

On the estimation rate of Bayesian PINN for inverse problems

YI SUN^{1,a}, DEBARGHYA MUKHERJEE^{??,a} and YVES ATCHADÉ^{??,a}

¹*Department of Mathematics and Statistics, Boston University, Boston, USA, ^aatchade@bu.edu*

Solving partial differential equations (PDEs) and their inverse problems using Physics-informed neural networks (PINNs) is a rapidly growing approach in the physics and machine learning community. Although several architectures exist for PINNs that work remarkably in practice, our theoretical understanding of their performances is somewhat limited. In this work, we study the behavior of a Bayesian PINN estimator of the solution of a PDE from n independent noisy measurement of the solution. We focus on a class of equations that are linear in their parameters (with unknown coefficients θ_\star). We show that when the partial differential equation admits a classical solution (say u_\star), differentiable to order β , the mean square error of the Bayesian posterior mean is at least of order $n^{-2\beta/(2\beta+d)}$. Furthermore, we establish a convergence rate of the linear coefficients of θ_\star depending on the order of the underlying differential operator. Last but not least, our theoretical results are validated through extensive simulations.

Keywords: partial differential equations; inverse problems; physics-informed neural network; nonparametric regression; Bayesian deep learning; posterior contraction

1. Introduction

We consider a class of inverse problems that consists in estimating the parameters of a partial differential equation from noisy measurements of the solution. Specifically, let $\Omega \subset \mathbb{R}^m$ be a bounded open domain with a smooth boundary equipped with a probability measure ν . For $j \geq 1$, let $L^2(\Omega, \mathbb{R}^j, \nu)$ denote the L^2 -space of \mathbb{R}^j -valued functions on Ω , i.e.

$$L^2(\Omega, \mathbb{R}^j, \nu) \stackrel{\text{def}}{=} \left\{ f : \Omega \mapsto \mathbb{R}^j : \int_{\Omega} \|f(x)\|_2^2 \nu(dx) < \infty \right\}.$$

We consider a differential equation with parameter $\theta \in \mathbb{R}^d$ that aims to find a smooth function $u \in L^2(\Omega, \mathbb{R}, \nu)$ satisfying

$$\mathcal{H}_0 u + \theta^\top \mathcal{H}_1 u = f, \text{ with initial/boundary condition } \mathcal{B}u = \mathbf{0}, \quad (1.1)$$

for some known function $f \in L^2(\Omega, \mathbb{R}, \nu)$, and a pair of differential operators $\mathcal{H}_0 : L^2(\Omega, \mathbb{R}, \nu) \rightarrow L^2(\Omega, \mathbb{R}, \nu)$ and $\mathcal{H}_1 : L^2(\Omega, \mathbb{R}, \nu) \rightarrow L^2(\Omega, \mathbb{R}^d, \nu)$. The operator \mathcal{H}_0 is typically (although not necessarily) the "time" partial derivative, and the operator $\mathcal{B} : L^2(\Omega, \mathbb{R}, \nu) \rightarrow L^2(\Omega, \mathbb{R}^{d_1}, \nu)$ imposes the boundary/initial conditions. A large class of differential equations (linear and nonlinear) that are linear in their parameters can be written in this form. A classical example of such an inverse problem arises from the heat equation as elaborated below:

Example (Heat equation). Given $L > 0, T > 0$, and parameter $\theta \in \mathbb{R}$, consider the heat equation $u_t - \theta u_{xx} = 0$ on $(0, L)$, with time domain $(0, T)$, initial condition $u(0, \cdot) = g(\cdot)$, and boundary condition $u(\cdot, 0) = u(\cdot, L) = 0$. Here u_t (resp. u_{xx}) denotes the partial derivative of u with respect to t (resp. the second order partial derivative of u with respect to x). With $\Omega = (0, T) \times (0, L)$, this equation can be

written as $\mathcal{H}_0 u - \theta \mathcal{H}_1 u = f$, where $\mathcal{H}_0 u = u_t$, $\mathcal{H}_1 u = u_{xx}$, and $f = 0$. The boundary/initial condition operator is given by

$$\mathcal{B}u = \begin{pmatrix} u_{\{t=0\}} - g \\ u_{\{x=0\}} \\ u_{\{x=L\}} \end{pmatrix}.$$

where $u_{\{t=0\}}$ is the map $u_{\{t=0\}}(x) = u(0, x)$, with similar definition for $u_{\{x=0\}}$ and $u_{\{x=L\}}$. Here $d = 1$, $d_1 = 3$.

Let (θ_\star, u_\star) be a tuple that satisfies the PDE (1.1). We assume to observe some noisy measurements of u_\star as noted below:

H1. We have n i.i.d. random locations and observations $(\mathbf{s}_i, Y_i) \in \Omega \times \mathbb{R}$, where

$$Y_i \mid \mathbf{s}_i \stackrel{\text{ind.}}{\sim} \mathbf{N}(u_\star(\mathbf{s}_i), \sigma^2),$$

where $u_\star : \Omega \rightarrow \mathbb{R}$ is the unique solution of the pde (1.1) with $\theta = \theta_\star$, for some unknown vector $\theta_\star \in \mathbb{R}^d$. Here, $\mathbf{N}(m, v^2)$ denotes the univariate Gaussian distribution with mean m and variance v^2 . The variance parameter σ^2 is assumed to be known. Throughout we write \mathbb{P} for the joint distribution of $(\mathbf{s}_1, Y_1), \dots, (\mathbf{s}_n, Y_n)$, and \mathbb{P}_n for their corresponding empirical measure.

Our goal is to estimate (θ_\star, u_\star) using the noisy measurements. Over the last few years, physics-informed neural network (PINN) has taken the numerical pde literature by storm ((24, 30, 33, 40), and (8) for an extensive review of the literature). In the setting of H1, the approach consists of estimating u_\star by regression while explicitly using the information that the true function u_\star is the solution of a PDE. In this paper, as typically done in the PINN literature, we resort to a sieve-based approach for estimating u_\star , i.e., we consider a sequence of increasingly complex parametric models. Hence let $\mathcal{F} \stackrel{\text{def}}{=} \{u_W, W \in \mathbb{R}^q\}$ be a function class, where $u_W : \Omega \rightarrow \mathbb{R}$ is a function with parameter $W \in \mathbb{R}^q$. Due to their superior empirical performances, u_W is typically taken as a deep neural network, and we follow that practice, although our results can be applied more broadly. Given the data $\mathcal{D} \stackrel{\text{def}}{=} \{(\mathbf{s}_i, Y_i), 1 \leq i \leq n\}$, PINN in its frequentist formulation estimates jointly θ and W by minimizing the loss function

$$\frac{1}{2\sigma^2} \sum_{i=1}^n (Y_i - u_W(\mathbf{s}_i))^2 + \frac{\lambda}{2} \left\{ \alpha_1 \|f - \mathcal{H}_0 u_W - \theta^\top \mathcal{H}_1 u_W\|_2^2 + \alpha_2 \|\mathcal{B}u_W\|_2^2 \right\}, \quad (1.2)$$

for some regularization parameters $\lambda \geq 0$, and $\alpha_1, \alpha_2 \geq 0$, where $\|\cdot\|_2$ denotes the L^2 -norm on the appropriate function space. In this work, we approach the problem from a Bayesian perspective. Hence, starting from a standard Gaussian prior for $\theta \in \mathbb{R}^d$, and a sparsity inducing prior density Π_0 for $W \in \mathbb{R}^q$ (see Section 1.5 below for our choice of Π_0), we define the informative PINN prior distribution for (θ, W) as the probability measure on $\mathbb{R}^d \times \mathbb{R}^q$ with density proportional to

$$(\theta, W) \mapsto \Pi_0(W) \exp \left(-\frac{\lambda}{2} \left\{ \alpha_1 \|f - \mathcal{H}_0 u_W - \theta^\top \mathcal{H}_1 u_W\|_2^2 + \alpha_2 \|\mathcal{B}u_W\|_2^2 \right\} - \frac{1}{2} \|\theta\|_2^2 \right). \quad (1.3)$$

For appropriately large choices of λ , any realization (θ, W) from the PINN prior distribution (1.3) produces (θ, u_W) that is biased toward solving the pde (1.1). Rigorous general results (not specific to PINN) of this flavor can be found, for instance, in (15). When θ is known, realizations from the

conditional distribution of W given θ in (1.3) produce u_W that approximately solve the pde (1.1) for the given θ . This corresponds to the initial PINN methodology of (24).

For some convenience in the analysis, we will use the prior (1.3) with $\alpha_1 = 1$, and $\alpha_2 = 0$. Since the parameter θ does not appear in the boundary condition, the case $\alpha_2 > 0$ (although useful in practice) does not induce any fundamentally new behavior, and at the expense of more involved notations, our analysis can be easily modified to handle $\alpha_2 > 0$. Given the observed data \mathcal{D} , the data generating model postulated in (H1), and the PINN prior distribution (1.3), we thus consider the posterior distribution

$$\begin{aligned} \Pi(d\theta, dW|\mathcal{D}) \propto \\ \exp\left(-\frac{1}{2\sigma^2} \sum_{i=1}^n (Y_i - u_W(s_i))^2 - \frac{\lambda}{2} |f - \mathcal{H}_0 u_W - \theta^\top \mathcal{H}_1 u_W|_2^2 - \frac{\|\theta\|_2^2}{2}\right) \Pi_0(dW) d\theta. \end{aligned} \quad (1.4)$$

1.1. Main contributions

We study the behavior of the posterior distribution (1.4) as $n \rightarrow \infty$ (and $q \rightarrow \infty$). We consider the case where the operator $\mathcal{H} \stackrel{\text{def}}{=} (\mathcal{H}_0, \mathcal{H}_1)$ is a differential operator that involves derivatives up to order τ , for some $\tau > 0$. Under some additional regularity conditions, we establish in Theorem 3.1 that when the true pde solution u_\star is a classical solution that possesses derivatives up to order $\beta \geq \tau$, the mean square error of the posterior mean of θ under (1.4) satisfies (up to log terms that we ignore),

$$\mathbb{E} \left[\left\| \int_{\mathbb{R}^d} \vartheta \Pi^{(\theta)}(d\vartheta|\mathcal{D}) - \theta_\star \right\|_2^2 \right] \lesssim n^{-2(\beta-\tau)/(m+2\beta)}, \quad (1.5)$$

where $\Pi^{(\theta)}$ is the θ -marginal of (1.4). In fact we obtain this rate by showing that the 2-Wasserstein distance between $\Pi^{(\theta)}$ (the posterior distribution of θ) and the Gaussian distribution $\mathbf{N}(\theta_\star, \lambda^{-1}\Sigma_\star)$ converges to zero at the rate given above (see Section 2 for the definition of Σ_\star). Our simulation results suggest that, in general, this convergence does not hold in the total variation metric. We also show in Theorem 2.2 that the pde solution u_\star is estimated at a rate at least as fast as the nonparametric minimax optimal rate:

$$\mathbb{E} \left[\int_{\mathbb{R}^q} |u_W - u_\star|_2^2 \Pi^{(W)}(dW|\mathcal{D}) \right] \lesssim n^{-2\beta/(m+2\beta)},$$

where $\Pi^{(W)}$ denotes the W -marginal of (1.4). We conjecture that the PINN posterior distribution actually converges faster than the nonparametric minimax rate. Indeed, in a limiting case of an infinitely strong PINN prior, we construct an estimator which, although computationally intractable in general, achieves the parametric rate $n^{-1/2}$ (Theorem 2.3).

1.2. Related work

Despite their popularity, the theoretical aspects of PINN-related methods remain under-studied. Most of the existing theoretical literature on PINN deals with the forward problem of estimating u for a given θ , when the L^2 -norm $|\cdot|_2$ is replaced by a Monte Carlo approximation ((17, 20, 29)). The objective in this line of work is to derive the mean square error of the PINN solution as a function of the Monte Carlo sample size. More related to our work, (19) considers a PINN inverse problem where the parameter of interest appears in the boundary condition. They work mainly in a noiseless

(or small-magnitude noise) data regime. Furthermore, they assume a continuous space observation model, which is substantially different from the framework considered here.

Inverse problems as related to differential and partial differential equations have a long history that predates the current PINN literature ((18, 25, 36)). In particular, we note the close similarity between the PINN methodology and the approach by (36). The estimation rate in pde-driven inverse problems has been extensively studied in statistics in recent years ((9, 13, 21, 22, 31)). However, none of these results are directly applicable to our setting or PINN more generally, since in this literature, the forward map is typically assumed known and is not parametrized as in PINN.

A Bayesian version of PINN was proposed and studied in ((37)) from an empirical viewpoint, where the authors noted a certain robustness of Bayesian PINN to measurement noise. However, no theoretical analysis is proposed.

A related version of the inverse problem considered in this work assumes that θ is high-dimensional and sparse ((6, 7)). The goal is then to estimate θ under a sparsity assumption. These physics discovery problems can be viewed as high-dimensional versions of the problem investigated here. Thus, our work serves as a stepping-stone toward a theoretical understanding of high-dimensional physics discovery inverse problems.

1.3. Outline of the paper

We close the introduction with a description of the function class $\mathcal{F} = \{u_W, W \in \mathbb{R}^q\}$, the prior distribution Π_0 for W , and a description of the main notations that we use throughout the paper. Section 2 deals with the estimation rate for u_\star . A Bernstein-von Mises theorem for the marginal distribution of θ is discussed in Section 3. Section 4 presents our numerical illustrations of the theoretical findings.

1.4. The function class $\mathcal{F} = \{u_W\}$

Although our results apply widely, we follow the PINN literature and focus on the case where the function class \mathcal{F} is constructed using neural networks. Let $\xi > 0$ be the depth of the network. Let (p_ξ, \dots, p_0) be a sequence of integers representing the sizes of the layers of the network, with $p_0 = m$, and $p_\xi = 1$. For $1 \leq i \leq \xi$, let $A_i : \mathbb{R}^{p_i} \rightarrow \mathbb{R}^{p_i}$ be a component-wise application of a 1-Lipschitz function $a_i : \mathbb{R} \rightarrow \mathbb{R}$. For $B \in \mathbb{R}^{p_i \times p_{i-1}}$, and $b \in \mathbb{R}^{p_i}$, we set

$$\Psi_{B,b}^{(i)}(\mathbf{z}) \stackrel{\text{def}}{=} A_i(B\mathbf{z} + b), \quad \mathbf{z} \in \mathbb{R}^{p_{i-1}}. \quad (1.6)$$

We consider functions u_W of the form

$$u_W(\mathbf{x}) = \Psi_{W_\xi, w_\xi}^{(\xi)} \circ \dots \circ \Psi_{W_1, w_1}^{(1)}(\mathbf{x}), \quad \mathbf{x} \in \mathbb{R}^m, \quad (1.7)$$

with parameter $W = (W_\xi, w_\xi, \dots, W_1, w_1)$, where $W_i \in \mathbb{R}^{p_i \times p_{i-1}}$, and $w_i \in \mathbb{R}^{p_i}$, and where $f \circ g$ is the composition of f with g . For convenience, and by vectorization we view W as an element of \mathbb{R}^q , where

$$q \stackrel{\text{def}}{=} \sum_{i=1}^{\xi} p_i(1 + p_{i-1}).$$

We assume that the activation functions a_ℓ are of class C^∞ , which implies that u_W is of class C^∞ . The activation function a_ξ in the last layer is typically taken as the identity function.

1.5. The prior distribution Π_0

We use sparsity to control the complexity of the function class $\{u_W, W \in \mathbb{R}^q\}$. We use a version of the spike-and-slab prior distribution taken from (2). We set $\mathcal{S} \stackrel{\text{def}}{=} \{0, 1\}^q$. To construct the prior Π_0 , we first define a pair of random vectors $(\Lambda, W) \in \mathcal{S} \times \mathbb{R}^q$ as follows. Let $\Lambda_j \stackrel{i.i.d}{\sim} \mathbf{Ber}((1 + q^{u+1})^{-1})$, $1 \leq j \leq q$, for some sparsity parameter $u \geq 1$, where $\mathbf{Ber}(\alpha)$ denotes the Bernoulli distribution with parameter $\alpha \in (0, 1)$. Given $\Lambda \in \mathcal{S}$, the components of W are independent with joint density on \mathbb{R}^q given by

$$W \mapsto \prod_{j: \Lambda_j=1} \sqrt{\frac{1}{2\pi}} e^{-\frac{1}{2}W_j^2} \prod_{j: \Lambda_j=0} \sqrt{\frac{\rho_0}{2\pi}} e^{-\frac{\rho_0}{2}W_j^2},$$

for some given parameter $\rho_0 > 1$. We denote $\bar{\Pi}_0$ the joint distribution of (Λ, W) , and we let Π_0 (our prior on \mathbb{R}^q) be the distribution of $W \odot \Lambda$, where for $a, b \in \mathbb{R}^q$, $a \odot b$ denotes the component-wise product of a and b . By construction, for any measurable function $f: \mathbb{R}^q \rightarrow \mathbb{R}$, we have

$$\Pi_0(f) \stackrel{\text{def}}{=} \int_{\mathbb{R}^q} f(v) \Pi_0(dv) = \int_{\mathcal{S} \times \mathbb{R}^q} f(\Lambda \odot W) \bar{\Pi}_0(d\Lambda, dW).$$

1.6. Notation

For $\alpha \geq 1$, $\|\cdot\|_\alpha$ denotes the ℓ^α -norm on finite-dimensional Euclidean spaces ($\mathbb{R}^d, \mathbb{R}^n, \mathbb{R}^d$, etc., – which space should be clear from the context), and we use $a^\top b$ to denote the inner product between two finite dimensional vector a, b . As usual, we allow $\alpha = +\infty$ (resp. $\alpha = 0$), by defining $\|a\|_\infty \stackrel{\text{def}}{=} \max_{1 \leq i \leq d} |a_i|$ (resp. $|a|_0$ is the number of non-zero components of a). Given a matrix A , $\lambda_{\min}(A)$ (resp. $\|A\|_{\text{op}}$) denotes its smallest (resp. largest) singular value.

Let $\Omega \subset \mathbb{R}^m$ be as above with a probability measure ν (typically the uniform measure on Ω). Given an integer $j \geq 1$, $L^2(\Omega, \mathbb{R}^j, \nu)$ denotes the Hilbert L^2 -space of \mathbb{R}^j -valued functions with inner product

$$\langle f_1, f_2 \rangle \stackrel{\text{def}}{=} \int_{\Omega} f_1(\mathbf{x})^\top f_2(\mathbf{x}) \nu(d\mathbf{x}), \quad \text{and norm } \|f\|_2 \stackrel{\text{def}}{=} \sqrt{\langle f, f \rangle}.$$

For $f: \Omega \rightarrow \mathbb{R}^j$, we set $\|f\|_\infty \stackrel{\text{def}}{=} \sup_{x \in \Omega} \|f(x)\|_\infty$.

We use multi-index derivatives: given $\mathbf{k} = (k_1, \dots, k_m)$, where $k_i \geq 0$ is an integer,

$$D^{\mathbf{k}}u(x) \stackrel{\text{def}}{=} \frac{\partial^{|\mathbf{k}|} u(x)}{\partial x_1^{k_1} \dots \partial x_m^{k_m}},$$

where $|\mathbf{k}| = \sum_i k_i$. Given $\beta > 1$, we set

$$\|u\|_{C^\beta} \stackrel{\text{def}}{=} \sum_{\alpha: |\alpha| \leq \lfloor \beta \rfloor} \|D^\alpha u\|_\infty + \sum_{\alpha: |\alpha| = \lfloor \beta \rfloor} \sup_{x \neq y} \frac{|D^\alpha u(y) - D^\alpha u(x)|}{\|y - x\|_2^{\beta - \lfloor \beta \rfloor}},$$

where $\lfloor \beta \rfloor$ is the largest integer strictly smaller than β . $\|u\|_{C^\beta}$ is the Holder norm of u . We set $C^\beta(\Omega)$ (resp. $C^\beta(\Omega, b)$) as the set of all function $u: \Omega \rightarrow \mathbb{R}$ with finite Holder norm (resp. with Holder norm bounded by b).

Throughout the paper, we use $C_0, C_1, C_2, \text{etc.}$ to denote absolute constants and C to denote a generic absolute constant that depends on $C_0, C_1, \text{etc.}$. The actual value of C may change from one appearance

to the next. Similarly, $c_0, c_1, \text{etc.}$ denote problem-dependent constants (that do not depend on the sample size n or the model size q). Specifically, these constants typically depend on the noise parameter σ^2 , the dimension d , and the true parameter θ_\star . We will also use c to denote a generic constant that depends on $c_0, c_1, \text{etc.}$, and that we do not track. The actual value of c may change from one appearance to the next.

2. Estimation rate of u_\star

We maintain throughout the basic assumption that the operator $\mathcal{H} = (\mathcal{H}_0, \mathcal{H}_1)$ is a differential operator of order $\tau > 0$, in the sense that for all $\beta > \tau$, and for $u, v \in C^\beta(\Omega)$,

$$|\mathcal{H}u - \mathcal{H}v|_2 \leq C_0 \max_{\mathbf{k}: |\mathbf{k}| \leq \tau} |D^{\mathbf{k}}u - D^{\mathbf{k}}v|_\infty, \quad (2.1)$$

for some absolute constant C_0 .

Because the pde equation is linear in θ , and the log-prior is quadratic in θ , it is straightforward to integrate out θ from the posterior distribution (1.4). We first introduce some appropriate notations to do this. For $W \in \mathbb{R}^q$, we define $\bar{f}_W \stackrel{\text{def}}{=} f - \mathcal{H}_0 u_W$. Let $\Phi_W \in \mathbb{R}^d$ be the vector with j -th component given by

$$(\Phi_W)_j \stackrel{\text{def}}{=} \langle \bar{f}_W, (\mathcal{H}_1 u_W)_j \rangle,$$

and where $(\mathcal{H}_1 u_W)_j$ denotes the j -component of $\mathcal{H}_1 u_W$. We define $\Phi_\star \in \mathbb{R}^d$ similarly, replacing u_W by u_\star . Let $\Sigma_W \in \mathbb{R}^{d \times d}$ given by

$$(\Sigma_W)_{j,k} \stackrel{\text{def}}{=} \langle (\mathcal{H}_1 u_W)_j, (\mathcal{H}_1 u_W)_k \rangle + \frac{1}{\lambda} \mathbf{1}_{\{j=k\}}, \quad 1 \leq j, k \leq d.$$

We define Σ_\star similarly by replacing u_W by u_\star . Given $u \in L^2(\Omega, \mathbb{R}, \nu)$ we define

$$\mathcal{J}(u) \stackrel{\text{def}}{=} \min_{\theta \in \mathbb{R}^d} \left[|f - \mathcal{H}_0 u - \theta^\top \mathcal{H}_1 u|_2^2 + \frac{1}{\lambda} \|\theta\|_2^2 \right]. \quad (2.2)$$

It is easy to check that $\mathcal{J}(u_W) = |\bar{f}_W|_2^2 - \Phi_W^\top \Sigma_W^{-1} \Phi_W$, and we can integrate out θ from the posterior distribution (1.4) to obtain the marginal distribution of W given by

$$\Pi(W|\mathcal{D}) \propto \frac{\Pi_0(W)}{\sqrt{\det(\Sigma_W)}} \exp \left(-\frac{1}{2\sigma^2} \sum_{i=1}^n (Y_i - u_W(\mathbf{s}_i))^2 - \frac{\lambda}{2} \mathcal{J}(u_W) \right). \quad (2.3)$$

At times, when the distinction is needed we will write $\Pi^{(W)}$ to denote (2.3). We make the following basic assumption on the function class $\mathcal{F} \stackrel{\text{def}}{=} \{u_W, W \in \mathbb{R}^q\}$.

H2. 1. For all $W_0 \in \mathbb{R}^q$, and $r > 0$, we can find $L_{W_0, r} \geq 1$ such that if $W_1, W_2 \in \mathbb{R}^q$ satisfy $\max(\|W_1 - W_0\|_2, \|W_2 - W_0\|_2) \leq r$, then it holds

$$|u_{W_1} - u_{W_2}|_\infty \leq L_{W_0, r} \|W_1 - W_2\|_2.$$

2. For all $W \in \mathbb{R}^q$, u_W has derivatives to the order τ , and there exists $\kappa > 0$ and a constant c_1 such that for all $W_1, W_2 \in \mathbb{R}^q$, such that $|u_{W_1} - u_{W_2}|_\infty \leq \varepsilon$, we have

$$\max_{\mathbf{k}: |\mathbf{k}| \leq \tau} |D^{\mathbf{k}} u_{W_1} - D^{\mathbf{k}} u_{W_2}|_\infty \leq c_1 \varepsilon^\kappa.$$

The local Lipschitz condition imposed in H2-(1) is known to hold for most deep neural net architectures with Lipschitz activation functions (see e.g. Proposition 6 of (34)). Specifically, if u_W is a deep neural network as constructed in Section 1.4 with 1-Lipschitz activation functions, then with $W_1, W_2 \in \mathbb{R}^q$ satisfying $\max(\|W_1 - W_0\|_2, \|W_2 - W_0\|_2) \leq r$, Proposition 6 of (34) and some easy calculations yield

$$|u_{W_1} - u_{W_2}|_\infty \leq L_{W_0, r} \|W_1 - W_2\|_2,$$

with

$$L_{W_0, r} = \left(\sup_{\mathbf{x} \in \Omega} \|\mathbf{x}\|_2 \right) \sqrt{\xi} \left(1 + \frac{(\|W_0\|_2 + r)^2}{\xi} \right)^\xi, \quad (2.4)$$

where ξ is the depth of the DL function.

Since we aim to solve pdes, it is natural to require the function u_W to be smooth. H2-(2) further requires a Holder-type assumption on the differential operator $D^{\mathbf{k}}$ over the function class \mathcal{F} . This assumption is less common in the literature, and we check its validity in Theorem 2.4.

Let $\epsilon : \mathbb{N} \rightarrow \mathbb{R}$ be a nonincreasing function such that for all $s \geq 1$

$$\min \{ |u_W - u_\star|_\infty, W \in \mathbb{R}^q : \|W\|_0 \leq s \} \leq \epsilon(s). \quad (2.5)$$

$\epsilon(s)$ defines the approximation skills of the function class $\{u_W\}$ at the sparsity level s , and is a nonincreasing function of s . Of importance are sparsity levels $s > 1$ at which the approximation error $\epsilon(s)$ matches the statistical error:

$$\epsilon(s) \leq \sigma \sqrt{\frac{us \log(q)}{n}}. \quad (2.6)$$

We make the following assumption.

H3. There exists $\epsilon : \mathbb{N} \rightarrow \mathbb{R}$ such that (2.5) holds. Given n and q , let $s_0 \geq 2$ be the smallest integer satisfying (2.6), and $\epsilon_0 \stackrel{\text{def}}{=} \epsilon(s_0)$. We assume that there exists $W_0 \in \mathbb{R}^q$, with $\|W_0\|_0 \leq s_0$ that satisfies $|u_{W_0} - u_\star|_\infty \leq \epsilon_0$, and

$$\min(\lambda_{\min}(\Sigma_{W_0}), \lambda_{\min}(\Sigma_\star)) \geq C_2, \quad (2.7)$$

for some absolute constant C_2 . Furthermore, we assume that $\lim_{n \rightarrow \infty} \epsilon_0 = 0$, and

$$(\|W_0\|_\infty + 1)^2 + \log(Ls_0^{1/2}/\epsilon_0) \leq C_3 u \log(q), \quad (2.8)$$

for some absolute constant C_3 , where $L \stackrel{\text{def}}{=} L_{W_0, 1}$ is as in H2.

We recall that we aim to estimate u_\star by the sieve method, i.e., we approximate u_\star by a collection of functions $\mathcal{F} \equiv \mathcal{F}_n$, where the *size* of the function class typically grows with the sample size n . The assumption that $\epsilon_0 \rightarrow 0$ as $n \rightarrow \infty$ in H3 highlights the relationship between the sample size n and the model size: given n , we assume that the model \mathcal{F} is chosen such that its approximation error at some sparsity level s can match (or be smaller than) the statistical error at sparsity level s .

There has been a flurry of research activities in recent years to derive precise estimates on $\epsilon(s)$ for various DL function classes, using various norms ((4, 16, 27, 38)). For example for piecewise polynomial activation functions, and u_W as defined in Section 1.4, it is known (see e.g. (16, 39)) that at the sparsity level s , and depth $\xi = O(\log(s))$, \mathcal{F} can approximate a function $u_\star \in C^\beta(\Omega)$ with precision $(1/s)^{\beta/m}$. This yields

$$\epsilon(s) \lesssim (1/s)^{\beta/m}.$$

In that case, solving for s_0 in (2.6) yields

$$s_0 \sim \left(\frac{n}{\log(q)} \right)^{\frac{m}{m+2\beta}}, \quad \text{and} \quad \epsilon_0 \sim \left(\frac{\log(q)}{n} \right)^{\frac{\beta}{m+2\beta}}.$$

In the high-dimensional setting considered here where q is typically much larger than n , (2.8) is typically satisfied.

Ultimately, we need some assumptions on the stability of the pde operator \mathcal{H} as presented below:

H4. *With κ as in H2-(2), there exists an absolute constant C_4 such that for all $\varepsilon > 0$, if $|u_W - u_\star|_2 \leq \varepsilon$, then $|\mathcal{H}u_W - \mathcal{H}u_\star|_2 \leq C_4\varepsilon^\kappa$.*

H4 is similar to H2-(2) and both can be checked in a similar way if u_\star is a classical solution. Indeed, if the pde solution u_\star is a classical solution, and $u_\star \in C^\beta(\Omega)$ for some $\beta > \tau$, then by (2.1), and Theorem 2.4 below, H4 holds with $\kappa = 1 - \tau/\beta$.

Given $s \geq 0$, and constant C , we set $\tau_s \stackrel{\text{def}}{=} \sqrt{(2+u)(1+s)\log(q)}$, and

$$\mathcal{F}_s = \mathcal{F}_{s,C} \stackrel{\text{def}}{=} \left\{ u_W \in \mathcal{F} : \|W\|_0 \leq s, \|W\|_\infty \leq \tau_s, \mathcal{J}(u_W) \leq \frac{Cn\epsilon_0^2}{\lambda\sigma^2} \right\},$$

where \mathcal{J} is as defined in (2.2). We show below that the PINN posterior distribution puts high probabilities on the sets \mathcal{F}_s . Specifically, the following theorem is proved as Lemma A.3 in Section A.2.

Theorem 2.1. *Assume H1-H4. Let ϵ_0, W_0 be as in H3. We can find an absolute constant C such that with $s = C\|W_0\|_0$, it holds,*

$$\mathbb{E} \left[\Pi(\mathcal{F}_{s,C} | \mathcal{D}) \right] \geq 1 - C_0 e^{-\frac{n\epsilon_0^2}{2\sigma^2}},$$

for some absolute constant C_0 .

To highlight this point, we note that if a function $u \in L^2(\Omega, \mathbb{R}, \nu)$ solves the pde equation (1.1) (without the initial/boundary condition) for some parameter θ , say, then

$$\mathcal{J}(u) = \frac{1}{\lambda} \theta^\top \left(\mathbb{1} - \frac{1}{\lambda} \Sigma_u^{-1} \right) \theta \leq \frac{\|\theta\|_2^2}{\lambda}.$$

Hence with the size of θ remaining bounded, and for λ large, functions that approximately solve the pde equation (1.1) satisfies $\mathcal{J}(u) \lesssim C/\lambda$ for some constant C . Theorem 2.1 thus implies that for λ large, the PINN posterior distribution inherits the inductive bias of the prior and restricts the search of a solution for the nonparametric regression problem on approximate solutions of the pde (1.1). The rate of convergence of PINN, therefore, depends mainly on the complexity of the function class \mathcal{F}_s , as measured, for instance, by its covering number. The next assumption models the metric entropy of \mathcal{F}_s .

H5. Given $s \geq 2$, let $3 \leq b < \infty$ be such that

$$\sup_{u \in \mathcal{F}_s} |u - u_\star|_\infty \leq b. \quad (2.9)$$

There exists $V_1 = V_1(s)$ and $V_2 = V_2(s) \geq 6b$ such that for all $0 < \varepsilon \leq \tau_s$,

$$\log \mathcal{N}(\varepsilon, \mathcal{F}_s, \|\cdot\|_\infty) \leq V_1 \log \left(\frac{V_2}{\varepsilon} \right), \quad (2.10)$$

where $\mathcal{N}(\varepsilon, \mathcal{F}_s, \|\cdot\|_\infty)$ denotes the ε -covering number of \mathcal{F}_s in the L_∞ norm.

The following result establishes the rate of convergence of the PINN posterior and is proved in Section A. We assume that the prior parameter $\lambda > 0$ is taken such that

$$\|\theta_\star\|_2^2 \lambda \leq C_5 n \epsilon_0^{2(1-\kappa)} \quad (2.11)$$

for some absolute constant C_5 . We also impose the following mild technical condition: there exists an absolute constant C_6 such that

$$d \left(\log(n \|\Sigma_\star\|_{\text{op}}) + \|\theta_\star\|_\infty^2 \right) \leq \frac{C_6 n \epsilon_0^2}{2\sigma^2}. \quad (2.12)$$

Theorem 2.2. Assume H1-H5, (2.11), and (2.12). Then we can find absolute constants C and $M \geq 1$ such that with

$$s = C s_0, \quad \text{and } r \stackrel{\text{def}}{=} 2(b + \sigma) \sqrt{\frac{V_1 \log(V_2 \sqrt{n})}{n}}, \quad (2.13)$$

where $V_1 = V_1(s)$, $V_2 = V_2(s)$ are as defined in H5, it holds for all n large enough,

$$\mathbb{E} [\Pi(\{W \in \mathbb{R}^q : |u_W - u_\star|_2 > Mr\} \mid \mathcal{D})] \leq c \left(e^{-n \epsilon_0^2 / (2\sigma^2)} + e^{-V_1 \log(V_2 \sqrt{n}) / C} \right),$$

where c is a constant that depends only on c_1, c_2 and σ^2 .

2.1. Minimax nonparametric rate

Bounding the metric entropy of the function class \mathcal{F}_s in a way that leverages the pde information $\mathcal{J}(u_W) \leq Cn\epsilon_0^2/(\lambda\sigma^2)$ has proved challenging. However, we note that $\mathcal{F}_s \subseteq \{u_W \in \mathcal{F} : \|W\|_0 \leq s, \|W\|_\infty \leq \tau_s\}$, and the metric entropy of this latter set is straightforward to bound using H2-(1). See also ((16, 27)). Indeed, for all W, W' in the ball $\{W \in \mathbb{R}^q : \|W\|_0 \leq s, \|W\|_\infty \leq \tau_s\}$, $\max(\|W\|_2, \|W'\|_2) \leq s^{1/2}\tau_s$. Hence by H2-(1), $|u_W - u_{W'}|_\infty \leq Ls^{1/2}\|W - W'\|_\infty$, where $L = L_{\mathbf{0}, s^{1/2}\tau_s}$. The size of all the $(\epsilon' = \epsilon/Ls^{1/2})$ -cover of all the τ_s -balls of all the s -sparse subspaces of \mathbb{R}^q in the infinity-norm is bounded from above by

$$\binom{q}{s} \left(1 + \frac{2r_s}{\epsilon'}\right)^s.$$

It follows that

$$\log \mathcal{N}(\epsilon, \mathcal{F}_s, \|\cdot\|_\infty) \leq s \log(q) + s \log\left(1 + \frac{2s^{1/2}\tau_s L}{\epsilon}\right),$$

and H5 holds with

$$V_1(s) = s, \quad \text{and} \quad V_2(s) = qr_s \left(1 + 2s^{1/2}L\right) \leq c\sqrt{\xi}q \left(1 + s^2 \log(q)\right)^{\xi+1/2},$$

for some constant c that depends only on u , and $\sup_{x \in \Omega} \|\mathbf{x}\|_2$, where ξ denotes the depth of the neural network class. Hence, with $s = Cs_0$,

$$V_1(s) = s = Cs_0, \quad \text{and} \quad \log(\sqrt{n}V_2(s)) \lesssim \log(n) + \xi \log(q),$$

where ξ is the depth of the neural network. Furthermore, we have seen in Remark 2 that if u_\star is a classical solution and $u_\star \in C^\beta(\Omega)$, then for piecewise polynomial activation at depth $\xi = O(\log(p))$, we have

$$s_0 \sim \left(\frac{n}{\log(q)}\right)^{\frac{m}{m+2\beta}}.$$

As a result, up to log terms that we ignore, we obtain the estimation rate

$$r \lesssim \sqrt{\frac{s_0}{n}} \lesssim \left(\frac{1}{n}\right)^{\frac{\beta}{m+2\beta}}.$$

Hence, Bayesian PINN estimates u_\star at least at the nonparametric minimax rate $n^{-\beta/(m+2\beta)}$, that is, at the optimal rate achievable when u_\star is known only to be β -smooth ((32)) without any PDE constraint. We note that our metric entropy calculations of the function class \mathcal{F}_s are rather crude, as they ignore the inductive bias induced by the PINN prior. Thus, the PINN posterior distributions likely enjoy a faster convergence rate than the abovementioned nonparametric minimax rate. To understand this, consider the limiting case when $\lambda = +\infty$, i.e., the prior put all the mass on $\{f \in L^2(\Omega, \mathbb{R}, \nu) : \mathcal{J}(u) = 0\}$, that is, the set of functions that are solutions of the pde equation $\mathcal{H}_0 u + \theta^\top \mathcal{H}_1 u = f$ for some $\theta \in \mathbb{R}^d$. In the subsequent section, we show that estimating u_\star at a parametric rate (up to a log factor) is possible in this limiting case.

2.2. Estimation at the parametric rate

We assume that θ_\star belongs to a compact set $\Theta \subset \mathbb{R}^d$, and for each $\theta \in \Theta$ the pde equation (1.1) (including the initial/boundary condition) admits a unique solution $u_\theta \in C^\beta(\Omega)$ for some given and fixed β . Furthermore we assume that the map $\Gamma : \Theta \rightarrow C^\beta(\Omega)$ that maps $\theta \rightarrow u_\theta$ is Lipschitz. Specifically,

H6. *There exists a constant $L > 0$ such that:*

$$|\Gamma(\theta_1) - \Gamma(\theta_2)|_\infty \leq L \|\theta_1 - \theta_2\|_2 \quad \forall \theta_1, \theta_2 \in \Theta.$$

Let $\Gamma(\Theta) \stackrel{\text{def}}{=} \{\Gamma(\theta), \theta \in \Theta\}$. Under the data generating process assumed in H1, we consider the constrained empirical risk minimization estimator

$$\hat{u} = \arg \min_{u \in \Gamma(\Theta)} \frac{1}{n} \sum_{i=1}^n (Y_i - u(\mathbf{s}_i))^2. \quad (2.14)$$

The following theorem is proved in Section A.7.

Theorem 2.3. *Assume H1, and H6. Then any solution \hat{u} of (2.14) satisfies*

$$|\hat{u} - u_\star|_2^2 = O_p \left(\frac{d}{n} \log \left(\frac{n}{d} \right) \right).$$

It is immediate from the above theorem that $|\hat{u} - u_\star|_2 = O_p(\sqrt{\frac{d}{n}})$ (up to a log factor) which is the standard parametric rate for estimating a d dimensional parameter. The estimator \hat{u} is typically not computable since it requires solving the pde for each $\theta \in \Theta$. But the rate in Theorem 2.3, in light of Theorem 2.1, suggests that the PINN posterior distribution may actually have a convergence rate faster than the nonparametric minimax rate. More research is needed on this issue.

Regarding Assumption 6, one may replace the L_2 norm in the assumption by any L_p norm without hurting the rate, but to translate the complexity of Θ to $\Gamma(\Theta)$, we believe that some sort of smoothness assumption is necessary. Whether one can do entirely without Assumption 6 is an intriguing question that we leave for future research.

2.3. Checking Assumption H2-(2)

We end this section with a result showing that sufficiently smooth functions typically satisfy H2-(2). The proof is given in Section A.6. We recall that Ω is a bounded open subset of \mathbb{R}^m , and we write $\mathbf{B}(x, \alpha)$ to denote the Euclidean ball of \mathbb{R}^m centered at x and with radius α . We define $\text{int}_\alpha(\Omega) \stackrel{\text{def}}{=} \{x \in \Omega : \mathbf{B}(x, \alpha) \subset \Omega\}$. We recall also that $C^\beta(\Omega, b)$ denotes the ball with radius b in the Holder space $C^\beta(\Omega)$.

Theorem 2.4. *Given $\varepsilon > 0$, $\beta > 1$, integer $\tau \in [0, \beta)$, and $M > 0$, define*

$$\bar{\alpha} = \left(\frac{\tau \varepsilon \lfloor \beta \rfloor!}{2M(\beta - \tau)} \right)^{1/\beta}.$$

For all $u, \tilde{u} \in C^\beta(\Omega, M)$ such that $\sup_{x \in \Omega} |u(x) - \tilde{u}(x)|_\infty \leq \varepsilon$, and for all multi-index $\mathbf{k} = (k_1, \dots, k_m)$, with $|\mathbf{k}| \leq \tau$, we have

$$\sup_{x \in \text{int}_{\bar{\alpha}}(\Omega)} |D^{\mathbf{k}}u(x) - D^{\mathbf{k}}\tilde{u}(x)|_\infty \leq CM^{\frac{\tau}{\beta}} \varepsilon^{\frac{\beta-\tau}{\beta}},$$

for some constant C that depends only on τ and β .

We note that $\bar{\alpha} \downarrow 0$ as $M \uparrow \infty$. Hence if $u, \tilde{u} \in C^\beta(\bar{\Omega}, M_0)$ for some well-chosen $\bar{\Omega} \supset \Omega$, and $\sup_{x \in \bar{\Omega}} |u(x) - \tilde{u}(x)|_\infty \leq \varepsilon$ then by taking $M \geq M_0$ large enough such that $\Omega \subset \text{int}_{\bar{\alpha}}(\bar{\Omega})$, we get

$$\sup_{x \in \Omega} |D^{\mathbf{k}}u(x) - D^{\mathbf{k}}\tilde{u}(x)|_\infty \leq CM^{\frac{\tau}{\beta}} \varepsilon^{\frac{\beta-\tau}{\beta}}.$$

3. A Bernstein-von-Mises theorem for the θ -marginal

Given $\mu \in \mathbb{R}^d$, and $\Lambda \in \mathbb{R}^{d \times d}$ symmetric and positive definite, we write $\mathbf{N}(\mu, \Lambda)(\cdot)$ as the probability measure of the Gaussian distribution $\mathbf{N}(\mu, \Lambda)$. Given $W \in \mathbb{R}^q$, let

$$\hat{\theta}_W \stackrel{\text{def}}{=} \Sigma_W^{-1} \Phi_W. \quad (3.1)$$

It is then easy to see from (1.4) that the conditional distribution of θ given W is $\mathbf{N}(\hat{\theta}_W, \Sigma_W^{-1}/\lambda)$. For clarity's sake we will write $\Pi^{(\theta)}$ (resp. $\Pi^{(W)}$) to denote the marginal distribution of θ (resp. W) under (1.4). We have

$$\Pi^{(\theta)}(\cdot | \mathcal{D}) = \int_{\mathbb{R}^q} \left[\mathbf{N}\left(\hat{\theta}_W, \frac{1}{\lambda} \Sigma_W^{-1}\right)(\cdot) \right] \Pi^{(W)}(dW | \mathcal{D}).$$

We set

$$\Pi_\star^{(\theta)}(\cdot) \stackrel{\text{def}}{=} \mathbf{N}\left(\theta_\star, \frac{1}{\lambda} \Sigma_\star^{-1}\right)(\cdot). \quad (3.2)$$

We investigate the proximity between $\Pi^{(\theta)}$ and $\Pi_\star^{(\theta)}$ for n large, using the 2-Wasserstein metric that we denote W_2 . The following result is established in Section A.5.

Theorem 3.1. *Assume H1-H5 and the notations of Theorem 2.2. Then, for all n large enough, we have*

$$\mathbb{E} \left[W_2^2(\Pi^{(\theta)}, \Pi_\star^{(\theta)}) \right] \leq c \left(r^{2\kappa} + \lambda e^{-n\epsilon_0^2/(2\sigma^2)} + \lambda e^{-V_1 \log(V_2 \sqrt{n})/C} \right),$$

for some constant c . In particular,

$$\mathbb{E} \left[\left| \int_{\mathbb{R}^d} u \Pi^{(\theta)}(du | \mathcal{D}) - \theta_\star \right|^2 \right] \leq c \left(r^{2\kappa} + \lambda e^{-n\epsilon_0^2/(2\sigma^2)} + \lambda e^{-V_1 \log(V_2 \sqrt{n})/C} \right).$$

For DL functions with sufficiently smooth piecewise polynomial activation functions, and for $u_\star \in C^\beta(\Omega)$, we have seen above that $r \lesssim n^{-\beta/(m+2\beta)}$, and by Theorem 2.4, $\kappa = (\beta - \tau)/\beta$. With these we conclude from Theorem 3.1 that the estimation rate of θ_\star afforded by B-PINN is at least $n^{-2(\beta-\tau)/(m+2\beta)}$. This estimation rate corresponds to the optimal plug-in rate. Recall that our pde

equation is linear in θ . Therefore, the bottleneck in estimating θ is the difficulty of estimating u and its derivatives up to order τ . As shown above, for a function $u \in C^\beta(\Omega)$, PINN estimates these derivatives at the nonparametric minimax rate $n^{-2(\beta-\tau)/(m+2\beta)}$. Hence, the obtained estimation rate for θ_\star . As conjectured above, PINN likely estimates u_\star at a rate faster than the nonparametric minimax rate. If true, this would translate to a faster estimation rate for θ_\star . More research is needed on this point.

4. Numerical illustration

We illustrate the results above using the one-dimensional heat equation of Example 1. Extensive illustrations of PINN and Bayesian PINN can be found in the literature ((8, 24, 37)). We focus here on comparing PINN and an approach that does not directly exploit the PDE structure of the problem. We consider the heat equation with $L = \pi$, and $T = 1$. Hence $\Omega = (0, 1) \times (0, \pi)$. We use the boundary condition $u(t, 0) = u(t, \pi) = 0$, and the initial condition $u(0, x) = \sin(x)$. Given θ the heat equation has a unique solution $u_\theta(t, x) = \sin(x)\exp(-\theta t)$. We set the true value of θ to $\theta_\star = 0.5$. The observed data \mathcal{D} is generated as follows: n sensor locations $\{\mathbf{s}_i = (t_i, x_i), 1 \leq i \leq n\}$ are evenly distributed in Ω , yielding measurements

$$Y_i = u_{\theta_\star}(\mathbf{s}_i) + \sigma \epsilon_i, \text{ where } \epsilon_i \sim \mathbf{N}(0, 1),$$

for some noise parameter σ that we control. For the function class $\mathcal{F} = \{u_W\}$, $u_W : \mathbb{R}^2 \rightarrow \mathbb{R}$ is taken to be a fully connected neural network with depth 4, width 64, and \tanh activation function. The resulting posterior distribution is

$$\Pi(d\theta, dW|\mathcal{D}) \propto \exp\left(-\frac{1}{2\sigma^2} \sum_{i=1}^n (Y_i - u_W(\mathbf{s}_i))^2 - \frac{\lambda}{2} \ell_0(\theta, W) - \frac{1}{2} \|\theta\|_2^2\right) \Pi_0(dW) d\theta,$$

where

$$\ell_0(\theta, W) = |(u_W)_t - \theta(u_W)_{xx}|_2^2 + |u_W(\cdot, 0)|_2^2 + |u_W(\cdot, \pi)|_2^2 + |u_W(0, \cdot) - \sin(\cdot)|_2^2, \quad (4.1)$$

where $|\cdot|_2$ denotes the function space L^2 -norm. Observe that $\lambda = 0$ in the above posterior corresponds to the estimation of W without the physics-informed prior (non-PINN estimation). Before presenting the numerical results, we describe briefly our MCMC sampling method.

4.1. Approximation and MCMC sampling

In general, the L^2 -norm in the loss (4.1) is intractable and is typically replaced by a Monte Carlo or a numerical quadrature approximation. Here we use Monte Carlo by drawing $N = 10,000$ interior points $(\mathbf{s}_1^{(i)}, \dots, \mathbf{s}_N^{(i)})$ uniformly in Ω to estimate the L^2 norm on Ω . For the L^2 norm along the boundaries we draw two sets of $B = 128$ points $(\mathbf{s}_1^{(b_i)}, \dots, \mathbf{s}_B^{(b_i)})$ for $i = 1, 2$ uniformly on the time and space boundaries respectively. Hence, we consider the approximate posterior distribution

$$\widehat{\Pi}(\theta, W|\mathcal{D}) \propto \Pi_0(W) \exp\left(-\frac{1}{2\sigma^2} \sum_{i=1}^n (y_i - u_W(\mathbf{s}_i))^2 - \frac{\lambda}{2} \widehat{\ell}_0(\theta, W) - \frac{\rho}{2} \|\theta\|_2^2\right). \quad (4.2)$$

where the loss (4.1) is replaced by

$$\begin{aligned} \widehat{\ell}_0(\theta, W) &= \frac{1}{N} \sum_{k=1}^N \left((u_W)_t(\mathbf{s}_k^{(i)}) - \theta \times (u_W)_{xx}(\mathbf{s}_k^{(i)}) \right)^2 \\ &+ \frac{1}{B} \sum_{k=1}^B \left(u_W(\mathbf{s}_k^{(b_1)}, 0) \right)^2 + \frac{1}{B} \sum_{k=1}^B \left(u_W(\mathbf{s}_k^{(b_1)}, \pi) \right)^2 + \frac{1}{B} \sum_{k=1}^B \left(u_W(0, \mathbf{s}_k^{(b_2)}) - \sin(\mathbf{s}_k^{(b_2)}) \right)^2. \end{aligned} \quad (4.3)$$

Given W , and as seen in Section 3, the posterior conditional distribution of θ given W is

$$\theta|W, \mathcal{D} \sim \mathbf{N} \left(\Sigma_W^{-1} \Phi_W, \frac{1}{\lambda} \Sigma_W^{-1} \right) \quad (4.4)$$

where in the particular case of this example

$$\Sigma_W = \frac{1}{N} \sum_{k=1}^N (u_W)_{xx}(\mathbf{s}_k^{(i)})^2 + \frac{\rho}{\lambda}, \quad \text{and} \quad \Phi_W = \frac{1}{N} \sum_{k=1}^N (u_W)_t(\mathbf{s}_k^{(i)}) \times (u_W)_{xx}(\mathbf{s}_k^{(i)})$$

We sample from (4.2) using the approximate asynchronous sampler of (3). The algorithm is a data-augmentation Metropolis-with-Gibbs sampler where the update of the sparsity support Λ given θ, W uses asynchronous sampling, and the update of W given θ, Λ is a sparse stochastic gradient Langevin dynamics. Then θ given W, Λ is sample from Gaussian $\mathbf{N} \left(\Sigma_{W_\Lambda}^{-1} \Phi_{W_\Lambda}, \frac{1}{\lambda} \Sigma_{W_\Lambda}^{-1} \right)$ using (4.4) where W_Λ is the component-wise product of W and Λ (a sparse neural network weight). A PyTorch implementation is available from the GitHub page <https://github.com/xliu-522/SA-cSGLD>.

Throughout the experiment, we use $\theta^* = 0.5$, $\lambda = n$ and $\rho = 1$. For each instance of the posterior distribution (4.2) under consideration, we run the aforementioned MCMC sampler until convergence and keep running another 200,000 iterations. We then record every 20th sample, resulting in a total of $K = 10,000$ samples, denoted by $\{\hat{\theta}^{(k)}\}_{k=1}^K$. The mean and standard deviations of the marginal posterior distribution of θ are then approximated respectively by

$$\mu_{\hat{\theta}} \stackrel{\text{def}}{=} \frac{1}{K} \sum_{k=1}^K \hat{\theta}^{(k)}, \quad \text{and} \quad \sigma_{\hat{\theta}} \stackrel{\text{def}}{=} \sqrt{\frac{1}{K} \sum_{k=1}^K (\hat{\theta}^{(k)} - \mu_{\hat{\theta}})^2}.$$

In this section, assuming that the MCMC sampler has converged, we take the distribution of the samples $\{\hat{\theta}^{(k)}\}_{k=1}^K$ to be $\Pi^{(\theta)}$ and we denote its normal approximation by $\tilde{\Pi}^{(\theta)} \stackrel{\text{def}}{=} \mathbf{N}(\mu_{\hat{\theta}}, \sigma_{\hat{\theta}}^2)$. Since the exact solution $u_\star(t, x) = \sin(x) \exp(-0.5t)$ is known, we can compute the variance of the limiting Gaussian distribution in Bernstein-von Mises theorem (Theorem 3.1)

$$\Sigma_* = \int_0^1 \int_0^\pi ((u_\star)_{xx}(t, x))^2 dx dt + \frac{1}{n} = \frac{\pi(1 - \frac{1}{e})}{2} + \frac{1}{n} \approx 0.993 + \frac{1}{n}$$

Hence the limiting distribution in Theorem 3.1 denoted $\Pi_\star^{(\theta)}$ is

$$\Pi_\star^{(\theta)} \stackrel{\text{def}}{=} \mathbf{N}(\theta_\star, \frac{1}{n} \Sigma_\star^{-1}). \quad (4.5)$$

4.2. PINN versus non-PINN comparison

In this section, we compare PINN to a two-step approach (denoted by non-PINN) that does not directly use the PDE structure of u . Specifically, in the non-PINN approach, we fit the same DL model for

estimating u but without the PINN prior. The resulting posterior is given by

$$\widehat{\Pi}(W|\mathcal{D}) \propto \Pi_0(W) \exp\left(-\frac{1}{2\sigma^2} \sum_{i=1}^n (y_i - u_W(\mathbf{s}_i))^2\right).$$

We sample from this posterior distribution using the same MCMC algorithm described above. For each MCMC draw W , we then subsequently solve for θ by sampling from the Gaussian distribution $\mathbf{N}\left(\Sigma_W^{-1}\Phi_W, \frac{1}{n}\Sigma_W^{-1}\right)$ from (4.4). Hence, in the two-step approach, the PDE information is not used in the first step for estimating u , but is used in the second step to recover θ by the linear regression model derived from (1.1). By contrast, PINN uses the PDE information and estimates (u, θ) jointly. We compare the two approaches under varying sample sizes and noise levels, as it allows us to observe the contribution of the PINN prior $\ell_0(\theta, W)$.

Figure 1 provides a comparative analysis of the posterior distribution of θ in the two approaches (PINN and non-PINN) at a noise level of 10% ($\sigma \approx 0.025$), across different sample sizes. Both methods exhibit improved parameter estimates as the sample size n increases, as indicated by the decreasing bias and interquartile range (IQR). However, training with the PDE term generally yields more concentrated distributions of θ with narrower IQRs, indicating greater stability and reliability in parameter estimation. This demonstrates that incorporating the PDE term results in more accurate and robust parameter estimates.

To further investigate the role of the PDE term in the recovery of θ , the root mean square error (rMSE) and the Wasserstein-2 distance (W_2) are computed and summarized in Table 1 and Figure 2. The rMSE is defined as $\sqrt{\frac{1}{K} \sum_{k=1}^K (\hat{\theta}^{(k)} - \theta^*)^2}$, whereas the W_2 metric is computed using the `ot.emd2_1d` function from the POT library by (11) with the parameter $p = 2$. These metrics provide a comprehensive comparison of the accuracy (rMSE) and distribution similarity (W_2) between sample θ and the true value θ^* .

Table 1. Comparison of PDE and No-PDE under different noise levels

		rMSE		$W_2(\Pi^{(\theta)}, \Pi_*^{(\theta)})$	
Noise Level	n	PDE	No-PDE	PDE	No-PDE
1%	50	0.287	0.352	0.0379	0.0699
1%	500	0.083	0.088	0.0018	0.0024
1%	1000	0.058	0.060	0.0008	0.0010
1%	5000	0.027	0.028	0.0002	0.0003
10%	500	0.103	0.140	0.0052	0.0131
10%	1000	0.080	0.116	0.0036	0.0078
10%	5000	0.031	0.040	0.0004	0.0008
10%	10000	0.023	0.031	0.0002	0.0007
25%	1000	0.094	0.377	0.0055	0.1387
25%	5000	0.042	0.104	0.0011	0.0090
50%	5000	0.129	0.474	0.0147	0.2246

Figure 2 and Table 1 illustrate the performance comparison of models trained with and without PDE constraints across various noise levels and sample sizes. As the number of samples (n) increases or the

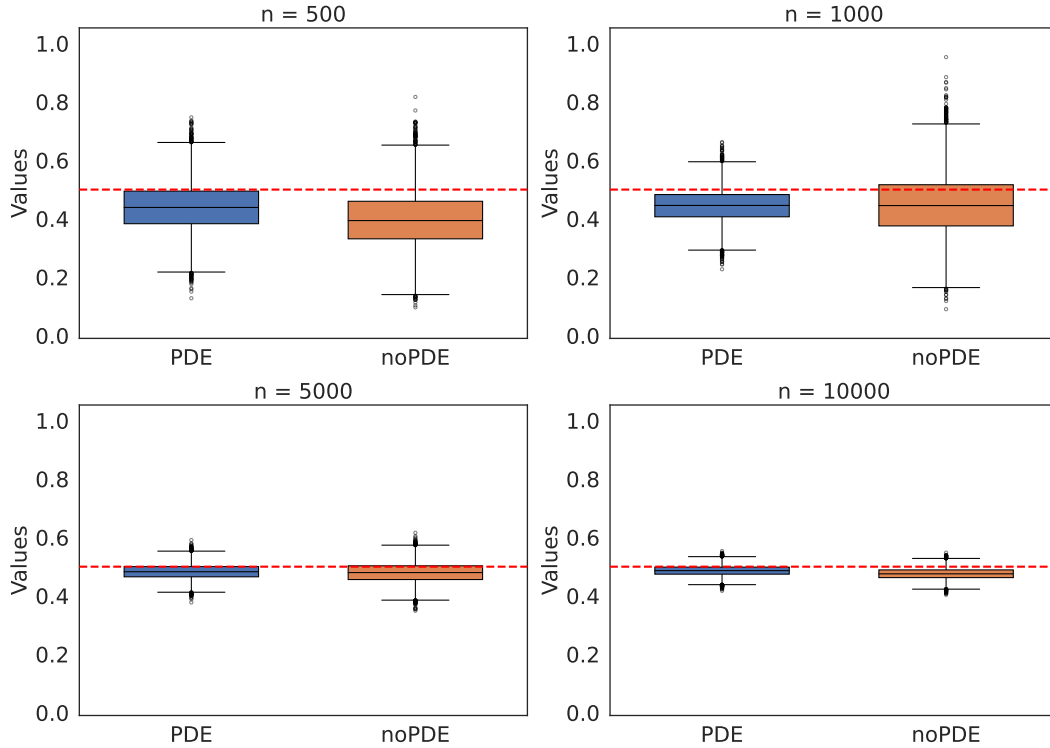


Figure 1. Comparison between boxplots of sampled θ from training with/without pde with different sample sizes n at noise level = 10%. True $\theta^* = 0.5$ in the dotted line.

noise level decreases, the values of rMSE and W_2 distance decrease for both training methods, with and without the PDE constraint. However, the PDE-constrained models perform better, demonstrating lower rMSE and W_2 distances than corresponding models without PDE constraints. The difference in performance becomes more noticeable at low sample sizes or high noise levels. This suggests that the PDE constraint can help the model learn more effectively when there is insufficient information from the data alone.

4.3. Posterior contraction Behavior

From the previous section and Table 1, we observed that the W_2 distance between $\Pi^{(\theta)}$ and the limiting distribution $\Pi_{\star}^{(\theta)}$ given in (4.5) decays significantly as the sample size n increases, which is consistent with the conclusion of Theorem 3.1. To further study this contraction behavior, we plot the histogram of $\Pi^{(\theta)}$ (using samples from the MCMC sampler), its Gaussian approximation $\tilde{\Pi}^{(\theta)} = \mathbf{N}(\mu_{\hat{\theta}}, \sigma_{\hat{\theta}}^2)$ where $\mu_{\hat{\theta}}$ and $\sigma_{\hat{\theta}}$ are computed from the MCMC samples, and the limiting distribution $\Pi_{\star}^{(\theta)}$, for various sample sizes ($n = 500, 1000, 5000, 10000$) at a noise level of 10% ($\sigma \approx 0.025$). The perfect match between $\Pi^{(\theta)}$ and $\tilde{\Pi}^{(\theta)}$ in Figure 3 suggests that the posterior distribution of θ is approximately Gaussian. However the persistent discrepancy between $\Pi^{(\theta)}$ and $\Pi_{\star}^{(\theta)}$ suggests that the W_2 convergence




Figure 2. Comparison between training with/without pde term. The number in each grid is the rMSE between sampled θ and θ^* . The color represents W_2 distance $W_2(\Pi^{(\theta)}, \Pi_\star^{(\theta)})$

to zero between $\Pi^{(\theta)}$ and $\Pi_\star^{(\theta)}$ (as established in Theorem 3.1 and illustrated above), likely do not hold in total variation. For two probability measures P, Q with densities p, q respectively, their total variation distance is $\text{TVD}(P, Q) \stackrel{\text{def}}{=} \frac{1}{2} \int |p(x) - q(x)| dx$. To further illustrate this point we compute in Table 2 a lower and an upper on the total variation distance between $\tilde{\Pi}^{(\theta)}$ and $\Pi_\star^{(\theta)}$. We compute these bounds by noting (see (10)) that for any two probability measures P, Q with densities p, q respectively we have

$$H^2(P, Q) \leq \text{TVD}(P, Q) \leq \min\left(1, \sqrt{KL(P\|Q)/2}, \sqrt{KL(Q\|P)/2}\right), \quad (4.6)$$

where in the univariate Gaussian case where $P \sim \mathcal{N}(\mu_1, \sigma_1^2)$ with density $p(x)$ and $Q \sim \mathcal{N}(\mu_2, \sigma_2^2)$ with density $q(x)$, the Hellinger distance $H(P, Q)$ is (see e.g. (23))

$$H^2(P, Q) \stackrel{\text{def}}{=} \frac{1}{2} \int \left(\sqrt{p(x)} - \sqrt{q(x)}\right)^2 dx = 1 - \sqrt{\frac{2\sigma_1\sigma_2}{\sigma_1^2 + \sigma_2^2}} \exp\left(-\frac{(\mu_1 - \mu_2)^2}{4(\sigma_1^2 + \sigma_2^2)}\right),$$

and the KL -divergence is (see e.g. (26))

$$D_{KL}(P\|Q) \stackrel{\text{def}}{=} \int p(x) \log \frac{p(x)}{q(x)} dx = \log\left(\frac{\sigma_2}{\sigma_1}\right) + \frac{\sigma_1^2 + (\mu_1 - \mu_2)^2}{2\sigma_2^2} - \frac{1}{2}.$$

Using these formulas, the third and fourth column of Table 2 shows the left-hand side and right-hand side of (4.6) respectively. From this table, we see, for instance, that at 10% noise level, the W_2 distance decreases with the sample size, whereas the TVD lower and upper bounds do not. This lack of TV convergence is because the posterior mean does not converge to θ_\star fast enough. Indeed, given two univariate normal distributions $\mathcal{N}(\mu_{1,n}, \frac{1}{n})$ and $\mathcal{N}(\mu_{2,n}, \frac{1}{n})$, as $n \rightarrow \infty$, their W_2 distance converges to

0 as soon as $|\mu_{1,n} - \mu_{2,n}| \rightarrow 0$, whereas their Hellinger distance (using the formula above) converges to zero if and only if $|\mu_{1,n} - \mu_{2,n}|$ converges to zero faster than $1/\sqrt{n}$.

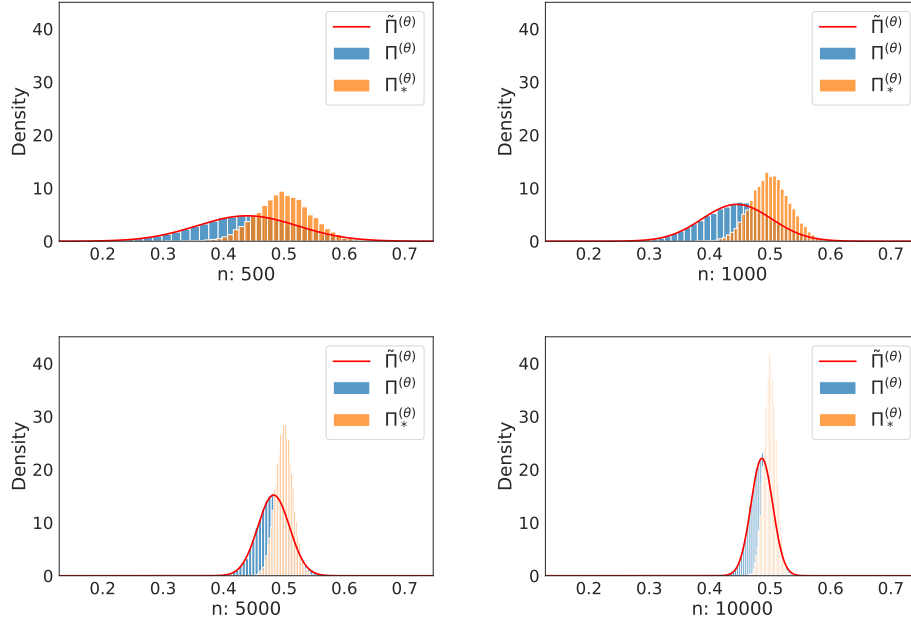


Figure 3. Histograms of sampled θ , denoted $\Pi^{(\theta)}$, in blue with its normal approximation $\tilde{\Pi}^{(\theta)}$ in red line and target Gaussian distribution $\Pi_*^{(\theta)}$ in orange with increasing sample size n at noise level 10%

Table 2. W_2 and upper and lower bound of TVD under different sample sizes and noise levels

Noise Level	n	W_2	TVD Lower	TVD Upper
1%	50	0.0377	0.149	0.482
1%	500	0.0018	0.090	0.372
1%	1000	0.0008	0.084	0.360
1%	5000	0.0002	0.099	0.387
10%	500	0.0051	0.176	0.515
10%	1000	0.0037	0.228	0.594
10%	5000	0.0004	0.156	0.485
10%	10000	0.0002	0.173	0.512
25%	1000	0.0054	0.270	0.637
25%	5000	0.0011	0.287	0.665
50%	5000	0.0148	0.616	1.000

Appendix A: Proof of Theorem 2.2

We define

$$\Delta_W(y, \mathbf{s}) \stackrel{\text{def}}{=} (y - u_W(\mathbf{s}))^2, \quad \Delta_\star(y, \mathbf{s}) \stackrel{\text{def}}{=} (y - u_\star(\mathbf{s}))^2, \quad (y, \mathbf{s}) \in \mathbb{R} \times \Omega.$$

We recall that $\mathbb{P}_n(f) = n^{-1} \sum_{i=1}^n f(Y_i, \mathbf{s}_i)$. It follows from (2.3) that for any set $C \subseteq \mathbb{R}^q$, we have

$$\Pi(C|\mathcal{D}) = \frac{\int_C \exp\left(-\frac{n}{2\sigma^2} \mathbb{P}_n(\Delta_W - \Delta_\star) + \mathcal{R}_W\right) \Pi_0(dW)}{\int_{\mathbb{R}^q} \exp\left(-\frac{n}{2\sigma^2} \mathbb{P}_n(\Delta_W - \Delta_\star) + \mathcal{R}_W\right) \Pi_0(dW)}, \quad (\text{A.1})$$

where

$$\mathcal{R}_W \stackrel{\text{def}}{=} \mathcal{R}_W^{(1)} + \mathcal{R}_W^{(2)}, \quad \text{with } \mathcal{R}_W^{(1)} \stackrel{\text{def}}{=} -\frac{1}{2} (\log \det(\Sigma_W) - \log \det(\Sigma_\star)),$$

$$\text{and } \mathcal{R}_W^{(2)} \stackrel{\text{def}}{=} -\frac{\lambda}{2} \mathcal{J}(u_W).$$

We break the proof into three parts. First, in Section A.1, we give a lower bound on the normalizing constant of the posterior distribution as given in (A.1). Then we show in Section A.2 that the prior distribution Π_0 has a good inductive bias, as it puts a high probability on W that is sparse. The third part of the proof in Section A.3 establishes some deviation bounds for the empirical process of the log-likelihood ratio. Then, we put all the pieces together in Section A.4.

A.1. Lower bound on the normalizing constant

Lemma A.1. *Assume H1-H4. Let ϵ_0, s_0, W_0 be as in H3. Then for all n large enough, with probability at least $1 - 2e^{-n\epsilon_0^2/(2\sigma^2)}$, we have*

$$\int_{\mathbb{R}^q} \exp\left(-\frac{n}{2\sigma^2} \mathbb{P}_n(\Delta_W - \Delta_\star) + \mathcal{R}_W\right) \Pi_0(dW) \geq C \times \exp\left(-\frac{Cn\epsilon_0^2}{\sigma^2}\right),$$

for some absolute constant C .

Proof. By definition of W_0 , $|u_{W_0} - u_\star|_\infty \leq \epsilon_0$. Let $L \stackrel{\text{def}}{=} L_{W_0,1} \geq 1$ be as in H2. Let $\Lambda_0 \in \{0, 1\}^q$ denote the sparsity support of W_0 , and let

$$\eta \stackrel{\text{def}}{=} \frac{\epsilon_0}{Ls_0^{1/2}}, \quad \text{and } \mathcal{V} \stackrel{\text{def}}{=} \{W \in \mathbb{R}^q : \text{supp}(W) = \Lambda_0, \|W - W_0\|_\infty \leq \eta\}. \quad (\text{A.2})$$

Since $\eta s_0^{1/2} \leq 1$, for all $W \in \mathcal{V}$, we have

$$\|W - W_0\|_2 \leq s_0^{1/2} \|W - W_0\|_\infty \leq s_0^{1/2} \eta \leq 1. \quad (\text{A.3})$$

Using (A.3), and the Lipschitz assumption imposed in H2, we deduce that for all $W \in \mathcal{V}$,

$$|u_W - u_{W_0}|_\infty \leq L \|W - W_0\|_2 \leq Ls_0^{1/2} \eta \leq \epsilon_0. \quad (\text{A.4})$$

From the above display, and appealing to (2.1) and H2-(3) we further deduce that for all $W \in \mathcal{V}$,

$$|\mathcal{H}u_W - \mathcal{H}u_{W_0}|_2 \leq C_0 \max_{\mathbf{k}: |\mathbf{k}| \leq \tau} |D^{\mathbf{k}}u_W - D^{\mathbf{k}}u_{W_0}|_\infty \leq C_0 c_1 \epsilon_0^K. \quad (\text{A.5})$$

In view of H4, we also deduce that for all $W \in \mathcal{V}$,

$$|u_W - u_\star|_\infty \leq 2\epsilon_0, \quad \text{and} \quad |\mathcal{H}u_W - \mathcal{H}u_\star|_2 \leq 2^K C_4 \epsilon_0^K. \quad (\text{A.6})$$

A.1.0.1. Step 1. First we show that with probability at least $1 - 2e^{-n\epsilon_0^2/(2\sigma^2)}$, it holds

$$\int_{\mathbb{R}^q} \exp\left(-\frac{n}{2\sigma^2} \mathbb{P}_n(\Delta_W - \Delta_\star) + \mathcal{R}_W\right) \Pi_0(dW) \geq \frac{1}{2} e^{-4n\epsilon_0^2/\sigma^2} \int_{\mathcal{V}} e^{\mathcal{R}_W} \Pi_0(dW). \quad (\text{A.7})$$

To establish this, we recall that with $\xi_i = (Y_i - u_\star(\mathbf{s}_i))/\sigma$, we have

$$-\frac{n}{2\sigma^2} \mathbb{P}_n(\Delta_W - \Delta_\star) = -\frac{1}{2\sigma^2} \sum_{i=1}^n (u_W(\mathbf{s}_i) - u_\star(\mathbf{s}_i))^2 + \frac{1}{\sigma} \sum_{i=1}^n \xi_i (u_W(\mathbf{s}_i) - u_\star(\mathbf{s}_i)).$$

Using (A.6), we note that for $W \in \mathcal{V}$,

$$\frac{1}{\sigma^2} \sum_{i=1}^n (u_W(\mathbf{s}_i) - u_\star(\mathbf{s}_i))^2 \leq A,$$

where

$$A \stackrel{\text{def}}{=} \frac{4n\epsilon_0^2}{\sigma^2}.$$

By Gaussian tail bounds, for all $W \in \mathcal{V}$,

$$\mathbb{P}\left(\frac{1}{\sigma} \sum_{i=1}^n \xi_i (u_W(\mathbf{s}_i) - u_\star(\mathbf{s}_i)) > \frac{A}{2} \mid \mathbf{s}_{1:n}\right) \leq \exp\left(-\frac{A^2}{8A}\right) \leq e^{-A/8}.$$

Therefore, with $\mathcal{E}_W \stackrel{\text{def}}{=} \{\mathcal{D} : \frac{1}{\sigma} \sum_{i=1}^n \xi_i (u_W(\mathbf{s}_i) - u_\star(\mathbf{s}_i)) \leq A/2\}$,

$$\int_{\mathbb{R}^q} \exp\left(-\frac{n}{2\sigma^2} \mathbb{P}_n(\Delta_W - \Delta_\star) + \mathcal{R}_W\right) \Pi_0(dW) \geq e^{-A} \int_{\mathcal{V}} \mathbf{1}_{\mathcal{E}_W}(\mathcal{D}) e^{\mathcal{R}_W} \Pi_0(dW),$$

and with \mathcal{E}_W^c denoting the complement of \mathcal{E}_W , we obtain

$$\begin{aligned} & \mathbb{P}\left[\int_{\mathbb{R}^q} \exp\left(-\frac{n}{2\sigma^2} \mathbb{P}_n(\Delta_W - \Delta_\star) + \mathcal{R}_W\right) \Pi_0(dW) < \frac{e^{-A}}{2} \int_{\mathcal{V}} e^{\mathcal{R}_W} \Pi_0(dW) \mid \mathbf{s}_{1:n}\right] \\ & \leq \mathbb{P}\left[\int_{\mathcal{V}} \mathbf{1}_{\mathcal{E}_W}(\mathcal{D}) e^{\mathcal{R}_W} \Pi_0(dW) < \frac{1}{2} \int_{\mathcal{V}} e^{\mathcal{R}_W} \Pi_0(dW) \mid \mathbf{x}_{1:n}\right] \\ & = \mathbb{P}\left[\int_{\mathcal{V}} \mathbf{1}_{\mathcal{E}_W^c}(\mathcal{D}) e^{\mathcal{R}_W} \Pi_0(dW) > \frac{1}{2} \int_{\mathcal{V}} e^{\mathcal{R}_W} \Pi_0(dW) \mid \mathbf{x}_{1:n}\right] \end{aligned}$$

$$\leq \frac{2}{\int_{\mathcal{V}} e^{\mathcal{R}_W} \Pi_0(dW)} \int_{\mathcal{V}} \mathbb{P} \left(\frac{1}{\sigma} \sum_{i=1}^n \xi_i(u_W(\mathbf{s}_i) - u_{\star}(\mathbf{s}_i)) > \frac{A}{2} \mid \mathbf{s}_{1:n} \right) e^{\mathcal{R}_W} \Pi_0(dW) \leq 2e^{-A/8},$$

which is (A.7).

A.1.0.2. Step 2. We now show that

$$\int_{\mathcal{V}} e^{\mathcal{R}_W} \Pi_0(dW) \geq c \times \exp \left(-C \frac{n\epsilon_0^2}{\sigma^2} \right). \quad (\text{A.8})$$

Using the definition of Σ_W , we first note that for all W_1, W_2

$$\begin{aligned} \|\Sigma_{W_1} - \Sigma_{W_2}\|_{\text{op}} &\leq \sqrt{2} (|\mathcal{H}u_{W_1}|_2 + |\mathcal{H}u_{W_2}|_2) |\mathcal{H}u_{W_1} - \mathcal{H}u_{W_2}|_2 \\ &\leq \sqrt{2} (2|\mathcal{H}u_{W_1}|_2 + |\mathcal{H}u_{W_2} - \mathcal{H}u_{W_1}|_2) |\mathcal{H}u_{W_1} - \mathcal{H}u_{W_2}|_2. \end{aligned} \quad (\text{A.9})$$

We combine this with (A.5) to obtain that for all $W \in \mathcal{V}$,

$$\|\Sigma_W - \Sigma_{W_0}\|_{\text{op}} \leq 2C_0c_1 (2|\mathcal{H}u_{W_0}|_2 + C_0c_1\epsilon_0^\kappa) \epsilon_0^\kappa \leq \frac{C_2}{2},$$

for all n large enough, since $\epsilon_0 \rightarrow 0$, as $n \rightarrow \infty$. Therefore, Weyl's inequality, and $\lambda_{\min}(\Sigma_{W_0}) \geq C_2$ imply that for all $W \in \mathcal{V}$, and n large enough

$$\lambda_{\min}(\Sigma_W) \geq \frac{C_2}{2}.$$

As a result of the last display, we can use a first order Taylor expansion of the log det to conclude that for all $W \in \mathcal{V}$,

$$\begin{aligned} |\mathcal{R}_W^{(1)}| &= \left| \frac{1}{2} (\log \det(\Sigma_W) - \log \det(\Sigma_{\star})) \right| \leq \frac{2d^{1/2}}{C_2} \|\Sigma_W - \Sigma_{\star}\|_{\text{F}} \leq \frac{2d}{C_2} \|\Sigma_W - \Sigma_{\star}\|_{\text{op}} \\ &\leq \frac{4d}{C_2} (2|\mathcal{H}u_{\star}|_2 + |\mathcal{H}u_W - \mathcal{H}u_{\star}|_2) |\mathcal{H}u_W - \mathcal{H}u_{\star}|_2 \\ &\leq 4 \frac{2^\kappa C_4 d}{C_2} (2|\mathcal{H}u_{\star}|_2 + 2^\kappa C_4 \epsilon_0^\kappa) \epsilon_0^\kappa \leq 1, \end{aligned}$$

again, for all n large enough. For $W \in \mathcal{V}$, let J_W be the L^2 projector on the linear space spanned by the function $\{(\mathcal{H}u_W)_i, 1 \leq i \leq d\}$ in $L^2(\Omega, \mathbb{R}, \nu)$. Note that, since $\lambda_{\min}(\Sigma_W) \geq C_2/2$, that sub-space is isomorphic to \mathbb{R}^d , and by expressing the calculation in \mathbb{R}^d , it easily follows that

$$\left| \Phi_W^\top \Sigma_W^{-1} \Phi_W - |J_W \bar{f}_W|_2^2 \right| \leq \frac{2}{\lambda C_2},$$

where we recall that $\bar{f}_W = f - \mathcal{H}_0 u_W$. As a result, for $W \in \mathcal{V}$,

$$\left| \Phi_W^\top \Sigma_W^{-1} \Phi_W - |\bar{f}_W|_2^2 \right| \leq \frac{2}{\lambda C_2} + \left| |J_W \bar{f}_W|_2^2 - |\bar{f}_W|_2^2 \right| = \frac{2}{\lambda C_2} + |J_W \bar{f}_W - \bar{f}_W|_2^2,$$

where the equality uses the fact that for all $u \in L^2(\Omega, \mathbb{R}, \nu)$, $|u|_2^2 = |J_W u|_2^2 + |(J_W - \mathbb{1})u|_2^2$. By the definition of the projector as closest element, and since $\tilde{f}_W = f - \mathcal{H}_0 u_W = (\mathcal{H}_0 u_\star - \mathcal{H}_0 u_W) + \theta_\star^\top \mathcal{H} u_\star$,

$$\begin{aligned} |J_W \tilde{f}_W - \tilde{f}_W|_2^2 &\leq |\theta_\star^\top (\mathcal{H} u_W) - \tilde{f}_W|_2^2 \leq 2 |\mathcal{H}_0 u_\star - \mathcal{H}_0 u_W|_2^2 + 2 |\theta_\star^\top (\mathcal{H} u_W) - \theta_\star^\top (\mathcal{H} u_\star)|_2^2 \\ &\leq C \left(1 + \|\theta_\star\|_2^2\right) \epsilon_0^{2\kappa}, \end{aligned}$$

for some absolute constant C that depends only on κ and C_4 . We conclude that for $W \in \mathcal{V}$,

$$|\mathcal{R}_W^{(2)}| = \frac{\lambda}{2} \left(|f|_2^2 - \Phi_W^\top \Sigma_W^{-1} \Phi_W \right) \leq \frac{2}{C_2} + C(1 + \|\theta_\star\|_2^2) \lambda \epsilon_0^{2\kappa} \leq \frac{1}{C_2} + \frac{C n \epsilon_0^2}{2\sigma^2},$$

where in the last inequality we have used (2.11). Hence, there exists an absolute constant C , such that for all n large enough,

$$\int_{\mathcal{V}} e^{\mathcal{R}_W} \Pi_0(dW) \geq C e^{-\frac{C n \epsilon_0^2}{2\sigma^2}} \Pi_0(\mathcal{V}). \quad (\text{A.10})$$

Lemma A.2 below gives the lower bound

$$\Pi_0(\mathcal{V}) \geq \frac{1}{2} \exp \left(-s_0(u+2) \log(q) - \frac{s_0}{2} (\|W_0\|_\infty + 1)^2 - s_0 \log \left(\frac{L s_0^{1/2}}{\epsilon_0} \right) \right),$$

which together with assumption (2.8) and the inequality in (A.10) yields

$$\int_{\mathcal{V}} e^{\mathcal{R}_W} \Pi_0(dW) \geq c \times \exp \left(-C \left[\frac{n \epsilon_0^2}{\sigma^2} + s_0(u+2) \log(q) \right] \right).$$

From the definition of s_0 in (2.6), we have

$$s_0(u+2) \log(q) \leq 6u(s_0 - 1) \log(q) \leq 6 \frac{n \epsilon_0^2}{\sigma^2},$$

which then yields (A.8). The lemma follows from (A.7) and (A.8). \square

We show here that the prior Π_0 has good contraction properties.

Lemma A.2. 1. Given $s, r \geq 0$, define

$$\mathcal{W}(s, r) \stackrel{\text{def}}{=} \{W \in \mathbb{R}^q : \|W\|_0 \leq s, \text{ and } \|W\|_\infty \leq r\}.$$

If $r \geq \sqrt{(1+s)(2+u) \log(q)}$, we have

$$\Pi_0(\mathcal{W}(s, r)) \geq 1 - \frac{4}{q^{u(1+s)}}.$$

2. Assume $q \geq \sqrt{2\pi}$. Fix $W_0 \in \mathbb{R}^q$ with sparsity support Λ_0 . Given $r > 0$, let

$$\mathcal{V}(W_0, r) \stackrel{\text{def}}{=} \{W \in \mathbb{R}^q : \text{supp}(W) = \Lambda_0, \|W - W_0\|_\infty \leq r\}.$$

$$\Pi_0(\mathcal{V}(W_0, r)) \geq \frac{1}{2} \exp\left(-s_0(u+2)\log(q) - \frac{s_0}{2}(\|W_0\|_\infty + r)^2 + s_0\log(r)\right),$$

where $s_0 = \|W_0\|_0$.

Proof. 1. Let \mathcal{W}^c be a short for $\mathbb{R}^q \setminus W(s, r)$. Let $\bar{\Pi}_0$ denote the joint of (Λ, W) in the definition of Π_0 (see Section 1.5). Then

$$\Pi_0(\mathcal{W}^c) = \Pi_0(\|\Lambda\|_0 > s) + \sum_{\Lambda: \|\Lambda\|_0 \leq s} \bar{\Pi}_0(\Lambda) \times \bar{\Pi}_0(\|\Lambda \odot W\|_\infty > r | \Lambda).$$

If $(\Lambda, W) \sim \bar{\Pi}_0$, then Λ is an ensemble of iid random variables drawn from the Bernoulli distribution with success probability $(1 + q^{u+1})^{-1}$. Hence

$$\begin{aligned} \Pi_0(\|\Lambda\|_0 > s) &\leq \sum_{j>s} \binom{q}{j} \left(\frac{1}{1+q^{u+1}}\right)^j \left(\frac{q^{u+1}}{1+q^{u+1}}\right)^{q-j} \\ &\leq \sum_{j>s} \binom{q}{j} \left(\frac{1}{q^{u+1}}\right)^j \leq 2 \left(\frac{1}{q^u}\right)^{s+1}, \end{aligned}$$

where we use $\binom{q}{j} \leq q^j$, and $q^u \geq 2$. Given $\Lambda_d = 1$, $W_d \sim \mathbf{N}(0, 1)$. Therefore, $\mathbb{P}(|W_d| > t) \leq 2e^{-t^2/2}$ for all $t \geq 0$. Hence by union bound, for $\|\Lambda\|_0 \leq s$, and since $r \geq \sqrt{(1+s)(2+u)\log(q)}$, we obtain

$$\Pi_0(\|\Lambda \odot W\|_\infty > r | \Lambda) \leq 2e^{-r^2/2+\log(s)} \leq \frac{2}{q^{u(1+s)}}.$$

We conclude that

$$\Pi_0(\mathcal{W}^c) \leq \frac{4}{q^{u(1+s)}}. \quad (\text{A.11})$$

2. We write \mathcal{V} as a short for $\mathcal{V}(W_0, r)$. We have

$$\Pi_0(\mathcal{V}) = \bar{\Pi}_0(\Lambda_0) \bar{\Pi}_0(\|\Lambda \odot W - W_0\|_\infty \leq r | \Lambda = \Lambda_0).$$

Since $\log(1-x) \geq -2x$ for all $0 \leq x \leq 1/2$, for $q^u \geq 2/\log(2)$, we have

$$\begin{aligned} \bar{\Pi}_0(\Lambda_0) &= \left(\frac{1}{1+q^{u+1}}\right)^{\|\Lambda_0\|_0} \left(1 - \frac{1}{1+q^{u+1}}\right)^{q-\|\Lambda_0\|_0} \\ &= \left(\frac{1}{q^{u+1}}\right)^{\|\Lambda_0\|_0} \exp\left(q \log\left(1 - \frac{1}{1+q^{u+1}}\right)\right) \\ &\geq \left(\frac{1}{q^{u+1}}\right)^{\|\Lambda_0\|_0} \exp\left(-\frac{2q}{1+q^{u+1}}\right) \geq \frac{1}{2} \left(\frac{1}{q^{u+1}}\right)^{\|\Lambda_0\|_0} = \frac{1}{2} \left(\frac{1}{q^{u+1}}\right)^{s_0}. \end{aligned}$$

If $U \sim \mathbf{N}(0, 1)$, and $t \geq 0$, then for all a , and $c \geq |a|$,

$$P(|U - a| \leq t) \geq P(c \leq U \leq c+t) = \Phi(c+t) - \Phi(c) \geq e^{-(c+t)^2/2} \frac{t}{\sqrt{2\pi}},$$

where Φ is the cdf of the standard normal distribution. We use this inequality with $c = \|W_0\|_\infty$, and we the assumption $q \geq \sqrt{2\pi}$, we deduce that

$$\bar{\Pi}_0(\|\Lambda \odot W - W_0\|_\infty \leq r | \Lambda = \Lambda_0) \geq \frac{1}{q^{s_0}} \exp\left(-\frac{s_0}{2}(\|W_0\|_\infty + r)^2 + s_0 \log(r)\right).$$

Hence

$$\bar{\Pi}_0(\Lambda_0) \geq \frac{1}{2} \left(\frac{1}{q^{u+2}}\right)^{s_0} \exp\left(-\frac{s_0}{2}(\|W_0\|_\infty + r)^2 + s_0 \log(r)\right),$$

as claimed. □

A.2. Ignorability of unsuitable weights

Given integer $s \geq 0$, $\tau_s \stackrel{\text{def}}{=} \sqrt{(2+u)(1+s) \log(q)}$, and a constant C , we define that

$$\mathcal{W}_0(s, C) \stackrel{\text{def}}{=} \left\{ W \in \mathbb{R}^q : \|W\|_0 \leq s, \text{ and } \|W\|_\infty \leq \tau_s, \text{ and } \mathcal{J}(u_W) \leq \frac{Cn\epsilon_0^2}{\lambda\sigma^2} \right\}.$$

Our next result shows that the prior Π_0 puts most of its probability mass on $\mathcal{W}_0(s, C)$.

Lemma A.3. *Assume H1-H4. Let s_0, ϵ_0, W_0 be as in H3. We can find an absolute constant C such that with*

$$s = Cs_0, \tag{A.12}$$

it holds,

$$\mathbb{E}[\Pi(\mathcal{W}_0(s, C) | \mathcal{D})] \geq 1 - C_0 e^{-\frac{n\epsilon_0^2}{2\sigma^2}},$$

for some absolute constant C_0 .

Proof. For any measurable set $A \subseteq \mathbb{R}^q$, we have

$$\Pi(A | \mathcal{D}) = \frac{\int_A \exp\left(-\frac{n}{2\sigma^2} \mathbb{P}_n(\Delta_W - \Delta_\star) + \mathcal{R}_W\right) \Pi_0(dW)}{\int_{\mathbb{R}^q} \exp\left(-\frac{n}{2\sigma^2} \mathbb{P}_n(\Delta_W - \Delta_\star) + \mathcal{R}_W\right) \Pi_0(dW)}.$$

By Lemma A.1, we can find an absolute constant C_0 , such that for all n large enough, and with

$$\beta \stackrel{\text{def}}{=} \frac{n\epsilon_0^2}{2\sigma^2},$$

$$\mathbb{E}[\Pi(A | \mathcal{D})] \leq 2e^{-n\epsilon_0^2/(2\sigma^2)} + \frac{e^{C_0\beta}}{C_0} \mathbb{E}\left[\int_A \exp\left(-\frac{n}{2\sigma^2} \mathbb{P}_n(\Delta_W - \Delta_\star) + \mathcal{R}_W\right) \Pi_0(dW)\right].$$

By Fubini's theorem, the expectation on the right hand side of the last display is

$$\int_A e^{\mathcal{R}_W} \mathbb{E} \left[\exp \left(-\frac{n}{2\sigma^2} \mathbb{P}_n(\Delta_W - \Delta_\star) \right) \right] \Pi_0(dW).$$

By conditioning on $\mathbf{s}_{1:n}$ we see that the expectation inside the last integral is equal to 1 for all W . Hence,

$$\mathbb{E} [\Pi(A|\mathcal{D})] \leq 2e^{-n\epsilon_0^2/(2\sigma^2)} + \frac{e^{C_0\beta}}{C_0} \int_A e^{\mathcal{R}_W} \Pi_0(dW). \quad (\text{A.13})$$

By definition, we have $\mathcal{R}_W = -\frac{1}{2}(\log \det(\Sigma_W) - \log \det(\Sigma_\star)) - \lambda \mathcal{J}(u_W)/2$. For all $W \in \mathbb{R}^q$, we have

$$-\frac{1}{2}(\log \det(\Sigma_W) - \log \det(\Sigma_\star)) \leq \frac{d}{2} \log(n \|\Sigma_\star\|_{\text{op}}) \leq C_1\beta, \quad (\text{A.14})$$

where the second inequality uses (2.12). We note also that $\mathcal{J}(u_W) \geq 0$. As a result (A.13) becomes

$$\mathbb{E} [\Pi(A|\mathcal{D})] \leq 2e^{-n\epsilon_0^2/(2\sigma^2)} + \frac{e^{C_0\beta}}{C_0} \int_A e^{-\frac{\lambda}{2} \mathcal{J}(u_W)} \Pi_0(dW), \quad (\text{A.15})$$

for some possibly different absolute constant C_0 . We apply this with $A = \mathcal{W}_0(Cs_0, C)^c = \{W \in \mathbb{R}^q : \mathcal{J}(u_W) > 2C\beta/\lambda\} \cup [\mathcal{W}_0(s)^c \cap \{W \in \mathbb{R}^q : \mathcal{J}(u_W) \leq 2C\beta/\lambda\}]$. For $C \geq 1 + C_0$,

$$\mathbb{E} [\Pi(\mathcal{W}(s)^c|\mathcal{D})] \leq 2e^{-n\epsilon_0^2/(2\sigma^2)} + \frac{1}{C_0} e^{-n\epsilon_0^2/(2\sigma^2)} + \frac{e^{C_0\beta}}{C_0} \int_{\mathcal{W}(s)^c} \Pi_0(dW),$$

where $\mathcal{W}(s) = \{W \in \mathbb{R}^q : \|W\|_0 \leq s, \text{ and } \|W\|_\infty \leq \tau_s\}$. Lemma A.2-(1) shows that $\Pi_0(\mathcal{W}(s)^c) \leq 4\exp(-us \log(q))$. As a result, by taking $s = Cs_0$ with $C \geq C_0 + 1$, we have

$$us \log(q) = Cus_0 \log(q) \geq \frac{Cn\epsilon_0^2}{2\sigma^2} \geq \frac{(C_0 + 1)n\epsilon_0^2}{2\sigma^2},$$

and we conclude that

$$\mathbb{E} [\Pi(\mathcal{W}_0(s)^c|\mathcal{D})] \leq 2e^{-n\epsilon_0^2/(2\sigma^2)} + \frac{1}{C_0} e^{-n\epsilon_0^2/(2\sigma^2)} \leq C \times e^{-n\epsilon_0^2/(2\sigma^2)},$$

for some absolute constant C . □

A.3. Deviation bounds

The last ingredient of the proof is a concentration inequality for the empirical process of the log-likelihood ratio that we derive next. The proof is based on Theorem 5.11 of (12) that we first present. For any random variable X , its κ -Bernstein norm ρ_κ is defined as:

$$\rho_\kappa(X) \stackrel{\text{def}}{=} \sqrt{2\kappa^2 \mathbb{E} \left[e^{\frac{|X|}{\kappa}} - 1 - \frac{|X|}{\kappa} \right]}. \quad (\text{A.16})$$

We recall that we say a random variable satisfies Bernstein condition with parameter (K, R) if:

$$\mathbb{E}[|X|^m] \leq \frac{1}{2} m! K^{m-2} R^2 \quad \forall m = 2, 3, \dots$$

It is immediate that if a random variable X satisfies Bernstein condition with parameters (K, R) then $\rho_{2K}(X) \leq \sqrt{2}R$. Let $X, X_{1:n} \stackrel{i.i.d.}{\sim} P, \mathbb{P}_n$ their empirical measure. Let \mathcal{G} be a collection of real-valued functions, where $g(X)$ satisfies the Bernstein condition with parameters (K, R) for all $g \in \mathcal{G}$. Let $H_{B,2K}(u, \mathcal{G}, P)$ denote the bracketing entropy of \mathcal{G} with respect to the pseudo-metric ρ_{2K} .

Theorem A.4 (Theorem 5.11 of (12)). *Suppose that $g(X)$ satisfies the Bernstein condition with parameters (K, R) for all $g \in \mathcal{G}$. Then there exists a universal constant C such that for any a, C_0, C_1 satisfying:*

$$C_0 \left(\max \left\{ \int_0^R \sqrt{H_{B,2K}(u, \mathcal{G}, P)} du, R \right\} \right) \leq a \leq \min \left\{ \frac{C_1 \sqrt{n} R^2}{2K}, 8\sqrt{n}R \right\} \quad (\text{A.17})$$

and $C^2 \leq C_0^2 / (C_1 + 1)$:

$$\mathbb{P} \left(\sup_{g \in \mathcal{G}} |\sqrt{n}(\mathbb{P}_n - P)g| \geq a \right) \leq C e^{-\frac{a^2}{C^2(C_1+1)R^2}}.$$

Given $s \geq 0$, we recall the definition

$$\mathcal{W}(s) = \{W \in \mathcal{W} : \|W\|_0 \leq s, \|W\|_\infty \leq \tau_s\}, \quad \text{where } \tau_s = \sqrt{(2+u)(1+s) \log(q)}.$$

We also recall that

$$\Delta_W(y, \mathbf{s}) \stackrel{\text{def}}{=} (y - u_W(\mathbf{s}))^2, \quad \Delta_\star(y, \mathbf{s}) \stackrel{\text{def}}{=} (y - u_\star(\mathbf{s}))^2, \quad (y, \mathbf{s}) \in \mathbb{R} \times \Omega.$$

We use Theorem A.4 to obtain the following.

Lemma A.5. *Assume H1-H2. Fix $s \geq 1$, and let $\mathcal{F}_s, V_1 = V_1(s)$ and $V_2 = V_2(s) \geq 6b$ be as in H2. There exists an absolute constant $M_0 \geq 1$ such that the following holds. For all $M \geq M_0, A_M \stackrel{\text{def}}{=} Mb\sqrt{V_1 \log(V_2 \sqrt{n})}$, and $\delta > 0$ such that*

$$\frac{2\sqrt{2}}{\sqrt{n}} \leq \frac{A_M}{\sqrt{n}} \leq \delta \leq 2b,$$

we have

$$\mathbb{P} \left[\sup_{W \in \mathcal{W}(s): |u_W - u_\star|_2 \leq \delta} |\sqrt{n}(\mathbb{P}_n - \mathbb{P})(\Delta_W - \Delta_\star)| > \frac{A_M \delta}{2} \left(1 + \frac{\sigma}{b}\right) \right] \leq C e^{-\frac{A_M^2}{cb^2}},$$

for some absolute constant C .

Proof. We set

$$Z_n(\delta) \stackrel{\text{def}}{=} \sup_{W \in \mathcal{W}(s): |u_W - u_\star|_2 \leq \delta} |\sqrt{n}(\mathbb{P}_n - \mathbb{P})(\Delta_W - \Delta_\star)|.$$

Recall that, by definition of Y , we have $Y_i = u_\star(\mathbf{s}_i) + \sigma \xi_i$ where $\xi_i \sim \mathbb{N}(0, 1)$, (ξ_i, \mathbf{s}_i) are independent, and $(\xi_i, \mathbf{s}_i) \sim \mathbb{P}$. Therefore, we can simplify the difference $(\Delta_W - \Delta_\star)(\xi, \mathbf{s})$ as:

$$(\Delta_W - \Delta_\star)(\xi, \mathbf{s}) = (u_\star(\mathbf{s}) - u_W(\mathbf{s}))^2 + 2\sigma \xi (u_\star(\mathbf{s}) - u_W(\mathbf{s})),$$

and consequently, we can write:

$$\begin{aligned} Z_n(\delta) &\leq \sup_{W \in \mathcal{W}(s): |u_W - u_\star|_2 \leq \delta} \left| \sqrt{n}(\mathbb{P}_n - \mathbb{P})(u_\star(\mathbf{s}) - u_W(\mathbf{s})) \right| \\ &\quad + 2\sigma \sup_{W \in \mathcal{W}(s): |u_W - u_\star|_2 \leq \delta} \left| \sqrt{n}(\mathbb{P}_n - \mathbb{P})(\xi u_\star(\mathbf{s}) - \xi u_W(\mathbf{s})) \right| \\ &= Z_{n,1}(\delta) + 2\sigma Z_{n,2}(\delta), \end{aligned}$$

where $Z_{n,1}(\delta)$ and $Z_{n,2}(\delta)$ can also be written as

$$Z_{n,1}(\delta) \stackrel{\text{def}}{=} \sup_{h \in \mathcal{H}_\delta} \left| \frac{1}{\sqrt{n}} \sum_{i=1}^n (h(\mathbf{s}_i) - \mathbb{P}(h)) \right|, \quad \text{and} \quad Z_{n,2}(\delta) \stackrel{\text{def}}{=} \sup_{f \in \bar{\mathcal{F}}_\delta} \left| \frac{1}{\sqrt{n}} \sum_{i=1}^n f(\xi_i, X_i) \right|,$$

where $\mathcal{H}_\delta \stackrel{\text{def}}{=} \{f^2, f \in \mathcal{F}_\delta\}$, and where $\mathcal{F}_\delta = \{u_W - u_\star : W \in \mathcal{W}(s), |u_W - u_\star|_2 \leq \delta\}$, and $\bar{\mathcal{F}}_\delta = \{(\xi, x) \mapsto \xi f(x), f \in \mathcal{F}_\delta\}$.

A.3.0.1. (A) Deviation bound for $Z_{n,1}(\delta)$. By H2-(1), $|h|_\infty \leq b^2$ for all $h \in \mathcal{H}_\delta$. Therefore for $f \in \mathcal{F}_\delta$, and $m \geq 2$, we have:

$$\mathbb{E}[|f^2(X)|^m] \leq (b^2)^{m-2} \mathbb{E}[|f^2(X)|^2] \leq (b^2)^{m-2} (b\delta)^2 \leq \frac{1}{2} m! (b^2)^{m-2} (b\delta)^2.$$

Therefore the function class \mathcal{H}_δ satisfies the Bernstein condition with parameter $(b^2, b\delta)$. The bracketing number of \mathcal{H}_δ satisfies

$$\mathcal{H}_{B,2b^2}(u\sqrt{2}, \mathcal{H}_\delta, \mathbb{P}) \leq \log \mathcal{N}(u/2b, \mathcal{F}_\delta, L_\infty), \quad (\text{A.18})$$

for all $0 < u < b^2$. To see this fix some $u > 0$ and let $\{f_j\}_{1 \leq j \leq N}$ be $(u/2b)$ -cover of \mathcal{F}_δ with respect to L_∞ norm. Therefore, for any $f \in \mathcal{F}_\delta$ there exists $1 \leq j \leq N$ such that $\|f - f_j\|_\infty \leq u/2b$. Consequently, we have $\|f^2 - f_j^2\|_\infty \leq 2b(u/2b) = u$. Now consider the collection of brackets $\{(f_j^2 - u/2, f_j^2 + u/2)\}_{1 \leq j \leq N}$. For this collection we have, $\rho_{2b^2}^2 \left(f_j^2 + \frac{u}{2} - f_j^2 + \frac{u}{2} \right) = \rho_{2b^2}^2(u) \leq (\sqrt{2}u)^2$, for $u \leq b^2$. Hence we have established (A.18).

We can then apply Theorem A.4 with $a = A_M \delta/2$, $R = b\delta$, $K = b^2$, $C_1 = 1$, $C_0^2 = 2C$, where C is the absolute constant in Theorem A.4. The condition $a \geq C_0 R$ is satisfied with $M \geq M_0 \geq 2C_0$. For $\delta \leq 2b$, using H2-(4),

$$\begin{aligned} \int_0^{b\delta} \sqrt{\mathcal{H}_{B,2b^2}(u, \mathcal{H}_\delta, \mathbb{P})} du &= \sqrt{2} \int_0^{b\delta/\sqrt{2}} \sqrt{\mathcal{H}_{B,2b^2}(u\sqrt{2}, \mathcal{H}_\delta, \mathbb{P})} du \\ &\leq \sqrt{2} \int_0^{b\delta/\sqrt{2}} \sqrt{\log \mathcal{N}(u/2b, \mathcal{F}_\delta, L_\infty)} du \end{aligned}$$

$$\begin{aligned} &\leq 2b\sqrt{2V_1} \int_0^{\delta/2\sqrt{2}} \sqrt{\log\left(\frac{V_2}{u}\right)} du \\ &\leq 2b\delta\sqrt{V_1} \sqrt{\log\left(\frac{2\sqrt{2}V_2}{\delta}\right)} \leq 2b\delta\sqrt{V_1 \log(V_2\sqrt{n})}, \end{aligned}$$

where the inequality before last follows from the following fact:

$$\int_0^c \sqrt{\log\left(\frac{C}{u}\right)} du \leq 2c\sqrt{\log\left(\frac{C}{c}\right)} \text{ as soon as } \log(C/c) \geq 2,$$

which in our case is implied by $V_2 \geq 6b$. The last inequality uses $\delta \geq 2\sqrt{2}/\sqrt{n}$.

Therefore the condition $a \geq C_0 \int_0^{b\delta} \sqrt{\mathcal{H}_{B,2b^2}(u, \mathcal{F}_\delta, \mathbb{P})} du$ is satisfied with $M \geq M_0 \geq 4C_0$. The condition $a \leq C_1\sqrt{n}R^2/(2K)$ boils down to $A_M/\sqrt{n} \leq \delta$, whereas the condition $a \leq 8R\sqrt{n}$ boils down to $A_M/\sqrt{n} \leq 16b$ which holds by assumption. Hence by Theorem A.4,

$$\mathbb{P}\left(Z_{n,1}(\delta) > \frac{A_M\delta}{2}\right) \leq Ce^{-\frac{A_M^2}{Cb^2}},$$

for some absolute constant C .

A.3.0.2. (B) Deviation bound for $Z_{n,2}(\delta)$. The argument is similar. Note that, for any $u_W - u_\star \in \mathcal{F}_\delta$, $|u_W - u_\star|_2 \leq \delta$, and since the function $u_W - u_\star$ are bounded by b as assumed in H2, we have:

$$\begin{aligned} \mathbb{E}[|\xi f(X)|^m] &= \mathbb{E}[|\xi|^m] \mathbb{E}[|u_W(X) - u_\star(X)|^m] \\ &\leq \mathbb{E}[|\xi|^m] b^{m-2} \delta^2 \\ &\leq \frac{1}{2} m! b^{m-2} \delta^2, \end{aligned}$$

where the last inequality follows from the centered absolute moment of Gaussian random variables. Therefore, all $f \in \tilde{\mathcal{F}}_\delta$ satisfies the Bernstein condition with parameters (b, δ) .

The bracketing entropy satisfies,

$$\mathcal{H}_{B,2b}(u\sqrt{2}, \tilde{\mathcal{F}}_\delta, \mathbb{P}) \leq \log \mathcal{N}(u/2, \mathcal{F}_\delta, L_\infty) \quad 0 < u \leq \delta. \quad (\text{A.19})$$

To establish (A.19), fix $u \in (0, b]$, and let $\{f_1, \dots, f_N\}$ be a $u/2$ cover of \mathcal{F}_δ with respect to L_∞ norm, i.e.

$$\sup_{f \in \mathcal{F}_\delta} \min_{1 \leq j \leq N} \|f - f_j\|_\infty \leq \frac{u}{2}.$$

Now consider the brackets $\{(\xi, x) \mapsto (\xi f_j(x) - u|\xi|/2, \xi f_j(x) + u|\xi|/2), 1 \leq j \leq N\}$. For any f , there exists f_j such that $\|f - f_j\|_\infty \leq u/2$ due to the property of the covering set, and it follows that

$$\xi f_j(x) - \frac{u|\xi|}{2} \leq \xi f(x) \leq \xi f_j(x) + \frac{u|\xi|}{2} \quad \forall x.$$

Furthermore, for any $1 \leq j \leq N$, since $\mathbb{E}(|u\xi|^m) \leq u^m m! / 2 \leq b^{m-2} u^2 m! / 2$,

$$\rho_{2b}^2 (\xi f_j(x) + u|\xi|/2 - \xi f_j(x) + u|\xi|/2) = \rho_{2b}^2 (u|\xi|) \leq (\sqrt{2}u)^2,$$

which implies (A.19).

Hence we can apply Theorem A.4 with $a = A_M \delta / 2b$, $R = \delta$, $K = b$, $C_1 = 1$, $C_0^2 = 2C$, where C is the absolute constant in Theorem A.4. All the conditions of Theorem A.4 can be checked as we did in bounding $Z_{n,1}(\delta)$. We conclude that

$$\mathbb{P} \left(Z_{n,2}(\delta) > \frac{A_M \delta}{2b} \right) \leq C e^{-\frac{A_M^2}{Cb^2}},$$

for some absolute constant C . Combining the two bounds, we get

$$\mathbb{P} \left(Z_n(\delta) > \frac{A_M \delta}{2} \left(1 + \frac{\sigma}{b} \right) \right) \leq C e^{-\frac{A_M^2}{Cb^2}},$$

for some absolute constant C , which implies the stated result. \square

A.4. Finishing the proof

Proof. Let s_0, ϵ_0 , and W_0 be as in H3. Let $s = Cs_0$ as in (A.12). We set $\tau_s = \sqrt{(2+u)(1+s) \log(q)}$,

$$\mathcal{W}_0(s) \stackrel{\text{def}}{=} \left\{ W \in \mathbb{R}^q : \|W\|_0 \leq s, \|W\|_\infty \leq \tau_s, \mathcal{J}(u_W) \leq \frac{Cn\epsilon_0^2}{\lambda\sigma^2} \right\}.$$

We set $r \stackrel{\text{def}}{=} 2M_0(b + \sigma)\sqrt{V_1 \log(V_2 \sqrt{n})/n}$, where M_0 is as in Lemma A.5. For $j \geq 1$, we set $A_j \stackrel{\text{def}}{=} bM_0 j \sqrt{V_1 \log(V_2 \sqrt{n})}$, $r_j \stackrel{\text{def}}{=} jr$. With C as in Lemma A.1, we set

$$\beta \stackrel{\text{def}}{=} \frac{C}{\sigma^2} n \epsilon_0^2,$$

and define

$$\mathcal{E} \stackrel{\text{def}}{=} \left\{ \mathcal{D} : \int_{\mathbb{R}^q} \exp \left(-\frac{n}{2\sigma^2} \mathbb{P}_n(\Delta_W - \Delta_\star) + \mathcal{R}_W \right) \Pi_0(dW) \leq C e^{-\beta}, \text{ or} \right. \\ \left. \sup_{W \in \mathcal{W}(s): r_j < \|u_W - u_\star\|_2 \leq r_{j+1}} |\sqrt{n}(\mathbb{P}_n - \mathbb{P})(\Delta_W - \Delta_\star)| > \frac{A_{j+1} r_{j+1}}{2} \left(1 + \frac{\sigma}{b} \right) \text{ for some } j \geq 1 \right\}.$$

By Lemma A.1 and Lemma A.5, and for all n large enough

$$\mathbb{P}(\mathcal{D} \in \mathcal{E}) \leq 2e^{-n\epsilon_0^2/(2\sigma^2)} + C \sum_{j \geq 1} e^{-A_j^2/(Cb^2)} \leq 2e^{-n\epsilon_0^2/(2\sigma^2)} + C e^{-V_1 \log(V_2 \sqrt{n})/C}, \quad (\text{A.20})$$

for some absolute constant C .

We set $B \stackrel{\text{def}}{=} \{W \in \mathbb{R}^q : |u_W - u_\star|_2 > r\}$. We can write $B \cap \mathcal{W}(s)$ as $\cup_{j \geq 1} B_j$, where

$$B_j \stackrel{\text{def}}{=} \{W \in \mathcal{W}(s) : Mjr < |u_W - u_\star|_2 \leq M(j+1)r\}.$$

We should point out that the union $\cup_{j \geq 1} B_j$ is over a finite number of terms since $B_j = \emptyset$ for $Mjr \geq b$. Since $\Pi(B|\mathcal{D}) \leq \Pi(\mathcal{W}_0(s)^c|\mathcal{D}) + \mathbf{1}_{\mathcal{E}}(\mathcal{D}) + \mathbf{1}_{\mathcal{E}^c}(\mathcal{D})\Pi(B \cap \mathcal{W}_0(s)|\mathcal{D})$, taking expectation on both sides, and using Lemma A.3, the definition of \mathcal{E} and (A.20), yields for all n large enough,

$$\begin{aligned} \mathbb{E}[\Pi(B|\mathcal{D})] &\leq Ce^{-n\epsilon_0^2/(2\sigma^2)} + Ce^{-V_1 \log(V_2 \sqrt{n})/C} \\ &\quad + \frac{e^\beta}{C} \sum_{j \geq 1} \mathbb{E} \left[\mathbf{1}_{\mathcal{E}^c}(\mathcal{D}) \int_{B_j} \exp\left(-\frac{n}{2\sigma^2} \mathbb{P}_n(\Delta_W - \Delta_\star) + \mathcal{R}_W\right) \Pi_0(dW) \right], \end{aligned}$$

for some absolute constant C . We have

$$\frac{n}{2\sigma^2} \mathbb{P}_n(\Delta_W - \Delta_\star) + \mathcal{R}_W = -\frac{n}{2\sigma^2} \varrho(u_W, u_\star) - \frac{n}{2\sigma^2} (\mathbb{P}_n - \mathbb{P})(\Delta_W - \Delta_\star) + \mathcal{R}'_W,$$

where

$$\mathcal{R}'_W = -\frac{1}{2} (\log \det \Sigma_\star - \log \det \Sigma_W),$$

and satisfies, as we show in the proof of Lemma A.3

$$|\mathcal{R}'_W| \leq \frac{d}{2} \log \left(\frac{n \|\Sigma_\star\|_{\text{op}}}{\sigma^2} \right) \leq C_0 \log(q),$$

where the second inequality is our assumption (2.12). Hence, we can find an absolute constant C such that for all n large enough

$$\begin{aligned} \mathbb{E}[\Pi(B|\mathcal{D})] &\leq Ce^{-n\epsilon_0^2/(2\sigma^2)} + Ce^{-V_1 \log(V_2 \sqrt{n})/C} \\ &\quad + \frac{e^\beta}{C} \sum_{j \geq 1} \mathbb{E} \left[\mathbf{1}_{\mathcal{E}^c}(\mathcal{D}) \int_{B_j} \exp\left(-\frac{n}{2\sigma^2} \varrho(u_W, u_\star) - \frac{n}{2\sigma^2} (\mathbb{P}_n - \mathbb{P})(\Delta_W - \Delta_\star)\right) \Pi_0(dW) \right]. \quad (\text{A.21}) \end{aligned}$$

For $W \in B_j$, and $\mathcal{D} \notin \mathcal{E}$, the expression inside the exponential in the last display is bounded from above by

$$-\frac{nr_j^2}{2\sigma^2} + \frac{\sqrt{n}A_{j+1}r_{j+1}}{2\sigma^2} \left(1 + \frac{\sigma}{b}\right) \leq -\frac{nr_j^2}{4\sigma^2},$$

with the choice of r . Since

$$\sum_{j \geq 1} e^{-\frac{nr_j^2}{4\sigma^2}} \leq Ce^{-\frac{nr^2}{4\sigma^2}},$$

for some absolute constant C , we conclude that for all n large enough,

$$\mathbb{E}[\Pi(B|\mathcal{D})] \leq C \left(e^{-n\epsilon_0^2/(2\sigma^2)} + e^{-V_1 \log(V_2 \sqrt{n})/C} + e^{-nr^2/(4\sigma^2)} \right).$$

Hence the theorem. \square

A.5. Proof of Theorem 3.1

Proof. First, a simple coupling argument shows that

$$W_2^2(\Pi^{(\theta)}, \Pi_\star^{(\theta)}) \leq \int_{\mathbb{R}^q} W_2^2\left(\mathbf{N}(\hat{\theta}_W, \frac{1}{\lambda}\Sigma_W^{-1}), \mathbf{N}(\theta_\star, \frac{1}{\lambda}\Sigma_\star^{-1})\right) \Pi^{(W)}(dW|\mathcal{D}). \quad (\text{A.22})$$

We recall (see e.g. (5)) that for symmetric and positive definite matrices Λ_1, Λ_2 ,

$$\begin{aligned} W_2^2(\mathbf{N}(\mu_1, \Lambda_1), \mathbf{N}(\mu_2, \Lambda_2)) &= \|\mu_2 - \mu_1\|_2^2 + \text{Tr}\left(\Lambda_1 + \Lambda_2 - 2\left(\Lambda_1^{1/2}\Lambda_2\Lambda_1^{1/2}\right)^{1/2}\right) \\ &\leq \|\mu_2 - \mu_1\|_2^2 + \|\Lambda_1^{1/2} - \Lambda_2^{1/2}\|_F^2, \end{aligned}$$

where the inequality follows from Theorem 1 of ((5)). By the Hemmen-Ando inequality ((1) Proposition 2.1),

$$\|\Lambda_1^{1/2} - \Lambda_2^{1/2}\|_F \leq \frac{\|\Lambda_1 - \Lambda_2\|_F}{\sqrt{\lambda_{\min}(\Lambda_1)} + \sqrt{\lambda_{\min}(\Lambda_2)}}.$$

Hence

$$W_2^2(\mathbf{N}(\mu_1, \Lambda_1), \mathbf{N}(\mu_2, \Lambda_2)) \leq \|\mu_2 - \mu_1\|_2^2 + \frac{\|\Lambda_1 - \Lambda_2\|_F^2}{\left(\sqrt{\lambda_{\min}(\Lambda_1)} + \sqrt{\lambda_{\min}(\Lambda_2)}\right)^2}. \quad (\text{A.23})$$

We apply this bound, to conclude that there exists a constant c (that we can take as $d/\lambda_{\max}(\Sigma_\star)$ such that for all $W \in \mathbb{R}^q$,

$$W_2^2\left(\mathbf{N}(\hat{\theta}_W, \frac{1}{\lambda}\Sigma_W^{-1}), \mathbf{N}(\theta_\star, \frac{1}{\lambda}\Sigma_\star^{-1})\right) \leq \|\hat{\theta}_W - \theta_\star\|_2^2 + \frac{c}{\lambda}\|\Sigma_W^{-1} - \Sigma_\star^{-1}\|_{\text{op}}^2.$$

We consider the integrand of (A.22) under two scenarios.

A.5.0.1. Case 1: W is such that $|u_W - u_\star|_2 \leq Mr$. By (A.9), and for W such that $|u_W - u_\star|_2 \leq Mr$, for some absolute onstant C we have

$$\|\Sigma_W - \Sigma_\star\|_{\text{op}} \leq C(|\mathcal{H}u_\star|_2 + (Mr)^\kappa)(Mr)^\kappa,$$

for all n large enough. So, since $\lambda_{\min}(\Sigma_\star) > C_2$, we conclude that for all such W under consideration, and for all n large enough, $\lambda_{\min}(\Sigma_W) > C_2/2$.

For all W ,

$$\begin{aligned} (\hat{\theta}_W - \theta_\star)^\top \mathcal{H}_1 u_W &= \hat{\theta}_W^\top \mathcal{H}_1 u_W - \theta_\star^\top \mathcal{H}_1 u_\star - \theta_\star^\top (\mathcal{H}_1 u_W - \mathcal{H}_1 u_\star) \\ &= \hat{\theta}_W^\top \mathcal{H}_1 u_W - \mathcal{H}_0 u_W - f + (\mathcal{H}_0 u_W - \mathcal{H}_0 u_\star) - \theta_\star^\top (\mathcal{H}_1 u_W - \mathcal{H}_1 u_\star). \end{aligned}$$

Therefore

$$|(\hat{\theta}_W - \theta_\star)^\top \mathcal{H}_1 u_W|_2^2 \leq 2|f - \mathcal{H}_0 u_W - \hat{\theta}_W^\top \mathcal{H}_1 u_W|_2^2 + c \times |\mathcal{H}u_W - \mathcal{H}u_\star|_2^2.$$

Since $\hat{\theta}_W$ minimizes $\theta \mapsto |f - \mathcal{H}_0 u_W - \theta^\top \mathcal{H}_1 u_W|_2 + \|\theta\|_2^2/\lambda$, we have

$$\begin{aligned} |f - \mathcal{H}_0 u_W - \hat{\theta}_W^\top \mathcal{H}_1 u_W|_2^2 &\leq |f - \mathcal{H}_0 u_W - \theta_\star^\top \mathcal{H}_1 u_W|_2^2 + \frac{\|\theta_\star\|_2^2}{\lambda} \\ &\leq c \times |\mathcal{H} u_W - \mathcal{H} u_\star|_2^2 + \frac{\|\theta_\star\|_2^2}{\lambda}. \end{aligned}$$

Using the above, and H4, we conclude that we can find a constant c such that for W such that $|u_W - u_\star|_2 \leq Mr$, it holds

$$|(\hat{\theta}_W - \theta_\star)^\top \mathcal{H}_1 u_W|_2^2 \leq cr^{2\kappa}.$$

However, $|(\hat{\theta}_W - \theta_\star)^\top \mathcal{H}_1 u_W|_2^2 = (\hat{\theta}_W - \theta_\star)^\top (\Sigma_W - (1/\lambda)\mathbf{1}_d)(\hat{\theta}_W - \theta_\star)$. And since $\lambda_{\min}(\Sigma_W) > C/2$ as seen above, and since $\lambda \rightarrow \infty$, as $n \rightarrow \infty$, for all n large enough we have

$$|(\hat{\theta}_W - \theta_\star)^\top \mathcal{H}_1 u_W|_2^2 \geq C\|\hat{\theta}_W - \theta_\star\|_2^2,$$

for some absolute constant C . In conclusion, for W such that $|u_W - u_\star|_2 \leq Mr$, and for all n large enough,

$$\|\hat{\theta}_W - \theta_\star\|_2^2 \leq cr^{2\kappa},$$

and

$$W_2^2\left(\mathbf{N}(\hat{\theta}_W, \frac{1}{\lambda}\Sigma_W^{-1}), \mathbf{N}(\theta_\star, \frac{1}{\lambda}\Sigma_\star^{-1})\right) \leq cr^{2\kappa}.$$

A.5.0.2. Case 2: W is such that $|u_W - u_\star|_2 > Mr_n$. Since the smallest eigenvalue of Σ_W is at least $1/\lambda$, there exists a constant c such that for all $W \in \mathbb{R}^q$, we have

$$\|\Sigma_W^{-1} - \Sigma_\star^{-1}\|_{\text{op}} \leq c\lambda.$$

Similarly,

$$\|\hat{\theta}_W - \theta_\star\|_2 \leq \|\hat{\theta}_W\|_2 + \|\theta_\star\|_2,$$

and

$$\|\hat{\theta}_W\|_2^2 = \hat{\theta}_W^\top \Sigma_W^{1/2} \Sigma_W^{-1} \Sigma_W^{1/2} \hat{\theta}_W \leq \lambda \hat{\theta}_W^\top \Sigma_W \hat{\theta}_W = \lambda \Phi_W^\top \Sigma_W^{-1} \Phi_W \leq \lambda |\bar{f}_W|_2^2 \leq 2\lambda(|f|_2^2 + |\mathcal{H}_0 u_W|_2^2).$$

We deduce that for all W such that $|u_W - u_\star|_2 > Mr$, we have

$$W_2^2\left(\mathbf{N}(\hat{\theta}_W, \frac{\sigma^2}{\lambda}\Sigma_W^{-1}), \mathbf{N}(\theta_\star, \frac{\sigma^2}{\lambda}\Sigma_\star^{-1})\right) \leq c\lambda \left(1 + |\mathcal{H}_0 u_W|_2^2\right),$$

for some constant c . Given the behaviors of the W_2 distance between $\mathbf{N}(\hat{\theta}_W, \frac{1}{\lambda}\Sigma_W^{-1})$ and $\mathbf{N}(\theta_\star, \frac{1}{\lambda}\Sigma_\star^{-1})$ obtained in the cases above, we return to (A.22) to write

$$W_2^2(\Pi^{(\theta)}, \Pi_\star^{(\theta)}) \leq cr^{2\kappa} + c\lambda \int_{\mathbb{R}^q} \left(1 + |\mathcal{H}_0 u_W|_2^2\right) \mathbf{1}_{\{|u_W - u_\star|_2 > Mr_n\}} \Pi^{(W)}(dW|\mathcal{D}).$$

Taking the expectation on both sides and with a similar argument as in the proof of Theorem 2.2 yields

$$\mathbb{E}\left[W_2^2(\Pi^{(\theta)}, \Pi_\star^{(\theta)})\right] \leq c\left(r^{2\kappa} + \lambda e^{-n\epsilon_0^2/(2\sigma^2)} + \lambda e^{-V_1 \log(V_2 \sqrt{n})/C}\right),$$

which is the stated bound. \square

A.6. Proof of Theorem 2.4

Proof. The result is based on a generalization of Markov's polynomial inequality due to (14) Theorem 1. The result shows that differential operators are bounded operators when restricted to polynomials, and an upper on their norms is given.

Let X, Y be a real Banach space with norms $\|\cdot\|_X$, and $\|\cdot\|_Y$ respectively. A function $P : X \rightarrow Y$ is a homogeneous polynomial of degree $d \geq 0$ if $P(x) = L(x, \dots, x)$, where L is a linear symmetric map from $\underbrace{X \times \dots \times X}_d$ to Y (for $d = 0$, these are constant functions on X). A function $P : X \rightarrow Y$ is

a polynomial of degree d is $P = \sum_{j=0}^d P_j$, where P_j is a homogeneous polynomial of degree j . We let $\wp_d(X, Y)$ be the set of all polynomials $P : X \rightarrow Y$ with degree at most d . When $Y = \mathbb{R}$, we write $\wp_d(X)$. If $f : X \rightarrow Y$ has Frechet derivatives to order k we write $\nabla^{(k)} f(x)$ to denote its k -th order Frechet derivative, and for $z \in X$, $\nabla^{(k)} f(x) \cdot z^k \stackrel{\text{def}}{=} \nabla^{(k)} f(x)(z, \dots, z)$. Furthermore, we set

$$\|\nabla^{(k)} f(x)\| \stackrel{\text{def}}{=} \sup_{z \in X: \|z\|_X \leq 1} \|\nabla^{(k)} f(x)(z, \dots, z)\|_Y.$$

Let $T_d(t) = \cos(d \arccos(t))$, $t \in [-1, 1]$ denote the Chebyshev polynomial of degree d , and $T_d^{(j)}(t)$ its j -th order derivative. The following lemma is due to (14) Theorem 1.

Lemma A.6. For $P \in \wp_d(\mathbb{R}^m)$, and $k \geq 1$, let $P^{(k)}$ denote the k -th order derivative of P . We have

$$\sup_{x: \|x\|_2 \leq 1} \|P^{(k)}(x)\| \leq T_d^{(k)}(1) \times \sup_{x: \|x\|_2 \leq 1} |P(x)|,$$

where $\|P^{(k)}(x)\| \stackrel{\text{def}}{=} \sup_{y: \|y\|_2 \leq 1} |P^{(k)}(x)(y, \dots, y)|$.

We turn to the proof of Theorem 2.4. Let $\alpha = \bar{\alpha}$ as defined in the statement of the theorem, and fix $x_0 \in \text{int}_\alpha(\Omega)$. The function u has derivatives to the order $\lfloor \beta \rfloor$ at x_0 . By Taylor approximation we have, for all $z \in \mathbb{R}^m$, with $\|z\|_2 \leq 1$, setting $r \stackrel{\text{def}}{=} \lfloor \beta \rfloor - 1$,

$$\begin{aligned} u(x_0 + \alpha z) &= \\ u(x_0) + \alpha \nabla u(x_0) \cdot z + \dots + \frac{\alpha^r}{r!} \nabla^{(r)} u(x_0) \cdot z^r + \frac{\alpha^{\lfloor \beta \rfloor}}{r!} \int_0^1 (1-t)^r \nabla^{(\lfloor \beta \rfloor)} u(x_0 + \alpha t z) \cdot z^{\lfloor \beta \rfloor} dt \\ &= u(x_0) + \alpha \nabla u(x_0) \cdot z + \dots + \frac{\alpha^{\lfloor \beta \rfloor}}{\lfloor \beta \rfloor!} \nabla^{(\lfloor \beta \rfloor)} u(x_0) \cdot z^{\lfloor \beta \rfloor} + R_u(z), \end{aligned}$$

where, using the fact that $u \in C^\beta(\Omega, M)$, the remainder satisfies

$$\begin{aligned} |R_u(z)| &= \frac{\alpha^{\lfloor \beta \rfloor}}{r!} \left| \int_0^1 (1-t)^r \left(\nabla^{(\lfloor \beta \rfloor)} u(x_0 + \alpha t z) - \nabla^{(\lfloor \beta \rfloor)} u(x_0) \right) \cdot z^{\lfloor \beta \rfloor} dt \right| \\ &\leq \frac{\alpha^{\lfloor \beta \rfloor}}{r!} \int_0^1 (1-t)^r M(\alpha t)^{\beta - \lfloor \beta \rfloor} \|z\|_2^{\beta - \lfloor \beta \rfloor} \|z\|_\infty^{\lfloor \beta \rfloor} dt \leq \frac{M \alpha^\beta}{\lfloor \beta \rfloor!}. \end{aligned}$$

A similar expansion holds for \tilde{u} . And since $x_0 + \alpha z \in \Omega$, we have $|u(x_0 + \alpha z) - \tilde{u}(x_0 + \alpha z)| \leq \varepsilon$. Therefore, setting

$$P(z) \stackrel{\text{def}}{=} u(x_0) - \tilde{u}(x_0) + \alpha(\nabla u(x_0) - \nabla \tilde{u}(x_0)) \cdot z \\ + \cdots + \frac{\alpha^{\lfloor \beta \rfloor}}{\lfloor \beta \rfloor!} \left(\nabla^{(\lfloor \beta \rfloor)} u(x_0) - \nabla^{(\lfloor \beta \rfloor)} \tilde{u}(x_0) \right) \cdot z^{\lfloor \beta \rfloor},$$

we conclude that

$$\sup_{z \in \mathbb{R}^m: \|z\|_2 \leq 1} |P(z)| \leq \varepsilon + \frac{2M\alpha^\beta}{\lfloor \beta \rfloor!} \quad (\text{A.24})$$

We note that the function $z \mapsto P(z)$ is infinitely differentiable polynomial of degree $\lfloor \beta \rfloor$ on \mathbb{R}^m , and for $0 \leq \tau \leq \lfloor \beta \rfloor$, and $y \in \mathbb{R}^m$,

$$\nabla^{(\tau)} P(z) \cdot y^\tau = \sum_{j=\tau}^{\lfloor \beta \rfloor} \frac{\alpha^j}{(j-\tau)!} \left(\nabla u^{(j)}(x_0) - \nabla \tilde{u}^{(j)}(x_0) \right) \cdot \left(z^{j-\tau}, y^\tau \right). \quad (\text{A.25})$$

On the other hand, the function $x \mapsto \nabla^{(\tau)} u(x) : \Omega \mapsto \mathcal{L}(\mathbb{R}^m \times \cdots \times \mathbb{R}^m; \mathbb{R})$ has derivatives to order $\lfloor \beta \rfloor - \tau$ and $(\nabla^{(j)} \nabla^{(\tau)} u(x_0) \cdot y^j) \cdot x^\tau = \nabla^{(\tau+j)} u(x_0) \cdot (y^j, x^\tau)$. Setting $k = \lfloor \beta \rfloor - \tau$, its Taylor expansion at x_0 yields

$$\nabla^{(\tau)} u(x_0 + \alpha z) = \underbrace{\nabla^{(\tau)} u(x_0) + \alpha \nabla^{(\tau+1)} u(x_0) \cdot z + \cdots + \frac{\alpha^k}{k!} \nabla^{(\lfloor \beta \rfloor)} u(x_0) \cdot z^k}_{Q(z)} + R'(z),$$

and

$$\nabla^{(\tau)} \tilde{u}(x_0 + \alpha z) = \underbrace{\nabla^{(\tau)} \tilde{u}(x_0) + \alpha \nabla^{(\tau+1)} \tilde{u}(x_0) \cdot z + \cdots + \frac{\alpha^k}{k!} \nabla^{(\lfloor \beta \rfloor)} \tilde{u}(x_0) \cdot z^k}_{\tilde{Q}(z)} + \tilde{R}'(z).$$

Using the same calculations as above, we check that the remainders R' and \tilde{R}' satisfy

$$\sup_{z \in \mathbb{R}^d: \|z\|_2 \leq 1} \|R'(z)\| + \|\tilde{R}'(z)\| \leq \frac{2M\alpha^{\beta-\tau}}{(\lfloor \beta \rfloor - \tau)!}.$$

Using the last display, and noting that $\alpha^\tau(Q(z) - \tilde{Q}(z)) = \nabla^{(\tau)} P(z)$ given in (A.25), it follows that for all z in the unit ball,

$$\|\nabla^{(\tau)} u(x_0 + \alpha z) - \nabla^{(\tau)} \tilde{u}(x_0 + \alpha z)\| \leq \frac{1}{\alpha^\tau} \|\nabla^{(\tau)} P(z)\| + \frac{2M\alpha^{\beta-\tau}}{(\lfloor \beta \rfloor - \tau)!}.$$

Therefore, by Markov's polynomial inequality (Lemma A.6), for all z in the unit ball

$$\begin{aligned} \|\nabla^{(\tau)} u(x_0 + \alpha z) - \nabla^{(\tau)} \tilde{u}(x_0 + \alpha z)\| &\leq \frac{T_{\lfloor \beta \rfloor}^{(\tau)}(1)}{\alpha^\tau} \left(\varepsilon + \frac{2M\alpha^\beta}{\lfloor \beta \rfloor!} \right) + \frac{2M\alpha^{\beta-\tau}}{(\lfloor \beta \rfloor - \tau)!} \\ &= \frac{A}{\alpha^\tau} + B\alpha^{\beta-\tau}, \end{aligned}$$

where

$$A = \varepsilon T_{\lfloor \beta \rfloor}^{(\tau)}(1), \quad B = \frac{2M \times T_{\lfloor \beta \rfloor}^{(\tau)}(1)}{\lfloor \beta \rfloor!} + \frac{2M}{(\lfloor \beta \rfloor - \tau)!}.$$

With the choice

$$\alpha = \left(\frac{\tau A}{(\beta - \tau)B} \right)^{1/\beta} \leq \left(\frac{\tau \varepsilon \lfloor \beta \rfloor!}{2M(\beta - \tau)} \right)^{1/\beta},$$

we get

$$\|\nabla^{(\tau)} u(x_0) - \nabla^{(\tau)} \tilde{u}(x_0)\| \leq CM^{\frac{\tau}{\beta}} \varepsilon^{\frac{\beta-\tau}{\beta}},$$

for some constant C that depends only on τ and β . In fact, C can be taken as

$$C = \frac{\beta T_{\lfloor \beta \rfloor}^{(\tau)}(1)}{\beta - \tau} \left(\frac{\beta - \tau}{\tau} \right)^{\tau/\beta} \leq \frac{\beta}{\tau} T_{\lfloor \beta \rfloor}^{(\tau)}(1).$$

□

A.7. Proof of Theorem 2.3

We start with a lemma that establishes that the covering number of $\Gamma(\Theta)$ grows polynomially near 0, which consequently implies that $\Gamma(\Theta)$ has a finite VC dimension. This effectively proves that $\Gamma(\Theta)$ is a finite-dimensional subset of $C^\beta(\Omega)$, serving as a cornerstone of our main theorem.

Lemma A.7. *Given any $\varepsilon > 0$, we have:*

$$\mathcal{N}(\varepsilon, \Gamma(\Theta), L_\infty) \leq C \left(\frac{c}{\varepsilon} \right)^d,$$

for some constants $C, c > 0$.

Proof. Recall that Θ is a compact subset of \mathbb{R}^d , which implies $\text{vol}(\Theta) < \infty$. Therefore, from the standard covering number calculation, we have:

$$\mathcal{N}(\varepsilon, \Theta, L_2) \leq \frac{\text{vol}(\Theta)}{\text{vol}(B)} \left(\frac{3}{\varepsilon} \right)^d \triangleq C \left(\frac{3}{\varepsilon} \right)^d,$$

where B is the unit ball (centered at the origin) in \mathbb{R}^d (e.g., see Theorem 14.2 of Yihong Wu's lecture notes, ECE598, Spring 2016). We next claim that:

$$\mathcal{N}(\varepsilon, \Gamma(\Theta), L_\infty) \leq \mathcal{N}(\varepsilon/L, \Theta, L_2).$$

which will complete the proof. To show this, suppose $\{\theta_1, \dots, \theta_n\}$ are the center of (ϵ/L) -covers of Θ , i.e.

$$\sup_{\theta \in \Theta} \min_{1 \leq j \leq n} \|\theta_j - \theta\|_2 \leq \epsilon/L.$$

where $n = \mathcal{N}(\epsilon/L, \Theta, L_2)$. Now consider the images corresponding to these θ 's, meaning, $\{\gamma_1, \dots, \gamma_n\}$ where $\gamma_i = \Gamma(\theta_i)$. Then we have:

$$\sup_{\theta \in \Theta} \min_{1 \leq j \leq n} \|\Gamma(\theta_j) - \Gamma(\theta)\|_\infty \leq L \sup_{\theta \in \Theta} \min_{1 \leq j \leq n} \|\theta_j - \theta\|_2 \leq L(\epsilon/L) = \epsilon.$$

Hence, the proof of the lemma is complete with $c = 3L$. \square

Proof of Theorem 2.3. We use the standard rate theorem (e.g., see 3.2.5 of (35)) to establish the convergence rate of $|\hat{u} - u_\star|_2^2$. Define, for simplicity, $M(u) = \mathbb{E}[(Y - u(\mathbf{s}))^2]$ and $M_n(u) = (1/n) \sum_i (Y_i - u(\mathbf{s}_i))^2$. Therefore, from (2.14) we have \hat{u} is the minimizer of $M_n(u)$ over $\Gamma(\Theta)$. Furthermore, we have:

$$M(u) - M(u_\star) = |u - u_\star|_2^2.$$

To establish the rate of convergence, we need to find the modulus of continuity function $\phi_n(\delta)$ that satisfies:

$$\mathbb{E} \left[\sup_{|u - u_\star|_2 \leq \delta} |(M_n(u) - M(u)) - (M_n(u_\star) - M(u_\star))| \right] \leq \frac{\phi_n(\delta)}{\sqrt{n}}.$$

Let $\xi_i \stackrel{\text{def}}{=} (Y_i - u_\star(\mathbf{s}_i))/\sigma$. Some simple algebra yields:

$$\begin{aligned} & |(M_n(u) - M(u)) - (M_n(u_\star) - M(u_\star))| \\ &= \left| (\mathbb{P}_n - \mathbb{P})(u(\mathbf{s}) - u_\star(\mathbf{s}))^2 + \sigma(\mathbb{P}_n - \mathbb{P})\xi(u(\mathbf{s}) - u_\star(\mathbf{s})) \right| \\ &\leq \left| (\mathbb{P}_n - \mathbb{P})(u(\mathbf{s}) - u_\star(\mathbf{s}))^2 \right| + \sigma \left| (\mathbb{P}_n - \mathbb{P})\xi(u(\mathbf{s}) - u_\star(\mathbf{s})) \right| \end{aligned}$$

Therefore, by symmetrization we have (using (η_1, \dots, η_n) i.i.d Rademacher random variables):

$$\begin{aligned} & \mathbb{E} \left[\sup_{|u - u_\star|_2 \leq \delta} |(M_n(u) - M(u)) - (M_n(u_\star) - M(u_\star))| \right] \\ &\leq \mathbb{E} \left[\sup_{|u - u_\star|_2 \leq \delta} \left| (\mathbb{P}_n - \mathbb{P})(u(\mathbf{s}) - u_\star(\mathbf{s}))^2 \right| \right] + \sigma \mathbb{E} \left[\sup_{|u - u_\star|_2 \leq \delta} |(\mathbb{P}_n - \mathbb{P})\xi(u(\mathbf{s}) - u_\star(\mathbf{s}))| \right] \\ &\leq 2\mathbb{E} \left[\sup_{|u - u_\star|_2 \leq \delta} \left| \frac{1}{n} \sum_i \eta_i (u(\mathbf{s}_i) - u_\star(\mathbf{s}_i))^2 \right| \right] + \sigma \mathbb{E} \left[\sup_{|u - u_\star|_2 \leq \delta} \left| \frac{1}{n} \sum_i \xi_i (u(\mathbf{s}_i) - u_\star(\mathbf{s}_i)) \right| \right] \\ &\leq 2B\mathbb{E} \left[\sup_{|u - u_\star|_2 \leq \delta} \left| \frac{1}{n} \sum_i \eta_i (u(\mathbf{s}_i) - u_\star(\mathbf{s}_i)) \right| \right] + \sigma \mathbb{E} \left[\sup_{|u - u_\star|_2 \leq \delta} \left| \frac{1}{n} \sum_i \xi_i (u(\mathbf{s}_i) - u_\star(\mathbf{s}_i)) \right| \right] \\ &\leq C_1 \left(\delta \sqrt{\frac{d}{n} \log \left(\frac{C_2}{\delta} \right)} + \frac{d}{n} \log \left(\frac{C_2}{\delta} \right) \right) \end{aligned}$$

Here, in the second last inequality, we use Ledoux-Talagrand contraction inequality. The last line follows Dudley's chaining bound and some standard calculations (e.g., proof of Theorem 8.7 of (28)). Therefore, a valid choice of $\phi_n(\delta)$ is:

$$\phi_n(\delta) = C_1 \left(\delta \sqrt{d \log \left(\frac{C_2}{\delta} \right)} + \frac{d}{\sqrt{n}} \log \left(\frac{C_2}{\delta} \right) \right).$$

Therefore, we conclude that $\|\hat{u} - u_\star\|_2 = O_p(\delta_n)$ where δ_n satisfies:

$$\sqrt{n} \delta_n^2 \geq \phi_n(\delta_n).$$

Some simple algebra yields that a valid choice for δ_n is $\sqrt{(d/n) \log(n/d)}$, which completes the proof. \square

References

- [1] ANDO, T. and VAN HEMMEN, J. L. (1980). An inequality for trace ideals. *Communications in Mathematical Physics* **76** 143 – 148.
- [2] ATCHADE, Y. and BHATTACHARYYA, A. (2019). An approach to large-scale Quasi-Bayesian inference with spike-and-slab priors.
- [3] ATCHADE, Y. and WANG, L. (2023). A fast asynchronous Markov chain Monte Carlo sampler for sparse Bayesian inference. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **85** 1492-1516. <https://doi.org/10.1093/jrssi/bqad078>
- [4] BELOMESTNY, D., NAUMOV, A., PUCHKIN, N. and SAMSONOV, S. (2022). Simultaneous approximation of a smooth function and its derivatives by deep neural networks with piecewise-polynomial activations.
- [5] BHATIA, R., JAIN, T. and LIM, Y. (2019). On the Bures-Wasserstein distance between positive definite matrices. *Expositiones Mathematicae* **37** 165-191. <https://doi.org/10.1016/j.exmath.2018.01.002>
- [6] BOTH, G.-J., CHOUDHURY, S., SENS, P. and KUSTERS, R. (2021). DeepMoD: Deep learning for model discovery in noisy data. *Journal of Computational Physics* **428** 109985. <https://doi.org/10.1016/j.jcp.2020.109985>
- [7] CHEN, Z., LIU, Y. and SUN, H. (2020). Physics-informed learning of governing equations from scarce data. *Nature Communications* **12**.
- [8] CUOMO, S., COLA, V. S. D., GIAMPAOLO, F., ROZZA, G., RAISSI, M. and PICCIALI, F. (2022). Scientific Machine Learning Through Physics-Informed Neural Networks: Where we are and What's Next. *Journal of Scientific Computing* **92**.
- [9] DE HOOP, M. V., KOVACHKI, N. B., NELSEN, N. H. and STUART, A. M. (2023). Convergence Rates for Learning Linear Operators from Noisy Data. *SIAM/ASA Journal on Uncertainty Quantification* **11** 480-513.
- [10] DEVROYE, L., MEHRABIAN, A. and REDDAD, T. (2023). The total variation distance between high-dimensional Gaussians with the same mean.
- [11] FLAMARY, R., COURTY, N., GRAMFORT, A., ALAYA, M. Z., BOISBUNON, A., CHAMBON, S., CHAPEL, L., CORENFLOS, A., FATRAS, K., FOURNIER, N., GAUTHERON, L., GAYRAUD, N. T. H., JANATI, H., RAKOTOMAMONJY, A., REDKO, I., ROLET, A., SCHUTZ, A., SEGUY, V., SUTHERLAND, D. J., TAVENARD, R., TONG, A. and VAYER, T. (2021). POT: Python Optimal Transport. *Journal of Machine Learning Research* **22** 1-8.
- [12] GEER, S. A. (2000). *Empirical Processes in M-estimation* **6**. Cambridge university press.
- [13] GIORDANO, M. and KEKKONEN, H. (2020). Bernstein–von Mises Theorems and Uncertainty Quantification for Linear Inverse Problems. *SIAM/ASA Journal on Uncertainty Quantification* **8** 342-373. <https://doi.org/10.1137/18M1226269>

-
- [14] HARRIS, L. (2008). Multivariate Markov polynomial inequalities and Chebyshev nodes. *Journal of Mathematical Analysis and Applications* **338** 350-357. <https://doi.org/10.1016/j.jmaa.2007.05.044>
- [15] HWANG, C.-R. (1980). Laplace's Method Revisited: Weak Convergence of Probability Measures. *The Annals of Probability* **8** 1177 – 1182. <https://doi.org/10.1214/aop/1176994579>
- [16] LU, J., SHEN, Z., YANG, H. and ZHANG, S. (2021). Deep Network Approximation for Smooth Functions. *SIAM Journal on Mathematical Analysis* **53** 5465-5506. <https://doi.org/10.1137/20M134695X>
- [17] LU, Y., CHEN, H., LU, J., YING, L. and BLANCHET, J. (2022). Machine Learning For Elliptic PDEs: Fast Rate Generalization Bound, Neural Scaling Law and Minimax Optimality. In *International Conference on Learning Representations*.
- [18] MCGOFF, K., MUKHERJEE, S. and PILLAI, N. (2015). Statistical inference for dynamical systems: A review. *Statistics Surveys* **9** 209 – 252.
- [19] MISHRA, S. and MOLINARO, R. (2021). Estimates on the generalization error of physics-informed neural networks for approximating a class of inverse problems for PDEs. *IMA Journal of Numerical Analysis* **42** 981-1022.
- [20] MISHRA, S. and MOLINARO, R. (2022). Estimates on the generalization error of physics-informed neural networks for approximating PDEs. *IMA Journal of Numerical Analysis* **43** 1-43.
- [21] MONARD, F., NICKL, R. and PATERNAIN, G. P. (2021). Statistical guarantees for Bayesian uncertainty quantification in nonlinear inverse problems with Gaussian process priors. *The Annals of Statistics* **49** 3255 – 3298. <https://doi.org/10.1214/21-AOS2082>
- [22] NICKL, R., VAN DE GEER, S. and WANG, S. (2020). Convergence rates for penalized least squares estimators in PDE constrained regression problems. *SIAM/ASA Journal on Uncertainty Quantification* **8** 374–413.
- [23] PARDO, L. (2006). *Statistical Inference Based on Divergence Measures*. *Statistics: Textbooks and Monographs* **185**. Chapman & Hall/CRC, Boca Raton, FL.
- [24] RAISSI, M., PERDIKARIS, P. and KARNIADAKIS, G. E. (2019). 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. <https://doi.org/10.1016/j.jcp.2018.10.045>
- [25] RAMSAY, J. O., HOOKER, G., CAMPBELL, D. and CAO, J. (2007). Parameter Estimation for Differential Equations: A Generalized Smoothing Approach. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* **69** 741–796.
- [26] RASMUSSEN, C. E. and WILLIAMS, C. K. I. (2005). *Gaussian Processes for Machine Learning*. *Adaptive Computation and Machine Learning*. The MIT Press, Cambridge, MA. <https://doi.org/10.7551/mitpress/3206.001.0001>
- [27] SCHMIDT-HIEBER, J. (2020). Nonparametric regression using deep neural networks with ReLU activation function. *Annals of Statistics* **48**. <https://doi.org/10.1214/19-AOS1875>
- [28] SEN, B. (2018). A gentle introduction to empirical process theory and applications. *Lecture Notes, Columbia University* **11** 28–29.
- [29] SHIN, Y., ZHANG, Z. and KARNIADAKIS, G. E. (2023). ERROR ESTIMATES OF RESIDUAL MINIMIZATION USING NEURAL NETWORKS FOR LINEAR PDES. *Journal of Machine Learning for Modeling and Computing* **4** 73–101.
- [30] SIRIGNANO, J. and SPILIOPOULOS, K. (2018). DGM: A deep learning algorithm for solving partial differential equations. *Journal of computational physics* **375** 1339–1364.
- [31] SPOKOINY, V. (2020). Bayesian inference for nonlinear inverse problems.
- [32] STONE, C. J. (1982). Optimal Global Rates of Convergence for Nonparametric Regression. *The Annals of Statistics* **10** 1040 – 1053. <https://doi.org/10.1214/aos/1176345969>
- [33] SUN, L., GAO, H., PAN, S. and WANG, J.-X. (2020). Surrogate modeling for fluid flows based on physics-constrained deep learning without simulation data. *Computer Methods in Applied Mechanics and Engineering* **361** 112732.
- [34] TAHERI, M., XIE, F. and LEDERER, J. (2021). Statistical guarantees for regularized neural networks. *Neural Networks* **142** 148-161.
- [35] VAART, A. V. D. and WELLNER, J. A. (1997). Weak convergence and empirical processes with applications to statistics. *Journal of the Royal Statistical Society-Series A Statistics in Society* **160** 596–608.
- [36] XUN, X., CAO, J., MALLICK, B., MAITY, A. and CARROLL, R. J. (2013). Parameter Estimation of Partial Differential Equation Models. *Journal of the American Statistical Association* **108** 1009–1020.
-

- [37] YANG, L., MENG, X. and KARNIADAKIS, G. E. (2021). B-PINNs: Bayesian physics-informed neural networks for forward and inverse PDE problems with noisy data. *Journal of Computational Physics* **425** 109913.
 - [38] YAROTSKY, D. (2017). Error bounds for approximations with deep ReLU networks. *Neural Networks* **94** 103–114.
 - [39] YAROTSKY, D. and ZHEVNERCHUK, A. (2020). The phase diagram of approximation rates for deep neural networks. In *Proceedings of the 34th International Conference on Neural Information Processing Systems. NIPS'20*. Curran Associates Inc.
 - [40] YU, B. et al. (2018). The deep Ritz method: a deep learning-based numerical algorithm for solving variational problems. *Communications in Mathematics and Statistics* **6** 1–12.
-