theta \in \mathrm{NN}_2$ such that $\|u_\theta - \hat{f}\|_{W^{K,\infty}([0,1]^{d_1})} \leq \epsilon$. And, since $\tilde{f}$ is an extension of $\hat{f}$, $\tilde{f} \in C^\infty(\overline{\Omega})$ and it follows that $\|u_\theta - \tilde{f}\|_{C^K(\tau(\Omega))} \leq \epsilon$.

Let $\alpha$ be a multi-index such that $|\alpha| = m$. Then, $\partial^\alpha(\hat{f} \circ \tau) = A_\tau^m \times (\partial^\alpha \hat{f})(\tau)$ and $\partial^\alpha(u_\theta \circ \tau) = A_\tau^m \times (\partial^\alpha u_\theta)(\tau)$. Therefore, $\|u_\theta \circ \tau - \hat{f} \circ \tau\|_{C^K(\Omega)} \leq \epsilon \max(1, \|A_\tau^K\|)$, in other words,

$$\|u_\theta \circ \tau - f\|_{C^K(\Omega)} \leq \epsilon \max(1, \|A_\tau^K\|).$$

We note that $u_\theta \circ \tau \in \mathrm{NN}_2$ due to $\tau$ being an affine transformation. Hence, we have proven that for any neural network, there exists a function that is arbitrarily close.

**(Step 2) Generalization to any number of layers $H \geq 2$.** We will show that $\mathrm{NN}_H$ is dense in $(C^\infty(\overline{\Omega}, \mathbb{R}), \|\cdot\|_{C^K(\Omega)})$ for all $H \geq 2$. Firstly, let $f \in C^\infty(\overline{\Omega}, \mathbb{R})$ be an arbitrary function and let us denote by $v$ a function that is the composition of $H-2$ hyperbolic tangents

$$v(x_1, \ldots, x_{d_1}) := (\tanh^{\circ(H-2)}(x_1), \ldots, \tan^{\circ(H-2)}(x_{d_1}))$$

where $\tanh^{\circ(H-2)}$ is the composition of $H-2$ hyperbolic tangents. It is clear that $v$ is a neural network of $H-2$ layers. For any $u_\theta \in \mathrm{NN}_2$, $u_\theta \circ v \in \mathrm{NN}_H$ is a neural network with the first weight matrices $(W_\ell)_{1 \leq \ell \leq H-2}$ are the identity and the biases $(b_\ell)_{1 \leq \ell \leq H-2}$ are equal to zero. Since, $v$ is a differomorphism, $v(\Omega)$ is a bounded Lipschitz domain and $f \circ v^{-1} \in C(v(\Omega), \mathbb{R})$. Lemma 1.2, the composition $f \circ v^{-1}$ must belong to $C^\infty(\overline{v(\Omega)}, \mathbb{R})$. Owing to the previous paragraph, there exists a sequence $(\theta_n)_{n \in \mathbb{N}}$ of parameters such that

$$\|u_{\theta_n} - f \circ v^{-1}\|_{C^K(v(\Omega))} \xrightarrow{n \to \infty} 0.$$

In other words, $u_{\theta_n}$ approximates $f \circ v^{-1}$, but we want to prove that $u_{\theta_n} \circ v$ approximates $f$. Using Lemma 1.2 we let the first term be $u_{\theta_n} - f \circ v^{-1}$ and the second $v$. Hence,

$$\|(u_{\theta_n} - f \circ v^{-1}) \circ v\|_{C^K(\Omega)} = \|u_{\theta_n} \circ v - f\|_{C^K(\Omega)} \leq B_K \|u_{\theta_n} - f \circ v^{-1}\|_{C^K(\Omega)} (1 + \|v\|_{C^K(\Omega)})^K.$$

Since $\|v\|_{C^K(\Omega)})^K = \|\tanh^{\circ(H-2)}\|_{C^K(\Omega)}$ and Lemma 1.4 bounds $\|\tanh^{\circ(H-2)}\|_{C^K(\mathbb{R})} < \infty$. Therefore, we conclude that $\|(u_{\theta_n} - f \circ v^{-1}) \circ v\|_{C^K(\Omega)} < \infty$ with $u_{\theta_n} \circ v \in \mathrm{NN}_H$.

**(Step 3) Generalization to all output dimension $d_2$.** We have seen that $\mathrm{NN}_H$ is dense in $(C^\infty(\overline{\Omega}, \mathbb{R}), \|\cdot\|_{C^K(\Omega)})$. We will prove that $\mathrm{NN}_H$ is dense in $(C^\infty(\overline{\Omega}, \mathbb{R}^{d_2}), \|\cdot\|_{C^K(\Omega)})$.

Let $f = (f_1, \dots, f_{d_2}) \in C^\infty(\Omega, \mathbb{R}^{d_2})$ and let $(\theta_n^{(i)})_{n\in\mathbb{N}}$ for $i = 1, \dots, d_2$, be a sequence of neural networks such that $\|u_{\theta_n^{(i)}} - f_i\|_{C^K(\Omega)} \xrightarrow{n\to\infty} 0$. Therefore, the neural network that results of stacking all the parameters $u_{\theta_n} := (u_{\theta_n^{(1)}}, \dots, u_{\theta_n^{(1)}})$ satisfies that

$$\|u_{\theta_n} - f\|_{C^K(\Omega)} \leq \sum_{i=0}^{d_2} \|u_{\theta_n^{(i)}} - f_i\|_{C^K(\Omega)} \xrightarrow{n\to\infty} 0.$$

□

### 1.1.2. Bounding neural networks by its parameters

The next results bound the norm of the neural network by its parameters. Beware that the results shown are only applicable to neural networks with hyperbolic tangent as actiavtion function. However, we can prove the same results for other differentiable activation functions.

**Lemma 1.6** (Bounding the norm of a neural network by the norm of its parameters)**.** Consider the class of neural networks $\mathrm{NN}_H^D = \{u_\theta, \theta \in \Theta_{H,D}\}$. Let $K \in \mathbb{N}$. Then, there exists a constant $C_{K,H}$ (depending only on $K$ and $H$), such that

$$\|u_\theta\|_{C^K(\mathbb{R}^{d_1})} \leq C_{K,H}(D+1)^{1+KH}(1+\|\theta\|_2)^{KH}\|\theta\|_2 \qquad \forall \theta \in \Theta_{H,D}.$$

*Proof.* Recall, that $u_\theta$ is function from $\mathbb{R}^{d_1}$ to $\mathbb{R}^{d_2}$, defined by $u_\theta = \mathcal{A}_{H+1} \circ (\phi \circ \mathcal{A}_H) \circ \cdots \circ (\phi \circ \mathcal{A}_1)$ where $\phi$ is the *activation function* and each $\mathcal{A}_k : \mathbb{R}^{L_{k-1}} \to \mathbb{R}^{L_k}$ is an affine function with $W_k \in M_{L_{k-1}\times L_k}(\mathbb{R})$ and $b_k \in \mathbb{R}^{L_k}$. We establish $L_0 = d_1, L_1 = \dots, L_H = D, L_{H+1} = d_2$. We let $\|W\|_\infty$ be the maximum of the absolute value of each entry in the matrix, i.e., we view it as a vector. Note that for each layer $k$, $\|W_k x\|_\infty \leq L_{k-1}\|W_k\|_\infty\|x\|_\infty$. Moreover, $\max_{1\leq k\leq H+1}(\|W_k\|_\infty, \|b_k\|_\infty) \leq \|\theta\|_\infty \leq \|\theta\|_2$.
The constant $C_{K,H}$ is defined recursively over the number of layers $H \geq 1$ and the regularity $K \geq 0$. Firstly, assume that $C_{0,H} = 1$, $C_{K,1} = 2^{K-1}(K+2)!$, and

$$C_{K,H+1} = B_K 2^{K-1}(K+2)! \max_{\substack{i_1,\dots,i_K\in\mathbb{N} \\ i_1+2i_2+\cdots+Ki_K=K}} \prod_{1\leq \ell\leq K} C_{\ell,H} \tag{1.3}$$

where $B_K$ is the $K$-th Bell number, defined in (1.1). Let us start considering the case $H = 1$. As we know $u_\theta = W_2 \tanh(\mathcal{A}_1) + b_2$, then

$$\|u_\theta\|_\infty \leq \|W_2 \tanh(\mathcal{A}_1)\|_\infty + \|b_2\|_\infty \leq \|W_2\|_\infty D + \|b_2\|_\infty \leq (D+1)\|\theta\|_2. \tag{1.4}$$

Next, for any multi-index $\alpha = (\alpha_1, \dots, \alpha_{d_1})$ such that $|\alpha| \geq 1$,

$$\partial^\alpha u_\theta(x) = W_2 \begin{pmatrix} (W_1)_{1,1}^{\alpha_1} \times \cdots \times (W_1)_{1,d_1}^{\alpha_{d_1}} \times \tanh^{(\alpha)}(\mathcal{A}_1(x))_1 \\ \vdots \\ (W_1)_{D,1}^{\alpha_1} \times \cdots \times (W_1)_{D,d_1}^{\alpha_{d_1}} \times \tanh^{(\alpha)}(\mathcal{A}_1(x))_D \end{pmatrix} \in \mathbb{R}^{d_1}. \tag{1.5}$$

We can note that $|(W_1)_{i,j}| \leq \|\theta\|_\infty$, we see that

$$\|\partial^\alpha u_\theta\|_\infty \leq D\|W_2\|_\infty \|\theta\|_\infty^{|\alpha|} \|\tanh^{(|\alpha|)}\|_\infty \leq D\|\theta\|_2^{1+|\alpha|} \|\tanh^{(|\alpha|)}\|_\infty. \tag{1.6}$$

Combining (1.4) and (1.6), we obtain that for any $K \geq 1$,

$$\begin{aligned}\|u_\theta\|_{C^K(\mathbb{R}^{d_1})} &\leq (D+1)\max_{k\leq K}\|\tanh^{(k)}\|_\infty (1+\|\theta\|_2)^K \|\theta\|_2 \\ &\leq \underbrace{2^{K-1}(K+2)!}_{C_{K,1}}(D+1)(1+\|\theta\|_2)^K\|\theta\|_2.\end{aligned}$$

By induction, assume that for a given $H \geq 1$, for any neural network $v_\theta \in \mathrm{NN}_H^D$ and any $K \geq 0$,

$$\|v_\theta\|_{C^K(\mathbb{R}^{d_1})} \leq C_{K,H}(D+1)^{1+KH}(1+\|\theta\|_2)^{KH}\|\theta\|_2. \tag{1.7}$$

We want to show that for any $u_\theta \in \mathrm{NN}_{H+1}^D$,

$$\|u_\theta\|_{C^K(\mathbb{R}^{d_1})} \leq C_{K,H+1}(D+1)^{1+K(1+H)}(1+\|\theta\|_2)^{K(H+1)}\|\theta\|_2.$$

For a $u_\theta \in \mathrm{NN}_H^D$, we have, by definition, $u_\theta = \mathcal{A}_{H+2} \circ \tanh \circ v_\theta$ where $v_{\theta'} \in \mathrm{NN}_H^D$. In fact, the parameter of $v_{\theta'}$ is $\theta' = (W_1, b_1, \ldots, W_{H+1}, b_{H+1})$ whereas $\theta = (W_1, b_1, \ldots, W_{H+2}, b_{H+2})$, so $\|\theta'\|_2 \leq \|\theta\|_2$. Therefore,

$$\|u_\theta\|_\infty \leq \|W_{H+2}\tanh(v_\theta) + b_{H+2}\|_\infty \leq \|W_{H+2}\|_\infty D + \|b_{H+2}\|_\infty \leq (D+1)\|\theta\|_2.$$

Furthermore, for any multi-index $\alpha = (\alpha_1, \ldots, \alpha_{d_1})$ such that $|\alpha| \geq 1$,

$$\partial^\alpha u_\theta(x) = W_{H+2}\begin{pmatrix}\partial^\alpha(\tanh\circ(v_\theta(x))_1) \\ \vdots \\ \partial^\alpha(\tanh\circ(v_\theta(x))_D)\end{pmatrix}.$$

Therefore,

$$\|\partial^\alpha u_\theta\|_\infty \leq D\|W_{H+2}\|_\infty \max_{1\leq j\leq D}\|(\tanh(v_\theta))_j\|_{C^K(\mathbb{R}^{d_1})}. \tag{1.8}$$

Using the bound for partial derivatives of composition of functions (1.2), we obtain

$$\|\tanh(v_\theta)_j\|_{C^K(\mathbb{R}^{d_1})} \leq B_K\|\tanh\|_{C^K(\mathbb{R})} \max_{i_1+2i_2+\cdots+Ki_k=K} \prod_{1\leq\ell\leq K}\|(v_\theta)_j\|_{C^\ell(\mathbb{R}^{d_1})}^{i_\ell}.$$

We notice that $(v_\theta)_j$ belongs to $\mathrm{NN}_H^D$ where $d_2 = 1$. Using Lemma 1.3 and inequality (1.7),

$$\begin{aligned}&\|\tanh(v_\theta)_j\|_{C^K(\mathbb{R}^{d_1})} \\ &\leq B_K 2^{K-1}(K+2)! \max_{i_1+2i_2+\cdots+Ki_k=K}\prod_{1\leq\ell\leq K}\left(C_{\ell,H}(D+1)^{1+\ell H}(1+\|\theta\|_2)^{\ell H}\|\theta\|_2\right)^{i_\ell} \\ &\leq C_{K,H+1}\max_{i_1+2i_2+\cdots+Ki_k=K}\prod_{1\leq\ell\leq K}\left((D+1)^{1+\ell H}(1+\|\theta\|_2)^{\ell H}\|\theta\|_2\right)^{i_\ell} \\ &\leq C_{K,H+1}(D+1)^{K+KH}(1+\|\theta\|_2)^{KH}(1+\|\theta\|_2)^K \\ &= C_{K,H+1}(D+1)^{K(1+H)}(1+\|\theta\|_2)^{K(1+H)}.\end{aligned} \tag{1.9}$$

Hence, we can combine the previous equation with (1.8),

$$\begin{aligned}\|\partial^\alpha u_\theta\|_\infty &\le D\|W_{H+2}\|_\infty C_{K,H+1}(D+1)^{K(1+H)}(1+\|\theta\|_2)^{K(1+H)}\\ &\le C_{K,H+1}(D+1)^{1+K(1+H)}(1+\|\theta\|_2)^{K(1+H)}\|\theta\|_2.\end{aligned}$$

□

### 1.1.3. Lipschitz dependence of Neural Networks

**Lemma 1.7** (Lipschitz dependence of parameters of the Hölder norm in NN)**.** Consider the class $\mathrm{NN}_H^D = \{u_\theta, \theta \in \Theta_{H,D}\}$. Fix $K \in \mathbb{N}$. Then, there exists a constant $\tilde{C}_{K,H} > 0$, depending only on the regularity $K$ and the number of hidden layers $H$, such that

$$\|u_\theta - u_{\theta'}\|_{C^K(\Omega)} \le \tilde{C}_{K,H}(1 + d_1 M(\Omega))(D+1)^{H+KH^2}(1+\|\theta\|_2)^{H+KH^2}\|\theta-\theta'\|_2$$

where $M(\Omega) := \sup_{x\in\Omega}\|x\|_\infty$

*Proof.* We define recursively the constants $\tilde{C}_{K,H}$, starting by $\tilde{C}_{K,1} := (K+2)2^{2K-1}(K+2)!(K+3)!$, and

$$\tilde{C}_{K,H+1} := C_{K,H+1}[1 + (K+1)B_K 2^{2K-1}(K+3)!(K+2)!\tilde{C}_{K,H}]$$

where $C_{K,H}$ is the constant defined in previous lemma. Before embarking on the proof, we need to show the following remark. Firstly, we note that for two sequences $(a_i)_{1\le i\le n}$ and $(b_i)_{1\le i\le n}$,

$$\prod_{i=1}^n a_i - \prod_{i=1}^n b_i = \sum_{i=1}^n (a_i - b_i)\left(\prod_{j=i+1}^n a_j\right)\left(\prod_{j=1}^{i-1} b_j\right).$$

This is true because it is a telescopic series, i.e., for the terms $i$ and $i+1$ of the series,

$$\begin{array}{rll} i: & (a_i - b_i)\prod_{j=i+1}^n a_j \prod_{j=1}^{i-1} b_j & = \prod_{j=i}^n a_j \prod_{j=1}^{i-1} b_j - \cancel{\prod_{j=i+1}^n a_j \prod_{j=1}^{i} b_j}\\ i+1: & (a_{i+1} - b_{i+1})\prod_{j=i+2}^n a_j \prod_{j=1}^{i} b_j & = \cancel{\prod_{j=i+1}^n a_j \prod_{j=1}^{i} b_j} - \prod_{j=i+2}^n a_j \prod_{j=1}^{i+1} b_j \end{array}$$

Therefore, we can bound the difference of the products as follows

$$\left|\prod_{i=1}^n a_i - \prod_{i=1}^n b_i\right| \le n \max_{1\le i\le n}\{|a_i - b_i|\}\prod_{i=1}^n \max(|a_i|,|b_i|). \tag{1.10}$$

Owing to (1.2), for all $u, v \in C^K(\Omega, \mathbb{R}^D)$ for all $1 \le i \le D$,

$$\begin{aligned}&\partial^\alpha(\tanh(u_i) - \tanh(v_i))\\ &\quad= \sum_{P\in\Pi(K)}[\tanh^{(|P|)}(u_i)]\prod_{S\in P}\partial^{\alpha(S)}(u_i) - \sum_{P\in\Pi(K)}[\tanh^{(|P|)}(v_i)]\prod_{S\in P}\partial^{\alpha(S)}(v_i)\\ &\qquad= \sum_{P\in\Pi(K)}\left([\tanh^{(|P|)}(u_i)]\prod_{S\in P}\partial^{\alpha(S)}(u_i) - [\tanh^{(|P|)}(v_i)]\prod_{S\in P}\partial^{\alpha(S)}(v_i)\right).\end{aligned} \tag{1.11}$$

Let us denote the summation term $A(P)$, i.e., $\partial^\alpha(\tanh(u_i) - \tanh(v_i)) = \sum_{P\in\Pi(K)} A(P)$. We notice that $A(P)$ is the difference of a product of $|P|+1$ factors. We can use the previous result (1.10) to bound it,

$$\begin{aligned}\|A(P)\|_{\infty,\Omega} \leq (|P|+1)(\|\tanh^{(|P|)}(u_i) - \tanh^{(|P|)}(v_i)\| + \max_{S\in P}\|\partial^{\alpha(S)}u_i - \partial^{\alpha(S)}\|)\\ \|\tanh^{(|P|)}\|_\infty \prod_{S\in P}\max(\|\partial^{\alpha(S)}u\|_{\infty,\Omega}, \|\partial^{\alpha(S)}v\|_{\infty,\Omega})\\ \leq (|P|+1)(\|\tanh^{(|P|)}\|_{\text{Lip}}\|u-v\|_{\infty,\Omega} + \|u-v\|_{C^K(\Omega)})\\ \|\tanh^{(|P|)}\|_\infty \prod_{S\in P}\max(\|\partial^{\alpha(S)}u\|_{\infty,\Omega}, \|\partial^{\alpha(S)}v\|_{\infty,\Omega}).\end{aligned} \tag{1.12}$$

We notice that the Lipschitz constant is the norm of the derivative, i.e., $\|\tanh^{(|P|)}\|_{\text{Lip}} = \|\tanh^{(|P|+1)}\|_\infty$.

With this result, we are equipped to prove the statement by induction on $H$. We will begin with the base case $H=1$. Assume that $K=0$ and then we will generalize to all $K\geq 1$. Let $u_\theta = \mathcal{A}_2 \circ \tanh \mathcal{A}_1$ and $u_{\theta'} = \mathcal{A}_2' \circ \tanh \circ \mathcal{A}_1'$. Firstly, we notice that

$$\|\mathcal{A}_1 - \mathcal{A}_1'\|_{\infty,\Omega} \leq \|b_1 - b_1'\|_\infty + d_1 M(\Omega)\|W_1 - W_1'\|_\infty \leq \|\theta-\theta'\|_2(1 + d_1 M(\Omega)), \tag{1.13}$$

where $M(\Omega) := \max_{x\in\Omega}\|x\|_\infty$. Since $1 = \|\tanh'\|_\infty = \|\tanh\|_{\text{Lip}}$, implies that $\|\tanh\circ\mathcal{A}_1 - \tanh\circ\mathcal{A}_1'\|_\infty \leq \|\theta-\theta'\|_2(1+d_1M(\Omega))$. Subsequently, $\|\mathcal{A}_2 - \mathcal{A}_2'\|_{\infty,B(1,\|\cdot\|_\infty)} \leq \|\theta-\theta'\|_2(1+D)$. Hence,

$$\begin{aligned}\|u_\theta - u_{\theta'}\|_{\infty,\Omega} &\leq \|\mathcal{A}_2\circ\tanh\circ\mathcal{A}_1 - \mathcal{A}_2'\circ\tanh\circ\mathcal{A}_1'\|_{\infty,\Omega}\\ &\leq \|(\mathcal{A}_2 - \mathcal{A}_2')\circ\tanh\circ\mathcal{A}_1\|_{\infty,\Omega} + \|\mathcal{A}_2'\circ\tanh\circ\mathcal{A}_1 - \mathcal{A}_2'\circ\tanh\circ\mathcal{A}_1'\|_{\infty,\Omega}\\ &\leq \|\mathcal{A}_2 - \mathcal{A}_2'\|_{\infty,B(1,\|\cdot\|_\infty)} + D\|W_2'\|_\infty\|\tanh\circ\mathcal{A}_1 - \tanh\circ\mathcal{A}_1'\|_{\infty,\Omega}\\ &\leq \|\theta-\theta'\|_2(1 + D + D\|\theta\|_2(1+d_1M(\Omega)))\\ &\leq (1+d_1M(\Omega))(D+1)(1+\|\theta\|_2)\|\theta-\theta'\|_2.\end{aligned} \tag{1.14}$$

Therefore, it is true for $\tilde{C}_{0,1} = 6$. We have proven that it is true for $H=1$ and $K=0$. We will address when $K\geq 1$. Let $\alpha$ be a multi-index such that $|\alpha| = K$. It follows that

$$\|\partial^\alpha(u_\theta - u_{\theta'})\|_{\infty,\Omega} \leq \|(W_2 - W_2')\partial^\alpha(\tanh\circ\mathcal{A}_1)\|_{\infty,\Omega} + \|W_2'\partial^\alpha(\tanh\circ\mathcal{A}_1 - \tanh\circ\mathcal{A}_1')\|_{\infty,\Omega}. \tag{1.15}$$

Owing to Lemma 1.3, and following a similar argument as in (1.8), we have

$$\begin{aligned}\|(W_2 - W_2')\partial^\alpha(\tanh\circ\mathcal{A}_1)\|_{\infty,\Omega} &\leq (D+1)\|\theta-\theta'\|_2\|\theta\|_2^K\|\tanh\|_{C^K(\mathbb{R})}\\ &\leq 2^{K-1}(K+2)!(D+1)\|\theta-\theta'\|_2\|\theta\|_2^K.\end{aligned} \tag{1.16}$$

Now we want to bound the right part of (1.15), we use inequality (1.12) and substituting $u, v$ by $\mathcal{A}_1, \mathcal{A}_1'$, respectively. In this case, the maximum $S\in P$ occurs when $P = \{\{1\},\ldots,\{K\}\}$ because $\mathcal{A}_1$ is a linear function and, hence, higher order derivatives annuls. On the other hand, by (1.13), recall that $\|\mathcal{A}_1 - \mathcal{A}_1'\|_{\infty,\Omega} \le \|\theta - \theta'\|(1 + d_1 M(\Omega))$. Moreover, we notice that if $|\alpha| = 1$, $\|\partial^\alpha(\mathcal{A}_1 - \mathcal{A}_1')\|_{\infty,\Omega} \le \|W_1 - W_1'\|_\infty \le \|\theta - \theta'\|_2$. We notice that $\|\mathcal{A} - \mathcal{A}_1'\|_{C^K(\Omega)} = \|\mathcal{A} - \mathcal{A}_1'\|_{C^1(\Omega)} \le \|\theta - \theta'\|(1 + d_1 M(\Omega))$. Finally, we notice that $\|\partial^{\alpha(i)}\mathcal{A}_1\|_{\infty,\Omega} \le \|\theta\|_2$ for all $i = 1, \dots, K$. Therefore, we count bound (1.11) as follows

$$
\begin{aligned}
\|\partial^\alpha(\tanh \circ \mathcal{A}_1 - \tanh \circ \mathcal{A}_1')\|_{\infty,\Omega} &\le (K+1)(\|\tanh^{(K+1)}\|_\infty \|\mathcal{A}_1 - \mathcal{A}_1'\|_{\infty,\Omega} + \|\mathcal{A}_1 - \mathcal{A}_1'\|_{C^K(\Omega)}) \\
&\quad \prod_{S \in \{\{1\},\dots,\{K\}\}} \max(\|\partial^{\alpha(B)}\mathcal{A}_1\|_{\infty,\Omega}, \|\partial^{\alpha(B)}\mathcal{A}_1'\|_{\infty,\Omega}) \\
&\le (K+1)(2^K(K+3)!\|\theta - \theta'\|_2(1 + d_1 M(\Omega)) + \|\theta - \theta'\|_2(1 + d_1 M(\Omega))) \\
&\quad \max(\|\theta\|_2, \|\theta'\|_2)^K \\
&\le (K+1)\|\theta - \theta'\|_2(1 + d_1 M(\Omega))2^{2K-1}(K+2)!(K+3)! \max(\|\theta\|_2, \|\theta'\|_2)^K.
\end{aligned}
$$

Therefore, we can combine the previous equation with (1.15) and (1.16) and obtain

$$
\|\partial^\alpha(u_\theta - u_{\theta'})\|_{\infty,\Omega} \le \tilde{C}_{K,1}(1 + d_1 M(\Omega))(D+1)(1 + \max(\|\theta\|_2, \|\theta'\|_2))^{K+1}\|\theta - \theta\|_2
$$

with $\tilde{C}_{K,1} = (K+2)2^{2K-1}(K+2)!(K+3)!$ Hence, it satisfies the hypothesis for $\|u_\theta - u_{\theta'}\|_{C^K(\Omega)}$.

By induction, we assume that for a fixed $H \ge 1$, all $v_\theta, v_{\theta'} \in \mathrm{NN}_H^D$

$$
\|v_\theta - v_{\theta'}\|_{C^K(\Omega)} \le \tilde{C}_{K,H}(1 + d_1 M(\Omega))(D+1)^{H+KH^2}(1 + \max(\|\theta\|_2, \|\theta'\|_2)^{H+KH^2}\|\theta - \theta'\|_2. \tag{1.17}
$$

Let $u_\theta, u_{\theta'} \in \mathrm{NN}_{H+1}^D$. In other words, $u_\theta = \mathcal{A}_{H+2} \circ \tanh \circ v_\theta$ and $u_{\theta'} = \mathcal{A}_2' \circ \tanh \circ v_{\theta'}$. Let us proceed as in aforementioned lemma,

$$
\begin{aligned}
&\|\partial^\alpha(u_\theta - u_{\theta'})\|_{\infty,\Omega} \\
&\quad \le \|(W_{H+2} - W_{H+2}')\partial^\alpha(\tanh \circ v_\theta)\|_{\infty,\Omega} + \|W_{H+2}'\partial^\alpha(\tanh \circ v_\theta - \tanh \circ v_{\theta'})\|_{\infty,\Omega} \qquad (1.18) \\
&\quad \le D\|\theta - \theta'\|_2\|\partial^\alpha(\tanh \circ v_\theta)\|_{\infty,\Omega} + D\|\theta'\|_2\|\partial^\alpha(\tanh \circ v_\theta - \tanh \circ v_{\theta'})\|_{\infty,\Omega}.
\end{aligned}
$$

We will bound both summation terms. The first one can be bounded by (1.9) in Lemma 1.6,

$$
\|\partial^\alpha(\tanh \circ v_\theta)\|_{\infty,\Omega} \le C_{K,H+1}(D+1)^{1+K(H+1)}(1 + \|\theta\|_2)^{K(H+1)}\|\theta\|_2.
$$

The right term can be bounded using (1.11) as follows

$$\partial^\alpha(\tanh(v_\theta) - \tanh(v_{\theta'})) \le \sum_{P\in\Pi(K)} (|P|+1)(\|\tanh^{(|P|)}\|_{\mathrm{Lip}}\|v_\theta - v_{\theta'}\|_{\infty,\Omega} + \|u-v\|_{C^K(\Omega)})$$
$$\|\tanh^{(|P|)}\|_\infty \prod_{S\in P} \max(\|\partial^{\alpha(S)}v_\theta\|_{\infty,\Omega}, \|\partial^{\alpha(S)}v_{\theta'}\|_{\infty,\Omega})$$
$$\le (K+1)2^K(K+3)!2^{K-1}(K+2)!2\|v_\theta - v_{\theta'}\|_{C^K(\Omega)}$$
$$\sum_{P\in\Pi(K)}\prod_{S\in P} \max(\|\partial^{\alpha(S)}v_\theta\|_{\infty,\Omega}, \|\partial^{\alpha(S)}v_{\theta'}\|_{\infty,\Omega})$$
$$\le (K+1)2^{2K}(K+3)!(K+2)!\|v_\theta - v_{\theta'}\|_{C^K(\Omega)}B_k$$
$$\max_{i_1+2i_2+Ki_K=K}\prod_{1\le\ell\le K}[C_{\ell,H}(D+1)^{1+\ell H}\max\{(1+\|\theta\|_2)^{\ell H}\|\theta\|_2, (1+\|\theta'\|_2)^{\ell H}\|\theta'\|_2\}]^{i_\ell}$$
$$\le (K+1)2^{2K-1}(K+3)!(K+2)!\|v_\theta - v_{\theta'}\|_{C^K(\Omega)}$$
$$C_{K,H+1}(D+1)^{K+KH}(1+\max(\|\theta\|_2, \|\theta'\|_2)^{K+KH}.$$

We can bound the term $\|v_\theta - v_{\theta'}\|_{C^K(\Omega)}$ using the induction hypothesis (1.17). Hence, we can combine the previous equations inside (1.18) and obtain the desired result,

$$\|u_\theta - u_{\theta'}\|_{C^K(\Omega)} \le \tilde{C}_{K,H+1}(1+d_1M(\Omega))(D+1)^{(H+1)+K(H+1)^2}$$
$$(1+\max(\|\theta\|_2, \|\theta'\|_2))^{(H+1)+K(H+1)^2}\|\theta-\theta'\|_2.$$

□

**Lemma 1.8** (Lipschitz dependence of input of the Hölder norm in NN)**.** Consider the class $\mathrm{NN}_H^D = \{u_\theta, \theta\in\Theta_{H,D}\}$. Fix $K\in\mathbb{N}$. Then, there exists a constant $\tilde{C}_{K,H} > 0$, depending only on the regularity $K$ and the number of hidden layers $H$, such that

$$\|u_\theta(x) - u_\theta(y)\|_2 \le \tilde{C}_{H,\theta}\|x-y\|_2 \qquad \forall x,y\in\Omega.$$

where $\tilde{C}_{H,\theta}$ is a constant that depends on the number of layers $H$ and the parameters $\theta\in\Theta$.

*Proof.* We know that the Lipschitz constant of the composition of two functions, e.g., $f\circ g$ with Lipschitz constants $L_f$ and $L_g$ respectively is given by $L_fL_g$ due to

$$\|f(g(x)) - f(g(y))\|_2 \le L_f\|g(x)-g(y)\|_2 \le L_fL_g.$$

Since, we have defined the neural network as a composition of funtions $u_\theta = \mathcal{A}_{H+1}\circ(\phi\circ\mathcal{A}_H)\circ\cdots\circ(\phi\circ\mathcal{A}_1)$. Then,

$$\|u_\theta(x) - u_\theta(y)\|_2 \le \|\tanh\|_{\mathrm{Lip}}\prod_{k=1}^{H+1}\|\mathcal{A}_k\|_{\mathrm{Lip}}\|x-y\|_2.$$

We can easily see that for each affine function $\mathcal{A}_i$ the Lipschitz function is given by the norm of the matrix

$$\|\mathcal{A}_k(x) - \mathcal{A}_k(y)\|_2 = \|W_k x + b_k - (W_k y + b_k)\| \leq \|W_k\|_2 \|x - y\|_2 \leq \|\theta\|_2 \|x - y\|_2.$$

On the other hand, tha activation function applied element-wise verifies that

$$|\tanh(x_j) - \tanh(y_j)| \leq \sup_z |\tanh'(z)||x_j - y_j|.$$

Since the derivative of tanh is maximized at $z = 0$ with value 1, then $\|\phi(x) - \phi(y)\|_2 \leq 1 \cdot \|x - y\|_2$. Hence we obtain the intended result

$$\|u_\theta(x) - u_\theta(y)\|_2 \leq \|\theta\|_2^{H+1} \|x - y\|_2.$$

where we define $\tilde{C}_{H,\theta} := \|\theta\|_2^{H+1}$. □

## 1.2. Physics-Informed Neural Networks (PINNs)

While many machine learning models are purely data-driven, their performance can be significantly enhanced by incorporating domain knowledge. This is particularly true in scientific applications where data may be scarce, but the underlying physical principles are well understood. For instance, a system might be governed by known differential equations, be subject to physical invariants, or have other structural constraints. Ignoring such valuable information is inefficient and can harm the model's predictive accuracy. This prior knowledge can be integrated into a model in several ways, often categorized as different types of biases Watson et al. [2024], Hao et al. [2022].

*Observational bias* is introduced through the data itself. By collecting data from a system that inherently follows certain physical laws, a model trained on this data will learn to mimic the underlying structure.

*Inductive bias* refers to the prior assumptions built into a model's architecture. Specific architectures are chosen because they are guaranteed to respect certain properties. For example: Convolutional Neural Networks are designed with translation equivariance, making them ideal for image processing; Graph Neural Networks incorporate permutation equivariance, suitable for unstructured data like molecular graphs; Neural Ordinary Differential Equations use an architecture that inherently respects the structure of continuous-time dynamical systems.

*Learning bias* is embedded during training, typically by customizing the loss function to penalize solutions that violate known physical properties. A prominent example of this approach is the Physics-Informed Neural Network (PINN). PINNs are neural networks that are trained not only to fit observed data but also to obey the laws of physics described by a set of differential equations. This is achieved by adding a term to the loss function that measures how well the network's output satisfies these equations.

In this section, we define the learning framework for Physics-Informed Neural Networks (PINNs). We present the associated error formulation and examine some pathological behaviors that motivate the use of alternative loss functions.

**Notation** The PINN framework is typically formulated as an empirical risk minimization problem constrained by a partial differential equation (PDE). Throughout this manuscript, let $\Omega \subset \mathbb{R}^{d_1}$ denote a bounded Lipschitz domain with boundary $\partial\Omega$. Let $(X, Y) \in \Omega \times \mathbb{R}^{d_2}$ be a pair of random variables. A function $h : E \subset \mathbb{R}^{d_1} \to \mathbb{R}^{d_2}$ encodes the initial and/or boundary conditions of the PDE. The differential operator $\mathfrak{F} : C^\infty(\Omega, \mathbb{R}^{d_2}) \times \Omega \to \mathbb{R}$ characterizes the PDE itself.

**Example 1.9** (Poisson equation)**.** Let $\Omega = B(0, 1) \subset \mathbb{R}^2$, $f : \Omega \to \mathbb{R}$, $h : E = \partial\Omega \to \mathbb{R}$. Then, our goal is to find $u : \Omega \to \mathbb{R}$ such that

$$\begin{cases} -\Delta u = f & \text{in } \Omega \\ u = h & \text{in } \partial\Omega. \end{cases}$$

In this case, the differential operator is defined as $\mathfrak{F}(u, \mathbf{x}) := \Delta u(\mathbf{x}) + f(\mathbf{x}) \equiv \partial^2_{1,1} u(\mathbf{x}) + \partial^2_{2,2} u(\mathbf{x}) + f(\mathbf{x})$.

In hybrid modeling, we build upon a classical regression framework to estimate an unknown function $u^\star$ The underlying statistical model is given by:

$$Y = u^*(X) + \epsilon$$

where $Y$ is the response variable, $X$ is the vector of input variables and $\epsilon$ is a random noise term with a conditional mean of zero, i.e., $\mathbb{E}[\epsilon|X] = 0$. The novelty of PINNs lies in constraining the function $u^\star$ with a prior physical or theoretical knowledge. This knowledge is incorporated as a set of $M > 1$ PDE constaints of order at most $K$. The function $u^\star$ is assumed to satisfy these constraints, at leas approximataly

$$\mathfrak{F}_k(u^\star, x) \approx 0 \qquad \text{for } k = 1, \dots, M.$$

A direct consequence of this is that the function $u^\star$ must be at least $K$ times differentiable.

Furthermore, the function is subject to initial or boundary conditions. On a given subset of the domain's boundary, $E \subseteq \partial\Omega$, the function's values are approximated by a known function $h : E \to \mathbb{R}$. This is expressed as:

$$u^\star(x) \approx h(x) \qquad \text{for all } x \in E.$$

The next section will focus on measuring the error of a solution to a problem like the Poisson equation. We want to quantify how well a proposed solution performs. A challenge with the true loss function is that it often cannot be calculated in practice. Therefore, instead of the true error function, we will use an empirical version.

## 1.3. Empirical Risk Function

One of the most common ways to measure error is with the mean squared error (MSE) function or its variants. The MSE is widely used because it's convex and differentiable, allowing for gradient-based optimization methods like gradient descent. Furthermore, both the function and its derivative are computationally efficient.

To define the empirical risk for a partial differential equation (PDE) problem, we aim to minimize a loss function composed of three parts: (i) Data error: This component ensures the solution fits any available measurement data. (ii) Boundary error: This enforces the solution's adherence to the specified boundary conditions. (iii) PDE residual error: This component penalizes deviations from the PDE itself, ensuring the solution satisfies the governing equation. Hereafter we define in detail each one of the compontents:

(i) A collection of $n_1 \in \mathbb{N}$ i.i.d. random variables $(X_1, Y_1), \ldots, (X_{n_1}, Y_{n_1})$ distributed as $(X, Y) \in \Omega \times \mathbb{R}^{d_2}$, where the distribution is unknown. The partial risk function is defined by

$$\widehat{L}_1(u_\theta) := \frac{1}{n_1} \sum_{i=1}^{n_1} \|u_\theta(X_i) - Y_i\|_2^2. \qquad \text{(Data Conditions)}$$

(ii) A collection of $n_2 \in \mathbb{N}$ i.i.d. random variables $X_1^{(2)}, \ldots, X_{n_2}^{(2)}$ distributed according to some known distribution $\mu_E$ on $E$ where is the set of the boundary/initial conditions. Using the Dirichlet boundary condition, the partial risk function is defined by

$$\widehat{L}_2(u_\theta) := \frac{1}{n_2} \sum_{i=1}^{n_2} \|u_\theta(X_j^{(2)}) - h(X_j^{(2)})\|_2^2. \quad \text{(Initial/Boundary Conditions Type I)}$$

Likewise, if we consider the Neumann boundary condition

$$\widehat{L}_2(u_\theta) := \frac{1}{n_2} \sum_{i=1}^{n_2} \left\| \frac{\partial u_\theta}{\partial \mathbf{n}}(X_j^{(2)}) - h(X_j^{(2)}) \right\|_2^2 \quad \text{(Initial/Boundary Conditions Type II)}$$

(iii) A sample of $n_3 \in \mathbb{N}$ i.i.d. random variables $X_1^{(3)}, \ldots, X_{n_3}^{(3)}$ uniformly distributed on $\Omega$. The partial risk function is defined by

$$\widehat{L}_3(u_\theta) := \frac{1}{n_3} \sum_{k=1}^{M} \sum_{\ell=1}^{n_3} \mathfrak{F}_k(u_\theta, X_\ell^{(3)})^2. \qquad \text{(PDEs Equation)}$$

The solution $u$ is estimated by minimizing the empirical risk function, which is a linear combination of the individual risk functions defined above.

$$\widehat{R}_{n_1,n_2,n_3}(u_\theta) := \lambda_1 \widehat{L}_1(u_\theta) + \lambda_2 \widehat{L}_2(u_\theta) + \lambda_3 \widehat{L}_3(u_\theta). \tag{1.19}$$

Here, $\lambda_1, \lambda_2, \lambda_3$ establishes a trade-off between the three terms. In practice, we usually set $\lambda_1 = 1$ to avoid redudancy. And sometimes, one may encounter the case where $\lambda_3 = 0$ (data + PDEs) or $\lambda_2 = 0$ (PDEs + initial/boundary condition). Our goal is to find a minimizer of the empirical risk function 1.19, however, it does not necessarily exist. Instead, we use a minimizing sequence, which we denote $\{\hat{\theta}(n, n_1, n_2, n_3, D)\}_{n\in\mathbb{N}} \subseteq \Theta_{H,D}$, that is,

$$\lim_{n\to\infty} \widehat{R}_{n_1,n_2,n_3}(u_{\hat{\theta}(n,n_1,n_2,n_3,D)}) = \inf_{\theta\in\Theta_{H,D}} \widehat{R}_{n_1,n_2,n_3}(u_\theta).$$

# 1.4. Theoretical Risk Function

**Theoretical Risk Function** We use the theoretical risk function to measure the generalization of the solution. This function does not use samples, instead, it uses the continuous counterpart, e.g., the expected value and the integral.

(i) The theoretical risk function for the data is defined similarly to the empirical risk function counterpart

$$L_1(u_\theta) := \frac{1}{n_1}\sum_{i=1}^{n_1} \|u_\theta(X_i) - Y_i\|_2^2. \qquad \text{(Data Conditions)}$$

The rationale for holding $n_1$ fixed is that we usually have a limited ability to measure the number of samples.

(ii) The theoretical risk function for the initial/boundary condition uses the expectation value instead and is defined by

$$L_2(u_\theta) := \mathbb{E}_{X\sim\mu_E}\|u_\theta(X) - h(X)\|_2^2. \qquad \text{(Initial/Boundary Conditions Type I)}$$

This regime corresponds to letting $n_2 \to \infty$, which can usually be sampled freely up to computational resources. Likewise, the Neumann boundary condition is expressed as

$$L_2(u_\theta) := \mathbb{E}_{X\sim\mu_E}\left\|\frac{\partial u_\theta}{\partial \mathbf{n}}(X) - h(X)\right\|_2^2. \qquad \text{(Initial/Boundary Conditions Type II)}$$

(iii) The theoretical risk function for the PDE equation is defined by

$$L_3(u_\theta) := \frac{1}{|\Omega|}\sum_{k=1}^{M}\int_\Omega \mathfrak{F}_k(u_\theta, x)^2 dx. \qquad \text{(PDEs Equation)}$$

Similarly, we let $n_3 \to \infty$ and, because it is uniformly distributed, we use the integral.

The solution $u$ is estimated by minimizing the empirical risk function, which is a linear combination of the individual risk functions defined above.

$$R_{n_1}(u_\theta) := \lambda_1 L_1(u_\theta) + \lambda_2 L_2(u_\theta) + \lambda_3 L_3(u_\theta). \tag{1.20}$$

**Remark 1.10.** We can note that for a fixed $u_\theta \in \mathrm{NN}_{H,D}$, the expectation of empirical risk is equal to the expectation of the theoretical risk function, i.e.,

$$\mathbb{E}[\widehat{R}_{n_1,n_2,n_3}(u_\theta)] = \mathbb{E}[R_{n_1}(u_\theta)].$$

*Proof.* We will see that each one of the components verifies that the expectation is centered

$$\begin{aligned}
&\mathbb{E}[\widehat{L}_1(u_\theta)] = \mathbb{E}[L_1(u_\theta)] \\
&\mathbb{E}[\widehat{L}_2(u_\theta)] = \mathbb{E}_{X\sim\mu_E}\|u_\theta(X) - h(X)\|_2^2 = L_2(u_\theta) = \mathbb{E}[L_2(u_\theta)] \\
&\mathbb{E}[\widehat{L}_3(u_\theta)] = \sum_{k=1}^{M} \mathbb{E}_{X\sim U(\Omega)}[\mathfrak{F}_k(u_\theta, X)] = \sum_{k=1}^{M} \frac{1}{\Omega} \int_\Omega \mathfrak{F}_k(u_\theta, x)dx = L_3(u_\theta) = \mathbb{E}[L_3(u_\theta)].
\end{aligned}$$

□

The cornerstone in PINNs that differs from other learning frameworks is the error function.

# CHAPTER 2
# Stochastic Processes

In this section, we present several results used to establish the consistency of the regularized loss for PINNs. Most of these results are drawn from Van Handel [2014] and Doumèche et al. [2024].

Let us recall the definitions of some key concepts and clarify the notation. We denote $(\Omega, \mathcal{A}, \mathbb{P})$ as a probability space. A *random* or *stochastic process* is a collection of random variables $\{X_t : t \in T\}$ indexed by some set $T$. In most cases, the index set $T$ is some subset of the real line, such as the natural numbers or an interval, giving the set $T$ the interpretation of time. However, $T$ can take more general forms, such as a metric space. In this manuscript, we will consider both cases. When $T$ is an ordered set, we say that the random process $\{X_t\}_{t\in T}$ is adapted to a filtration $\{\mathcal{F}_t\}_{t\in T}$ when that the random variable $X_t$ is measurable with respect to $\mathcal{F}_t$ where the filtration is formed by a monotone chain of sub-$\sigma-$algebras, i.e., $\mathcal{F}_s \subseteq \mathcal{F}_t \subseteq \mathcal{A}$ when $s \leq t$. The *natural* filtration is defined as the $\sigma-$algebra generated the previous random variables, i.e., $\mathcal{F}_t = \sigma(\{X_s\}_{s\leq t})$.

We will study the family of random variables whose tail behavior and concentration properties resemble those of a Gaussian distribution—namely, the *sub-gaussian random variables*. The moment-generating function of a centered normal distribution $X \sim \mathcal{N}(0, \sigma^2)$ is

$$M_X(\lambda) = \mathbb{E}[e^{\lambda X}] = e^{\sigma^2\lambda^2/2}.$$

Subgaussian random variables are the class of random variables whose moment-generating function is bounded by that of a normal distribution. This leads to the following definition.

## 2.1. Subgaussian Random Variables

**Definition 2.1** (Subgaussian random variable)**.** Let $X$ be a random variable, we define the *log-moment generating function* $\psi$ of $X$ as

$$\psi(\lambda) := \log \mathbb{E}[e^{\lambda(X-\mathbb{E}[X])}]$$

and we say that $X$ random variable is $\sigma^2-$subgaussian if its log-moment generating function satisfies $\psi(\lambda) \leq \lambda^2\sigma^2/2$ for all $\lambda \in \mathbb{R}$, and the smallest $\sigma^2$ for which that holds is called the variance proxy of $X$.

The following result is a variant of the Markov's inequality written in terms of the log-moment generating function.

**Lemma 2.2** (Chernoff bound)**.** Let $X$ be a $\sigma^2-$subgaussian random variable. Then, the following inequality holds

$$\mathbb{P}[X - \mathbb{E}X \geq a] \leq \exp\left\{\frac{-a^2}{2\sigma^2}\right\}.$$

In particular, $\mathbb{P}[|X - \mathbb{E}X| \geq a] \leq 2\exp\{-a^2/2\sigma^2\}$.

*Proof.* We exponentiate inside the probability before appliying Markov's inequality. For any $\lambda \geq 0$, we have

$$\mathbb{P}[X - \mathbb{E}X \geq a] = \mathbb{P}[e^{\lambda(X-\mathbb{E}X)} \geq e^{\lambda a}] \leq e^{-\lambda a}\mathbb{E}[e^{\lambda(X-\mathbb{E}X)}] = e^{\psi(\lambda)-\lambda a}.$$

Since this holds for every $\lambda \in \mathbb{R}$, we can choose $\lambda$ to obtain the best bound. The optimal choice is $\lambda = \frac{a}{\sigma^2}$. Substituting this into the inequality and using the bound $\psi(\lambda) \leq \lambda^2\sigma^2/2$, we obtain

$$\mathbb{P}[X - \mathbb{E}X \geq a] \leq \exp\left\{\left(\frac{a}{\sigma^2}\right)^2\sigma^2/2 - \left(\frac{a}{\sigma^2}\right)a\right\} = \exp\left\{\frac{-a^2}{2\sigma^2}\right\}.$$

□

**Lemma 2.3** (Maximal inequality)**.** Let $\{X_k\}_{1\leq k\leq n}$ be a collection of $\sigma^2$-subgaussian random variables with the same variance, satisfying $\mathbb{E}[X_k] = 0$ for all $k = 1, \ldots, n$. Then

$$\mathbb{E}\left[\sup_{1\leq k\leq n} X_k\right] \leq \sqrt{2\sigma^2\log n}.$$

*Proof.* Since $-\log x$ is convex, by Jensen's inequality, we have for any $\lambda > 0$

$$\begin{aligned}\mathbb{E}\left[\sup_k X_k\right] &= \mathbb{E}\left[\frac{1}{\lambda}\log(e^{\lambda\sup_k X_k})\right] \leq \frac{1}{\lambda}\log\mathbb{E}[e^{\lambda\sup_k X_k}] \\ &\leq \frac{1}{\lambda}\log\sum_{1\leq k\leq n}\mathbb{E}[e^{\lambda X_k}] \leq \frac{1}{\lambda}\log(ne^{\psi(\lambda)}) \\ &= \frac{\log n + \psi(\lambda)}{\lambda} \leq \frac{\log n + \lambda^2\sigma^2/2}{\lambda}.\end{aligned}$$

Since this holds for every $\lambda > 0$, we can now optimize over $\lambda$ on the right hand side. Differentiating and setting it equal to zero, we obtain the minimum at $\lambda = \frac{\sqrt{2\log n}}{\sigma}$. Substituting this value into the equation, we obtain the desired result

$$\mathbb{E}\left[\sup_{1\leq k\leq n} X_k\right] \leq \sqrt{2\sigma^2\log n}.$$

□

**Lemma 2.4** (Hoeffding lemma)**.** Let $X$ be a random variable such that $a \leq X \leq b$ a.s. for some $a, b \in \mathbb{R}$. Then,

$$\mathbb{E}[e^{\lambda(X-\mathbb{E}[X])}] \leq e^{\lambda^2(b-a)^2/8}.$$

In other words, $X$ is a $(b-a)^2/4$-subgaussian random variable.

*Proof.* Let us define the mean-centered random variable $Y := X - \mathbb{E}[X]$, then $\tilde{a} \leq Y \leq \tilde{b}$ where $\tilde{a} := a - \mathbb{E}[X]$ and $\tilde{b} := b - \mathbb{E}[X]$. By definition, log-moment generating function of $X$ is given by $\psi(\lambda) = \log \mathbb{E}[e^{\lambda Y}]$, along with its derivatives

$$\psi'(\lambda) = \frac{\mathbb{E}[Ye^{\lambda Y}]}{\mathbb{E}[e^{\lambda Y}]}, \qquad \psi''(\lambda) = \frac{\mathbb{E}[Y^2 e^{\lambda Y}]}{\mathbb{E}[e^{\lambda Y}]} - \left(\frac{\mathbb{E}[Ye^{\lambda Y}]}{\mathbb{E}[e^{\lambda Y}]}\right)^2.$$

We can interpret $\psi''(\lambda)$ as the variance of the random variable $Y$ under another probability measure. We define anew probability measure $\mathbb{Q}$ by setting $d\mathbb{Q} := \frac{e^{\lambda Y}}{\mathbb{E}[e^{\lambda Y}]} d\mathbb{P}$. The definition is well-posed because the Radon-Nikodym derivative $\frac{e^{\lambda Y}}{\mathbb{E}[e^{\lambda Y}]}$ is positive and integrates to 1. From this, we deduce that

$$\mathbb{E}_{\mathbb{Q}}[g(Y)] = \frac{\mathbb{E}[g(Y)e^{\lambda Y}]}{\mathbb{E}[e^{\lambda Y}]}.$$

Therefore,

$$\text{Var}_{\mathbb{Q}}(Y) = \mathbb{E}_{\mathbb{Q}}[X^2] - (\mathbb{E}_{\mathbb{Q}}[Y])^2 = \frac{\mathbb{E}[Y^2 e^{\lambda Y}]}{\mathbb{E}[e^{\lambda Y}]} - \left(\frac{\mathbb{E}[Ye^{\lambda Y}]}{\mathbb{E}[e^{\lambda Y}]}\right)^2 = \psi''(\lambda).$$

We can bound the variance of $Y$ as follows

$$\text{Var}_{\mathbb{Q}}(Y) = \text{Var}_{\mathbb{Q}}(Y - z) \leq \mathbb{E}_{\mathbb{Q}}[(Y-z)^2].$$

Since $z$ can be any real number, we let $z = (\tilde{a} + \tilde{b})/2$. Given that $\tilde{a} \leq Y \leq \tilde{b}$ a.s., we obtain

$$\mathbb{E}_{\mathbb{Q}}[(Y-z)^2] = \frac{1}{4}\mathbb{E}_{\mathbb{Q}}[(2Y - \tilde{a} - \tilde{b})^2] \leq \frac{(\tilde{b} - \tilde{a})^2}{4} = \frac{(b-a)^2}{4}.$$

Next, since $\psi(\lambda) \leq \lambda^2\sigma^2/2$, we observe that $\psi'(0) = \psi(0) = 0$. Using the bound that $\psi''(\lambda) \leq \frac{(b-a)^2}{4}$ along with the Fundamental Theorem of Calculus, we obtain

$$\psi(\lambda) = \int_0^\lambda \int_0^\mu \psi''(\rho) d\rho d\mu \leq \frac{\lambda^2(b-a)^2}{8}.$$

Moreover, $X$ is a $\sigma^2-$sbugaussian variable with $\sigma^2 = (b-a)^2/4$. □

**Lemma 2.5** (Azuma)**.** Let $\{X_k\}_{1\leq k\leq n}$ be a stochastic process adapted to the natural filtration $\{\mathcal{F}_k\}_{1\leq k\leq n}$, i.e., $\mathcal{F}_k = \sigma(X_1, \ldots, X_k)$. Assume that the random variables satisfy:

$$\mathbb{E}[e^{\lambda X_k}|\mathcal{F}_{k-1}] \leq e^{\lambda^2\sigma_k^2/2} \quad a.s. \qquad k = 1, \ldots, n \tag{2.1}$$

where $\sigma_k^2$ is the variance of $X_k$. Then, $\mathbb{E}[X_k] = 0$ and the sum $\sum_{k=1}^n X_k$ is a $\sigma^2-$subgaussian with variance proxy $\sigma^2 := \sum_{k=1}^n \sigma_k^2$.

*Proof.* The tower property of expectation assures that $\mathbb{E}[X] = \mathbb{E}[\mathbb{E}[Y|\mathcal{F}_{k-1}]]$. For any $1 \le k \le n$ the random variables $X_1, \ldots, X_{k-1}$ are $\mathcal{F}_{k-1}-$measurables. Then, it follows

$$\mathbb{E}[e^{\lambda \sum_{i=1}^{k} X_i}] = \mathbb{E}[\mathbb{E}[e^{\lambda \sum_{i=1}^{k-1} X_i} e^{\lambda X_k}|\mathcal{F}_{k-1}]] = \mathbb{E}[e^{\lambda \sum_{i=1}^{k-1} X_i}\mathbb{E}[e^{\lambda X_k}|\mathcal{F}_{k-1}]] \overset{\text{eq. (2.1)}}{\le} e^{\lambda^2\sigma_k^2/2}\mathbb{E}[e^{\lambda \sum_{i=1}^{k-1} X_i}].$$

By induction, it yields $\mathbb{E}[e^{\lambda \sum_{i=1}^{n} X_i}] \le e^{\lambda^2 \sum_{i=1}^{n} \sigma_i^2/2} = e^{\lambda^2\sigma^2/2}$. Therefore, assuming that $\mathbb{E}[X_k] = 0$, it follows

$$\psi(\lambda) = \log \mathbb{E}[e^{\lambda \sum_{i=1}^{n} X_i}] \le \lambda^2\sigma^2/2.$$

Let us prove that $\mathbb{E}[X_k|\mathcal{F}_{k-1}] = 0$ (which, by the tower property, implies that $\mathbb{E}[X_k] = 0$). We observe that equality holds in (2.1) when $\lambda = 0$. Differentiating both sides with respect to $\lambda$ and evaluating at $\lambda = 0$ must yield the same result. This leads to the condition:

$$\frac{d}{d\lambda}\Big|_{\lambda=0} \mathbb{E}\left[e^{\lambda X_k}|\mathcal{F}_{k-1}\right] \overset{DCT}{=} \mathbb{E}\left[\frac{d}{d\lambda}\Big|_{\lambda=0} e^{\lambda X_k}|\mathcal{F}_{k-1}\right] = \frac{d}{d\lambda}\Big|_{\lambda=0} e^{\lambda^2\sigma_k^2/2} \Rightarrow \mathbb{E}[X_k|\mathcal{F}_{k-1}] = 0.$$

□

**Corollary 2.6** (Azuma-Hoeffding inequality)**.** Let $\{\mathcal{F}_k\}_{1\le k\le n}$ be a filtration, and let $X_k, A_k, B_k$ be stochastic processes satisfying that $A_k, B_k$ are $\mathcal{F}_{k-1}-$measurable and $A_k \le X_k \le B_k$ a.s. for $k = 1, \ldots, n$. Then, $\sum_{k=1}^{n} X_k$ is $\sigma^2-$subgaussian with $\sigma^2 = \frac{1}{4}\sum_{k=1}^{n} \|B_k - A_k\|_\infty^2$. In particular, for every $t \ge 0$ the tail bound

$$\mathbb{P}\left[\sum_{k=1}^{n} X_k \ge t\right] \le \exp\left(-\frac{2t^2}{\sum_{k=1}^{n} \|B_k - A_k\|_\infty^2}\right).$$

*Proof.* Applying Hoeffding's Lemma 2.4 to $X_k$ conditionally on $\mathcal{F}_{k-1}$ yields

$$\mathbb{E}[e^{\lambda X_k}|\mathcal{F}_{k-1}] \le e^{\lambda^2(B_k-A_k)^2/8}.$$

In other words, $X_k$ is $\sigma_k^2-$subgaussian with $\sigma_k^2 = \frac{1}{4}(B_k - A_k)^2$. This follows because $A_k$ and $B_k$ are $\mathcal{F}_{k-1}$ measurable, allowing us to treat them as constants. On the other hand, applying Azuma's Lemma 2.5 to $X_k$ implies that $\mathbb{E}[X_k] = 0$ and that $\sum_{k=1}^{n} X_k$ is $\sigma^2-$subgaussian with $\sigma^2 = \sum_{k=1}^{n} \sigma_k^2 = \sum_{k=1}^{n} \frac{1}{4}(B_k - A_k)^2$.

To establish the second part, note that $\mathbb{E}X_k = 0$, therefore, invoking Chernoff Bound Lemma 2.2 leads to the desired tail bound. □

## 2.2. Covering and Packing number

We will define the covering and packing number of a metric space $(E, d)$. Intuitively, the *covering number* is the minimum number of balls required to completely cover the space, while the *packing number* is the maximum number of disjoint balls that can be placed within the space without overlapping (see Figure 2.1).

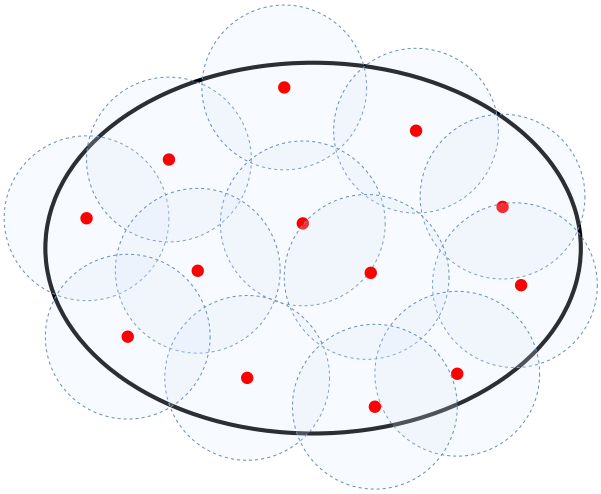

(a) The ellipse represents the $T$ and the dots in red are the the elements in $N$, i.e., the $\epsilon-$net. The circles around the dots are the balls of $\epsilon$ radius that should cover the metric space.

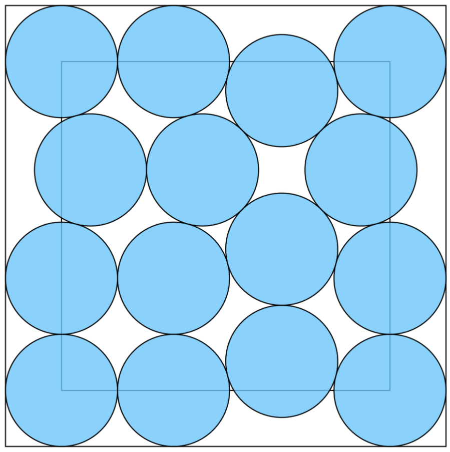

(b) The optimal packing problem in a square using the blue balls.

Figure 2.1: Packing vs Covering number

**Definition 2.7** ($\epsilon-$net and covering number)**.** A set $N \subseteq E$ is called a $\epsilon-$*net* for $(E, d)$ if for every $x \in E$, there exists $\pi(x) \in N$ such that $d(x, \pi(x)) \leq \epsilon$. The smallest cardinality of an $\epsilon-$net for $(E, d)$ is the called the *covering number*

$$N(E, d, \epsilon) := \inf\{|N| : N \text{ is an } \epsilon-\text{net for } (E, d)\}.$$

**Definition 2.8** ($\epsilon-$packing and packing number)**.** A set $N \subseteq E$ is called an $\epsilon-$*packing* of $(E, d)$ if $d(x, x') > \epsilon$ for every $x, x' \in N$, $x \neq x'$. The largest cardinality of an $\epsilon-$packing of $(E, d)$ is called the *packing number*

$$D(E, d, \epsilon) := \sup\{|N| : N \text{ is an } \epsilon-\text{packing of } (E, d).\}.$$

In the literature, the binary logarithm of the covering number of a set $E$ is commonly referred to as the $\boldsymbol{\epsilon}$**−entropy** of $E$, denoted by

$$H_\epsilon(E) := \log_2 N(E, \|\cdot\|, \epsilon).$$

Similary, the binary logarithm of the packing number is typically called the $\boldsymbol{\epsilon}$**−capacity** of $E$, defined as

$$C_\epsilon(E) := \log_2 D(E, \|\cdot\|, \epsilon).$$

The logarithm is usually taken with base 2 because it used in the context of Computer Science to measure the transmission properties of certain channel.

The following lemma illustrates the relationship between the covering number and the packing number.

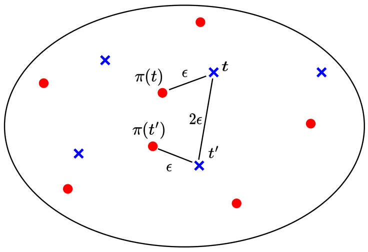


Figure 2.2: distance between $x$ and $x'$.

**Lemma 2.9** (Duality between covering and packing number)**.** For every $\epsilon > 0$

$$D(E, d, 2\epsilon) \leq N(E, d, \epsilon) \leq D(E, d, \epsilon).$$

*Proof.* (1) Let $D$ be a $2\epsilon-$packing and let $N$ be an $\epsilon-$net. For every $x \in D$, choose $\pi(x) \in N$ such that $d(x, \pi(x)) \leq \epsilon$. Then, for every $x' \in D$ such that $x \neq x'$, we have

$$2\epsilon < d(x, x') \leq d(x, \pi(x)) + d(\pi(x), \pi(x')) + d(\pi(x'), x') \leq 2\epsilon + d(\pi(x), \pi(x')),$$

which implies that $\pi(x) \neq \pi(x')$ (see Figure 2.2). Thus, the function $\pi : D \to N$ is injective, and thus, $|D| \leq |N|$. In other words, $D(E, d, 2\epsilon) \leq N(E, d, \epsilon)$.

(2) Let $D$ be a *maximal* $\epsilon-$packing with $|D| = D(E, d, \epsilon)$. We claim that $D$ is necessarily an $\epsilon-$net. Indeed, suppose for contradiction that there exists a point $x \in E$ such that $d(x, x') > \epsilon$ for every $x' \in D$. This would imply that $D \cup \{x\}$ is a *larger* $\epsilon-$packing, contradicting the maximality of $D$. Therefore, every point in $E$ must be within a distance at most of $\epsilon$ from some point in $D$, confirming that $D$ is an $\epsilon-$net. □

We are now ready to establish an upper bound on the covering number of the Euclidean ball $B_2^n$ with respect to the Euclidean distance. The proof of this fundamental result employs a clever technique known as a *volume argument*.

**Lemma 2.10.** Let $B_2^n$ be the $n-$dimensional euclidean ball centered at zero with radius 1, i.e., $B_2^n := \{x \in \mathbb{R}^n : \|x\|_2 < 1\}$. Then, we have that $N(B_2^n, \|\cdot\|, \epsilon) = 1$ for $\epsilon \geq 1$ and

$$\left(\frac{1}{\epsilon}\right)^n \leq N(B_2^n, \|\cdot\|, \epsilon) \leq \left(\frac{3}{\epsilon}\right)^n \qquad \text{for } 0 < \epsilon < 1.$$

*Proof.* Case $\epsilon \geq 1$: It is clear that for $\epsilon \geq 1$, $N(B_2^n, \|\cdot\|, \epsilon) = 1$ since $N := \{0\}$ is a $\epsilon-$net, i.e., $\|x - 0\| < 1 \leq \epsilon$ for every $x \in B_2^n$.

Case $0 < \epsilon < 1$: Let us begin with the upper bound. Let $D$ be a $2\epsilon-$packing of $B_2^n$. Since $d(x, x') > 2\epsilon$ for all $x \neq x'$ in $D$, the balls $\{B(x, \epsilon) : x \in D\}$ must be disjoint. On the other hand, every ball $B(x, \epsilon)$ for $x \in B_2^n$ must be cointained in a larger ball centered in zero with radius $1 + \epsilon$, i.e., $B(x, \epsilon) \subseteq B(0, 1 + \epsilon)$. Therefore, the sum of volumes of the balls satisfy

$$\sum_{x \in D} \lambda(B(x, \epsilon)) = \lambda\left(\bigcup_{x \in D} B(x, \epsilon)\right) \leq \lambda(B(0, 1 + \epsilon))$$

where $\lambda$ denotes the Lebesgue measure on $\mathbb{R}^n$. Using that the Lebesgue measure is homogeneous and that $\lambda(B(0,\alpha)) = \lambda(\alpha B(0,1)) = \alpha^n \lambda(B(0,1))$, then

$$\epsilon^n |D| B(0,1) = |D| B(0,\epsilon) = \sum_{x\in D} \lambda(B(x,\epsilon)) \leq \lambda(B(0,1+\epsilon)) = (1+\epsilon)^n B(0,1).$$

Therefore, dividing both sides by $B(0,1)$, we obtain that

$$|D| \leq \left(\frac{1+\epsilon}{\epsilon}\right)^n .$$

We have established that for every $2\epsilon-$packing $D$, we obtain an upper bound for $D(E,d,2\epsilon)$. This leads to the following chain of inequalities

$$N(B_2^n, \|\cdot\|, 2\epsilon) \overset{\text{Lemma 2.9}}{\leq} D(B_2^n, \|\cdot\|, 2\epsilon) \leq \left(1+\frac{1}{\epsilon}\right)^n \leq \left(\frac{3}{2\epsilon}\right)^n \qquad \text{for } 2\epsilon < 1.$$

Relabeling $2\epsilon$ as $\epsilon$ completes the proof.

We proceed similarly to obtain the lower bound. Let $N$ be an $\epsilon-$net for $B_2^n$. Then,

$$\lambda(B_2^n) \leq \lambda\left(\bigcup_{x\in N} B(x,\epsilon)\right) \leq \sum_{x\in N} \lambda(B(x,\epsilon)) = |N|\epsilon^n \lambda(B_2^n).$$

Hence, dividing by $\lambda(B_2^n)$, we obtain

$$|N| \geq \frac{\lambda(B_2^n)}{\lambda(B(0,\epsilon))} = \left(\frac{1}{\epsilon}\right)^n .$$

This inequality holds for every $\epsilon-$net $N$, so we conclude that $N(B_2^n, \|\cdot\|, \epsilon) \geq (1/\epsilon)^n$. □

The following Lemma finds a bound for the growing rate of the $\epsilon-$entropy of the Sobolev space $H^{m+1}$. The proof can be found in [Nickl and Pötscher, 2007, Corollary 4].

**Lemma 2.11** ($\epsilon$-Entropy of $H^{m+1}(\Omega, \mathbb{R}^{d_2})$)**.** Let $\Omega \subseteq \mathbb{R}^{d_2}$ be a Lipschitz domain. For $m \geq 1$, one has

$$\log N(B_{H^{m+1}(\Omega)}(1), \|\cdot\|_{H^{m+1}(\Omega)}, \epsilon) = \mathcal{O}_{\epsilon\to 0}(\epsilon^{-d_1/(m+1)}).$$

## 2.3. Subgaussian Process

**Definition 2.12** (Subgaussian process)**.** A random process $\{X_t\}_{t\in T}$ on a metric space $(T,d)$ is called *subgaussian* if $\mathbb{E}[X_t] = 0$ for all $t \in T$ and

$$\mathbb{E}[e^{\lambda(X_t - X_s)}] \leq e^{\lambda^2 d(t,s)^2/2} \qquad \text{for all } t,s \in T, \lambda \geq 0. \tag{2.2}$$

Notice that for every $s, t \in T$, the random variable $X_t - X_s$ is a $d(t,s)^2$-subgaussian random variable.

**Definition 2.13** (Separable process)**.** A random process $\{X_t\}_{t\in T}$ is called separable if there exists a countable set $T_0 \subseteq T$ such that

$$X_t \in \lim_{\substack{s\to t\\ s\in T_0}} X_s \qquad \text{for all } t \in T \quad a.s.$$

Here, $\lim_{\substack{s\to t\\ s\in T_0}} X_s \equiv \{X_t : \exists\{s_n\}_{n\in\mathbb{N}} \text{ such that } X_{s_n} \xrightarrow{n\to\infty} X_t \text{ a.s.}\}$

**Remark 2.14.** The assumption of separability is almost always satisfied. For example, assuming that $t \to X_t$ is continuous and $T$ is a separable metric space (as is the case in this manuscript), then we can take $T_0$ to be any countable dense subset of $T$, allowing us to verify the separability assumption.

For the next theorem, we must prove that $\sup_{t\in T} X_t$ is measurable. In fact, if $T$ is uncountable, the supremum is not necessarily measurable. However, under the assumption of separability, we have $\sup_{t\in T} X_t = \sup_{t\in T_0} X_t$. Since the supremum of a measurable function is always measurable, it follows that $\sup_{t\in T} X_t$ must be measurable.

**Theorem 2.15** (Dudley)**.** Let $\{X_t\}_{t\in T}$ be a separable subgaussian process in the metric space $(T, d)$. Then,

$$\mathbb{E}\left[\sup_{t\in T} X_t\right] \le 6 \sum_{k\in\mathbb{Z}} 2^{-k}\sqrt{\log N(T, d, 2^{-k})}.$$

*Proof.* We begin by proving the result for the case where $|T| < \infty$. Let $k_0 \in \mathbb{Z}$ be the largest integer such that $2^{-k_0} \ge \operatorname{diam}(T)$. It is clear that for every $t_0 \in T$, the set $N_0 := \{t_0\}$ forms a $2^{-k_0}-$net and $\pi_0(t) \equiv t_0$.

For $k > k_0$, let $N_k$ be a $2^{-k}-$net such that $|N_k| = N(T, d, 2^{-k})$. We denote $\pi_k(t)$ as the element in $N_k$ that satisfies $d(t, \pi_k(t)) \le 2^{-k}$. Using a chaining argument up to the scale $2^{-n}$, we proceed as follows

$$\begin{aligned}\mathbb{E}\left[\sup_{t\in T} X_t\right] &= \mathbb{E}\left[\sup_{t\in T}\left\{X_{\pi_0(t)} + \left(\sum_{k=k_0+1}^{n} X_{\pi_k(t)} - X_{\pi_{k-1}(t)}\right) + X_t - X_{\pi_n}(t)\right\}\right]\\ &\le \mathbb{E}[X_{t_0}] + \sum_{k=k_0+1}^{n} \mathbb{E}\left[\sup_{t\in T}\{X_{\pi_k(t)} - X_{\pi_{k-1}(t)}\}\right] + \mathbb{E}\left[\sup_{t\in T}\{X_t - X_{\pi_n(t)}\}\right].\end{aligned}$$

By definition of subgaussian process, $\mathbb{E}[X_{t_0}] = 0$. Since $|T| < \infty$, we can choose $n$ sufficiently large so that $N_n = T$, and hence $\pi_n(t) = t$, meaning that the third term vanishes. Next, we bound the second term. By definition, $X_{\pi_k(t)} - X_{\pi_{k-1}}(t)$ is a $d(\pi_k(t), \pi_{k-1}(t))$-subgaussian random variable. We can readily estimate the variance,

$$d(\pi_k(t), \pi_{k-1}(t)) \le d(\pi_k(t), t) + d(t, \pi_{k-1}(t)) \le 2^{-k} + 2^{-(k-1)} = 3 \times 2^{-k}.$$

Moreover, we can control the number of terms in the sum, note that $\{X_{\pi_k(t)} - X_{\pi_{k-1}(t)} : t \in T\}$ contains at most $|N_k||N_{k-1}|$ which is bounded by $|N_k|^2$ terms. Applying Maximal Inequality Lemma 2.3 to these terms, we obtain

$$\mathbb{E}\left[\sup_{t\in T} X_t\right] \leq \sum_{k=k_0+1}^{n} \sqrt{2d(\pi_k(t), \pi_{k-1}(t))^2 \log |N_k|^2} \leq 6 \sum_{k=k_0+1}^{n} 2^{-k}\sqrt{\log |N_k|}$$
$$\leq 6 \sum_{k=k_0+1}^{n} \sqrt{\log N(T, d, 2^{-k})}.$$

To prove the result when $|T|$ is infinite, we can use separability to approximate the general case by the finite case. Indeed, by separability, there is a countable set $T \subseteq T'$ such that $\sup_{t\in T} X_t = \sup_{t\in T'} X_t$ a.s. Let us denote $T'_k$ the first $k$ elements of $T'$. Then, by monote convergence of expectations

$$\mathbb{E}\left[\sup_{t\in T} X_t\right] = \mathbb{E}\left[\sup_{t\in T'} X_t\right] = \mathbb{E}\left[\sup_{k\geq 1}\sup_{t\in T'_k} X_t\right] = \sup_{k\geq 1}\mathbb{E}\left[\sup_{t\in T'_k} X_t\right].$$

And, applying the chaining inequality to each finite maximum and using that $N(T_k, d, \epsilon) \leq N(T, d, \epsilon)$ yields the same result. □

**Corollary 2.16** (Entropy integral)**.** Let $\{X_t\}_{t\in T}$ be a separable subgaussian process on the metric space $(T, d)$. Then,

$$\mathbb{E}\left[\sup_{t\in T} X_t\right] \leq 12 \int_0^\infty \sqrt{\log N(T, d, \epsilon)}d\epsilon.$$

*Proof.* Since $N(T, d, \cdot)$ is decreasing, we obtain the following chains of inequalities

$$\begin{aligned}\sum_{k\in\mathbb{Z}} 2^{-k}\sqrt{\log N(T, d, 2^{-k})} &= 2\sum_{k\in\mathbb{Z}} \int_{2^{-k-1}}^{2^{-k}} \sqrt{\log N(T, d, 2^{-k})}d\epsilon \\ &\leq 2\sum_{k\in\mathbb{Z}} \int_{2^{-k-1}}^{2^{-k}} \sqrt{\log N(T, d, \epsilon)}d\epsilon \\ &= 2\int_0^\infty \sqrt{\log N(T, d, \epsilon)}d\epsilon.\end{aligned}$$

□

**Definition 2.17** (Discrete derivative)**.** Let $f \in C(\mathbb{R}^n, \mathbb{R})$. We define the *discrete derivative* of $f$ with respect to variable $x_k$ at the point $x \in \mathbb{R}$ as follows:

$$\mathfrak{D}_k f(x) := \sup_z f(x_1, \ldots, x_{k-1}, z, x_{k+1}, \ldots, x_n) - \inf_z f(x_1, \ldots, x_{k-1}, z, x_{k+1}, \ldots, x_n).$$

**Theorem 2.18** (McDiarmid)**.** Let $X_1, \ldots, X_n$ be independent random variables, $f \in C(\mathbb{R}^n, \mathbb{R})$. Then, $f(X_1, \ldots, X_n)$ is $\sigma^2-$subgaussian with $\sigma^2 = \frac{1}{4}\sum_{k=1}^n \|\mathfrak{D}_k f\|_\infty^2$ where $\mathfrak{D}_k f$ is the discrete derivative (see Definition 2.17).

*Proof.* Let us define the following random processes for $k = 1, \ldots, n$

$$Y_k := \mathbb{E}[f(X_1, \ldots, X_n)|\mathcal{F}_k] - \mathbb{E}[f(X_1, \ldots, X_n)|\mathcal{F}_{k-1}].$$

where $\mathcal{F}_k$ is the natural filtration. We observe that it satisfies $\mathbb{E}[Y_k|\mathcal{F}_{k-1}] = 0$ due to the tower property. Moreover, we see that

$$\sum_{k=1}^n Y_k = f(X_1, \ldots, X_n) - \mathbb{E}f(X_1, \ldots, X_n). \tag{2.3}$$

On the other hand, let us denote $g_k(z) := f(X_1, \ldots, X_{k-1}, z, X_{k+1}, \ldots, X_n)$. Since $X_k$ is indpendent from $X_1, \ldots, X_{k-1}, X_{k+1}, \ldots, X_n$, it follows that

$$\mathbb{E}[\inf_z g_k(z)|\mathcal{F}_{k-1}] = \mathbb{E}[\inf_z g_k(z)|\mathcal{F}_k] \leq \mathbb{E}[g_k(X_k)|\mathcal{F}_k] \leq \mathbb{E}[\sup_z g_k(z)|\mathcal{F}_k] = \mathbb{E}[\sup_z g_k(z)|\mathcal{F}_{k-1}].$$

Hence, $A_k \leq Y_k \leq B_k$ a.s. where

$$\begin{aligned} A_k &:= \mathbb{E}[\inf_z g_k(z) - g_k(X_k)|\mathcal{F}_{k-1}], \\ B_k &:= \mathbb{E}[\sup_z g_k(z) - g_k(X_k)|\mathcal{F}_{k-1}]. \end{aligned}$$

We notice that $A_k, B_k$ are $\mathcal{F}_{k-1}-$measurables. Then, the Azuma-Hoeffding inequality (Corollary 2.6) assures us that $\sum_{k=1}^n Y_k$ is $\sigma^2-$subgaussian with $\sigma^2 = \frac{1}{4}\sum_{k=1}^n \|B_k - A_k\|_\infty^2$. Moreover, by the above identity (2.3), $f(X_1, \ldots, X_n)$ is also a $\sigma^2-$subgaussian. We note that

$$|B_k - A_k| = |\mathbb{E}[\mathfrak{D}_k f(X)|X_1, \ldots, X_{k-1}]| \leq \|\mathfrak{D}_k f\|_\infty.$$

Therefore, $f(X_1, \ldots, X_n)$ is $\sigma^2-$subgaussian with $\sigma^2 = \frac{1}{4}\sum_{k=1}^n \|D_k f\|_\infty^2$.

□

## 2.4. Uniform Approximation of Integrals

**Theorem 2.19** (Uniform approximation of integrals)**.** Let $\Omega \subseteq \mathbb{R}^{d_1}$ be a bounded Lipschitz domain, let $\alpha_1 > 0$ and let $\{X_k\}_{k\in\mathbb{N}}$ be a sequence of i.i.d. random variables in $\bar{\Omega}$, with distribution $\mu$. Let $f : C^\infty(\bar{\Omega}, \mathbb{R}^{d_2}) \times \bar{\Omega} \to \mathbb{R}^{d_2}$ be an operator, and we assume that the following two requirements are satisfied:

(i) there exists $C_1 > 0$ and $\beta_1 \in (0, 1/2)$ such that for all $n \geq 1$ and all $\theta, \theta' \in B_2(0, n^{\alpha_1})$,

$$\|f(u_\theta, \cdot) - f(u_{\theta'}, \cdot)\|_{\infty,\bar{\Omega}} \leq C_1 n^{\beta_1}\|\theta - \theta'\|_2. \tag{H.1}$$

(ii) there exist $C_2 > 0$ and $\beta_2 \in (0, 1/2)$ satisfying $\beta_2 > \alpha_1 + \beta_1$ such that for all $n \geq 1$, and all $\theta \in B_2(0, n^{\alpha_1})$,

$$\|f(u_\theta, \cdot)\|_{\infty,\bar{\Omega}} \leq C_2 n^{\beta_2}. \tag{H.2}$$

Then, almost surely, there exists $N \in \mathbb{N}$ such that for all $n \geq N$,

$$\sup_{\theta \in B_2(0,n^{\alpha_1})} \left\| \frac{1}{n} \sum_{i=1}^{n} f(u_\theta, X_i) - \int_{\bar{\Omega}} f(u_\theta, \cdot) d\mu \right\|_2 \leq \log^2(n) n^{\beta_2 - 1/2}.$$

Notice that $N$ is itself a random variable because it depends on $X_1, \ldots, X_n$. Moreover, note that for $\beta_2 \in (0, 1/2)$, $\lim_{n\to\infty} \log^2(n) n^{\beta_2 - 1/2} = 0$. This theorem assures us that the empirical mean approximates the integral under certain assumptions over the operator $f$.

*Proof.* Let us first analyze the particular case where $d_2 = 1$ before generalizing. Consider $\theta \in B_2(0, n^{\alpha_1})$, and we let

$$Z_{n,\theta} := \frac{1}{n} \sum_{i=1}^{n} f(u_\theta, X_i) - \int_{\bar{\Omega}} f(u_\theta, \cdot) d\mu.$$

Our objective is to establish an upper bound for the random variable

$$Z_n := \sup_{\theta \in B_2(0,n^{\alpha_1})} |Z_{n,\theta}|.$$

Let $M(\Omega) := \max_{x\in\bar{\Omega}} \|x\|_2$. For every $k = 1, \ldots, n$ we define the random variable $W_{k,n}$ as follows

$$W_{k,n} := \frac{1}{n}\left( f(u_\theta, X_k) - \int_{\bar{\Omega}} f(u_\theta, \cdot) d\mu \right) - \frac{1}{n}\left( f(u_{\theta'}, X_k) - \int_{\bar{\Omega}} f(u_{\theta'}, \cdot) d\mu \right).$$

Then, using hypthesis (H.1), we obtain that for any $\theta, \theta' \in B_2(0, n^{\alpha_1})$,

$$|W_{i,n}| \leq \frac{1}{n}\left( C_1 n^{\beta_1} \|\theta - \theta'\|_2 + \int_{\bar{\Omega}} C_1 n^{\beta_1} \|\theta - \theta'\|_2 d\mu \right) \overset{\int_{\bar{\Omega}} d\mu = 1}{=} 2C_1 n^{\beta_1 - 1} \|\theta - \theta'\|_2.$$

In other words, we have found $a, b \in \mathbb{R}$ such that $a \leq W_{k,n} \leq b$. By Hoeffding's Lemma 2.4, we know that $W_{k,n}$ is a $\sigma_k^2$−subgaussian random variable with

$$\sigma_k^2 = \frac{1}{4}(b-a)^2 = \frac{1}{4}(4C_1 n^{\beta_1 - 1}\|\theta - \theta'\|_2)^2 = 4C_1^2 n^{2\beta_1 - 2} \|\theta - \theta'\|_2^2.$$

On the other hand, we note that $Z_{n,\theta} - Z_{n,\theta'} = \sum_{k=1}^{n} W_{k,n}$ and since the random variables $\{W_{k,n}\}_{k\leq n}$ are independent, we can invoke Azuma's Lemma 2.5 to deduce that $Z_{n,\theta} - Z_{n,\theta'}$ is a subgaussian random variable with parameter $\sum_{k=1}^{n} \sigma_k^2 = n\sigma_k^2 = 4C_1^2 n^{2\beta_1 - 1} \|\theta - \theta'\|_2^2$. We observe that $\mathbb{E}[Z_{n,\theta}] = 0$. Therefore, by definition, for all $n \geq 1$, the family $\{Z_{n,\theta}\}_{\theta\in T}$ is a subgaussian process on the metric space $(T, d)$, where

$$T = B_2(0, n^{\alpha_1}), \qquad d(\theta, \theta') = 2C_1 n^{\beta_1 - 1/2} \|\theta - \theta'\|_2.$$

Furthermore, the mapping $\theta \mapsto Z_{n,\theta}$ is continuous with respect to the topology induced by the metric $d$. Consequently, by Remark 2.14, the process $(Z_{n,\theta})_{\theta \in B_2(0,n^{\alpha_1})}$ is separable. Applying Dudley's Corollary 2.16, we obtain

$$\mathbb{E}\left[\sup_{\theta \in B_2(0,n^{\alpha_1})} Z_{n,\theta}\right] \equiv \mathbb{E}[Z_n] \leq 12 \int_0^\infty \sqrt{\log N(T,d,r)}dr. \tag{2.4}$$

We can now establish the following relation

$$\begin{aligned} N(T,d,r) &= N(B_2(0,n^{\alpha_1}), d(\theta,\theta'), r) \\ &= N(B_2(0,n^{\alpha_1}), \|\cdot\|_2, n^{1/2-\beta_1} r/(2C_1)) \\ &= N(B_2(0,1), \|\cdot\|_2, n^{1/2-\beta_1-\alpha_1} r/(2C_1)). \end{aligned} \tag{2.5}$$

Combining (2.4) and (2.5), we obtain

$$\begin{aligned} \mathbb{E}[Z_n] &\leq 12 \int_0^\infty \sqrt{\log N(B_2(0,n^{\alpha_1}), \|\cdot\|_2, n^{1/2-\beta_1} r/(2C_1))}dr \\ &= 24 C_1 n^{\alpha_1+\beta_1-1/2} \int_0^\infty \sqrt{\log N(B_2(0,1), \|\cdot\|_2, s)}ds \\ &= 24 C_1 n^{\alpha_1+\beta_1-1/2} \int_0^1 \sqrt{\log N(B_2^n, \|\cdot\|_2, s)}ds. \end{aligned}$$

In the second term, we performed a change of variable given by $s = n^{1/2-\beta_1} r/(2C_1)$, while the third term follows from the fact that $N(B_2, \|\cdot\|_2, r) = 1 \ \forall r \geq 1$. Since $\theta$ belongs to the parameter space $\Theta_H^D := \mathbb{R}^{(d_1+1)D+(H-1)D(D+1)+(D+1)d_2}$, we can apply Lemma 2.10 to bound the covering number as follows

$$\log N(B_2^n, \|\cdot\|_2, s) \leq [(d_1+1)D + (H-1)D(D+1) + (D+1)d_2] \log(3/s).$$

Observing that $\int_0^1 \sqrt{\log(3/s)}ds \leq 3/2$ and noting that, by definition, $\alpha_1 + \beta_1 < \beta_2 < 1/2$, we observe that

$$\mathbb{E}[Z_n] \leq 36 C_1 n^{\alpha_1+\beta_1-1/2} \sqrt{(d_1+1)D + (H-1)D(D+1) + (D+1)d_2} \xrightarrow{n\to\infty} 0. \tag{2.6}$$

Note that $Z_n = Z_n(X_1, \dots, X_n)$, so applying the definition of the discrete derivative of $Z_n$ (Definition 2.17) along with the hypothesis (H.2), we obtain

$$\begin{aligned} \mathfrak{D}_i Z_n &:= \sup_{x_i \in \mathbb{R}^{d_1}} Z_n(X_1, \dots, X_{i-1}, x_i, X_{i+1}, \dots, X_n) - \inf_{x_i \in \mathbb{R}^{d_1}} Z_n(X_1, \dots, X_{i-1}, x_i, X_{i+1}, \dots, X_n) \\ &\leq 2n^{-1} \sup_{\theta \in B_2(0,n^{\alpha_1})} \|f(u_\theta, \cdot)\|_\infty \leq 2C_2 n^{\beta_2 - 1}. \end{aligned}$$

Applying McDiarmid's Theorem 2.18, we conclude that $Z_n(X_1, \dots, X_n)$ is $\sigma_z^2-$subgaussian variable with $\sigma_z^2 = \frac{1}{4} n (2C_2 n^{\beta_2-1})^2 = C_2^2 n^{2\beta_2 - 1}$. In particular, by the Chernoff bound (Lemma 2.2), for all $t_n \geq 0$, we obtain

$$\mathbb{P}(|Z_n - \mathbb{E}Z_n| \geq t_n) \leq 2e^{-t_n^2/2\sigma_z^2} = 2\exp\{-t_n^2 n^{1-2\beta_2}/2C_2^2\}.$$

We will now show that the events $E_n := \{|Z_n - \mathbb{E}Z_n| \le t_n\}$ satisfy the Borel-Cantelli Lemma when the sequence $\{t_n\}_n$ is given by $t_n := 2C_2 n^{\beta_2 - 1/2}\log^{1/2}(n)$. Indeed,

$$\sum_{n\in\mathbb{N}} \mathbb{P}(E_n) \le \sum_{n\in\mathbb{N}} 2\exp\{-(2C_2 n^{\beta_2-1/2}\log^{1/2}(n^2))^2 n^{1-2\beta_2}/2C_2^2\} < \infty.$$

Therefore, using the fact that $\mathbb{E}[Z_n] \to 0$ (2.6) and applying Borel-Cantelli Lemma yields the desired result

$$0 \le Z_n \le \log^{1/2}(n) n^{\beta_2 - 1/2} \le \log^2(n) n^{\beta_2-1/2} \quad a.s. \qquad \forall n \ge N.$$

for some sufficiently large $N$.

The generalization to the case $d_2 > 1$ is straightforward. Let $f = (f_1, \ldots, f_{d_2})$. Since $\|\cdot\|_2 \le \sqrt{d_2}\|\cdot\|_\infty$, then

$$\sup_{\theta\in B_2(0,n^{\alpha_1})} \left\| \frac{1}{n}\sum_{i=1}^n f(u_\theta, X_i) - \int_\Omega f(u_\theta,\cdot)d\mu \right\|_2$$
$$\le \sqrt{d_2} \max_{1\le j\le n} \sup_{\theta\in B_2(0,n^{\alpha_1})} \left\| \frac{1}{n}\sum_{i=1}^n f_j(u_\theta, X_i) - \int_\Omega f_j(u_\theta,\cdot)d\mu \right\|_2 .$$

□

## 2.5. Empirical Processes

**Lemma 2.20** (Empirical Process)**.** Let $X_1, \ldots, X_n, \epsilon_1, \ldots, \epsilon_n$ be independent random variables, such that $X_i$ is distributed along $\mu_X$ and $\epsilon_i$ is distributed along $\mu_\epsilon$ with $\mathbb{E}[\epsilon_i] = 0$. Then, there exists a constant $C_\Omega > 0$, depending on $\Omega$ such that

$$\mathbb{E}\left[\sup_{\|u\|_{H^{m+1}(\Omega)}\le 1} \frac{1}{n}\sum_{j=1}^n \langle \tilde{\Pi}(u)(X_j) - \mathbb{E}[\tilde{\Pi}(u)(X)], \epsilon_j\rangle\right]^2 \le \frac{d_2\mathbb{E}\|\epsilon\|_2^2}{n} C_\Omega,$$

where $\tilde{\Pi}$ is the Sobolev embedding in Theorem A.1.

*Proof.* Firstly, we note that $H^{m+1}(\Omega)$ is separable and the function

$$(x_1, \ldots, x_n, e_1, \ldots, e_n) \mapsto \frac{1}{n}\sum_{j=1}^n \langle \tilde{\Pi}(u)(x_j) - \mathbb{E}[\tilde{\Pi}(u)(X)], e_j\rangle$$

is continuous for all $u \in H^{m+1}(\Omega, \mathbb{R}^{d_2})$. Then, we can define the supremum $Z := \sup_{\|u\|_{H^{m+1}(\Omega)}\le 1} \frac{1}{n}\sum_{j=1}^n \langle \tilde{\Pi}(u)(X_j) - \mathbb{E}[\tilde{\Pi}(u)(X)], \epsilon_j\rangle$ is a random variable. Moreover, using

Theorem A.1 that assures us that $\|\tilde{\Pi}(u)\|_{\infty,\Omega} \le C_\Omega \|u\|_{H^m(\Omega)}$, we obtain

$$\begin{aligned}
|Z| &\le \frac{1}{n}\sum_{j=1}^{n} \sup_{\|u\|_{H^{m+1}(\Omega)}\le 1} \|\tilde{\Pi}(u)(X_j) - \mathbb{E}[\tilde{\Pi}(u)(X)]\| \cdot \|\epsilon_j\| \\
&\le \frac{1}{n}\sum_{j=1}^{n} \sup_{\|u\|_{H^{m+1}(\Omega)}\le 1} \sqrt{d_2}(\|\tilde{\Pi}(u)\|_{\infty,\Omega} + \mathbb{E}\|\tilde{\Pi}(u)\|_{\infty,\Omega}) \cdot \|\epsilon_j\| \\
&\le \frac{\sqrt{d_2}}{n}\sum_{j=1}^{n} \sup_{\|u\|_{H^{m+1}(\Omega)}\le 1} (C_\Omega\|u\|_{H^m(\Omega)} + \mathbb{E}[C_\Omega\|u\|_{H^m(\Omega)}]) \cdot \|\epsilon_j\| \\
&\le \frac{2C_\Omega\sqrt{d_2}}{n}\sum_{j=1}^{n} \|\epsilon_j\|_2.
\end{aligned}$$

If $\epsilon$ has finite second moment, then $\mathbb{E}[Z^2] < \infty$. Define, for any $u \in H^{m+1}(\Omega, \mathbb{R}^{d_2})$,

$$\begin{aligned}
Z_{n,u} &:= \frac{1}{n}\sum_{j=1}^{n} \langle \tilde{\Pi}(u)(X_j) - \mathbb{E}[\tilde{\Pi}(u)(X)], \epsilon_j \rangle \\
Z_n &:= \sup_{\|u\|_{H^{m+1}(\Omega)}\le 1} Z_{n,u}.
\end{aligned}$$

For any $u, v \in H^{m+1}(\Omega, \mathbb{R}^{d_2})$, we have

$$\begin{aligned}
&\left| \frac{1}{n}\langle \tilde{\Pi}(X_i) - \mathbb{E}[\tilde{\Pi}(u)(X)], \epsilon_i\rangle - \frac{1}{n}\langle \tilde{\Pi}(v)(X_i) - \mathbb{E}[\tilde{\Pi}(u)(X)], \epsilon_i\rangle \right| \\
&\qquad = \frac{1}{n}\left| \langle \tilde{\Pi}(u-v)(X_i) - \mathbb{E}[\tilde{\Pi}(u-v)(X)], \epsilon_j\rangle \right| \le \frac{2C_\Omega\sqrt{d_2}}{n}\|u-v\|_{H^{m+1}(\Omega)}\|\epsilon_i\|_2.
\end{aligned}$$

Using that $\epsilon$ is independent of $X$, we can apply Hoeffding's, Azuma's and Dudley's Theorem as in the previous proof to show that

$$\mathbb{E}[Z_n|\epsilon_1,\dots,\epsilon_n] \le \frac{24C_\Omega\sqrt{d_2}}{n}\sqrt{\sum_{i=1}^{n}\|\epsilon_i\|_2^2}\int_0^\infty [\log N(B_{H^{m+1}(\Omega)}(1), \|\cdot\|_{H^{m+1}(\Omega)}, r)]^{1/2} dr.$$

Hence, according to Lemma 2.11, there exists a constant $C'_\Omega > 0$, depending only on $\Omega$ such that $\mathbb{E}[Z_n|\epsilon_1,\dots,\epsilon_n] \le C'_\Omega n^{-1}\sqrt{d_2\sum_{i=1}^n \|\epsilon_i\|_2^2}$. Hence, we obtain

$$\mathbb{E}[Z_n] \le C'_\Omega\sqrt{d_2\frac{\mathbb{E}\|\epsilon\|_2^2}{n}}$$

and

$$\mathrm{Var}(\mathbb{E}[Z_n|\epsilon_1,\dots,\epsilon_n]) \le \mathbb{E}[\mathbb{E}[Z_n|\epsilon_1,\dots,\epsilon_n]^2] \le (C'_\Omega)^2 d_2\frac{\mathbb{E}\|\epsilon\|_2^2}{n}.$$

Applying McDiarmid's inequality (Theorem 2.18) as in the previous theorem shows that

$$\mathrm{Var}(Z_n|\epsilon_1, \dots, \epsilon_n) \leq 16C_\Omega^2 d_2 \frac{1}{n^2} \sum_{i=1}^{n} \|\epsilon_i\|_2^2.$$

Therefore,

$$\mathrm{Var}(Z_n) = \mathrm{Var}(\mathbb{E}[Z_n|\epsilon_1, \dots, \epsilon_n]) + \mathbb{E}[\mathrm{Var}(Z_n|\epsilon_1, \dots, \epsilon_n)] \leq \frac{d_2 \mathbb{E}\|\epsilon\|_2^2}{n}((C'_\Omega)^2 + 16C_\Omega^2).$$

Since $\mathbb{E}[Z_n^2] \leq \mathrm{Var}(Z_n) + \mathbb{E}[Z_n]^2$, we obtain that

$$\mathbb{E}[Z_n^2] \leq \frac{d_2 \mathbb{E}\|\epsilon\|_2^2}{n}((C'_\Omega)^2 + 16C_\Omega^2).$$

□

# CHAPTER 3
# Consistency of PINNs

In this section, we show that standard PINNs can lead to overfitting problems. The ideas are mainly brought from Doumèche et al. [2023] and Mohri et al. [2012]. This means that while you can achieve a low empirical risk function, the theoretical risk function does not necessarily stay low. In fact, will give some examples where we can obtain zero empirical risk function, while the theoretical risk function tends to infinity.

## 3.1. Risk-Consistency

In machine learning, we minimized the empirical risk and hope that it translates to the theoretical risk. This concept is described in the following definition and it is usually named after *risk-consistency*.

**Definition 3.1.** Fix $n_1 \in \mathbb{N}$. Let $\{\hat{\theta}(n, n_2, n_3)\}_{n\in\mathbb{N}}$ be a minimizing sequence for the empirical risk, i.e.,

$$\lim_{n\to\infty} \widehat{R}_{n_1,n_2,n_3}(u_{\hat{\theta}_n(n_2,n_3)}) = \inf_{\theta\in\Theta_{H,D}} \widehat{R}_{n_1,n_2,n_3}(u_\theta).$$

We say that $\{\hat{\theta}_n(n_2, n_3)\}_{n\in\mathbb{N}}$ satisfies the *risk-consistency* with respect to the theoretical risk $R_{n_1}$ if

$$\lim_{n_2,n_3\to\infty} \lim_{n\to\infty} R_{n_1}(u_{\hat{\theta}_n(n_2,n_3,D)}) = \inf_{u\in \mathrm{NN}_H^D} R_{n_1}(u). \tag{3.1}$$

This condition means that the sequence minimizing the empirical risk also asymptotically minimizes the theoretical risk. Risk consistency is a desirable property, since training procedures typically minimize empirical risk. Without this property, there is no guarantee that minimizing the empirical risk results in a low theoretical risk. In fact, we will see a counterexample down below.

## 3.2. Counterexample

We will show theoretically and empirically that standard PINNs fail to be risk-consistent, though a counterexample.

**Example 3.2** (Dynamic friction model)**.** Considering the following ODE defined on the domain $\Omega := (0, T)$ by

$$\mathfrak{F}(u, x) = mu''(x) + \gamma u'(x), \tag{3.2}$$

where $u \in C^2(\overline{\Omega}, \mathbb{R})$ and $x \in \Omega$. In this case, we assume that there are not boundary samples, i.e., $n_2 = 0$. We can prove that, whenever $D \geq n - 1$, for all $n_1, n_3 \in \mathbb{N}$, there exists a minimizing sequence $\{u_{\hat{\theta}_n(n_1,n_3)}\}_{n\in\mathbb{N}} \subseteq \mathrm{NN}_H^D$ such that

$$\lim_{n\to\infty} R_{n_1,n_2,n_3}(u_{\hat{\theta}_n(n_1,n_3)}) = 0$$
$$\lim_{n\to\infty} \widehat{R}_{n_1}(u_{\hat{\theta}(n_1,n_3)}) = \infty.$$

Therefore, we can say that PINN estimator is not risk-consistent.

*Proof.* This example assumes that we have data availability and the equation (3.2). We assume that the observations $(X_1, Y_1), \dots, (X_{n_1}, Y_{n_1}) \in \mathbb{R}^2$ are ordered in increasing values of the $X_i$, i.e., $X_1 \leq \cdots \leq X_a$. Let $\mathcal{G}(n_1, n_3) := \{X_i, 1 \leq i \leq n_1\} \cup \{X_j^{(3)}, 1 \leq j \leq n_3\}$. Consider the neural network $u_{\hat{\theta}_n(c)} \in \mathrm{NN}_H^D$ defined by

$$u_\theta(x) = Y_1 + \sum_{i=1}^{n_1-1} \frac{Y_{i+1} - Y_i}{2} \left[\tanh_\theta^{\circ H}\left(x - X_i - \frac{\delta(n_1, n_3)}{2}\right) + 1\right],$$

where $\delta(n_1, n_3) = \min_{\substack{z_1,z_2\in\mathcal{G}(n_1,n_3)\\ z_1\neq z_1}} |z_1 - z_2|$ and where $\tanh_\theta^{\circ H} := \tanh_\theta \circ \overset{H}{\dots} \circ \tanh_\theta$, and $\tanh_\theta(x) := \tanh(x\theta)$. Let us fix $\epsilon := \delta(n_1, n_3)/4$ and define the set of discontinuities

$$G := \mathbb{R} \backslash \bigcup_{i=1}^{n_1} \left(X_i + \frac{1}{4}\delta(n_1, n_3), X_i + \frac{3}{4}\delta(n_1, n_3)\right).$$

By Lemma 1.6, for all $K$, $\lim_{\theta\to\infty} \|u_\theta - u_\infty\|_{C^K(G)} = 0$, where

$$\begin{aligned} u_\infty(x) &= Y_1 + \sum_{i=1}^{n-1} \frac{Y_{i+1} - Y_i}{2} \left[\mathbf{1}_{x>X_i+\delta(n_1,n_3)/2} - \mathbf{1}_{x<X_i+\delta(n_1,n_3)/2} + 1\right] \\ &= Y_1 + \sum_{i=1}^{n-1} [Y_{i+1} - Y_i] \mathbf{1}_{x>X_i+\delta(n_1,n_3)/2}. \end{aligned}$$

Therefore, $u_\infty(X_i) = Y_i$. Since $u'_\infty(x) = 0$ for all $x \in G$, and $X_j^{(3)} \in G$, we deduce that $u_\infty^{(K)}(X_j^{(3)}) = 0$. Thus, we conclude that $\lim_{n\to\infty} \widehat{R}_{n_1,0,n_3}(u_{\hat{\theta}_n}) = 0$ when $\hat{\theta}_n \xrightarrow{n\to\infty} \infty$.

We will prove that it is not risk-consistent, in particular, we will prove that $R_{n_1}(u_{\hat{\theta}_n}) \xrightarrow{n\to\infty} \infty$. Indeed,

$$\begin{aligned} R_{n_1}(u_{\hat{\theta}_n}) &\geq \frac{1}{T}\int_{[0,T]} \mathfrak{F}(u_{\hat{\theta}_n(n_3)}, x)^2 dx \\ &\geq \frac{1}{T}\sum_{i=1}^{n_1}\int_{X_i+\delta(n_1,n_3)/2-\epsilon}^{X_i+\delta(n_1,n_3)/2-\epsilon} \mathfrak{F}(u_{\hat{\theta}_n(n_3)}, x)^2 dx. \end{aligned}$$

We notice that

$$2\epsilon\int_{-\epsilon}^{\epsilon}(mf'' + \gamma f')^2 \overset{C-S}{\geq} \left(\int_{-\epsilon}^{\epsilon} mf'' + \gamma f'\right)^2 = [m(f'(\epsilon) - f'(-\epsilon)) + \gamma(f(\epsilon) - f(-\epsilon))]^2.$$

Therefore, using this inequality and the fact that $\epsilon < \delta(n_1, n_3)/4$,

$$\begin{aligned} \lim_{n\to\infty} R_{n_1}(u_{\hat{\theta}_n}) \geq \lim_{n\to\infty}\frac{1}{T}\sum_{i=1}^{n_1}\frac{1}{2\epsilon}[m(\cancel{u'_{\hat{\theta}_n}(X_i + \delta(n_1, n_3)/2 + \epsilon)} - \cancel{u'_{\hat{\theta}_n}(X_i + \delta(n_1, n_3)/2 - \epsilon)}) \\ \gamma(\underbrace{u_{\hat{\theta}_n}(X_i + \delta(n_1, n_3)/2 + \epsilon)}_{Y_{i+1}} - \underbrace{u_{\hat{\theta}_n}(X_i + \delta(n_1, n_3)/2 - \epsilon)}_{Y_i}))]^2 \\ \geq \frac{\gamma}{2\epsilon}\sum_{i=1}^{n_1-1}(Y_{i+1} - Y_i)^2. \end{aligned}$$

Since, we assume that $Y_{i+1} \neq Y_i$ and that for any $\epsilon < \delta(n_1, n_3)/4$, there exists a sequence of parameters such that $\lim_{n\to\infty} R_{n_1}(u_{\hat{\theta}_n}) = \infty$. □

As shown in the previous section, training PINNs is not straightforward, it can lead to unintended results when minimizing the risk in the form of overfitting. To avoid such problems, there is a popular approach among machine learning practitioners that consist of using ridge regularization. The regularization consists of the empirical risk by the $L^2$ norm of the parameters $\theta$. This technique has shown consistently that it improves the generalization and training process Mohri et al. [2012]. In this section, we will study and formalize the ridge PINNs framework and prove the risk-consistency.

## 3.3. Ridge PINNs

We will intordュce a regularization term in the empirical risk function to incentivize the solutions that are *less complex*. As we have seen in the second chapter, there is a relationship between the norm of the parameters and the Lipschitz constant of the neural networks. We want to avoid training a complex model that can essentially *memorize* the training data, including the noise and idiosyncrasies. This phenomenon, is known as *overfitting*. Intuitively, likewise the Occam's Razor principle, we are advocating for *simplicity*, favouring simpler models that are more likely to generalize well to new, unseen data, and in particular, that reduce the theoretical risk function. There are two common regularization techniques, $\ell^1$ and $\ell^2$ regularization, beware that they are also usually name $L_1$ and $L_2$ regularization. Despite been similar, they produce different results in most cases:

- $\ell^1$ **Regularization (Lasso)**: $\ell^1$ regularization adds a penalty proportional to the absolute value of the weights. A key feature of $\ell^1$ regularization is that it can force some of the weights to become exactly zero. This has the effect of *switching off* certain connections in the network, leading to a sparser, and therefore simpler, model.
- $\ell^2$ **Regularization (Ridge)**: This technique adds a penalty proportional to the square of the magnitude of the weights. This encourages the network to use smaller weights, effectively "smoothing out" the function it learns and preventing it from making sharp, complex decisions based on small fluctuations in the training data.

We will focus hereafter on the Ridge regularization. We begin by introducing a slightly modified version of the empirical risk function.

**Definition 3.3** (Ridge PINNs). The ridge empirical risk function is defined by

$$\widehat{R}^{(\text{ridge})}_{n_1,n_2,n_3}(u_\theta) := \widehat{R}_{n_1,n_2,n_3}(u_\theta) + \lambda_{(\text{ridge})}\|\theta\|_2^2, \tag{3.3}$$

where $\lambda_{(\text{ridge})} > 0$ is the ridge hyperparameter. We denote by $\{\hat{\theta}^{(\text{ridge})}_n(n_2, n_3, D)\}_{n\in\mathbb{N}}$ the minimizing sequence of the ridge risk, that is,

$$\lim_{n\to\infty} \widehat{R}^{(\text{ridge})}_{n_1,n_2,n_3}(u_{\hat{\theta}^{(\text{ridge})}_n(n_2,n_3,D)}) = \inf_{\theta\in\Theta_{H,D}} \widehat{R}^{(\text{ridge})}_{n_1,n_2,n_3}(u_\theta).$$

We will restrict the family of functions $\mathfrak{F}$ that can be used for the PDE. Polynomial operators are general enough that allow almost any physical problem to be described using this family of functions.

**Example 3.4** (Navier-Stokes equations). Let $\Omega = \Omega_1 \times (0, T)$, where $\Omega_1 \subseteq \mathbb{R}^3$ is bounded Lipschitz domain and $T \geq 0$ is finite time horizon. The incompressible Naiver-Stokes system of equations id defined for all $\mathbf{v} = (u, v, w, p) \in C^2(\overline{\Omega}, \mathbb{R}^4)$ by

$$\begin{cases} \nabla \cdot \mathbf{v} = 0, \\ \rho \frac{D\mathbf{v}}{Dt} = -\nabla p + \mu \nabla^2 \mathbf{v}. \end{cases}$$

This can be written using our notation as

$$\begin{cases} \mathfrak{F}_1(u, x) = \partial_x u + \partial_y v + \partial_z w \\ \mathfrak{F}_2(u, x) = \rho(\partial_t u + u\partial_x u + v\partial_y u + w\partial_z u) + \partial_x p - \mu(\partial^2_{x,x} u + \partial^2_{y,y} u + \partial^2_{z,z} u) - \rho g_x \\ \mathfrak{F}_3(u, x) = \rho(\partial_t v + u\partial_x v + v\partial_y v + w\partial_z v) + \partial_y p - \mu(\partial^2_{x,x} v + \partial^2_{y,y} v + \partial^2_{z,z} v) - \rho g_y \\ \mathfrak{F}_4(u, x) = \rho(\partial_t w + u\partial_x w + v\partial_y w + w\partial_z w) + \partial_z p - \mu(\partial^2_{x,x} w + \partial^2_{y,y} w + \partial^2_{z,z} w) - \rho g_z \end{cases}$$

where $\rho, \mu > 0$ and $g_1, g_2, g_3 \in C^\infty(\overline{\Omega}, \mathbb{R})$. We notice that $\mathfrak{F}_i$ for $i = 1, 2, 3, 4$ are polynomials in $u$ and its derivatives, with coefficients in $C^\infty(\overline{\Omega}, \mathbb{R})$. For example,

$$\mathfrak{F}_3(u, x) = P_3(u, v, w, \partial_t v, \partial_x v, \partial_y v, \partial_z v, \partial_y p, \partial^2_{x,x} v, \partial^2_{y,y} v, \partial^2_{z,z} v),$$

where the polynomial $P_3 \in C^\infty(\overline{\Omega}, \mathbb{R})[Z_1, \dots, Z_{11}]$ is defined by

$$P_3(Z_1, \dots, Z_{11}) = \rho(Z_4 + Z_1 Z_5 + Z_2 Z_6 + Z_3 Z_7) + Z_8 - \mu(Z_9 + Z_{10} + Z_{11}) - \rho g_y.$$

This example gives rise to the family of functions named polynomial operators.

**Definition 3.5** (Polynomial Operators)**.** Let $s \in \mathbb{N}$ and multi-indexes $(\alpha_{i,j})_{1 \le i \le d_2, 1 \le j \le s}$ where $\alpha_{i,j} \in \mathbb{N}^{d_1}$. An operator $\mathfrak{F} : C^K(\overline{\Omega}, \mathbb{R}^{d_2}) \times \Omega \to \mathbb{R}$ is a polynomial operator of order $K \in \mathbb{N}$ is a function that can be written as

$$\mathfrak{F}(u, \cdot) = P((\partial^{\alpha_{i,j}} u_i)_{1 \le i \le d_2, 1 \le j \le s}), \qquad \forall u = (u_1, \dots, u_{d_2}) \in C^K(\overline{\Omega}, \mathbb{R}^{d_2}),$$

where $P \in C^\infty(\overline{\Omega}, \mathbb{R})[Z_{1,1}, \dots, Z_{d_2,s}]$ is a polynomial with smooth coefficients.

This can be defined in other ways by

$$\mathfrak{F}(u, \mathbf{x}) = \sum_{k=1}^{N(P)} \phi_k \times \prod_{i=1}^{d_2} \prod_{j=1}^{s} (\partial^{\alpha_{i,j}} u_i(\mathbf{x}))^{I(i,j,k)},$$

where $N(P) \in \mathbb{N}$, $\phi_k \in C^\infty(\overline{\Omega}, \mathbb{R})$, and $I(i, j, k) \in \mathbb{N}$. This formula can be expressed as a polynomial

$$P(Z_{1,1}, \dots, Z_{d_2,s}) = \sum_{k=1}^{N(P)} \phi_k \times \prod_{i=1}^{d_2} \prod_{j=1}^{s} Z_{i,j}^{I(i,j,k)},$$

The *degree* of the polynomial operator $\mathfrak{F}$ is

$$\deg(\mathfrak{F}) := \max_{1 \le k \le N(P)} \sum_{i=1}^{d_2} \sum_{j=1}^{s} (1 + |\alpha_{i,j}|) I(i, j, k).$$

Using example 3.4, $\deg(\mathfrak{F}_3) = 3$ and its degree is reached in the terms $u_i \partial_j u_j$ and $\partial^2_{i,i} u_j$ for $i, j \in \{x, y, z\}$. It is noteworthy that this class of functions includes a large number of PDEs. This constraint allows us to establish the risk-consistency of the ridge PINNs when $n_2, n_3 \to \infty$.

## 3.4. Risk-consistency of ridge PINNs

**Proposition 3.6.** Let $\Omega$ be a bounded Lipschitz domain, let $E$ be a closed subset of $\partial\Omega$, and let $h \in \mathrm{Lip}(E, \mathbb{R}^{d_2})$. Then, the operator $\mathscr{H}(u, x) = \mathbb{1}_{x \in E} \|u(x) - h(x)\|^2$ satisfies inequalities (H.1) and (H.2) from Theorem 2.19 with $\alpha_1 < (3 + H)^{-1}/2$, $\beta_1 = (1 + H)\alpha_1$, and $1/2 > \beta_2 \ge (3 + H)\alpha_1$.

*Proof.* Since $\mathrm{Lip}(E,\mathbb{R}^{d_2}) \subseteq C^0(E,\mathbb{R}^{d_2})$, then $\|h\|_\infty < \infty$. Observe that for any $v, w \in \mathbb{R}^{d_2}$, $|\,\|v\|_2^2 - \|w\|_2^2\,| = |\langle v+w, v-w\rangle| \leq \|v+w\|_2\|v-w\|_2 \leq d_2\|v+w\|_\infty\|v-w\|_\infty$. Thus, we obtain, for all $\theta, \theta' \in B_2(0, n^{\alpha_1})$ and $x \in E$,

$$\begin{aligned}
|\mathscr{H}(u_\theta, x) - \mathscr{H}(u_{\theta'}, x)| &\leq (\|u_\theta(x)\|_2 + \|u_{\theta'}(x)\|_2 + 2\|h(x)\|_2)\|u_\theta(x) - u_{\theta'}(x)\|_2 \\
&\leq d_2(\|u_\theta\|_{\infty,\bar\Omega} + \|u_{\theta'}\|_{\infty,\bar\Omega} + 2\|h\|_\infty)\|u_\theta - u_{\theta'}\|_{\infty,\bar\Omega} \\
&\leq d_2(2(D+1)n^{\alpha_1} + 2\|h\|_\infty)\tilde{C}_{0,H}(1 + d_1 M(\Omega))(D+1)^H(1+n^{\alpha_1})^H\|\theta - \theta'\|_2 \\
&\leq \underbrace{2^{H+1} d_2(D + 1 + \|h\|_\infty)\tilde{C}_{0,H}(1 + d_1 M(\Omega))(D+1)^H}_{C_1} \overbrace{n^{\alpha_1(1+H)}}^{n^{\beta_1}} \|\theta - \theta'\|_2 \\
&\equiv C_1 n^{\beta_1}\|\theta - \theta'\|_2.
\end{aligned}$$

The third inequality is given by (1.4) in Lemma 1.6 and due to Lemma 1.7. Next, using again (1.4), for all $\theta \in B_2(0, n^{\alpha_1})$,

$$\|\mathscr{H}(u_\theta, \cdot)\|_{\infty,\bar\Omega} \leq d_2(\|u_\theta\|_{\infty,\bar\Omega} + \|h\|_\infty)^2 \leq d_2((D+1)n^{\alpha_1} + \|h\|_\infty)^2 \leq C_2 n^{2\alpha_1}.$$

Recall, that we must verify the inequality: $\alpha_1 + \beta_1 < \beta_2 < 1/2$. This is true for $\beta_2 := 2(2+H)\alpha_1$, which completes the proof. □

**Proposition 3.7.** Let $\Omega$ be a bounded Lipschitz domain, let $E$ be a closed subset of $\partial\Omega$, and let $\mathfrak{F} \in \mathscr{P}_{\mathrm{op}}$. Then, the operator $\mathbb{1}_{x\in\Omega}\mathfrak{F}(u_\theta, x)^2$ satisfies inequalities (H.1) and (H.2) from Theorem 2.19 with $\alpha_1 < [2+H(1+(2+H)\deg(\mathfrak{F}))]^{-1}$, $\beta_1 = H(1+(2+H)\deg(\mathfrak{F}))\alpha_1$, and $1/2 > \beta_2 \geq [2 + H(1 + (2+H)\deg(\mathscr{F}))]\alpha_1$.

*Proof.* Let $\mathfrak{F} \in \mathscr{P}_{op}$. By definition, there exists a polynomial $P \in C^\infty(\mathbb{R}^{d_1}, \mathbb{R})[Z_{1,1}, \ldots, Z_{d_2,s}]$ with degree $s \geq 1$, and a sequence $(\alpha_{i,j})_{1\leq i\leq d_2, 1\leq j\leq s}$ of multi-indexes such that, for any $u \in C^\infty(\bar\Omega, \mathbb{R}^{d_2})$, $\mathfrak{F}(u, \cdot) = P((\partial^{\alpha_{i,j}} u_i)_{1\leq i\leq d_2, 1\leq j\leq s})$. In other words, the polynomial can be expressed as $P(Z_{1,1}, \ldots, Z_{d_2,s}) = \sum_{k=1}^{N(P)} \phi_k \prod_{i=1}^{d_2}\prod_{j=1}^{s} Z_{i,j}^{I(i,j,k)}$ where $N(P) \in \mathbb{N}$. Recall, by Definition 3.5, that $\deg(\mathfrak{F}) := \max_k \sum_{i=1}^{d_2}\sum_{j=1}^{s}(1 + |\alpha_{ij}|)I(i,j,k)$.

According to Proposition 1.6, there exists a constant $C_{\deg(\mathfrak{F}),H}$ such that

$$\begin{aligned}
\|\mathfrak{F}(u_\theta, \cdot)^2\|_{\infty,\bar\Omega} &\leq \left[\sum_{k=1}^{N(P)} \|\phi_k\|_{\infty,\bar\Omega} \prod_{i=1}^{d_2}\prod_{j=1}^{s} \|\partial^{\alpha_{i,j}} u_\theta\|_{\infty,\bar\Omega}^{I(i,j,k)}\right]^2 \\
&\leq N^2(P)\left[\sup_{1\leq k\leq N(P)} \|\phi_k\|_{\infty,\bar\Omega} \prod_{i=1}^{d_2}\prod_{j=1}^{s} \|u_\theta\|_{C^{|\alpha_{ij}|}(\mathbb{R}^{d_1})}^{I(i,j,k)}\right]^2 \\
&\leq N^2(P)[\max_{1\leq k\leq N(P)} \|\phi\|_{\infty,\bar\Omega}]^2 C^2_{\deg\mathfrak{F},H}\,[(D+1)(1+\|\theta\|_2)]^{2H\sum_i\sum_j(1+|\alpha_{ij}|)I(i,j,k)} \\
&\leq \underbrace{N^2(P)[\max_{1\leq k\leq N(P)} \|\phi\|_{\infty,\bar\Omega}]^2 C^2_{\deg\mathfrak{F},H}(D+1)^{2H\deg\mathfrak{F}}}_{C_2} \overbrace{(1+\|\theta\|_2)^{2H\deg\mathfrak{F}}} \\
&\leq C_2\|\theta\|_2^{2H\deg\mathfrak{F}}.
\end{aligned}$$

Therefore, for any $\theta \in B_2(0, n^{\alpha_1})$, $\|\mathfrak{F}(u_\theta, \cdot)^2\|_{\infty,\bar{\Omega}} \leq C_2 n^{\beta_2}$ for any $\beta_2 \geq 2H\alpha_1 \deg \mathfrak{F}$ for some sufficiently large $n$.

On the other hand, notice that $|\, u^2 - v^2\,| = |(u+v)(u-v)| \leq |u+v||u-v|$ for any $u, v \in \mathbb{R}$. Therefore,

$$\begin{aligned} |\mathfrak{F}(u_\theta, x)^2 - \mathfrak{F}(u_{\theta'}, x)^2| &\leq (|\mathfrak{F}(u_\theta, x)| + |\mathfrak{F}(u_{\theta'}, x)|)|\mathfrak{F}(u_\theta, x) - \mathfrak{F}(u_{\theta'}, x)| \\ &\leq 2C_2^{1/2} n^{H\alpha_1 \deg \mathfrak{F}} |\mathfrak{F}(u_\theta, x) - \mathfrak{F}(u_{\theta'}, x)|. \end{aligned}$$

Let us bound the right hand side. We notice that $\prod_{i=1}^{d_2} \prod_{j=1}^{s} Z_{i,j}^{I(i,j,k)}$ has less than $\deg \mathfrak{F}$ terms different from 1. Using (1.10) over $\mathfrak{F}$ the right hand side yields

$$\begin{aligned} |\mathfrak{F}(u_\theta, x) - \mathfrak{F}(u_{\theta'}, x)| &\leq \left[ \sum_{k=1}^{N(P)} \phi_k \prod_{i=1}^{d_2} \prod_{j=1}^{s} (\partial^{\alpha_{i,j}} (u_\theta)_i)^{I(i,j,k)} - \sum_{k=1}^{N(P)} \phi_k \prod_{i=1}^{d_2} \prod_{j=1}^{s} (\partial^{\alpha_{i,j}} (u_{\theta'})_i)^{I(i,j,k)} \right] \\ &\leq N(P) \max_{1 \leq k \leq N(P)} \|\phi_k\|_{\infty,\bar{\Omega}} \left[ \prod_{i,j} (\partial^{\alpha_{i,j}} (u_\theta)_i)^{I(i,j,k)} - \prod_{i,j} (\partial^{\alpha_{i,j}} (u_{\theta'})_i)^{I(i,j,k)} \right] \\ &\leq N(P) \max_{1 \leq k \leq N(P)} \|\phi_k\|_{\infty,\bar{\Omega}} \deg(\mathfrak{F}) \|u_\theta - u_{\theta'}\|_{C^{\deg \mathfrak{F}}(\Omega)} \\ &\qquad \max_{1 \leq k \leq N(P)} \prod_{i,j} \max\{\|u_\theta\|_{C^{|\alpha_{i,j}|}(\Omega)}, \|u_{\theta'}\|_{C^{|\alpha_{i,j}|}(\Omega)}\}^{I(i,j,k)}. \end{aligned}$$

From Lemma 1.6, we deduce that

$$\begin{aligned} &\max_{1 \leq k \leq N(P)} \prod_{i,j} \max\{\|u_\theta\|_{C^{|\alpha_{i,j}|}(\Omega)}, \|u_{\theta'}\|_{C^{|\alpha_{i,j}|}(\Omega)}\}^{I(i,j,k)} \\ &\qquad \leq C_{\deg \mathfrak{F}} (D+1)^{H \deg \mathfrak{F}} (1 + \max(\|\theta\|_2, \|\theta'\|_2))^{H \deg \mathfrak{F}}. \end{aligned}$$

Combining the last two inequalities with Lemma 1.7 to bound $\|u_\theta - u_{\theta'}\|_{C^{\deg \mathfrak{F}}(\Omega)}$, we obtain

$$\begin{aligned} |\mathfrak{F}(u_\theta, x) - \mathfrak{F}(u_{\theta'}, x)| \leq N(P) &\max_{1 \leq k \leq N(P)} \|\phi_k\|_{\infty,\bar{\Omega}} \deg(\mathfrak{F}) \tilde{C}_{\deg \mathfrak{F}, H} (1 + d_1 M(\Omega)) \|\theta - \theta'\|_2 \\ &(D+1)^{H(1+(1+H)\deg \mathfrak{F})} (1 + \max(\|\theta\|_2, \|\theta'\|_2)^{H(1+(1+H)\deg \mathfrak{F})}. \end{aligned}$$

Hence, for all $\theta, \theta' \in B_2(0, n^{\alpha_1})$,

$$|\mathfrak{F}(u_\theta, x)^2 - \mathfrak{F}(u_{\theta'}, x)^2| \leq C_1 n^{\beta_1} \|\theta - \theta'\|_2,$$

where $\beta_1 = H(1 + (2+H)\deg \mathfrak{F})\alpha_1$.

Recall that it must verify that $\alpha_1 + \beta_1 < \beta_2 < 1/2$. And this is true for $\beta_2 = [2 + H(1 + (2+H)\deg \mathfrak{F})]\alpha_1$ and $\alpha_1 < [2 + H(1 + (2+H)\deg \mathfrak{F})]^{-1}/2$

$\square$

**Theorem 3.8** (Risk-consistency of ridge PINNs)**.** Let the number of hidden layers be $H \geq 2$. Assume that the initial/boundary condition $h$ is Lipschitz and that $\mathfrak{F}_1, \ldots, \mathfrak{F}_M$ are polynomial operators. Assuming that the ridge parameter is of the form

$$\lambda_{(\text{ridge})} := \frac{1}{\min(n_2, n_3)^{\kappa}}, \qquad \text{where} \qquad \kappa := \frac{1}{12 + 4H(1 + (2+H)\max_k \deg(\mathfrak{F}_k))}.$$

Then, almost surely,

$$\lim_{n_2, n_3 \to \infty} \lim_{n \to \infty} R_{n_1}(u_{\hat{\theta}_n^{(\text{ridge})}(n_2,n_3,D)}) = \inf_{u \in \mathrm{NN}_H^D} R_{n_1}(u).$$

In other words, this theorem assures us that minimizing the ridge empirical risk (3.3) over $\Theta_{H,D}$ amounts to minimizing the theoretical risk (1.20) over $\Theta_{H,D}$ in the asymptotic regime where $n_2, n_3$ tends to infinity. The following results ensure that the class of neural networks $\mathrm{NN}_H^D \subseteq C^\infty(\bar{\Omega}, \mathbb{R}^{d_2})$ does not introduce any asymptotic bias in the width $D$.

*Proof.* Since each $\mathfrak{F}_k$ is a polynomial operator (see Definiton 3.5), it takes the form

$$\mathfrak{F}_k(u, x) := \sum_{\ell=1}^{N(P_k)} \phi_{\ell,k} \prod_{i=1}^{d_2} \prod_{j=1}^{s_k} (\partial^{\alpha_{i,j,k}} u_i(x))^{I_k(i,j,\ell)}.$$

Let $u_0 = 0 \in \mathrm{NN}_H^D$ be a neural network with parameters $\theta = (0, \ldots, 0)$. Obviously, $\widehat{R}^{(\text{ridge})}_{n_1,n_2,n_3}(u_0) = \widehat{R}_{n_1,n_2,n_3}(u_0)$. Moreover,

$$\begin{aligned} \widehat{R}_{n_1,n_2,n_3}(u_0) &\leq \frac{\lambda_1}{n_1} \sum_{i=1}^{n_1} \|Y_i\|_2^2 + \lambda_2 \|h\|_\infty + \frac{\lambda_3}{n_3} \sum_{k=1}^{M} \sum_{\ell=1}^{n_3} \|\mathfrak{F}_k(u_\theta, X_\ell^{(3)})\|_2^2 \\ &\leq \frac{\lambda_1}{n_1} \sum_{i=1}^{n_1} \|Y_i\|_2^2 + \lambda_2 \|h\|_\infty + \frac{\lambda_3}{n_3} \sum_{k=1}^{M} \sum_{\ell=1}^{n_3} \|\phi_{\ell,k}\|_{\infty,\bar{\Omega}} := I. \end{aligned} \tag{3.4}$$

We notice that $I$ does not depend on $\lambda_{(ridge)}, n_2, n_3$.

Let $\{\hat{\theta}_n^{(\text{ridge})}(n_2, n_3, D)\}_{n \in \mathbb{N}}$ be any minimizing sequence of the empirical risk of the ridge PINN, i.e., $\lim_{n \to \infty} \widehat{R}^{(\text{ridge})}_{n_1,n_2,n_3}(u_{\hat{\theta}_n^{(\text{ridge})}(n_2,n_3,D)}) = \inf_{\theta \in \Theta_H^D} \widehat{R}^{(\text{ridge})}_{n_1,n_2,n_3}(u_\theta)$. Let $N_{2,3} := \min(n_2, n_3)$. Let us define the following three sets:

$$\begin{aligned} \mathcal{E}_1(N_{2,3}) &:= \{\theta \in \Theta_H^D : \|\theta\|_2 \geq N_{2,3}^{\kappa}\} \\ \mathcal{E}_2(N_{2,3}) &:= \{\theta \in \Theta_H^D : N_{2,3}^{\kappa/4} \leq \|\theta\|_2 \leq N_{2,3}^{\kappa}\} \\ \mathcal{E}_3(N_{2,3}) &:= \{\theta \in \Theta_H^D : N_{2,3}^{\kappa/4} \leq \|\theta\|_2\}. \end{aligned}$$

Notice that $\Theta_H^D = \mathcal{E}_1 \sqcup \mathcal{E}_2 \sqcup \mathcal{E}_2$. The proof relies on the argument that for any $n_2, n_3$, given a sufficiently large $n$, $\hat{\theta}_n^{(\text{ridge})}(n_2, n_3, D) \in \mathcal{E}_2 \cup \mathcal{E}_3$. Furhtermore, when $N_{2,3}$ is sufficiently large, then $\widehat{R}^{(\text{ridge})}_{n_1,n_2,n_3}$ is close to the theoretical risk $R$. Let us divide the proof in four steps.

**Step 1.** We start by observing that, for any $\theta \in \mathcal{E}_1(N_{2,3})$, we have $\widehat{R}^{(\text{ridge})}_{n_1,n_2,n_3}(u_\theta) \geq \lambda_{(\text{ridge})}\|\theta\|_2^2 \geq \lambda_{(\text{ridge})}N_{2,3}^{2\kappa} = N_{2,3}^{\kappa}$. Assuming that $N_{2,3}$ is sufficiently large to verify $N_{2,3}^{\kappa} - 1 \geq I \overset{(3.4)}{\geq} \widehat{R}_{n_1,n_2,n_3}(u_0)$. Therefore,

$$\inf_{\theta\in\mathcal{E}_3(N_{2,3})} \widehat{R}^{(\text{ridge})}_{n_1,n_2,n_3}(u_\theta) \leq \widehat{R}^{(\text{ridge})}_{n_1,n_2,n_3}(u_0) \leq N^{\kappa}_{n_1,n_2,n_3} - 1 \leq \inf_{\theta\in\mathcal{E}_1(N_{2,3})} \widehat{R}^{(\text{ridge})}_{n_1,n_2,n_3}(u_\theta) - 1.$$

This shows that, for all $N_{2,3}$ and $n$ large enough, $\hat{\theta}^{(\text{ridge})}_n(n_2, n_3, D) \notin \mathcal{E}_1(N_{2,3})$.

**Step 2.** Invoking Proposition 3.6 and Proposition 3.7 with $\alpha_1 = \kappa$ and $\beta_2 = (2 + H(1 + (2 + H)\max_k \deg \mathfrak{F}_k))\alpha_1$, then by Theorem 2.19, we know that, almost surely, there exist $N \in \mathbb{N}$ such that for all, $N_{2,3} \geq N$,

$$\begin{aligned}\sup_{\theta\in\mathcal{E}_2(N_{2,3})\cup\mathcal{E}_3(N_{2,3})} &\left|\frac{1}{n_2}\sum_{i=1}^{n_b} \|u_\theta(X_i^{(2)}) - h(X_i^{(2)})\|_2^2 - \mathbb{E}\|u_\theta(X^{(2)}) - h(X^{(2)})\|_2^2\right| \\ &\leq \log^2(N_{2,3})N_{2,3}^{\beta_2-1/2}\end{aligned} \tag{3.5}$$

and, for each $1 \leq k \leq M$,

$$\sup_{\theta\in\mathcal{E}_2(N_{2,3})\cup\mathcal{E}_3(N_{2,3})} \left|\frac{1}{c}\sum_{i=1}^{c} \mathfrak{F}_k(u_\theta, X_\ell^{(r)})^2 - \frac{1}{|\Omega|}\int_\Omega \mathfrak{F}_k(u_\theta, x)^2 dx\right| \leq \log^2(N_{2,3})N_{2,3}^{\beta_2-1/2}. \tag{3.6}$$

Then, almost surely, for all $N_{2,3}$ large enough and all $\theta \in \mathcal{E}_2(N_{2,3})$,

$$|\underbrace{\widehat{R}^{(\text{ridge})}_{n_1,n_2,n_3}(u_\theta) - \lambda_{(\text{ridge})}\|\theta\|_2^2}_{\widehat{R}_{n_1,n_2,n_3}(u_\theta)} - R(u_\theta)| \leq (M+1)\log^2(N_{2,3})N_{2,3}^{\beta_2-1/2} \tag{3.7}$$

But, for all $\theta \in \mathcal{E}_2(N_{2,3})$, $\lambda_{(\text{ridge})}\|\theta\|_2^2 \geq \lambda_{(\text{ridge})}N_{2,3}^{\kappa/2} = N_{2,3}^{-\kappa/2}$. We conclude that, for all $N_{2,3}$ large enough and for all $\theta \in \mathcal{E}_2(N_{2,3})$,

$$\begin{aligned}|R_{n_1}(u_\theta) - \widehat{R}_{n_1,n_2,n_3}(u_\theta)| &\leq (M+1)\log^2(N_{2,3})N_{2,3}^{\beta_2-1/2} \\ &\Rightarrow \widehat{R}_{n_1,n_2,n_3}(u_\theta) \geq R(u_\theta) - (M+1)\log^2(N_{2,3})N_{2,3}^{\beta_2-1/2} \\ &\Rightarrow \widehat{R}^{(\text{ridge})}_{n_1,n_2,n_3}(u_\theta) \geq R(u_\theta) + \lambda_{(\text{ridge})}\|\theta\|_2^2 - (M+1)\log^2(N_{2,3})N_{2,3}^{\beta_2-1/2} \\ &\Rightarrow \widehat{R}^{(\text{ridge})}_{n_1,n_2,n_3}(u_\theta) \geq R(u_\theta) + \underbrace{N_{2,3}^{-\kappa/2} - (M+1)\log^2(N_{2,3})N_{2,3}^{\beta_2-1/2}}.\end{aligned}$$

Assuming that $-\kappa/2 > \beta_2 - 1/2$, then the right hand side is positive for $N_{2,3}$ sufficiently large, and we obtain

$$\widehat{R}^{(\text{ridge})}_{n_1,n_2,n_3}(u_\theta) \geq R_{n_1}(u_\theta).$$

**Step 3.** For all $\theta \in \mathcal{E}_3(N_{2,3})$, $\lambda_{(\text{ridge})}\|\theta\|_2^2 \leq N_{2,3}^{-\kappa/2}$. Using inequalities (3.5) and (3.6), we deduce that, almost surely, for all $N_{2,3}$ large enough

$$\begin{aligned}|\widehat{R}^{(\text{ridge})}_{n_1,n_2,n_3}(u_\theta) - R_{n_1}(u_\theta)| &\leq |\widehat{R}^{(\text{ridge})}_{n_1,n_2,n_3}(u_\theta) - \lambda_{\text{ridge}}\|\theta\|_2^2 - R_{n_1}(u_\theta)| + \lambda_{(\text{ridge})}\|\theta\|_2^2 \\ &\leq (M+1)\log^2(N_{2,3})N_{2,3}^{\beta_2-1/2} + N_{2,3}^{-\kappa/2} \\ &\leq (M+2)\log^2(N_{2,3})N_{2,3}^{-\kappa/2}.\end{aligned}$$

**Step 4.** Fix $\epsilon > 0$. Let $(\theta_n)_{n\in\mathbb{N}}$ be a minimizing sequence of the theoretical risk function $R_{n_1}$, i.e., $\lim_{n\to\infty} R_{n_1}(u_{\theta_n}) = \inf_{\theta\in\Theta_H^D} R_{n_1}(u_\theta)$. By definition, there exists some $P_\epsilon \in \mathbb{N}$ such that $|R_{n_1}(u_{\theta_n}) - \inf_{\theta\in\Theta_H^D} R_{n_1}(u_\theta)| \leq \epsilon$.

For some fixed $N_{2,3}$, according to Step 1, we have for all $n$ large enough, $\hat{\theta}_n^{(\text{ridge})}(n_2, n_3, D) \in \mathcal{E}_2 \cup \mathcal{E}_3$. According to Step 2 and Step 3,

$$R(u_{\hat{\theta}_n^{(\text{ridge})}(n_2,n_3,D)}) \leq \widehat{R}^{(\text{ridge})}_{n_1,n_2,n_3}(u_{\hat{\theta}_n^{(\text{ridge})}(n_2,n_3,nD)}) + (M+2)\log^2(N_{2,3})N_{2,3}^{-\kappa/2}. \tag{3.8}$$

Now, by definition of the minimizing sequence $(\hat{\theta}_n^\lambda(a,b,c))_{n\in\mathbb{N}}$, for all $n$ large enough,

$$\widehat{R}^{(\text{ridge})}_{n_1,n_2,n_3}(u_{\hat{\theta}_n^{(\text{ridge})}(n_2,n_3,D)}) \leq \inf_{\theta\in\Theta_H^D} \widehat{R}^{(\text{ridge})}_{n_1,n_2,n_3}(u_\theta) + \epsilon. \tag{3.9}$$

Moreover, according to Step 3,

$$\begin{aligned}\inf_{\theta\in\mathcal{E}_2(N_{2,3})\cup\mathcal{E}_3(N_{2,3})} \widehat{R}^{(\text{ridge})}_{n_1,n_2,n_3}(u_\theta) &\leq \inf_{\theta\in\mathcal{E}_3(N_{2,3})} \widehat{R}^{(\text{ridge})}_{n_1,n_2,n_3}(u_\theta) \\ &\leq \inf_{\theta\in\mathcal{E}_3(N_{2,3})} R(u_\theta) + (M+2)\log^2(N_{2,3})N_{2,3}^{-\kappa/2}.\end{aligned} \tag{3.10}$$

Observe that for all $N_{2,3}$ large enough, $\theta_{P_\epsilon} \in \mathcal{E}_3(N_{2,3})$. Therefore, $\inf_{\theta\in\mathcal{E}_3(N_{2,3})} R_{n_1}(u_\theta) \leq R_{n_1}(u_{\theta_{P_\epsilon}})$. Combining the previous inequalities, we conclude that, almost surely, for all $N_{2,3}$ large enough for all $n$ large enough,

$$\begin{aligned}R_{n_1}(u_{\hat{\theta}_n^{(\text{ridge})}(n_2,n_3,D)}) &\overset{(3.8)}{\leq} \widehat{R}^{(\text{ridge})}_{n_1,n_2,n_3}(u_{\hat{\theta}_n^{(\text{ridge})}(n_2,n_3,D)}) + (M+2)\log^2(N_{2,3})N_{2,3}^{-\kappa/2} \\ &\overset{(3.9)}{\leq} \inf_{\theta\in\Theta_H^D} \widehat{R}^{(\text{ridge})}_{n_1,n_2,n_3}(u_\theta) + \epsilon + (M+2)\log^2(N_{2,3})N_{2,3}^{-\kappa/2} \\ &\overset{(3.10)}{\leq} \inf_{\theta\in\Theta_H^D} R_{n_1}(u_\theta) + \epsilon + 2(M+2)\log^2(N_{2,3})N_{2,3}^{-\kappa/2}.\end{aligned}$$

Since $\epsilon$ is arbitrary and $N_{2,3}$ is arbitrarily large, then, almost surely,

$$\lim_{n_2,n_3\to\infty}\lim_{n\to\infty} R_{n_1}(u_{\hat{\theta}_n^{(\text{ridge})}(n_2,n_3,D)}) = \inf_{\theta\in\Theta_H^D} R_{n_1}(u_\theta).$$

□

**Corollary 3.9.** The ridge is asymptotically unbiased almost surely,

$$\lim_{D\to\infty} \lim_{n_2,n_3\to\infty} \lim_{n\to\infty} R_{n_1}(u_{\hat{\theta}_n^{(\text{ridge})}(n_2,n_3,D)}) = \inf_{u\in C^\infty(\overline{\Omega},\mathbb{R}^{d_2})} R_{n_1}(u).$$

*Proof.* The result is a direct consequence of Theorem 2.19, Proposition 1.5 and of the continuity of $R$ with respect to the $C^K(\Omega)$ norm. □

## 3.5. Strong convergence of regularized PINNs

This section proves the strong consistency of Physics-Informed Neural Networks (PINNs). To this end, we augment the empirical risk function with a Sobolev norm to control the solution and prevent singularities.

**Definition 3.10** (Affine operator)**.** The operator $\mathfrak{F}$ is affine of order $K$ if there exists $A_\alpha \in C^\infty(\bar{\Omega},\mathbb{R}^{d_2})$ and $B \in C^\infty(\bar{\Omega},\mathbb{R})$ such that for all $x\in\Omega$ and all $u \in H^K(\Omega,\mathbb{R}^{d_2})$,

$$\mathfrak{F}(u,x) = \mathfrak{F}^{(lin)}(u,x) + B(x),$$

where $\mathfrak{F}^{(lin)}(u,x) = \sum_{|\alpha|\le K}\langle A_\alpha(x), \partial^\alpha u(x)\rangle$ is linear.

Note that affine operators of order $K$ are polynomial operators of degree $K+1$ that are extended from smooth functions to the whole Sobolev space $H^K(\Omega,\mathbb{R}^{d_2})$.

**Definition 3.11** (Regularized PINNs)**.** The regularized theoretical risk function is

$$R^{(reg)}_{n_1}(u) := R_{n_1}(u) + \lambda_{(\text{reg})}\|u\|^2_{H^{m+1}(\Omega)}, \tag{3.11}$$

where $R_a$ is the original theoretical risk as defined in (1.20), $\lambda_{(\text{reg})}\in\mathbb{R}^+$ is a importance factor, and $m\in\mathbb{N}$ is the regularity of the Sobolev norm. The correspoding empirical risk function is

$$\widehat{R}^{(reg)}_{n_1,n_2,n_3}(u_\theta) = \widehat{R}_{n_1,n_2,n_3}(u_\theta) + \lambda_{(\text{ridge})}\|\theta\|_2^2 + \frac{\lambda_{(\text{reg})}}{n_3}\sum_{\ell=1}^{n_3}\sum_{|\alpha|\le m+1}\|\partial^\alpha u_\theta(X_\ell^{(3)})\|_2^2, \tag{3.12}$$

where $\widehat{R}_{n_1,n_2,n_3}$ is the empirical risk function defined in (1.19).

Note that the regularization term intentionally uses the same variables as in $\widehat{L}_3$. This choice allows us to treat the regularization as an integral part of the objective function in the subsequent analysis.

The next proposition is used to characterize the unique minimizer of the theoretical risk as the unique minimizer that satisfies an equality of the type $\mathcal{A}(u,v) = \mathcal{B}(v)$ for all $v\in H^{m+1}$.

**Proposition 3.12** (Characterization of the unique minimizer of $R_{n_1}^{(reg)}$). Fix $n_1 \in \mathbb{N}$. Assume that $\mathfrak{F}_1, \dots, \mathfrak{F}_M$ are affine operators of order $K$, $\lambda_t > 0$ and $m \geq \max\{\lfloor d_1/2 \rfloor, K\}$. Then, the regularized theoretical risk $R_{n_1}^{(reg)}$ (defined in (3.11)) has a unique minimizer $\hat{u}_n$ over $H^{m+1}(\Omega, \mathbb{R}^{d_2})$. Moreover, $\hat{u}_{n_1}$ is the unique element of $H^{m+1}(\Omega, \mathbb{R}^{d_2})$ that satisfies

$$\forall v \in H^{m+1}(\Omega, \mathbb{R}^{d_2}), \quad \mathcal{A}_{n_1}(\hat{u}_{n_1}, v) = \mathcal{B}_{n_1}(v),$$

where

$$\mathcal{A}_{n_1}(\hat{u}_{n_1}, v) := \frac{\lambda_1}{n_1} \sum_{i=1}^{n_1} \langle \tilde{\Pi}(\hat{u}_{n_1})(X_i), \tilde{\Pi}(v)(X_i) \rangle + \lambda_2 \mathbb{E} \langle \tilde{\Pi}(\hat{u}_{n_1})(X^{(2)}), \tilde{\Pi}(v)(X^{(2)}) \rangle$$
$$+ \frac{\lambda_3}{|\Omega|} \sum_{k=1}^{M} \int_\Omega \mathfrak{F}_k^{(lin)}(\hat{u}_{n_1}, x) \mathfrak{F}_k^{(lin)}(v, x) dx + \frac{\lambda_{(\text{reg})}}{|\Omega|} \sum_{|\alpha| \leq m+1} \int_\Omega \langle \partial^\alpha \hat{u}_{n_1}(x), \partial^\alpha v(x) \rangle dx,$$

and

$$\mathcal{B}_{n_1}(v) := \frac{\lambda_1}{n_1} \sum_{i=1}^{n_1} \langle Y_i, \tilde{\Pi}(v)(X_i) \rangle + \lambda_2 \mathbb{E} \langle \tilde{\Pi}(v)(X^{(2)}), h(X^{(2)}) \rangle - \frac{\lambda_3}{|\Omega|} \sum_{k=1}^{M} \int_\Omega B_k(x) \mathfrak{F}_k^{(lin)}(v, x) dx,$$

where $\tilde{\Pi} : H^{m+1}(\Omega, \mathbb{R}^{d_2}) \to C(\Omega, \mathbb{R}^{d_2})$ is the Sobolev embedding (Theorem A.1) and $B_k$ is the *intercept* of the affine operator (Definition 3.10).

*Proof.* To define the pointwise evaluation $u(X_i)$ for a function $u \in H^{m+1}(\Omega, \mathbb{R}^{d_2})$, we leverage the Sobolev embedding theorem. As $\Omega$ is a bounded Lipschitz domain and $m + 1 > d_1/2$ holds, Theorem A.1 guarantees that $u$ has a unique continuous representative, denoted $\tilde{\Pi}(u)$. The pointwise evaluation of $u$ at $X_i$ is therefore defined as the evaluation of this continuous representative: $u(X_i) := \tilde{\Pi}(u)(X_i)$.

By assumption, $\mathfrak{F}_k(u, \cdot) = \mathfrak{F}_k^{(lin)}(u, \cdot) + B_k$ where $\mathfrak{F}_k^{(lin)}(u, \cdot) = \sum_{|\alpha| \leq K} \langle A_{k,\alpha}, \partial^\alpha u \rangle$ and $A_{k,\alpha} \in C^\infty(\bar{\Omega}, \mathbb{R}^{d_1})$. Next, consider the symmetric bilinear form $\mathcal{A}_{n_1}(u, v)$ defined for $u, v \in H^{m+1}(\Omega, \mathbb{R}^{d_2})$, along with the linear form $\mathcal{B}_{n_1}(u)$ defined for all $u \in H^{m+1}(\Omega, \mathbb{R}^{d_2})$. Observe that

$$\mathcal{A}_{n_1}(u, u) - 2\mathcal{B}_{n_1}(u) = \frac{\lambda_1}{n_1} \sum_{i=1}^{n_1} \langle u(X_i) - 2Y_i, u(X_i) \rangle + \lambda_2 \mathbb{E} \langle u(X^{(2)}), u(X^{(2)}) - 2h(X^{(2)}) \rangle$$
$$+ \frac{\lambda_3}{|\Omega|} \sum_{k=1}^{M} \int_\Omega \mathfrak{F}_k^{(lin)}(u, x)[\mathfrak{F}_k^{(lin)}(u, x) - 2B_k(x)] dx + \frac{\lambda_{(\text{reg})}}{|\Omega|} \sum_{|\lambda| \leq m+1} \int_\Omega \langle \partial^\lambda u(x), \partial^\lambda u(x) \rangle$$
$$= R_{n_1}^{(reg)}(u) - \frac{\lambda_1}{n_1} \sum_{i=1}^{n_1} \|Y_i\|_2^2 - \lambda_2 \mathbb{E} \|h(X^{(e)})\|_2^2 - \frac{1}{|\Omega|} \sum_{k=1}^{M} \int_\Omega B_k(x)^2 dx.$$

In addition, $\mathcal{A}_{n_1}(u, u) \geq \lambda_{(\text{reg})} \|u\|_{H^{m+1}(\Omega)}^2$, where $\lambda_{(\text{reg})} > 0$, so that $\mathcal{A}_{n_1}$ is coercive on the normed space $(H^{m+1}(\Omega), \|\cdot\|_{H^{m+1}(\Omega)})$. Now, let us prove that the operators $\mathcal{A}_n$ and $\mathcal{B}_n$ are continuous. Indeed, we can prove that they are bounded, since $m+1 > \max\{d_1/2, K\}$, one has that

$$|\mathcal{A}_{n_1}(u,v)| \leq \left((\lambda_1+\lambda_2)C_\Omega^2 + \lambda_3 \sum_{i=k}^{M}\left(\sum_{|\alpha|\leq K}\|A_{k,\alpha}\|_{\infty,\Omega}\right)^2 + \lambda_{(\mathrm{reg})}\right)\|u\|_{H^{m+1}(\Omega)}\|v\|_{H^{m+1}(\Omega)}$$

and

$$|\mathcal{B}_{n_1}(u)| \leq C_\Omega\left(\frac{\lambda_1}{n_1}\sum_{i=1}^{n_1}\|Y_i\|_2 + \lambda_2\|h\|_\infty + \lambda_3\sum_{k=1}^{M}\left(\|B_k\|_{\infty,\Omega}\sum_{|\alpha|\leq K}\|A_{k,\alpha}\|_{\infty,\Omega}\right)\right)\|u\|_{H^{m+1}(\Omega)}.$$

and due to linearity of the operators, they are continuous. Therefore, we can apply the Lax-Milgram theorem, which guarantees that there exists a unique $\hat{u} \in H^{m+1}(\Omega,\mathbb{R}^{d_2})$ such that $\mathcal{A}_{n_1}(\hat{u},\hat{u}) - 2\mathcal{B}_{n_1}(\hat{u}) = \min_{u\in H^{m+1}(\Omega,\mathbb{R}^{d_2})}\mathcal{A}_{n_1}(u,u) - 2\mathcal{B}_{n_1}(u)$. This, along with the above expression of $\mathcal{A}_{n_1}(u,u) - 2\mathcal{B}_{n_1}(u)$, directly implies that $\hat{u}$ is the unique minimizer of $R_{n_1}^{(reg)}$ over $H^{m+1}(\Omega,\mathbb{R}^{d_2})$. Futhermore, Lax-Milgram states that $\hat{u}$ is the unique element of $H^{m+1}(\Omega,\mathbb{R}^{d_2})$ such that $\mathcal{A}_{n_1}(\hat{u},v) = \mathcal{B}_{n_1}(v)$ for all $v \in H^{m+1}(\Omega,\mathbb{R}^{d_2})$. □

**Proposition 3.13.** Let $(u_n)_{n\in\mathbb{N}} \in C^\infty(\bar{\Omega},\mathbb{R}^{d_2})$ be a sequence of smooth functions such that $\lim_{n\to\infty} R_{n_1}^{(\mathrm{reg})}(u_n) = \inf_{u\in C^(\bar{\Omega},\mathbb{R}^{d_2})} R_{n_1}^{(\mathrm{reg})}(u)$ assuming $\lambda_{(\mathrm{reg})} > 0$ and $m \geq \max\{\lfloor d_1/2\rfloor, K\}$. Then, $\lim_{n\to\infty}\|u_n - \hat{u}\|_{H^m(\Omega)} = 0$, where $\hat{u}$ is the unique minimizer of $R_{n_1}^{(\mathrm{reg})}$ over $H^{m+1}(\Omega,\mathbb{R}^{d_2})$.

*Proof.* Let $\hat{u}$ be the unique minimizer of $R_{n_1}^{(\mathrm{reg})}$ over $H^{m+1}(\Omega,\mathbb{R}^{d_2})$ given by Proposition 3.12. We note that $C^\infty(\bar{\Omega},\mathbb{R}^{d_2})$ is dense in $H^{m+1}(\Omega,\mathbb{R}^{d_2})$, together with the continuity of the function $R_{n_1}^{(\mathrm{reg})} : H^{m+1}(\Omega,\mathbb{R}^{d_2}) \to \mathbb{R}$ with respect to the norm $H^{m+1}(\Omega)$ yields

$$\inf_{u\in C^\infty(\bar{\Omega},\mathbb{R}^{d_2})} R_{n_1}^{(\mathrm{reg})}(u) = \inf_{u\in H^{m+1}(\Omega,\mathbb{R}^{d_2})} R_{n_1}^{(\mathrm{reg})}(u) = R_{n_1}(\hat{u}). \tag{3.13}$$

Next, we will rewrite the theoretical regularized risk using a different formulation. Observe that

$$R_{n_1}^{(\mathrm{reg})}(u) = F(u) + \frac{\lambda_3}{|\Omega|}I(u),$$

where

$$F(u) := \frac{\lambda_1}{n_1}\sum_{i=1}^{n_1}\|\tilde{\Pi}(u)(X_i) - Y_i\|_2^2 + \lambda_2\mathbb{E}\|\tilde{\Pi}(u)(X^{(2)}) - h(X^{(2)}\|_2^2,$$

$$I(u) := \int_\Omega L((\partial_{i_1,\ldots,i_{m+1}}^{m+1}u(x))_{1\leq i_1,\ldots,i_{m+1}\leq d_1},\ldots,u(x),x)dx.$$

The function $L$ satisfies

$$L(x^{(m+1)},\ldots,x^{(0)},z) := \sum_{k=1}^{M}\left(B_k(z) + \sum_{|\alpha|\leq K}\langle A_{k,\alpha}(z), x_\alpha^{(|\alpha|)}\rangle\right)^2 + \lambda_{(\mathrm{reg})}\sum_{j=0}^{m+1}\|x^{(j)}\|_2^2.$$

Here, we concatenate all the derivatives of order $j$, i.e., $(\partial^j_{i_1,\dots,i_j} u(x))_{1\le i_1,\dots,i_j\le d_1}$ and denote it by $x^{(j)} \in \mathbb{R}^{\binom{d_1+j-1}{j-1}d_2}$. Clearly, $L \ge 0$, and since, $m+1 > K$, there is no dependence on $x^{(i+1)}$ hiding inside the square-sum part and hence $L$ is convex. According to Lemma A.5, the function $I$ must be weak lower-semi continuous on $H^{m+1}(\Omega, \mathbb{R}^{d_2})$ (see Definition A.4).

Let us assume that the oppositive and proceed by contradiction. Assume that there is a sequence $\{u_n\}_{n\in\mathbb{N}}$ of functions such that $u_n \in C^\infty(\bar{\Omega}, \mathbb{R}^{d_2})$, such that $\lim_{n\to\infty} R^{(\mathrm{reg})}_{n_1}(u_n) = R^{(\mathrm{reg})}_{n_1}(\hat{u})$ and that it does not converge to $\hat{u}$ with respect to the $H^m(\Omega)$ norm. Upong passing to a subsequence, there exist $\epsilon > 0$ such that for all $p \ge 0$ with $\|u_n - \hat{u}\|_{H^m(\Omega)} \ge \epsilon$.

Since $R^{(reg)}_{n_1}(u_n) \ge \lambda_{(\mathrm{reg})}\|u_n\|_{H^{m+1}(\Omega)}$, $\lambda_t > 0$ and $\{u_n\}_{p\in\mathbb{N}}$ is a minimizing sequence, then $\{u_n\}_{n\in\mathbb{N}}$ is bounded in $H^{m+1}(\Omega)$. Therefore, we can pass to a subsequence that converges to a limit, denoting it by $u_\infty$, both weakly in $H^{m+1}(\Omega, \mathbb{R}^{d_2})$ and with respect to the $H^m(\Omega)$ norm. Since $I$ is weakly lower-semi continuous, we deduce that

$$\lim_{n\to\infty} I(u_n) \ge I(u_\infty). \tag{3.14}$$

Recall the definition of $\tilde{\Pi}$ in Theorem A.1, we know that there exists a constant $C_\Omega > 0$ such that $\|u_n - \tilde{\Pi}(u_\infty)\|_{\infty,\Omega} = \|\tilde{\Pi}(u_p - u_\infty)\|_{\infty,\Omega} \le C_\Omega \|u_p - u_\infty\|_{H^m(\Omega)}$. We deduce that $\lim_{n\to\infty} F(u_n) = F(u_\infty)$. Therefore, combining result (3.13) and (3.14), we deduce that $\lim_{n\to\infty} R^{(reg)}_{n_1}(u_n) \ge R^{(\mathrm{reg})}_{n_1}(u_\infty)$. However, recalling that $\lim_{n\to\infty} R^{(\mathrm{reg})}_{n_1}(u_n) = R^{(\mathrm{reg})}_{n_1}(\hat{u})$ and that $\hat{u}$ is the unique minimizer of $R^{(\mathrm{reg})}_{n_1}$, we conclude that $u_\infty = \hat{u}$. This contradicts our assumption. □

The next theorem asserts that under some mild conditions, that our minimizing sequence of the regularized empirical risk converges to the real solution of the PDE.

**Theorem 3.14** (Strong convergence of regularized PINNs)**.** Assume that $\mathfrak{F}_1, \dots, \mathfrak{F}_M$ are affine operators of order $K$. Assume, in addition, that $\lambda_{(\mathrm{reg})} > 0$, $m \ge \max(\lfloor d_1/2 \rfloor, K)$ and the condition function $h$ is Lipschitz. Let $\{\hat{\theta}^{(reg)}_n(n_2, n_3, D)\}_{n\in\mathbb{N}}$ be a minimizing sequence of the regularized empirical risk function $\widehat{R}^{(\mathrm{reg})}_{n_1,n_2,n_3}(u_\theta)$ over the class $\mathrm{NN}^D_H = \{u_\theta : \theta \in \Theta_{H,D}\}$ where $H \ge 2$. Then,

$$\lambda_{(ridge)} = \min(n_2, n_3)^{-\kappa}, \qquad \text{where} \quad \kappa = \frac{1}{12 + 4H(1 + (2+H)(m+2))},$$

one has, almost surely,

$$\lim_{D\to\infty} \lim_{n_2,n_3\to\infty} \lim_{n\to\infty} \|u_{\hat{\theta}^{(\mathrm{reg})}_n(n_2,n_3,D)} - \hat{u}\|_{H^m(\Omega)} = 0,$$

where $\hat{u}$ is the unique minimizer of $R^{(\mathrm{reg})}_{n_1}$ over $H^{m+1}(\Omega, \mathbb{R}^{d_2})$.

In particular, if we consider the case where $n_1 = 0$ or $\lambda_1 = 0$ and assume that the PDE system admits a unique solution $u^\star \in H^{m+1}(\Omega, \mathbb{R}^{d_2})$. Then, almost surely, we have

$$\lim_{\lambda_{(\mathrm{reg})}\to 0} \lim_{D\to\infty} \lim_{n_2,n_3\to\infty} \lim_{n\to\infty} \|u_{\hat{\theta}^{(\mathrm{reg})}_n(n_2,n_3,D,\lambda_{(\mathrm{reg})})} - u^\star\|_{H^m(\Omega)} = 0.$$

*Proof.* The Sobolev regularization can be viewed as a specific form of PINN regularization where the penalty term is a polynomial differential operator of degree $m+2$. According to Theorem 2.19, the risk associated with a minimizer of the Sobolev-regularized loss converges to the minimum theoretical risk. Furthermore, Proposition 3.12 guarantees that the regularized risk $R_{n_1}^{(\mathrm{reg})}$, has a unique minimizer. Proposition 3.13 then assures that any sequence of minimizers of the empirical loss must converge to this unique minimizer. This convergence completes the proof of the theorem.

For the second part, we denote the regularized risk by $R^{(\mathrm{reg})}$ instead of $R_{n_1}^{(\mathrm{reg})}$ because there is no data involved. Moreover, we denote the unique minimizer $\hat{u}(\lambda_{(\mathrm{reg})})$ to make clear the dependence in the hyperparameter $\lambda_{(\mathrm{reg})}$. Notice that $u^\star$ due to being a solution of the PDE, satisfies $R(u^\star)=0$.

We proceed by contradiction, assume that $\lim_{\lambda_{(\mathrm{reg})}\to 0}\|\hat{u}(\lambda_{(\mathrm{reg})})-u^\star\|_{H^m(\Omega)}\neq 0$. Then, there exists a subsequence $\{\lambda_{(\mathrm{reg}),n}\}_{n\in\mathbb{N}}$ such that $\lim_{n\to\infty}\lambda_{(\mathrm{reg}),n}=0$ such that $n\geq 0$, $\|\hat{u}(\lambda_{(\mathrm{reg}),n})-u^\star\|_{H^m(\Omega)}\geq\epsilon$.

Note that $\|\hat{u}(\lambda_{(\mathrm{reg}),n})\|_{H^{n+1}(\Omega)}\leq R^{(\mathrm{reg})}(u^\star)/\lambda_{(\mathrm{reg}),n}=\|u^\star\|_{H^{m+1}(\Omega)}$. Rellich-Kondrachov compactness Theorem A.3 ensures that upon passing to a subsequence $\{\hat{u}(\lambda_{(\mathrm{reg}),n})\}_{n\in\mathbb{N}}$ converges with respect to the $H^m(\Omega)$ norm to a function $u_\infty\in H^{m+1}(\Omega,\mathbb{R}^{d_2})$. Since $m\geq K$, the theoretical risk $R$ is continuous with respect to the $H^m(\Omega)$ norm (Theorem A.1) and we have that $R(u_\infty)=\lim_{n\to\infty}R(\hat{u}(\lambda_{(\mathrm{reg}),n}))$. By the definition of $\hat{u}(\lambda_{(\mathrm{reg}),n})$ and since $R(u^\star)=0$, we have

$$R(u(\lambda_{(\mathrm{reg}),n}))+\lambda_{(\mathrm{reg}),n}\|\hat{u}(\lambda_{(\mathrm{reg}),n})\|_{H^{m+1}(\Omega)}\leq \cancel{R(u^\star)}+\lambda_{(\mathrm{reg}),n}\|u^\star\|_{H^{m+1}(\Omega)}$$

Therefore, taking the limit, $R(u_\infty)=0$ when $\lambda_{(\mathrm{reg}),n}\xrightarrow{n\to\infty}0$. Hence, by unicity $u_\infty=u^\star$. This contradicts the assumption that $\|u(\lambda_{(\mathrm{reg}),n})-u^\star\|_{H^m(\Omega)}\geq\epsilon$ for all $n\geq 0$. □

The following property measures how well the initial/boundary conditions, encoded by $h$, and the PDE system, encoded by $\mathfrak{F}_k$, describe the function $u$.

**Definition 3.15** (Physics inconsistency)**.** For any $u\in H^{m+1}(\Omega,\mathbb{R}^{d_2})$ the physics inconsistency of $u$ is defined by

$$\mathrm{PI}(u):=\lambda_2\mathbb{E}\|\tilde{\Pi}(u)(X^{(2)})-h(X^{(2)})\|_2^2+\frac{\lambda_3}{|\Omega|}\sum_{k=1}^{M}\int_\Omega\mathfrak{F}_k(u,x)^2dx.$$

In particular, $\mathrm{PI}(u)$ measures the modeling the error, i.e., the the better the model, the lower $\mathrm{PI}(u)$. On the other hand, note that $R_{n_1}(u)=\frac{\lambda_1}{n_1}\sum_{i=1}^{n_1}\|\tilde{\Pi}(u)(X_i)-Y_i\|_2^2+\mathrm{PI}(u)$.

**Proposition 3.16** (Strong convergence of hybrid modeling)**.** Assume that the conditions of Theorem 3.14 are satisfied. Then $\hat{u}_n\equiv\hat{u}_{n_1}(X_1,\ldots,X_{n_1},\epsilon_1,\ldots,\epsilon_{n_1})$ is a random variable such that $\mathbb{E}\|\hat{u}_{n_1}\|_{H^{m+1}(\Omega)}^2<\infty$.

Suppose that $u^\star\in H^{m+1}(\Omega,\mathbb{R}^{d_2})$ that the noise $\epsilon$ is independent from $X$ and that $\epsilon$ has the same distribution as $-\epsilon$. Then, there exists a constant $C_\Omega>0$ depending only on

$\Omega$ such that

$$\mathbb{E}\int_\Omega \|\tilde{\Pi}(\hat{u}_{n_1} - u^\star)\|_2^2 d\mu_X \le \frac{1}{\lambda_1}(\mathrm{PI}(u^\star) + \lambda_{(\mathrm{reg})}\|u^\star\|_{H^{m+1}(\Omega)}^2)$$

$$+ \frac{C_\Omega d_2^{1/2}}{n_1^{1/2}}\left(2\|u^\star\|_{H^{m+1}(\Omega)}^2 + \frac{\mathrm{PI}(u^\star)}{\lambda_{(\mathrm{reg})}}\right) + \frac{8\mathbb{E}\|\epsilon\|_2^2}{n_1}\left(1 + C_\Omega d_2^{3/2}\left(\frac{\lambda_1}{\lambda_{(\mathrm{reg})}} + \frac{\lambda_1^2}{\lambda_{(\mathrm{reg})}^2\sqrt{n_1}}\right)\right).$$

In particular, with the choice $\lambda_2 = 1, \lambda_3 = 1, \lambda_{(\mathrm{reg})} = (\log n_1)^{-1}$ and $\sqrt{n_1}/\log n_1$, one has

$$\mathbb{E}\int_\Omega \|\tilde{\Pi}(\hat{u}_{n_1} - u^\star)\|_2^2 d\mu_X \le \frac{\Lambda \log^2 n_1}{\sqrt{n_1}},$$

where $\Lambda = 24 d_2^{3/2} C_\Omega(\mathrm{PI}(u^\star) + \|u^\star\|_{H^{m+1}(\Omega)} + \mathbb{E}\|\epsilon\|_2^2)$.

**Theorem 3.17** (Strong convergence of regularized PINNs)**.** Under the same assumptions 3.14 and Proposition 3.16 with the choice $\lambda_2 = 1, \lambda_3 = 1$, and $\lambda_{(\mathrm{reg})} = (\log n_1)^{-1}$ and $\lambda_1 = \sqrt{n_1}/\log n_1$, one has

$$\lim_{D\to\infty}\lim_{n_2,n_3\to\infty}\lim_{n\to\infty}\mathbb{E}\int_\Omega \|u_{\hat{\theta}_n^{(\mathrm{reg})}(n_1,n_2,n_3,D)} - u^\star\|_2^2 d\mu_X \le \frac{\Lambda \log^2 n_1}{\sqrt{n_1}}.$$

Moreover,

$$\lim_{D\to\infty}\lim_{n_2,n_3\to\infty}\lim_{n\to\infty}\mathbb{E}[\mathrm{PI}(u_{\hat{\theta}_n^{(\mathrm{reg})}(n_1,n_2,n_3,D)})] \le \mathrm{PI}(u^\star) + o_{n_1\to\infty}(1).$$

*Proof.* The proof is divided into two parts. The first part proves the bound on the difference between the empirical error and the true solution of the PDE. The second part shows that the physical inconsistency of the minimizing sequence is bounded by the physical inconsistency of the solution $u^\star$, plus a diminishing constant.

**(Step 1)** Let us denote $\{u_{\hat{\theta}_n^{(\mathrm{reg})}(n_1,n_2,n_3,D)}\}_{n\in\mathbb{N}}$ be a minimizing sequence of $\widehat{R}^{(\mathrm{reg})}_{n_1,n_2,n_3}$. Since $u_{\hat{\theta}_n^{(\mathrm{reg})}(n_1,n_2,n_3,D)} \in C^\infty(\bar{\Omega}, \mathbb{R}^{d_2})$, then

$$\tilde{\Pi}(u_{\hat{\theta}_n^{(\mathrm{reg})}(n_1,n_2,n_3,D)}) = u_{\hat{\theta}_n^{(\mathrm{reg})}(n_1,n_2,n_3,D)}.$$

Since

$$\mathbb{E}\int_\Omega \|u_{\hat{\theta}_n^{(\mathrm{reg})}(n_1,n_2,n_3,D)} - u^\star\|_2^2 d\mu_X$$

$$\le \mathbb{E}\int_\Omega \|\hat{u}_{n_1} - u^\star\|_2^2 d\mu_X + \mathbb{E}\int_\Omega \|u_{\hat{\theta}_n^{(\mathrm{reg})}(n_1,n_2,n_3,D)} - \hat{u}_{n_1}\|_2^2 d\mu_X,$$

Proposition 3.16 assures us that the first part is less than $\frac{\Lambda \log^2 n_1}{\sqrt{n_1}}$ and Theorem 3.14 assures as that when we take the limit $\lim_{D\to\infty}\lim_{n_2,n_3\to\infty}\lim_{n\to\infty}$ it tends to zero the second part. And this concludes the first part of the theorem.

**(Step 2)** We will proceed proving that $\mathbb{E}[\mathrm{PI}(\hat{u}_{n_1})] \leq \mathrm{PI}(u^\star) + o_{n\to\infty}(1)$.

Recall that $\hat{u}_{n_1}$ minimizes $R_{n_1}^{(\mathrm{reg})}$ over $H^{m+1}(\Omega, \mathbb{R}^{d_2})$. So $R_{n_1}^{(\mathrm{reg})}(u^\star) \geq R_{n_1}^{(\mathrm{reg})}(\hat{u}_{n_1})$. One has

$$\begin{aligned}\|\tilde{\Pi}(\hat{u}_{n_1})(X_i) - Y_i\|_2^2 &= \|\tilde{\Pi}(\hat{u}_{n_1} - u^\star)(X_i) + \epsilon_i\|_2^2 \\ &= \|\tilde{\Pi}(\hat{u}_{n_1} - u^\star)(X_i)\|_2^2 - 2\langle\tilde{\Pi}(\hat{u}_{n_1} - u^\star)(X_i), \epsilon_i\rangle + \|\epsilon_i\|_2^2.\end{aligned}$$

On the other hand, the second term can be further expanded using the dual norm of the coefficient space

$$\begin{aligned}\langle\tilde{\Pi}(\hat{u}_{n_1} - u^\star)(X_i), \epsilon_i\rangle &= \langle\tilde{\Pi}(\hat{u}_{n_1} - u^\star)(X_i) - \mathbb{E}[\tilde{\Pi}(\hat{u}_{n_1} - u^\star)(X)], \epsilon_i\rangle + \langle\mathbb{E}[\tilde{\Pi}(\hat{u}_{n_1} - u^\star)(X)], \epsilon_i\rangle \\ &\leq \|\hat{u}_{n_1} - u^\star\|_{H^{m+1}(\Omega)} \sup_{\|u\|_{H^{m+1}(\Omega)}\leq 1} \langle\tilde{\Pi}(u)(X_i) - \mathbb{E}[\tilde{\Pi}(u)(X)], \epsilon_i\rangle + \langle\mathbb{E}[\tilde{\Pi}(\hat{u}_{n_1} - u^\star)(X)], \epsilon_i\rangle.\end{aligned}$$

Then, arranging all together, one has

$$\begin{aligned}&\frac{1}{n_1}\sum_{i=1}^{n_1}\|\tilde{\Pi}(\hat{u}_{n_1})(X_i) - Y_i\|_2^2 \\ &\quad\geq -2(\|\hat{u}_{n_1}\|_{H^{m+1}(\Omega)} + \|u^\star\|_{H^{m+1}(\Omega)}) \sup_{\|u\|_{H^{m+1}(\Omega)}\leq 1} \frac{1}{n_1}\sum_{j=1}^{n_1}\langle\tilde{\Pi}(u)(X_j) - \mathbb{E}(\tilde{\Pi}(u)(X)), \epsilon_j\rangle \\ &\qquad - 2\left\langle \int_\Omega \tilde{\Pi}(\hat{u}_{n_1} - u^\star)d\mu_X, \frac{1}{n_1}\sum_{i=1}^{n_1}\|\epsilon_i\|_2^2\right\rangle + \frac{1}{n_1}\sum_{i=1}^{n_1}\|\epsilon_i\|_2^2.\end{aligned}$$

In the proof of Theorem 3.16 (Steps 4 and 5), we obtain the following inequality

$$\begin{aligned}\mathbb{E}\|\hat{u}_{n_1}\|_{H^{m+1}(\Omega)}^2 &\leq 2\mathbb{E}\|\hat{u}_{n_1}^\star\|_{H^{m+1}(\Omega)}^2 + 2\mathbb{E}\|\hat{u}_{n_1}^{(\mathrm{noise})}\|_{H^{m+1}(\Omega)}^2 \\ &\leq 2\left(\frac{\mathrm{PI}(u^\star)}{\lambda_{(\mathrm{reg})}} + \|u^\star\|_{H^{m+1}(\Omega)}^2\right) + \frac{8\lambda^2}{n_1\lambda_{(\mathrm{reg})}^2} C_\Omega^2 d_2 \mathbb{E}\|\epsilon\|_2^2.\end{aligned}$$

Using Lemma 2.20 and the Cauchy-Schwarz inequality $\mathbb{E}(XY)^2 \leq \mathbb{E}(X^2)\mathbb{E}(Y^2)$, we show that

$$\mathbb{E}\left[\|\hat{u}_{n_1}\|_{H^{m+1}(\Omega)} \sup_{\|u\|_{H^{m+1}(\Omega)}\leq 1} \frac{1}{n}\sum_{j=1}^{n}\langle\tilde{\Pi}(u)(X_j) - \mathbb{E}(\tilde{\Pi}(u)(X)), \epsilon_j\rangle\right] = \mathcal{O}_{n_1\to\infty}\left(\frac{\lambda_1}{n_1\lambda_{(\mathrm{reg})}}\right).$$

By Proposition 3.16,

$$\mathbb{E}\left|\left\langle \int_\Omega \tilde{\Pi}(\hat{u}_{n_1} - u^\star)d\mu_X, \frac{1}{n_1}\sum_{j=1}^{n_1}\epsilon_j\right\rangle\right| \leq \sqrt{\mathbb{E}\|u^\star - \hat{u}_{n_1}\|_{L^2(\mu_X)}^2 \frac{\mathbb{E}\|\epsilon\|_2^2}{\sqrt{n_1}}} = \mathcal{O}_{n_1\to\infty}\left(\frac{\lambda_1}{n_1^2\lambda_{(\mathrm{reg})}}\right)^{1/2}.$$

Combining these result with the previous, we conclude that

$$\mathbb{E}\left[\frac{1}{n_1}\sum_{i=1}^{n_1}\|\tilde{\Pi}(\hat{u}_{n_1})(X_i) - Y_i\|_2^2\right] \geq \mathbb{E}\|\epsilon\|_2^2 + \mathcal{O}_{n\to\infty}\left(\frac{\lambda_1}{n_1\lambda_{(\mathrm{reg})}}\right).$$

By assumption, $\lambda_1^2/n_1\lambda_{(\mathrm{reg})} \to 0$ for $n_1 \to \infty$ and since $R_{n_1}^{(\mathrm{reg})}(\hat{u}_{n_1}) = \frac{\lambda_1}{n_1}\sum_{i=1}\|\tilde{\Pi}(\hat{u}_{n_1})(X_i) - Y_i\|_2^2 + \mathrm{PI}(\hat{u}_{n_1}) + \lambda_{(\mathrm{reg})}\|\hat{u}_{n_1}\|_{H^{m+1}(\Omega)}^2$,

$$\mathbb{E}[R_{n_1}^{(\mathrm{reg})}(\hat{u}_{n_1})] \geq \lambda_1\mathbb{E}\|\epsilon\|_2^2 + \mathbb{E}[\mathrm{PI}(\hat{u}_{n_1})] + o_{n_1\to\infty}(1).$$

As before, we obtain

$$\frac{1}{n_1}\sum_{i=1}^{n_1}\|\tilde{\Pi}(u^\star)(X_i) - Y_i\|_2^2 = \frac{1}{n_1}\sum_{i=1}^{n_1}\|\epsilon_i\|_2^2.$$

Hence,

$$\mathbb{E}[R_{n_1}^{(\mathrm{reg})}(u^\star)] = \lambda_1\mathbb{E}\|\epsilon\|_2^2 + \mathrm{PI}(u^\star) + \lambda_{(\mathrm{reg})}\|u^\star\|_{H^{m+1}(\Omega)}^2$$

Since $\mathbb{E}[R_{n_1}^{(\mathrm{reg})}(\hat{u}_{n_1})] \leq \mathbb{E}(R_{n_1}^{(\mathrm{reg})}(u^\star)$ and since $\lambda_{(\mathrm{reg})} \to 0$, we are lead to

$$\mathbb{E}[\mathrm{PI}(\hat{u}_{n_1})] \leq \mathrm{PI}(u^\star) + o_{n_1\to\infty}(1). \tag{3.15}$$

Due to the continuity of the Physics inconstistency (PI) and Theorem 3.14, we have that

$$\lim_{D\to\infty}\lim_{n_2,n_3\to\infty}\lim_{n\to\infty}\mathbb{E}[\mathrm{PI}(u_{\hat{\theta}_n^{(\mathrm{reg})}(n_1,n_2,n_3,D)})] = \mathbb{E}[\mathrm{PI}(\hat{u}_{n_1})].$$

Then, using (3.15), we obtain

$$\lim_{D\to\infty}\lim_{n_2,n_3\to\infty}\lim_{n\to\infty}\mathbb{E}[\mathrm{PI}(u_{\hat{\theta}_n^{(\mathrm{reg})}(n_1,n_2,n_3,D)})] \leq \mathrm{PI}(u^\star) + o_{n_1\to\infty}(1).$$

□

# CHAPTER 4

# Risk-Consistency of Regularized Neural ODEs

Neural Ordinary Differential Equations (NeuralODEs) are a class of deep learning models that generalize traditional neural networks by interpreting the transformation of hidden states as a continuous dynamical system governed by ordinary differential equations (ODEs). Introduced by Chen et al. [2019], NeuralODEs replace discrete layers in a neural network with an ODE solver, modeling the evolution of features over continuous time.

To formalize this concept, NeuralODEs define the system's dynamics using an initial value problem based on a neural network representation of the time derivative Chen et al. [2019], Karniadakis et al. [2021], Hao et al. [2022]. Specifically, the system is described by:

$$\frac{dy}{dt}(t) = u_\theta(t, y(t)), \qquad \text{with an initial state of } y(0) = y_0. \tag{4.1}$$

In this formulation, $(t, y_0) \in \Omega \subseteq \mathbb{R}^{d_1}$, $y(t)$ represents the hidden state of the system as a $d_2-$dimensional vector at time $t$. The function $u_\theta : \mathbb{R}^{d_1} \to \mathbb{R}^{d_2}$ is a neural network, wihth learnable parameters $\theta$, that outputs the time derivative of the state. Here, $d_2 = d_1 - 1$ and we assume the $u_\theta$ is the standard neural network defined in Section 1.1 with tanh activation function.

Here, we introduce the NeuralODE framework and establish original results regarding its risk-consistency of ridge NeuralODEs. Using the tools and results presented in previous sections, we establish findings analogous to those for PINNs.

Since the tanh is Lipschitz continuous, the resulting neural network $u_\theta$ is also Lipschitz continuous, and according to Picard-Linderlöf theorem, this property guarantees that for any given initial state $y_0$, a unique and differentiable solution $y(t)$ exists over a given time interval. This solution defines a continuous trajectory through the $d_2$-dimensional state space, which is shaped by the learned parameters $\theta$ of the neural network. We will denote this unique solution by $y(t, y_0; u_\theta) \in \mathbb{R}^{d_2}$ for $t \in [0, T]$ to highlight its dependence on these parameters.

Let us introduce the empirical and theoretical risk functions within the Neural Ordinary Differential Equation (Neural ODE) framework. These definitions are analogous to those in a classical regression problem Chen et al. [2019], Kidger [2022], Gwak et al. [2020], Massaroli et al. [2021].

**Theoretical Risk Function** The theoretical risk measures the expected prediction error over the entire data distribution. It is defined as the continuous version of the empirical risk function:

$$R(u_\theta) := \mathbb{E}_{X\sim\mu_X}\|y(X;u_\theta) - h(X)\|_2^2,$$

where $y(X; u_\theta)$ represents the output of the NeuralODE for a given input $X$, parametrized by $\theta$. Here, $h(X)$ is the true outcome of the NeuralODE defined above. The goal is to find an optimal set of parameters $\theta^\star$ that minimizes this risk, such that the resulting function $y(\cdot, u_\theta) = h(\cdot)$.

**Empirical Risk Function** In practice, the true data distribution $\mu$ is unknown. Therefore, we approximate the theoretical risk using a finite dataset. Given a collection of $n_1 \in \mathbb{N}$ independent and identically distributed data pairs $\{(X_i, Y_i)\}_{1\le i\le n_1}$ distributed as $(X, Y) \in \Omega \times \mathbb{R}^{d_2}$, where the distribution $\mu_X$ is unknown, the empirical risk function is defined as:

$$\widehat{R}_{n_1}(u_\theta) := \frac{1}{n_1}\sum_{i=1}^{n_1}\|y(X_i;u_\theta) - Y_i\|_2^2,$$

where $X_i := (t, y_0)_i$ can be considered the point at which we are evaluating and $y(X_i; u_\theta)$ is the unique solution of the NueralODE at a specific time $t$ and initial condition $y_0$ guaranteed by the Picard-Lindelöf theorem. We assume that the observed outcomes $Y_i$ are related to a true underlying function $h(X)$ corrupted by some noise $\epsilon_i$, such that $Y_i = h(X_i) + \epsilon_i$, where we assume that the noise is zero mean, i.e., $\mathbb{E}[\epsilon_i|X_i]$.

**Ridge Risk Function** Training NeuralODEs can be tricky, similar to PINNs, it can lead to the type of pathological situations highlited in previous Section. To avoid such an overfitting behaviour, we resort again to the ridge regularization applied to the empirical risk minimization where we penalize the $\ell^2$ norm of the parameter $\theta$

$$\widehat{R}^{(\text{ridge})}_{n_1}(u_\theta) := \widehat{R}_{n_1}(u_\theta) + \lambda_{(\text{ridge})}\|\theta\|_2^2. \tag{4.2}$$

Analogously to the PINNs framework, we define risk-consistency for the NeuralODE's risk function as the condition that a minimizing sequence of the empirical risk converges to the infimum of the theoretical risk.

**Definition 4.1** (Risk-Consistency of NeuralODEs)**.** Let $\{\hat\theta_n(n_1)\}_{n\in\mathbb{N}}$ be a minimizing sequence of the empirical risk, i.e.,

$$\lim_{n\to\infty}\widehat{R}_{n_1}(u_{\hat\theta(n,n_1)}) = \inf_{\theta\in\Theta_{H,D}}\widehat{R}_{n_1}(u_\theta).$$

We say that $\{\hat{\theta}_n(n_1)\}_{n\in\mathbb{N}}$ verifies the *risk-consistency* condition if

$$\lim_{n_1\to\infty}\lim_{n\to\infty} R(u_{\hat{\theta}_n(n_1)}) = \inf_{u\in \mathrm{NN}_H^D} R(u). \tag{4.3}$$

**Proposition 4.2** (Continuous Dependence on Neural ODEs)**.** Let $u_\theta$ and $u_{\theta'}$ be two NeuralODEs with the same initial value:

$$\begin{aligned}\frac{dy}{dt}(t) &= u_\theta(t, y(t)), \qquad y(0) = y_0\\ \frac{dz}{dt}(t) &= u_{\theta'}(t, z(t)), \qquad z(0) = y_0\end{aligned}$$

where $u_\theta, u_{\theta'} : \mathbb{R}\times\mathbb{R}^{d_2} \to \mathbb{R}^{d_2}$ are two different parametrized neural networks. Then for all $t \in [0, T]$

$$\|y(t) - z(t)\|_2 \le L\|\theta - \theta'\|_2 e^{\tilde{L}T}\int_0^T e^{-\tilde{L}s}ds \tag{4.4}$$

for some constants $L$ and $\tilde{L}$.

*Proof.* Let us denote the difference between the solutions $y(t)$ and $z(t)$ as $e(t) := y(t) - z(t)$. Lemma 1.7 assures that there exists a constant $L := \sqrt{d_2}\tilde{C}_{0,H}(1 + d_1 M(\Omega))(D + 1)^H(1 + n^{\alpha_1})^H$ that satisfies $\|u_\theta - u_{\theta'}\|_{2,\Omega} \le L\|\theta - \theta\|_2$. Lemma 1.8 assure us that there exists a constant $\tilde{L}$ that satisfies $\|u_\theta(x) - u_\theta(y)\|_2 \le \tilde{L}\|x - y\|_2$ for all $x, y \in \bar{\Omega}$. We can bound the derivative of $e(t)$ as follows

$$\begin{aligned}\left\|\frac{d}{dt}e(t)\right\|_2 &= \|u_\theta(t, y(t)) - u_{\theta'}(t, z(t))\|_2\\ &\le \|u_\theta(t, y(t)) - u_\theta(t, z(t))\|_2 + \|u_\theta(t, z(t)) - u_{\theta'}(t, z(t))\|_2\\ &\le \tilde{L}\|\underbrace{y(t) - z(t)}_{e(t)}\|_2 + L\|\theta - \theta'\|_2 \le \tilde{L}\|e(t)\|_2 + L\|\theta - \theta'\|_2\end{aligned}$$

Let $w(t) := \|e(t)\|_2$. Then, we obtain the following chain of inequalities

$$\begin{aligned}\frac{d}{dt}w(t) = \frac{d}{dt}\|e(t)\|_2 &= \left|\frac{e(t)^\intercal \dot{e}(t)}{\|e(t)\|_2}\right| \le \frac{1}{\|e(t)\|_2}\|e(t)\|_2\|\dot{e}(t)\|_2\\ &= \left\|\frac{d}{dt}e(t)\right\|_2 \le \tilde{L}w(t) + L\|\theta - \theta'\|_2.\end{aligned}$$

That is

$$\frac{d}{dt}w(t) - \tilde{L}w(t) \le L\|\theta - \theta'\|_2.$$

Multiplication by the integrating factor $e^{-\tilde{L}t}$ yields

$$\frac{d}{dt}(e^{-\tilde{L}t}w(t)) \le e^{\tilde{L}t}L\|\theta - \theta'\|_2.$$

Integrate from 0 to $T$ to get

$$e^{-\tilde{L}T}w(t) - \|y_0 - z_0\| \leq \int_0^T e^{-\tilde{L}s} L\|\theta - \theta'\|_2 ds.$$

That is equivalent to (4.4) and finishes the proof. □

**Proposition 4.3.** Let $\Omega$ be a bounded Lipschitz domain, and let $h \in \mathrm{Lip}(E, \mathbb{R}^{d_2})$. Then, the operator $\mathscr{H}(u, x) = \mathbb{1}_{x\in\bar{\Omega}}\|y(x;u) - h(x)\|_2^2$ satisfies inequalities (H.1) and (H.2) from Theorem 2.19, with $\alpha_1 < (3+H)^{-1}/2, \beta_1 = (1+H)\alpha_1$, and $1/2 > \beta_2 \geq (3+H)\alpha_1$.

*Proof.* First note, since $\mathrm{Lip}(\bar{\Omega}, \mathbb{R}^{d_2}) \subseteq C(\bar{\Omega}, \mathbb{R}^{d_2})$, then $\|h\|_\infty < \infty$. Observe also that for any $v, w \in \mathbb{R}^{d_2}$, $|\,\|v\|_2^2 - \|w\|_2^2\,| = |\langle v+w, v-w\rangle| \leq \|v+w\|_2\|v-w\|_2 \leq d_2\|v+w\|_2\|v-w\|_2$. Thus, we obtain that for all $\theta, \theta' \in B_2(0, n^{\alpha_1})$ and $x \in \bar{\Omega}$,

$$\begin{aligned}
|\mathscr{H}(u_\theta, x) - \mathscr{H}(u_{\theta'}, x)| &\leq (\ \|y(x;u_\theta)\|_2 + \|y(x;u_{\theta'})\|_2 + 2\|h\|_2)\|y(x;u_\theta) - y(x;u_{\theta'})\|_2 \\
&\leq d_2(\ \|y(x;u_\theta)\|_{\infty,\bar{\Omega}} + \|y(x;u_{\theta'})\|_{\infty,\bar{\Omega}} + 2\|h\|_{\infty,\bar{\Omega}})\|y(x;u_\theta) - y(x;u_{\theta'})\|_{\infty,\bar{\Omega}} \\
&\leq d_2(2(D+1)n^{\alpha_1}t + 2\|h\|_{\infty,\bar{\Omega}})C_3\|\theta - \theta'\|_2 \\
&\leq C_1 n^{\beta_1}\|\theta - \theta'\|_2.
\end{aligned}$$

where $\beta_1 = (1+H)\alpha_1$. Therefore it satisfies (H.1)

Next, for all $\theta \in B_2(0, n^{\alpha_1})$,

$$\|\mathscr{H}(u_\theta, \cdot)\|_{\infty,\bar{\Omega}} \leq d_2(\|y(\cdot;u_\theta)\|_{\infty,\bar{\Omega}} + \|h\|_\infty)^2 \leq d_2((D+1)M(\Omega)n^{\alpha_1} + \|h\|_\infty)^2 \leq C_2 n^{2\alpha_1}$$

Recall that for inequality (H.2), $\beta_2$ must satisfy $\alpha_1 + \beta_1 \leq \beta_2 < 1/2$. This is true for $\beta_2 = (3+H)\alpha_1$, which completes the proof. □

**Theorem 4.4** (Risk-consistency of ridge NeuralODEs)**.** Let the number of hidden layers be $H \geq 2$. Assuming that the ridge parameter is of the form

$$\lambda_{(\text{ridge})} := {n_1}^{-\kappa}, \qquad \text{where} \qquad \kappa := \frac{1}{3+H}.$$

Let $\{\hat{\theta}_n^{(\text{ridge})}(n_1)\}_{n\in\mathbb{N}}$ be a minizing of the ridge risk function (4.2). Then, almost surely,

$$\lim_{n_1\to\infty}\lim_{n\to\infty} R(u_{\hat{\theta}_n^{(\text{ridge})}(n_1)}) = \inf_{u\in \mathrm{NN}_H^D} R(u).$$

*Proof.* Let $u_0 = 0 \in \mathrm{NN}_H^D$ be a neural network with parameters $\theta = (0, \ldots, 0)$. Obviously, $\widehat{R}_{n_1}^{(\text{ridge})}(u_0) = \widehat{R}_{n_1}(u_0)$. Moreover,

$$\widehat{R}_{n_1}(u_0) = \frac{1}{n_1}\sum_{i=1}^{n_1}\|y(X_i;u_0) - h(X_i)\|_2^2 \leq \frac{1}{n_1}\sum_{i=1}^{n_1}\|X_i\|_2^2 + \|h\|_\infty := I. \tag{4.5}$$

We notice that $I$ does not depend on $\lambda_{(ridge)}$.

Let $\{\hat{\theta}_n^{(\text{ridge})}(n_1)\}_{n\in\mathbb{N}}$ be any minimizing sequence of the empirical risk of the ridge PINN, i.e., $\lim_{n\to\infty} \widehat{R}_{n_1}^{(\text{ridge})}(u_{\hat{\theta}_n^{(\text{ridge})}(n_1)}) = \inf_{\theta\in\Theta_H^D} \widehat{R}_{n_1}^{(\text{ridge})}(u_\theta)$. Let $N_1 := \min(n_1)$. Let us define the following three sets:

$$\begin{aligned}
\mathcal{E}_1(N_1) &:= \{\theta \in \Theta_H^D : \|\theta\|_2 \geq N_1^\kappa\} \\
\mathcal{E}_2(N_1) &:= \{\theta \in \Theta_H^D : N_1^{\kappa/4} \leq \|\theta\|_2 \leq N_1^\kappa\} \\
\mathcal{E}_3(N_1) &:= \{\theta \in \Theta_H^D : N_1^{\kappa/4} \leq \|\theta\|_2\}.
\end{aligned}$$

Notice that $\Theta_H^D = \mathcal{E}_1 \sqcup \mathcal{E}_2 \sqcup \mathcal{E}_2$. The proof relies on the argument that for any $n_1$, given a sufficiently large $n$, $\hat{\theta}_n^{(\text{ridge})}(n_1) \in \mathcal{E}_2 \cup \mathcal{E}_3$. Furhtermore, when $N_1$ is sufficiently large, then $\widehat{R}_{n_1}^{(\text{ridge})}$ is close to the theoretical risk $R$. Let us divide the proof in four steps.

**Step 1.** We start by observing that, for any $\theta \in \mathcal{E}_1(N_1)$,

$$\widehat{R}_{n_1}^{(\text{ridge})}(u_\theta) \geq \lambda_{(\text{ridge})}\|\theta\|_2^2 \geq \lambda_{(\text{ridge})} N_1^{2\kappa} = N_1^\kappa.$$

Assuming that $N_1$ is sufficiently large to verify $N_1^\kappa - 1 \geq I \overset{(4.5)}{\geq} \widehat{R}_{n_1}(u_0)$. Therefore,

$$\inf_{\theta\in\mathcal{E}_3(N_1)} \widehat{R}_{n_1}^{(\text{ridge})}(u_\theta) \leq \widehat{R}_{n_1}^{(\text{ridge})}(u_0) \leq N_1^\kappa - 1 \leq \inf_{\theta\in\mathcal{E}_1} \widehat{R}_{n_1}^{(\text{ridge})}(u_\theta) - 1.$$

This shows that, for all $N_1$ and $n$ large enough, $\hat{\theta}_n^{(\text{ridge})}(n_1) \notin \mathcal{E}_1(N_1)$.

**Step 2.** Invoking Proposition 4.3 with $\alpha_1 = \kappa$ and $\beta_2 = (3+H)\alpha_1$, then by Theorem 2.19, we know that, almost surely, there exist $N \in \mathbb{N}$ such that for all, $N_1 \geq N$,

$$\begin{aligned}
\sup_{\theta\in\mathcal{E}_2(N_1)\cup\mathcal{E}_3(N_a)} &\left| \frac{1}{n_1}\sum_{i=1}^{n_1} \|y(X_i; u_\theta) - h(X_i)\|_2^2 - \mathbb{E}_{X\sim\mu_X}\|u(X) - h(X)\|_2^2 \right| \\
&\leq \log^2(N_1) N_1^{\beta_2 - 1/2}.
\end{aligned} \tag{4.6}$$

Then, almost surely, for all $N_1$ large enough and all $\theta \in \mathcal{E}_2(N_a)$,

$$|\underbrace{\widehat{R}_{n_1}^{(\text{ridge})}(u_\theta) - \lambda_{(\text{ridge})}\|\theta\|_2^2}_{\widehat{R}_{n_1}(u_\theta)} - R(u_\theta)| \leq (M+1)\log^2(N_1) N_1^{\beta_2 - 1/2}. \tag{4.7}$$

But, for all $\theta \in \mathcal{E}_2(N_1)$, $\lambda_{(\text{ridge})}\|\theta\|_2^2 \geq \lambda_{(\text{ridge})} N_1^{\kappa/2} = N_1^{-\kappa/2}$. We conclude that, for all $N_1$ large enough and for all $\theta \in \mathcal{E}_2(N_1)$,

$$\begin{aligned}
|R(u_\theta) - \widehat{R}_{n_1}(u_\theta)| &\leq (M+1)\log^2(N_1) N_1^{\beta_2 - 1/2} \\
&\Rightarrow \widehat{R}_{n_1}(u_\theta) \geq R(u_\theta) - (M+1)\log^2(N_1) N_1^{\beta_2 - 1/2} \\
&\Rightarrow \widehat{R}_{n_1}^{(\text{ridge})}(u_\theta) \geq R(u_\theta) + \lambda_{(\text{ridge})}\|\theta\|_2^2 - (M+1)\log^2(N_1) N_1^{\beta_2 - 1/2} \\
&\Rightarrow \widehat{R}_{n_1}^{(\text{ridge})}(u_\theta) \geq R(u_\theta) + \underbrace{N_1^{-\kappa/2} - (M+1)\log^2(N_1) N_1^{\beta_2 - 1/2}}.
\end{aligned}$$

Assuming that $-\kappa/2 > \beta_2 - 1/2$, then the right hand side is positive for $N_1$ sufficiently large, and we obtain

$$\widehat{R}_{n_1}^{(\text{ridge})}(u_\theta) \geq R(u_\theta).$$

**Step 3.** For all $\theta \in \mathcal{E}_3(N_1), \lambda_{(\text{ridge})}\|\theta\|_2^2 \leq N_1^{-\kappa/2}$. Using inequalities (4.6), we deduce that, almost surely, for all $N_a$ large enough

$$\begin{aligned}|\widehat{R}_{n_1}^{(\text{ridge})}(u_\theta) - R(u_\theta)| &\leq |\widehat{R}_{n_1}^{(\text{ridge})}(u_\theta) - \lambda_{\text{ridge}}\|\theta\|_2^2 - R(u_\theta)| + \lambda_{(\text{ridge})}\|\theta\|_2^2 \\ &\leq (M+1)\log^2(N_1)N_1^{\beta_2-1/2} + N_1^{-\kappa/2} \\ &\leq (M+2)\log^2(N_1)N_1^{-\kappa/2}.\end{aligned}$$

**Step 4.** Fix $\epsilon > 0$. Let $(\theta_n)_{n\in\mathbb{N}}$ be a minimizing sequence of the theoretical risk function $R$, i.e., $\lim_{n\to\infty} R(u_{\theta_n}) = \inf_{\theta\in\Theta_H^D} R(u_\theta)$. By definition, there exists some $P_\epsilon \in \mathbb{N}$ such that $|R(u_{\theta_n}) - \inf_{\theta\in\Theta_H^D} R(u_\theta)| \leq \epsilon$.

For some fixed $N_1$, according to Step 1, we have for all $n$ large enough, $\hat{\theta}_n^{(\text{ridge})}(n_1) \in \mathcal{E}_2 \cup \mathcal{E}_3$. According to Step 2 and Step 3,

$$R(u_{\hat{\theta}_n^{(\text{ridge})}(n_1)}) \leq \widehat{R}_{n_1}^{(\text{ridge})}(u_{\hat{\theta}_n^{(\text{ridge})}(n_1)}) + (M+2)\log^2(N_1)N_1^{-\kappa/2}. \tag{4.8}$$

Now, by definition of the minimizing sequence $\{\hat{\theta}_n^{\lambda}(n_1)\}_{n\in\mathbb{N}}$, for all $n$ large enough,

$$\widehat{R}_{n_1}^{(\text{ridge})}(u_{\hat{\theta}_n^{(\text{ridge})}(n_1)}) \leq \inf_{\theta\in\Theta_H^D} \widehat{R}_{n_1}^{(\text{ridge})}(u_\theta) + \epsilon. \tag{4.9}$$

Moreover, according to Step 3,

$$\begin{aligned}\inf_{\theta\in\mathcal{E}_2(N_1)\cup\mathcal{E}_3(N_1)} \widehat{R}_{n_1}^{(\text{ridge})}(u_\theta) &\leq \inf_{\theta\in\mathcal{E}_3(N_1)} \widehat{R}_{n_1}^{(\text{ridge})}(u_\theta) \\ &\leq \inf_{\theta\in\mathcal{E}_3(N_1)} R(u_\theta) + (M+2)\log^2(N_1)N_1^{-\kappa/2}.\end{aligned} \tag{4.10}$$

Observe that for all $N_1$ large enough, $\theta_{P_\epsilon} \in \mathcal{E}_3(N_1)$. Therefore, $\inf_{\theta\in\mathcal{E}_3(N_a)} R(u_\theta) \leq R(u_{\theta_{P_\epsilon}})$. Combining the previous inequalities, we conclude that, almost surely, for all $N_1$ large enough for all $n$ large enough,

$$\begin{aligned}R(u_{\hat{\theta}_n^{(\text{ridge})}(n_1)}) &\overset{(4.8)}{\leq} \widehat{R}_{n_1}^{(\text{ridge})}(u_{\hat{\theta}_n^{(\text{ridge})}(n_1)}) + (M+2)\log^2(N_1)N_1^{-\kappa/2} \\ &\overset{(4.9)}{\leq} \inf_{\theta\in\Theta_H^D} \widehat{R}_{n_1}^{(\text{ridge})}(u_\theta) + \epsilon + (M+2)\log^2(N_1)N_1^{-\kappa/2} \\ &\overset{(4.10)}{\leq} \inf_{\theta\in\Theta_H^D} R(u_\theta) + \epsilon + 2(M+2)\log^2(N_1)N_1^{-\kappa/2}.\end{aligned}$$

Since $\epsilon$ is arbitrary and $N_1$ is arbitrarily large, then, almost surely,

$$\lim_{n_1\to\infty}\lim_{n\to\infty} R(u_{\hat{\theta}_n^{(\mathrm{ridge})}(n_1)}) = \inf_{\theta\in\Theta_H^D} R(u_\theta).$$

□

# CHAPTER 5
# Experiments

This chapter addresses the computational realization of Physics-Informed Neural Networks (PINNs). The practical implementation of the described algorithms introduces a variety of details that have hitherto been taken for granted. For instance, a foundational assumption in many of the theoretical results presented in this manuscript is the existence of a parameter sequence that minimizes the empirical risk function. In practice, however, the identification of this minimizing sequence is a non-trivial task.

To connect the preceding theory with practical application, this chapter begins with two illustrative examples designed to motivate the theoretical framework. We then proceed to a more demanding case study involving the Navier-Stokes equations (Example 3.4) to probe the limitations of the PINN approach. The resulting PINN solution is then systematically compared against a conventional numerical solution to evaluate its accuracy and effectiveness.

## 5.1. First-Order ODE

To begin, we will analyze a first-order Ordinary Differential Equation (ODE) with a known explicit solution to better understand the Physics-Informed Neural Network (PINN) framework. Consider the following problem:

$$u'(x) = 2(1 - u(x)^2). \tag{5.1}$$

We assume we are given a dataset of $n_1$ points, $\{(x_1, y_1), \dots, (x_{n_1}, y_{n_1})\} \subseteq \Omega \times \mathbb{R}$, representing the konwn values of the solution. For simplicity, we will not enforce any boundary or initial conditions, which corresponds to settings their associated loss weight, $\lambda_2$, to zero. Figure 5.1a illustrates a potential dataset.

We approximate the solution $u^\star(x)$ using a simple neural network, $u_\theta(x)$, with a single neuron:

$$u_\theta = \tanh(\omega x + \beta) \in \mathrm{NN}_H^D,$$

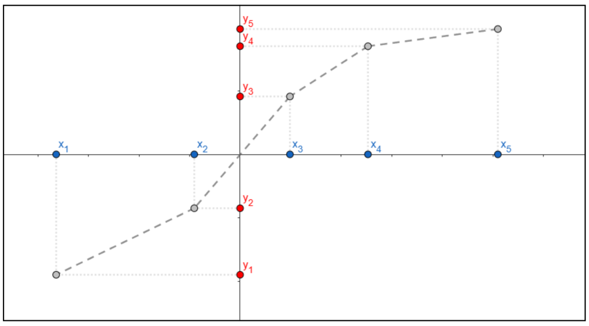


(a) A sample dataset consisting of five data points, with the curve representing a possible approximate solution.

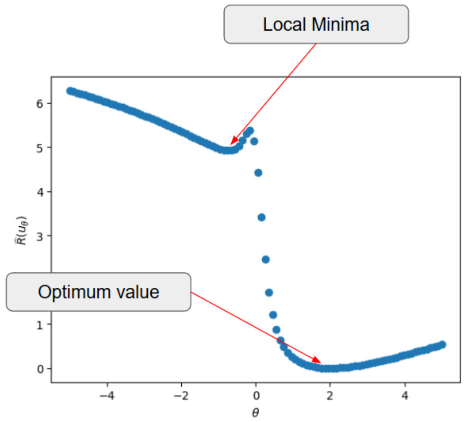


(b) The empirical risk function plotted against different parameter values, illustrating the optimization landscape.

Figure 5.1: Illustrative results for the ODE PINN.

where $\omega, \beta \in \mathbb{R}$ are the learnable parameters of the network.The empirical risk (or loss function) to be minimized is a weighted sum of the data fidelity loss and the PDE residual loss:

$$\widehat{R}_{n_1,n_3}(u_\theta) = \frac{\lambda_1}{n_1}\sum_{i=1}^{n_1}|u_\theta(x_i) - y_i|_2^2 + \frac{\lambda_3}{n_3}\sum_{\ell=1}^{n_3}|u'_\theta(x_i) - 2(1-u_\theta^2(x_i)|^2,$$

where $\{x_j^{(3)}\}_{j=1}^{n_3}$ is a set of $n_3$ collocation points sampled within the domain $\Omega$.

The corresponding theoretical risk is the expectation of this loss over the underlying probability distributions:

$$R_{n_1}(u_\theta) = \frac{\lambda_1}{n_1}\sum_{i=1}^{n_1}|u_\theta(x_i) - y_i|_2^2 + \frac{\lambda_3}{\Omega}\mathbb{E}|u'_\theta(X) - 2(1-u_\theta^2(X)|^2.$$

Finding the global minimum of the non-convex empirical risk function is a significant challenge. Optimization is typically performed using gradient-based methods, such as Gradient Descent or its variants. However, even for this simple problem, the algorithm can easily become trapped in a local minimum, as illustrated in Figure 5.1b.

The gradient descent method, when applied to the problem of Physics-Informed Neural Networks (PINNs), can be expressed as a consecutive loop:

$$\theta \leftarrow \theta - \mu\nabla_\theta(\widehat{R}_{n_1,n_3}),$$

where $\mu$ is the step size, or learning rate, of the method. Figure 5.2b illustrates the gradient descent method on a two dimensional parameter space. In this case, calculating the gradient with respect to both parameters is straightforward. However, when considering more complex risk functions and neural networks, such as those required for the

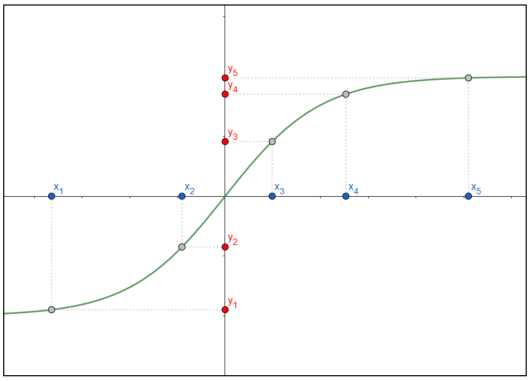


(a) The exact solution of the problem given by Eq. (5.1), which is $u^\star(x) = \tanh(2x)$.

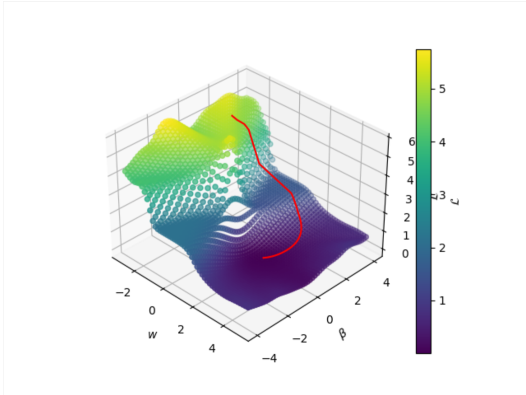


(b) The surface of the empirical risk function, illustrating the iterative steps of the gradient descent optimization.

Figure 5.2: Results for the ODE PINN.

Navier-Stokes problem discussed in the following section, this calculation can become computationally expensive.

To overcome this issue, software has been developed that automatically calculates the derivatives of functions, some examples include *PyTorch*[1], *Jax*[2], or *Tensorflow*[3]. This automatic differentiation software calculates the exact derivative of the risk function and the neural network at each point without any numerical approximation. Moreover, these libraries include advanced optimization algorithms that improve the numerical stability and convergence of the gradient descent method.

## 5.2. Counterexample: Dynamic friction model

In Example 3.2, we considered the differential operator $\mathfrak{F}$ on the domain $\Omega := (0, T)$, defined by

$$\mathfrak{F}(u, x) = mu''(x) + \gamma u'(x). \tag{5.2}$$

In this example, we demonstrated that a sequence of functions minimizing the empirical risk can exhibit a diverging theoretical risk. We will now provide experimental evidence that gradient descent can generate such a sequence, where the empirical risk is minimized, but the theoretical risk fails to decrease.

First, we consider the case where $m = 1$ and $\gamma = 1$. Here, the equation can be solved analytically, and the unique solution is determined by two boundary conditions. We use

[1]PyTorch: https://github.com/pytorch/pytorch
[2]Jax: https://github.com/jax-ml/jax
[3]Tensorflow: https://github.com/tensorflow/tensorflow

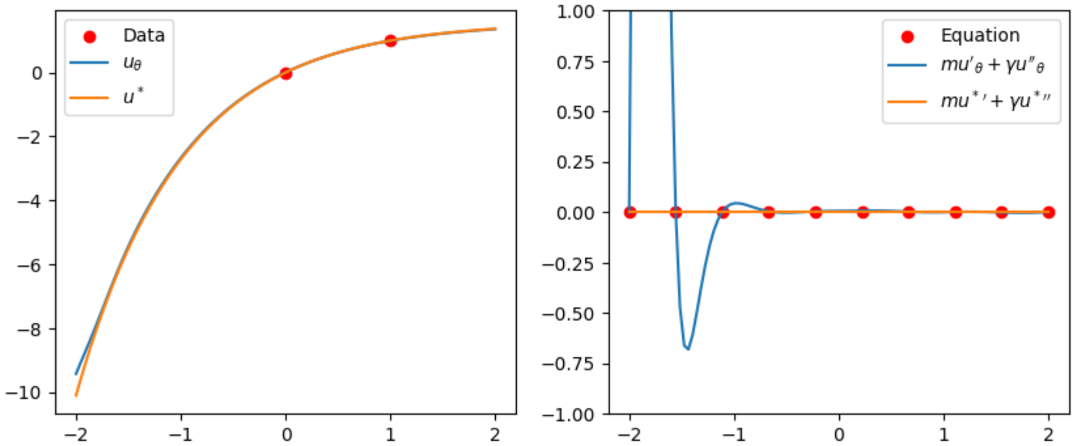


Figure 5.3: The left figure displays the solution $u_\theta$ versus the reference solution $u^\star$. The right figure shows the corresponding values of $\mathfrak{F}(u_\theta)$ vs $\mathfrak{F}(u^\star)$. The red dots represent the data used for training.

the conditions $\{(0,0),(1,1)\}$, which yiled the exact solution

$$u^\star(x) = \frac{e}{e-1}(1 - e^{-x}).$$

To train the network on the differential equation itself, we use a set of 10 collocation points equally spaced within the interval $[-2,2]$. By training a neural network with two hidden layers of five neurons each using gradient descent, we obtain a solution $u_\theta$ with an empirical risk approaching zero. Nonetheless, the solution generated by the neural network is far from the true solution $u^\star$. As shown in Figure 5.3, while the network's solution $u_\theta$ accurately approximates $u^\star$ on the training data, it fails to generalize. Outside this training set, $u_\theta$ deviates significantly from the true solution. The right panel of the

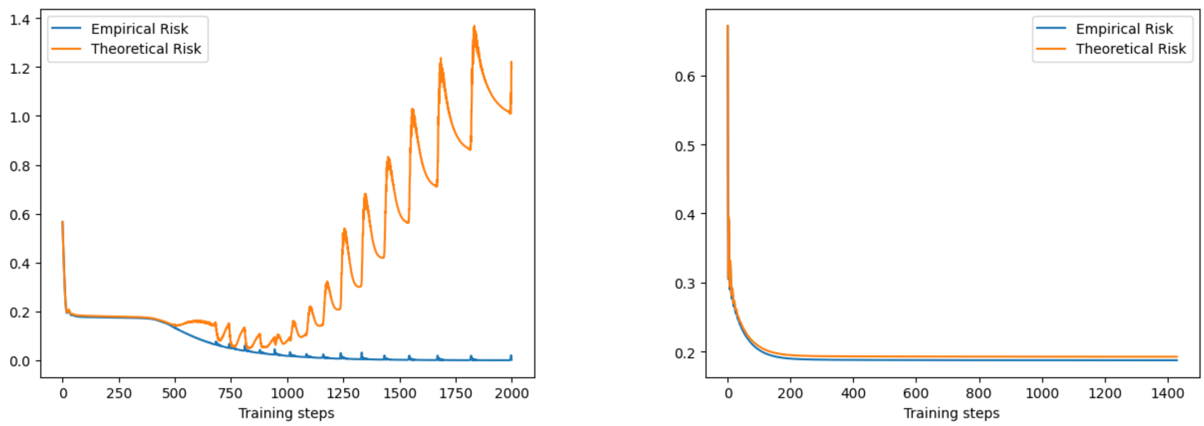


(a) Standard training.

(b) Training with ridge regularization.

Figure 5.4: Comparison of empirical and theoretical risk during the training of a PINN.

Figure 5.4 further illustrates this failure: the term $\mathfrak{F}(u_\theta)$, which should be zero everywhere, oscillates wildly outside the collocation points.

The divergence between the theoretical and empirical risks is also evident in Figure 5.4a. This figure shows that as the gradient descent method progresses, the empirical risk converges to zero, while the theoretical risk fails to converge. In contrast, when using the ridge risk (defined in Definition 3.3), both the empirical and theoretical ridge risks converge smoothly. Note that they do not converge to zero, which can be attributed to the limited representational capacity of the neural network. While a deeper network $D$ could better approximate any function, there is no guarantee of reaching the optimal solution using gradient descent.

## 5.3. Navier-Stokes: Lid-Driven Cavity Flow

This study investigates the use of Physics-Informed Neural Networks (PINNs) to solve the 2D Navier-Stokes equations under specific boundary conditions. The governing equations are identical to those presented in Example 3.4, adapted for a two-dimensional spatial domain.

The problem is defined on the spatio-temporal domai $\Omega = (0,2)^2 \times [0,5]$, where $(x,y) \in (0,2)^2$ is the spatial domain and $t \in [0,5]$ is the time interval. The objective is to find the state variables—velocity components $(u,v)$ and pressure $p$—as a continuous function $\mathbf{v} = (u,v,p) \in C(\Omega, \mathbb{R}^3)$ that satisfies the following initial and boundary conditions:

$$\begin{cases} u(x,y,0) = v(x,y,0) = p(x,y,0) = 0 & \forall x,y \in [0,2] \\ v(x,0,t) = v(x,2,t) = 0 & \forall x \in [0,2] \quad \forall t \in [0,5] \\ u(0,y,t) = 0 & \forall y \in [0,2] \quad \forall t \in [0,5] \\ u(2,y,t) = 1 & \forall y \in [0,2] \quad \forall t \in [0,5] \\ \frac{\partial p}{\partial \mathbf{n}}(x,0,t) = \frac{\partial p}{\partial \mathbf{n}}(x,2,t) = 0 & \forall x \in [0,2] \quad \forall t \in [0,5] \\ \frac{\partial p}{\partial \mathbf{n}}(0,y,t) = \frac{\partial p}{\partial \mathbf{n}}(2,y,t) = 0 & \forall y \in [0,2] \quad \forall t \in [0,5] \end{cases}$$

A PINN was trained without reliance on simulation data ($\lambda_1 = 0$). The network architecture consists of an input layer for spatio-temporal coordinates $(x,y,t)$, two hidden layers with 100 and 500 neurons, and an output layer predicts the solution $(u,v,p)$. This configuration contains a total of $52,403$ parameters.

Figure 5.5 presents a comparison between the PINN-derived solution and a conventional numerical simulation after 1000 training epochs. While the PINN accurately captures the qualitative behavior of the flow, a noticeable deviation from the numerical solution appears at $t = 5$ seconds. Furthermore, the model exhibits inaccuracies in the regions corresponding to the upper-left and upper-right corners of the domain.

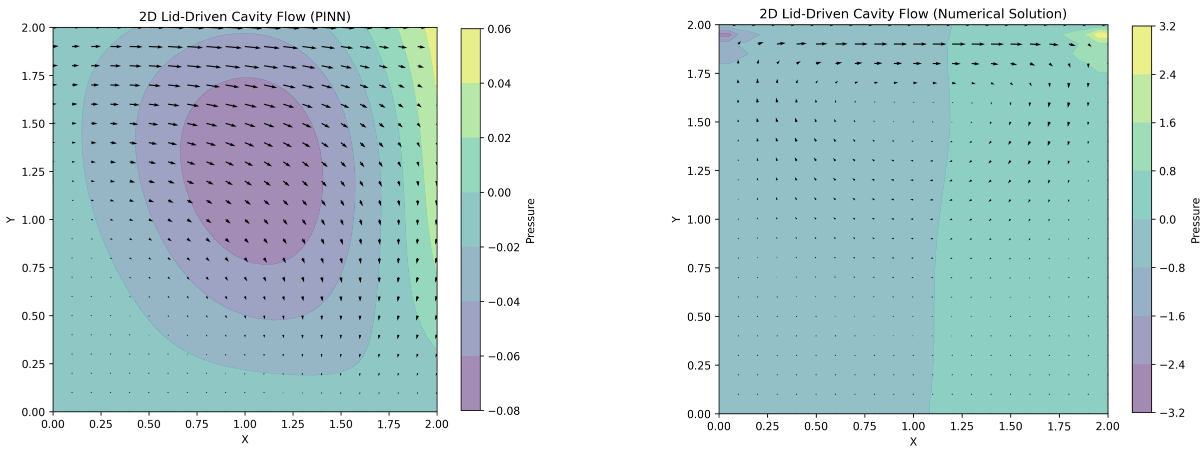


(a) Solution given by a PINN $u_\theta$ trained over 1000 epochs.

(b) Numerical Solution using finites differences and $n_x = n_y = 41$.

Figure 5.5: Comparison between methods for the Navier-Stokes equation.

# Conclusions

This manuscript began by introducing dense neural networks with tanh activation layers, establishing their fundamental properties. We demonstrated their denseness with respect to the Sobolev norm and their continuous Lipschitz dependence on both parameters and inputs. We then formally defined Physics-Informed Neural Networks (PINNs), detailing the empirical and theoretical risk functions that guide their training and evaluation.

After reviewing key results from stochastic processes, particularly concerning sub-Gaussian variables and covering and packing numbers, we investigated the critical property of risk-consistency. Through a counterexample, we proved that standard PINNs fail to satisfy risk-consistency. To address this limitation, we proposed a ridge regularization technique that constrains the minimization sequence, thereby ensuring risk-consistency.

Furthermore, we introduced a Sobolev regularizer that imposes stronger constraints on the solution space. This advanced regularization not only achieves strong consistency for PINNs but also establishes their order of convergence as the sample size increases. The applicability of these regularization techniques was then extended to the NeuralODE framework, where we similarly proved the risk-consistency of the ridge-regularized model.

Finally, the theoretical contributions were substantiated with numerical experiments, providing illustrative examples that corroborate our findings and highlight the practical implementation of the proposed methods.

# Bibliography


Arturs Berzins, Andreas Radler, Eric Volkmann, Sebastian Sanokowski, Sepp Hochreiter, and Johannes Brandstetter. Geometry-Informed Neural Networks, October 2024. URL http://arxiv.org/abs/2402.14009. arXiv:2402.14009.

Haim Brezis. *Functional Analysis, Sobolev Spaces and Partial Differential Equations*. Springer, New York, NY, 2011. ISBN 978-0-387-70913-0 978-0-387-70914-7. doi: 10.1007/978-0-387-70914-7. URL https://link.springer.com/10.1007/978-0-387-70914-7.

Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural Ordinary Differential Equations, December 2019. URL http://arxiv.org/abs/1806.07366. arXiv:1806.07366 [cs].

G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2(4):303–314, December 1989. ISSN 1435-568X. doi: 10.1007/BF02551274. URL https://doi.org/10.1007/BF02551274.

Tim De Ryck, Samuel Lanthaler, and Siddhartha Mishra. On the approximation of functions by tanh neural networks. *Neural Networks*, 143:732–750, November 2021. ISSN 0893-6080. doi: 10.1016/j.neunet.2021.08.015. URL https://www.sciencedirect.com/science/article/pii/S0893608021003208.

Nathan Doumèche, Gérard Biau, and Claire Boyer. Convergence and error analysis of PINNs, May 2023. URL http://arxiv.org/abs/2305.01240. arXiv:2305.01240 [math, stat].

Nathan Doumèche, Francis Bach, Gérard Biau, and Claire Boyer. Physics-informed kernel learning, September 2024. URL http://arxiv.org/abs/2409.13786. arXiv:2409.13786 [stat].

Lawrence C. Evans. *Partial Differential Equations*. American Mathematical Society, March 2022. ISBN 978-1-4704-6942-9. Google-Books-ID: Ott1EAAAQBAJ.

Charles Fefferman, Sanjoy Mitter, and Hariharan Narayanan. Testing the Manifold Hypothesis, December 2013. URL http://arxiv.org/abs/1310.0425. arXiv:1310.0425 [math, stat].

Takashi Furuya, Maarten V. de Hoop, and Gabriel Peyré. Transformers are Universal In-context Learners, August 2024. URL http://arxiv.org/abs/2408.01367. arXiv:2408.01367 [cs, stat].

Daehoon Gwak, Gyuhyeon Sim, Michael Poli, Stefano Massaroli, Jaegul Choo, and Edward Choi. Neural Ordinary Differential Equations for Intervention Modeling, October 2020. URL http://arxiv.org/abs/2010.08304. arXiv:2010.08304.

Zhongkai Hao, Songming Liu, Yichi Zhang, Chengyang Ying, Yao Feng, Hang Su, and Jun Zhu. Physics-Informed Machine Learning: A Survey on Problems, Methods and Applications, November 2022. URL https://arxiv.org/abs/2211.08064v2.

Michael Hardy. Combinatorics of Partial Derivatives, January 2006. URL http://arxiv.org/abs/math/0601149. arXiv:math/0601149.

George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, June 2021. ISSN 2522-5820. doi: 10.1038/s42254-021-00314-5. URL https://www.nature.com/articles/s42254-021-00314-5. Publisher: Nature Publishing Group.

Patrick Kidger. On Neural Differential Equations, February 2022. URL http://arxiv.org/abs/2202.02435. arXiv:2202.02435 [cs, math, stat].

Stefano Massaroli, Michael Poli, Jinkyoo Park, Atsushi Yamashita, and Hajime Asama. Dissecting Neural ODEs, January 2021. URL http://arxiv.org/abs/2002.08071. arXiv:2002.08071.

Marina Meilă and Hanyu Zhang. Manifold learning: what, how, and why, November 2023. URL http://arxiv.org/abs/2311.03757. arXiv:2311.03757 [cs, stat].

Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learning*. MIT Press, Boston, MA, 2012.

Richard Nickl and Benedikt M. Pötscher. Bracketing Metric Entropy Rates and Empirical Central Limit Theorems for Function Classes of Besov- and Sobolev-Type. *Journal of Theoretical Probability*, 20(2):177–199, June 2007. ISSN 1572-9230. doi: 10.1007/s10959-007-0058-1. URL https://doi.org/10.1007/s10959-007-0058-1.

M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, February 2019. ISSN 0021-9991. doi: 10.1016/j.jcp.2018.10.045. URL https://www.sciencedirect.com/science/article/pii/S0021999118307125.

Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Physics Informed Deep Learning (Part II): Data-driven Discovery of Nonlinear Partial Differential Equations, November 2017a. URL http://arxiv.org/abs/1711.10566. arXiv:1711.10566.

Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Physics Informed Deep Learning (Part I): Data-driven Solutions of Nonlinear Partial Differential Equations, November 2017b. URL https://arxiv.org/abs/1711.10561v1.

F. Rosenblatt. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6):386–408, 1958. ISSN 1939-1471. doi: 10.1037/h0042519. Place: US Publisher: American Psychological Association.

Domènec Ruiz-Balet and Enrique Zuazua. Neural ODE Control for Classification, Approximation, and Transport. *SIAM Review*, 65(3):735–773, August 2023. ISSN 0036-1445. doi: 10.1137/21M1411433. URL https://epubs.siam.org/doi/10.1137/21M1411433. Publisher: Society for Industrial and Applied Mathematics.

Yeonjong Shin, Jerome Darbon, and George Em Karniadakis. On the convergence of physics informed neural networks for linear second-order elliptic and parabolic type PDEs. *Communications in Computational Physics*, 28(5):2042–2074, January 2020. ISSN 1991-7120, 1815-2406. doi: 10.4208/cicp.OA-2020-0193. URL http://arxiv.org/abs/2004.01806. arXiv:2004.01806 [cs, math].

E.M. Stein. *Singular Integrals and Differentiability Properties of Functions*, volume 30 of *Princeton Mathematical Series*. Princeton University Press, Princeton, February 1971. URL https://press.princeton.edu/books/hardcover/9780691080796/singular-integrals-and-differentiability-properties-of-functions. ISBN: 9780691080796.

Ramon Van Handel. Probability in High Dimension:. Fort Belvoir, VA, June 2014. Defense Technical Information Center. doi: 10.21236/ADA623999. URL http://www.dtic.mil/docs/citations/ADA623999.

Joe Watson, Chen Song, Oliver Weeger, Theo Gruner, An T. Le, Kay Hansel, Ahmed Hendawy, Oleg Arenz, Will Trojak, Miles Cranmer, Carlo D'Eramo, Fabian Bülow, Tanmay Goyal, Jan Peters, and Martin W. Hoffman. Machine Learning with Physics Knowledge for Prediction: A Survey, August 2024. URL http://arxiv.org/abs/2408.09840. arXiv:2408.09840.

Dmitry Yarotsky. Error bounds for approximations with deep ReLU networks. *Neural Networks*, 94:103–114, October 2017. ISSN 0893-6080. doi: 10.1016/j.neunet.2017.07.002. URL https://www.sciencedirect.com/science/article/pii/S0893608017301545.

Dmitry Yarotsky. Optimal approximation of continuous functions by very deep ReLU networks. In *Proceedings of the 31st Conference On Learning Theory*, pages 639–649. PMLR, July 2018. URL https://proceedings.mlr.press/v75/yarotsky18a.html.

Dmitry Yarotsky. Universal Approximations of Invariant Maps by Neural Networks. *Constructive Approximation*, 55(1):407–474, February 2022. ISSN 1432-0940. doi: 10.1007/s00365-021-09546-1. URL https://doi.org/10.1007/s00365-021-09546-1.

# APPENDIX A
# Functional Analysis

In this section, we present several well-known results in functional analysis that will be used in future sections. Most of the results are gathered from Evans [2022], Stein [1971], Brezis [2011], Doumèche et al. [2024].

The Sobolev embedding of $H^m(\Omega)$ in $C(\Omega)$ [Evans, 2022, Chapter 5.6, Theorem 6]. It is used to operate in the continuous space and been able to do pointwise evaluations of functions.

**Theorem A.1** (Sobolev embedding)**.** Let $\Omega \subseteq \mathbb{R}^{d_1}$ be a bound Lipschitz domain and let $m \in \mathbb{N}$. If $m \geq d_1/2$, then there exists an operator $\tilde{\Pi} : H^m(\Omega, \mathbb{R}^{d_2}) \to C(\Omega, \mathbb{R}^{d_2})$ such that for any $u \in H^m(\Omega, \mathbb{R}^{d_2})$, $\tilde{\Pi}(u) = u$ almost everywhere. Moreover, there exists a constant $C_\Omega > 0$, depeding only on $\Omega$ such that $\|\tilde{\Pi}(u)\|_{\infty,\Omega} \leq C_\Omega \|u\|_{H^m(\Omega)}$.

The Lax-Milgram theorem can be found in [Brezis, 2011, Corollary 5.8].

**Theorem A.2** (Lax-Milgram)**.** Assume that $\mathcal{A}(u, v)$ is a continuous coercive bilinear for on $\mathcal{H}$ a Hilbert space. Then, given any $\varphi \in H^*$, there exists a unique element $u \in \mathcal{H}$ such that

$$\mathcal{A}(u, v) = \langle \varphi, v \rangle \quad \forall v \in \mathcal{H}.$$

Moreover, if $\mathcal{A}$ is symmetric, then $u \in \mathcal{H}$ is characterized by the property

$$\mathcal{A}(u, u) - 2\langle \varphi, u \rangle = \min_{v \in \mathcal{H}} \left\{ \mathcal{A}(v, v) - 2\langle \varphi, v \rangle \right\}.$$

The next theorem is used to find a convergent subsequence of a bounded sequence. The proof can be found in [Doumèche et al., 2023, Theorem B.4].

**Theorem A.3** (Rellich-Kondrachov compactness)**.** Let $\Omega \subseteq \mathbb{R}^{d_2}$ be a bounded Lipschitz domain and let $m \in \mathbb{N}$. Let $\{u_n\}_{n\in\mathbb{N}} \subseteq H^{m+1}(\Omega, \mathbb{R}^{d_2})$ be a sequence such that $\{\|u_n\|_{H^{m+1}(\Omega)}\}_{n\in\mathbb{N}}$ is bounded. Then, there exists a function $u_\infty \in H^{m+1}(\Omega, \mathbb{R}^{d_2})$ and a subsequence of $\{u_n\}_{n\in\mathbb{N}}$ that converges to $u_\infty$ both weakly in $H^{m+1}(\Omega, \mathbb{R}^{d_2})$ and with respect to the $H^m(\Omega)$ norm.

The next result is a generalization of [Evans, 2022, Theorem 1, Chapter 8.2] that can be found in [Doumèche et al., 2024, Lemma C.11]. Before introducing the theorem we will define the concept of weak lower semi-continuity.

**Definition A.4** (Weak lower semi-continuity)**.** A function $I : H^m(\Omega) \to \mathbb{R}$ is weakly lower semi-continuous on $H^m(\Omega)$ if for any sequence $\{u_n\}_{n\in\mathbb{N}} \subseteq H^m(\Omega)$ that weakly converges to $u_\infty \in H^m(\Omega)$, one has that $I(u_\infty) \leq \liminf_{n\to\infty} I(u_n)$.

**Lemma A.5** (Weak lower semi-continuity with convex Lagrangia)**.** Let the Lagrangian $L \in C^\infty(\mathbb{R}^{\binom{d_1+m}{m}d_2} \times \cdots \times \mathbb{R}^{d_2} \times \mathbb{R}^{d_1}, \mathbb{R})$ be such that for any $x^{(m)}, \ldots, x^{(0)}$, and $z$, the function $x^{(m+1)} \mapsto L(x^{(m+1)}, \ldots, x^{(0)}, z)$ is convez and nonnegative. Then, the function

$$I : u \mapsto \int_\Omega L((\partial^{m+1}_{i_1,\ldots,i_{m+1}} u(x))_{1\leq i_1, i_{m+1}\leq d_1}, \ldots, u(x), x)dx$$

is lower-semi continuous for the weak topology on $H^{m+1}(\Omega, \mathbb{R}^{d_2})$.

# APPENDIX B
# Generalization

**Theorem B.1** (Uniform approximation of integrals II)**.** Let $(T, d)$ be a metric space and let $\Omega \subseteq T$ be a doubling space and relatively compact set, let $\alpha_1 > 0$ and let $\{X_k\}_{k\in\mathbb{N}}$ be a sequence of i.i.d. random variables in $\bar{\Omega}$, with distribution $\mu$. Let $f : C(\bar{\Omega}, \mathbb{R}^{d_2}) \times \bar{\Omega} \to \mathbb{R}$ be an operator, and we assume that the following two requirements are satisfied:

(i) there exists $C_1 > 0$ and $\beta_1 \in (0, 1/2)$ such that for all $n \geq 1$ and all $u, v \in B(0, n^{\alpha_1}) \cap L\text{-Lip}(\Omega, \mathbb{R}^{d_2})$,

$$\|f(u, \cdot) - f(v, \cdot)\|_{\infty,\bar{\Omega}} \leq C_1 n^{\beta_1} \|u - v\|_\infty \tag{H.1}$$

(ii) there exist $C_2 > 0$ and $\beta_2 \in (0, 1/2)$ satisfying $\beta_2 > \alpha_1 + \beta_1$ such that for all $n \geq 1$, and all $u \in B_2(0, n^{\alpha_1}) \cap L\text{-Lip}(\Omega, \mathbb{R}^{d_2})$,

$$\|f(u, \cdot)\|_{\infty,\bar{\Omega}} \leq C_2 n^{\beta_2} \tag{H.2}$$

Then, almost surely, there exists $N \in \mathbb{N}$ such that for all $n \geq N$,

$$\sup_{u\in B(0,n^{\alpha_1})\cap L\text{-Lip}(\Omega,\mathbb{R}^{d_2})} \left\| \frac{1}{n} \sum_{i=1}^{n} f(u, X_i) - \int_{\bar{\Omega}} f(u, \cdot) d\mu \right\|_2 \leq \log^2(n) n^{\beta_2 - 1/2}$$

Notice that $N$ is itself a random variable because it depends on $X_1, \ldots, X_n$. Moreover, note that for $\beta_2 \in (0, 1/2)$, $\lim_{n\to\infty} \log^2(n) n^{\beta_2-1/2} = 0$. This theorem assures us that the empirical mean approximates the integral under certain assumptions over the operator $f$.

*Proof.* Let us first analyze the particular case where $d_2 = 1$ before generalizing. Consider $u \in B(0, n^{\alpha_1}) \cap L\text{-Lip}(\Omega, \mathbb{R}^{d_2})$, and we let

$$Z_{n,u} := \frac{1}{n} \sum_{i=1}^{n} f(u, X_i) - \int_{\bar{\Omega}} f(u, \cdot) d\mu$$

Our objective is to establish an upper bound for the random variable

$$Z_n := \sup_{u \in B(0,n^{\alpha_1}) \cap L\text{-Lip}(\Omega,\mathbb{R}^{d_2})} |Z_{n,u}|$$

Let $M(\Omega) := \max_{x\in\bar{\Omega}} \|x\|_2$. For every $k = 1, \dots, n$ we define the random variable $W_{k,n}$ as follows

$$W_{k,n} := \frac{1}{n}\left(f(u,X_k) - \int_{\bar{\Omega}} f(u,\cdot)d\mu\right) - \frac{1}{n}\left(f(v,X_k) - \int_{\bar{\Omega}} f(v,\cdot)d\mu\right)$$

Then, using hypthesis (H.1), we obtain that for any $u, v \in B_2(0, n^{\alpha_1}) \cap L\text{-Lip}(\Omega, \mathbb{R}^{d_2})$,

$$\begin{aligned} |W_{i,n}| &\le \frac{1}{n}\left(C_1 n^{\beta_1}\|u-v\|_\infty + \int_{\bar{\Omega}} C_1 n^{\beta_1}\|u-v\|_\infty d\mu\right) \overset{\int_{\bar{\Omega}} d\mu = 1}{=} 2C_1 n^{\beta_1 - 1}\|u-v\|_\infty \\ &\le 2C_1 n^{\beta_1-1}n^{\alpha_1} \end{aligned}$$

In other words, we have found $a, b \in \mathbb{R}$ such that $a \le W_{k,n} \le b$. By Hoeffding's Lemma 2.4, we know that $W_{k,n}$ is a $\sigma_k^2-$subgaussian random variable with

$$\sigma_k^2 = \frac{1}{4}(b-a)^2 = \frac{1}{4}(4C_1 n^{\beta_1+\alpha_1-1}\|u-v\|_\infty)^2 = 4C_1^2 n^{2\beta_1-2}\|u-v\|_\infty^2$$

On the other hand, we note that $Z_{n,u} - Z_{n,v} = \sum_{k=1}^n W_{k,n}$ and since the random variables $\{W_{k,n}\}_{k\le n}$ are independent, we can invoke Azuma's Lemma 2.5 to deduce that $Z_{n,u} - Z_{n,v}$ is a subgaussian random variable with parameter $\sum_{k=1}^n \sigma_k^2 = n\sigma_k^2 = 4C_1^2 n^{2\beta_1-1}\|u-v\|_\infty^2$. We observe that $\mathbb{E}[Z_{n,u}] = 0$. Therefore, by definition, for all $n \ge 1$, the family $\{Z_{n,u}\}_{u\in T}$ is a subgaussian process on the metric space $(\tilde{T}, \tilde{d})$, where

$$\tilde{T} = B_2(0,n^{\alpha_1}) \cap L\text{-Lip}(\Omega, \mathbb{R}^{d_2}), \qquad \tilde{d}(u,v) = 2C_1 n^{\beta_1 - 1/2}\|u-v\|_\infty$$

Furthermore, the mapping $u \mapsto Z_{n,u}$ is continuous with respect to the topology induced by the metric $\tilde{d}$. Consequently, by Remark 2.14, the process $(Z_{n,u})_{u\in B_2(0,n^{\alpha_1})\cap L\text{-Lip}(\Omega,\mathbb{R}^{d_2})}$ is separable. Applying Dudley's Corollary 2.16, we obtain

$$\mathbb{E}\left[\sup_{\theta\in B_2(0,n^{\alpha_1})\cap L\text{-Lip}(\Omega,\mathbb{R}^{d_2})} Z_{n,\theta}\right] \equiv \mathbb{E}[Z_n] \le 12\int_0^\infty \sqrt{\log N(\tilde{T},\tilde{d},r)}dr \tag{B.1}$$

We can now establish the following relation

$$\begin{aligned} N(\tilde{T},\tilde{d},r) &= N(B(0,n^{\alpha_1}) \cap L\text{-Lip}(\Omega,\mathbb{R}^{d_2}), \tilde{d}(u,v), r) \\ &= N(B(0,n^{\alpha_1}), \|\cdot\|_\infty, n^{1/2-\beta_1}r/(2C_1)) \\ &= N(B(0,1), \|\cdot\|_\infty, n^{1/2-\beta_1-\alpha_1}r/(2C_1)) \end{aligned} \tag{B.2}$$

Combining (B.1) and (B.2), we obtain

$$\begin{aligned} \mathbb{E}[Z_n] &\le 12\int_0^\infty \sqrt{\log N(B_2(0,n^{\alpha_1}), \|\cdot\|_\infty, n^{1/2-\beta_1}r/(2C_1))}dr \\ &= 24C_1 n^{\alpha_1+\beta_1-1/2}\int_0^\infty \sqrt{\log N(B_2(0,1),\|\cdot\|_\infty, s)}ds \\ &= 24C_1 n^{\alpha_1+\beta_1-1/2}\int_0^1 \sqrt{(16L/s)^{ddim(\Omega)}\log(4M(\Omega)/s)}ds \end{aligned}$$

In the second term, we performed a change of variable given by $s = n^{1/2-\beta_1} r/(2C_1)$, while the third term follows from the fact that $N(B_2, \|\cdot\|_2, r) = 1$ $\forall r \geq 1$. Since $\theta$ belongs to the parameter space $\Theta_H^D := \mathbb{R}^{(d_1+1)D+(H-1)D(D+1)+(D+1)d_2}$, we can apply Lemma 2.10 to bound the covering number as follows

$$\log N(B_2^n, \|\cdot\|_2, s) \leq [(d_1+1)D + (H-1)D(D+1) + (D+1)d_2] \log(3/s)$$

Observing that $\int_0^1 \sqrt{\log(3/s)}ds \leq 3/2$ and noting that, by definition, $\alpha_1 + \beta_1 < \beta_2 < 1/2$, we observe that

$$\mathbb{E}[Z_n] \leq 36 C_1 n^{\alpha_1+\beta_1-1/2} \sqrt{(d_1+1)D + (H-1)D(D+1) + (D+1)d_2} \xrightarrow{n\to\infty} 0 \qquad \text{(B.3)}$$

Note that $Z_n = Z_n(X_1, \ldots, X_n)$, so applying the definition of the discrete derivative of $Z_n$ (Definition 2.17) along with the hypothesis (H.2), we obtain

$$\begin{aligned}\mathfrak{D}_i Z_n &:= \sup_{x_i \in \mathbb{R}^{d_1}} Z_n(X_1, \ldots, X_{i-1}, x_i, X_{i+1}, \ldots, X_n) - \inf_{x_i \in \mathbb{R}^{d_1}} Z_n(X_1, \ldots, X_{i-1}, x_i, X_{i+1}, \ldots, X_n) \\ &\leq 2n^{-1} \sup_{\theta \in B_2(0, n^{\alpha_1})} \|f(u_\theta, \cdot)\|_\infty \leq 2C_2 n^{\beta_2 - 1}\end{aligned}$$

Applying McDiarmid's Theorem 2.18, we conclude that $Z_n(X_1, \ldots, X_n)$ is $\sigma_z^2-$subgaussian variable with $\sigma_z^2 = \frac{1}{4} n (2C_2 n^{\beta_2-1})^2 = C_2^2 n^{2\beta_2-1}$. In particular, by the Chernoff bound (Lemma 2.2), for all $t_n \geq 0$, we obtain

$$\mathbb{P}(|Z_n - \mathbb{E}Z_n| \geq t_n) \leq 2e^{-t_n^2/2\sigma_z^2} = 2\exp\{-t_n^2 n^{1-2\beta_2}/2C_2^2\}$$

We will now show that the events $E_n := \{|Z_n - \mathbb{E}Z_n| \leq t_n\}$ satisfy the Borel-Cantelli Lemma when the sequence $\{t_n\}_n$ is given by $t_n := 2C_2 n^{\beta_2-1/2} \log^{1/2}(n)$. Indeed,

$$\sum_{n\in\mathbb{N}} \mathbb{P}(E_n) \leq \sum_{n\in\mathbb{N}} 2\exp\{-(2C_2 n^{\beta_2-1/2} \log^{1/2}(n^2))^2 n^{1-2\beta_2}/2C_2^2\} < \infty$$

Therefore, using the fact that $\mathbb{E}[Z_n] \to 0$ (B.3) and applying Borel-Cantelli Lemma yields the desired result

$$0 \leq Z_n \leq \log^{1/2}(n) n^{\beta_2-1/2} \leq \log^2(n) n^{\beta_2-1/2} \quad a.s. \qquad \forall n \geq N$$

for some sufficiently large $N$.

The generalization to the case $d_2 > 1$ is straightforward. Let $f = (f_1, \ldots, f_{d_2})$. Since $\|\cdot\|_2 \leq \sqrt{d_2}\|\cdot\|_\infty$, then

$$\begin{aligned}&\sup_{\theta\in B_2(0,n^{\alpha_1})} \left\| \frac{1}{n}\sum_{i=1}^n f(u_\theta, X_i) - \int_\Omega f(u_\theta, \cdot)d\mu \right\|_2 \\ &\qquad \leq \sqrt{d_2} \max_{1\leq j\leq n} \sup_{\theta\in B_2(0,n^{\alpha_1})} \left\| \frac{1}{n}\sum_{i=1}^n f_j(u_\theta, X_i) - \int_\Omega f_j(u_\theta, \cdot)d\mu \right\|_2\end{aligned}$$

□