

Error Estimates for the Deep Ritz Method with Boundary Penalty

Johannes Müller

Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany

JMUELLER@MIS.MPG.DE

Marius Zeinhofer

Simula Research Laboratory, Oslo, Norway

MARIUSZ@SIMULA.NO

Editors: Bin Dong, Qianxiao Li, Lei Wang, Zhi-Qin John Xu

Abstract

We estimate the error of the Deep Ritz Method for linear elliptic equations. For Dirichlet boundary conditions, we estimate the error when the boundary values are imposed through the boundary penalty method. Our results apply to arbitrary sets of ansatz functions and estimate the error in dependence of the optimization accuracy, the approximation capabilities of the ansatz class and – in the case of Dirichlet boundary values – the penalization strength λ . To the best of our knowledge, our results are presently the only ones in the literature that treat the case of Dirichlet boundary conditions in full generality, i.e., without a lower order term that leads to coercivity on all of $H^1(\Omega)$. Further, we discuss the implications of our results for ansatz classes which are given through ReLU networks and the relation to existing estimates for finite element functions. For high dimensional problems our results show that the favourable approximation capabilities of neural networks for smooth functions are inherited by the Deep Ritz Method.

Keywords: Variational Problems, Deep Ritz Method, Boundary Penalty, High Dimensional PDEs

1. Introduction

Utilizing artificial neural networks as ansatz classes for the solution of PDEs has recently gained interest due to its potential for parametric families of PDEs [Li et al. \(2020\)](#), inverse or data enhanced problems [Zhang et al. \(2018\)](#); [Zhu et al. \(2019\)](#) and the solution of PDEs in high spatial dimensions [Han et al. \(2018, 2017\)](#); [Jentzen et al. \(2018\)](#). Popular approaches are physics informed neural networks [Raissi et al. \(2019\)](#), neural operator methods [Li et al. \(2020\)](#) and the Deep Ritz Method [E and Yu \(2018\)](#).

Classical numerical schemes for PDEs, such as FEM or FD, are usually designed to solve one PDE instance at a time. In contrast, neural network based approaches allow to simultaneously solve parametric problems. Examples range from parametric geometries to simultaneously solving equations for different material properties or boundary conditions. This allows for a very effective design space exploration and is of great benefit in engineering applications [Hennigh et al. \(2021\)](#).

Further, neural network based approaches to inverse problems, i.e., when certain parameters of the PDE are to be inferred from measured data, have shown promising results recently [Hennigh et al. \(2021\)](#); [Raissi et al. \(2019\)](#). For an effective software realization of the solution of the aforementioned problems we refer to Nvidia Modulus [Hennigh et al. \(2021\)](#), a software library built on TensorFlow [Abadi et al. \(2016\)](#) to enable seamless integration in engineering workflows.

Another surge of interest for neural network ansatz classes for the solution of PDEs stems from their ability to approximate high dimensional functions effectively, possibly breaking the curse of dimensionality [Weinan et al. \(2019\)](#); [Wojtowytsch et al. \(2020\)](#); [Jentzen et al. \(2018\)](#). Consequently, a line of research focuses on the capabilities of neural networks to approximate solutions of high dimensional PDEs breaking or mitigating the curse of dimensionality.

1.1. Main Contribution and Related Work

In this manuscript, we investigate the error of the Deep Ritz Method for linear elliptic equations in terms of the optimization quality, the expressiveness of the ansatz class and – in the case of Dirichlet boundary values – the penalty parameter. The class of linear elliptic PDEs is pervasive and estimating the error of the Deep Ritz Method for these PDEs constitutes an important first step before considering more complex situations, such as inverse and parametric problems.

As a model problem we consider the equation posed on $\Omega \subset \mathbb{R}^d$

$$\begin{aligned} -\operatorname{div}(A\nabla u) &= f && \text{in } \Omega \\ u &= 0 && \text{on } \partial\Omega, \end{aligned} \tag{1}$$

where $A \in L^\infty(\Omega, \mathbb{R}^{d \times d})$ is symmetric and elliptic and $f \in L^2(\Omega)$, although our method is general enough to handle a larger class of problems. In particular, our analysis covers also non-essential boundary value problems. Using a penalty parameter $\lambda > 0$ and neural network functions u_θ , where $\theta \in \Theta$ denotes the parameters of the network, the Deep Ritz formulation of (1) is the problem of (approximately) finding

$$\theta^* \in \operatorname{argmin}_{\theta \in \Theta} \left[\frac{1}{2} \int_{\Omega} A \nabla u_\theta \cdot \nabla u_\theta dx + \lambda \int_{\partial\Omega} u_\theta^2 ds - \int_{\Omega} f u_\theta dx \right]. \tag{2}$$

There are recent results in the literature that estimate the error of the Deep Ritz Method for elliptic equations in certain cases [Xu \(2020\)](#); [Jiao et al. \(2021\)](#); [Duan et al. \(2021a\)](#). Common in all of them is that they circumvent a major technical difficulty by not considering Dirichlet boundary conditions for general elliptic operators. Usually, Dirichlet boundary conditions are imposed by the boundary penalty method, compare to the formulation (2) and [E and Yu \(2018\)](#). The works [Xu \(2020\)](#); [Jiao et al. \(2021\)](#); [Duan et al. \(2021a\)](#) consider Dirichlet boundary conditions in the Deep Ritz Method via penalization only in the setting of a coercive operator on all of $H^1(\Omega)$ ¹. Thus, the effect of the boundary penalty in the manuscripts [Xu \(2020\)](#); [Jiao et al. \(2021\)](#); [Duan et al. \(2021a\)](#) can immediately be deduced by the results from the classical paper of [Maury \(2009\)](#). One of our main contributions (see [Theorem 9](#)) is to extend the results of [Maury \(2009\)](#) and to relax the coercivity assumption on the operator in the boundary penalty method, allowing operators that are coercive only on the space $H_0^1(\Omega)$. To the best of our knowledge, our result is presently the only one available in the literature that includes this case.

We then discuss the error decay rates for the boundary penalty method in dependence of the ansatz class and find that even for optimal penalization they might differ from the approximation rate (see [Section 4](#), [Theorem 19](#)). Finally note that in this manuscript, we do not take the numerical evaluations of the appearing integrals into account, however, this can be done exactly as in the paper of [Mishra and Molinaro \(2020\)](#).

With respect to high dimensional PDEs our results reveal – when combined with suitable approximation results such as [Gühring et al. \(2019\)](#), see also [Theorem 3](#) – that solving elliptic PDEs by the Deep Ritz Method retains the favorable approximation capabilities of neural networks for smooth (in the sense of Sobolev regularity), high dimensional functions. More precisely, the mitigation of the curse of dimensionality through sufficient smoothness of the approximating function can be transferred, also in the case of the boundary penalty method for Dirichlet boundary conditions, compare to [Remark 7](#) and [20](#). We stress that our results do not show a dimension independent

1. Note that this excludes the prototypical Poisson problem $-\Delta u = f$ with $u|_{\partial\Omega} = 0$.

error decay. However, this cannot be expected in this generality as we do not assume any special regularity (such as Barron regularity) or lower dimensionality conditions on the data of the PDEs. Note, that also in the important works of [Jentzen et al. \(2018\)](#) on breaking the curse of dimensionality for certain parabolic PDEs always low-dimensionality assumptions on the initial data and the PDE coefficients are present. In contrast, our results reveal that sufficient (dimension dependent) smoothness can serve as such an assumption for the application of the Deep Ritz Method to elliptic equations.

Elliptic Equations in high spatial dimensions are important as building blocks for parabolic equations which have well known applications in high dimensions, for instance in finance [Han et al. \(2017\)](#). For example, using a minimizing movement scheme, an elliptic problem (with an additional term corresponding to dissipation) needs to be solved for every time step [Hwang et al. \(2021\)](#) in order to solve a parabolic equation. Further, elliptic equations arise naturally as equilibrium states for parabolic equations.

2. Preliminaries

In this section we give an overview of the notations and concepts that we use in the remainder.

2.1. Sobolev spaces and Friedrich's inequality

We denote the space of functions on $\Omega \subseteq \mathbb{R}^d$ that are integrable in p -th power by $L^p(\Omega)$ and endow it with its canonical norm. We denote the subspace of $L^p(\Omega)$ of functions with weak derivatives up to order k in $L^p(\Omega)$ by $W^{k,p}(\Omega)$, which is a Banach space with the norm

$$\|u\|_{W^{k,p}(\Omega)}^p := \sum_{l=0}^k \|D^l u\|_{L^p(\Omega)}^p.$$

This space is called a *Sobolev space* and we denote its dual by $W^{k,p}(\Omega)^*$. The closure of all compactly supported smooth functions $C_c^\infty(\Omega)$ in $W^{k,p}(\Omega)$ is denoted by $W_0^{k,p}(\Omega)$. It is well known that for Lipschitz domains Ω , the operator restricting a Lipschitz continuous function defined on $\overline{\Omega}$ to the boundary admits a linear and bounded extension $\text{tr}: W^{1,p}(\Omega) \rightarrow L^p(\partial\Omega)$, called the *trace operator*. Further, we write $\|u\|_{L^p(\partial\Omega)}$ whenever we mean $\|\text{tr}(u)\|_{L^p(\partial\Omega)}$. In the following we mostly work with the case $p = 2$ and write $H_{(0)}^k(\Omega)$ instead of $W_{(0)}^{k,2}(\Omega)$.

Our analysis of the boundary penalty method relies on the Friedrich inequality which states that the $L^2(\Omega)$ norm of a function can be estimated by the norm of its gradient and boundary values.

Proposition 1 (Friedrich's inequality) *Let $\Omega \subseteq \mathbb{R}^d$ be a bounded and open set with Lipschitz boundary $\partial\Omega$. Then there exists a constant $c_F > 0$ such that*

$$\|u\|_{H^1(\Omega)} \leq c_F^2 \cdot \left(\|\nabla u\|_{L^2(\Omega)}^2 + \|u\|_{L^2(\partial\Omega)}^2 \right) \quad \text{for all } u \in H^1(\Omega). \quad (3)$$

In particular this implies that $(\|\nabla u\|_{L^2(\Omega)}^2 + \|u\|_{L^2(\partial\Omega)}^2)^{1/2}$ defines an equivalent norm on $H^1(\Omega)$. We refer to the optimal constant c_F as the *Friedrich constant* of Ω .

2.2. Preliminaries from neural network theory

Consider natural numbers d, m, L, N_0, \dots, N_L and let $\theta = ((A_1, b_1), \dots, (A_L, b_L))$ be a tuple of matrix-vector pairs where $A_l \in \mathbb{R}^{N_l \times N_{l-1}}, b_l \in \mathbb{R}^{N_l}$ and $N_0 = d, N_L = m$. Every matrix vector pair (A_l, b_l) induces an affine linear map $T_l: \mathbb{R}^{N_{l-1}} \rightarrow \mathbb{R}^{N_l}$. The *neural network function with parameters* θ and with respect to some *activation function* $\rho: \mathbb{R} \rightarrow \mathbb{R}$ is the function

$$u_\theta^\rho: \mathbb{R}^d \rightarrow \mathbb{R}^m, \quad x \mapsto T_L(\rho(T_{L-1}(\rho(\dots \rho(T_1(x)))))).$$

We refer to the parametric family of functions $\mathcal{F}_\Theta^\rho := \{u_\theta^\rho \mid \theta \in \Theta\}$ as a ρ -network, where Θ denotes a family of parameters and drop the superscript ρ from now on. If we have $f = u_\theta$ for some $\theta \in \Theta$ we say the function f can be *expressed* by the neural network. We call d the *input* and m the *output dimension*, L the *depth* and $W = \max_{l=0, \dots, L} N_l$ the *width* of the neural network. We call a network *shallow* if it has depth 2 and *deep* otherwise. The *number of parameters* and the *number of neurons* of such a network is given by $\dim(\Theta)$ and $\sum_{l=0}^L N_l$. In the remainder, we restrict ourselves to the case $m = 1$ since we only consider real valued functions.

The *rectified linear unit* or *ReLU activation function* is defined via $x \mapsto \max\{0, x\}$ and we call networks with this particular activation function *ReLU networks*. The class of ReLU networks coincides with the class of continuous and piecewise linear functions, which yields the following result [Arora et al. \(2016\)](#); [Dondl et al. \(2021\)](#).

Theorem 2 (Universal approximation with zero boundary values) *Consider an open set $\Omega \subseteq \mathbb{R}^d$ and $u \in W_0^{1,p}(\Omega)$ with $p \in [1, \infty)$. Then for all $\varepsilon > 0$ there exists $u_\varepsilon \in W_0^{1,p}(\Omega)$ that can be expressed by a ReLU network of depth $\lceil \log_2(d+1) \rceil + 1$ such that $\|u - u_\varepsilon\|_{W^{1,p}(\Omega)} \leq \varepsilon$.*

The difference of this result to other universal approximation results [Hornik \(1991\)](#); [Czarnecki et al. \(2017\)](#) is the approximating neural network function are guaranteed to have zero boundary values. This is a special property of the ReLU activation function and implies the consistency of the boundary penalty method for arbitrary penalization strengths as we will see later. In order to quantify the error that is being made by the variational training of ReLU networks with boundary penalty, we use the following result from [Gühring et al. \(2019\)](#), where other results on approximation bounds in Sobolev spaces have been obtained in [Gühring and Raslan \(2020\)](#); [Xu \(2020\)](#); [Siegel and Xu \(2020a,b\)](#); [Hon and Yang \(2021\)](#); [Duan et al. \(2021b\)](#); [De Ryck et al. \(2021\)](#).

Theorem 3 (Quantitative universal approximation) *Let $\Omega \subseteq \mathbb{R}^d$ be a bounded and open set with Lipschitz regular boundary², let $k \in (1, \infty)$, $p \in [1, \infty]$ and fix an arbitrary function $u \in W^{k,p}(\Omega)$. Then, for every $n \in \mathbb{N}$, there is a ReLU network u_n with $\mathcal{O}(\log_2^2(n^{k/d}) \cdot n)$ many parameters and neurons such that*

$$\|u - u_n\|_{W^{s,p}(\Omega)} \leq c(s) \cdot \|u\|_{W^{k,p}(\Omega)} \cdot n^{-(k-s)/d}$$

for every $s \in [0, 1]$.

3. Error estimates for the boundary penalty method

In this section we bound the distance of a function $v \in V$ to the solution u^* of the variational problem in dependence of the approximation capabilities of the ansatz class V , the optimization quality and the penalization strength. We begin by stating a general curvature based bound which can be seen as a nonlinear version of Céa's Lemma.

2. or more generally, that Ω is a Sobolev extension domain

3.1. A Céa Lemma

The following proof of Céa's Lemma is based on the curvature properties of a quadratic, coercive energy defined on a Hilbert space. Note that in the following proposition, V does not need to be a vector space.

Proposition 4 (Céa's Lemma) *Let X be a Hilbert space, $V \subseteq X$ any subset and $a: X \times X \rightarrow \mathbb{R}$ a symmetric, continuous and α -coercive bilinear form. For $f \in X^*$ define the quadratic energy $E(u) := \frac{1}{2}a(u, u) - f(u)$ and denote its unique minimiser by u^* . Then for every $v \in V$ it holds that*

$$\|v - u^*\|_X \leq \sqrt{\frac{2\delta}{\alpha} + \frac{1}{\alpha} \inf_{\tilde{v} \in V} \|\tilde{v} - u^*\|_a^2},$$

where $\delta = E(v) - \inf_{\tilde{v} \in V} E(\tilde{v})$ and $\|u\|_a^2 := a(u, u)$ is the norm induced by a .

Proof As E is quadratic it can be exactly expanded using Taylor's formula. Hence, for every $h \in X$ it holds that

$$E(u + h) = E(u^*) + \frac{1}{2}D^2E(u^*)(h, h) = E(u^*) + \frac{1}{2}a(h, h) = E(u^*) + \frac{1}{2}\|h\|_a^2,$$

where we used $DE(u^*) = 0$. Inserting $v - u^*$ for h we obtain

$$E(v) - E(u^*) = \frac{1}{2}a(v - u^*, v - u^*) \geq \frac{\alpha}{2}\|v - u^*\|_X^2.$$

On the other hand we compute

$$\begin{aligned} E(v) - E(u^*) &= E(v) - \inf_{\tilde{v} \in V} E(\tilde{v}) + \inf_{\tilde{v} \in V} (E(\tilde{v}) - E(u^*)) \\ &= \delta + \frac{1}{2} \inf_{\tilde{v} \in V} \|\tilde{v} - u^*\|_a^2. \end{aligned}$$

Combining the two estimates and rearranging terms yields the assertion. ■

Proposition 4 is all we need to derive error estimates for coercive problems with non-essential boundary conditions. We give an example.

Corollary 5 (Neumann Problem) *Let $\Omega \subset \mathbb{R}^d$ be a bounded Lipschitz domain and let f be a fixed member of $H^1(\Omega)^*$. Denote by $u \in H^1(\Omega)$ the weak solution to*

$$-\Delta u + u = f \quad \text{in } H^1(\Omega)^*.$$

Let Θ be the parameter set of a neural network architecture such that $u_\theta \in H^1(\Omega)$ for every $\theta \in \Theta$. Then for every $\theta \in \Theta$ it holds

$$\|u_\theta - u\|_{H^1(\Omega)} \leq \sqrt{2\delta + \inf_{\eta \in \Theta} \|u_\eta - u\|_{H^1(\Omega)}^2}$$

where

$$\delta = \|u_\theta\|_{H^1(\Omega)}^2 - f(u_\theta) - \inf_{\eta \in \Theta} \left[\|u_\eta\|_{H^1(\Omega)}^2 - f(u_\eta) \right].$$

Proof The bilinear form corresponding to the above Neumann problem is

$$a: H^1(\Omega) \times H^1(\Omega) \rightarrow \mathbb{R}, \quad a(u, v) = \int_{\Omega} \nabla u \nabla v + uv dx$$

and therefore its coercivity constant is $\alpha = 1$ and the associated norm $\|\cdot\|_a$ is the natural one on $H^1(\Omega)$. Employing Proposition 4 yields the assertion. \blacksquare

Remark 6 Corollary 5 yields $H^1(\Omega)$ convergence of the Deep Ritz Method provided the ansatz class possesses universal approximation properties with respect to the $H^1(\Omega)$ norm. This is of course also a necessary requirement and fulfilled by a wide class of network architectures and activation functions, see [Hornik \(1991\)](#); [Czarnecki et al. \(2017\)](#) or for approximation rates [Gühring et al. \(2019\)](#). We stress that any (quantitative) universal approximation theorem for Sobolev topologies can be combined with the above result, such as Theorem 3 for ReLU neural networks.

Furthermore, the form of the differential equation in the above corollary can easily be generalized. One can for example consider general second order elliptic PDEs in divergence form with non-essential boundary conditions as long as these are coercive and can be derived from a minimization principle.

Remark 7 (Dimension Dependence and Adaptation to Smoothness) Assume the solution u to the Neumann problem is a member of $H^k(\Omega)$ for some $k > 1$. Then applying the quantitative universal approximation Theorem 3 we estimate

$$\|u_{\theta} - u\|_{H^1(\Omega)} \lesssim \sqrt{2\delta + c \cdot \|u\|_{H^k(\Omega)} n^{-(k-1)/d}},$$

where $d \in \mathbb{N}$ is the spatial dimension. While this estimate is not dimension independent, it indicates how smoothness mitigates the deterioration of error decay rates for high dimensions. We see that the merit of neural networks to achieve approximation rates increasing with the smoothness of the target function carries over to the error decay in the deep Ritz method. In contrast, to achieve the approximation rate and an error decay rate of $(k-1)/d$ with finite elements one needs to for example use P^{k-1} elements [Ern and Guermond \(2004\)](#), which complicates the ansatz class and therefore the approach.

Remark 8 [Practical Realization of Rates] There is a gap between the theory and the practice of neural network based methods for the solution of PDEs. Error decay rates, as predicted by our results cannot be observed in practice due to the difficulties of computing minima of non-convex functions. In practice, one observes moderate errors that don't decrease beyond a certain accuracy when the numbers of the parameters of the neural network ansatz architecture are increased. We refer to [Lee et al. \(2021\)](#) and the references therein for a more detailed description of this phenomenon.

However, what one observes is that neural network based methods in general and the Deep Ritz Method in particular are well suited for problems in high spatial dimensions or a high dimensional parameter space. Practical evidence can already be found in the paper introducing the Deep Ritz Method, see [E and Yu \(2018\)](#) and further high dimensional examples – even of industrial scale – can be found in [Hennigh et al. \(2021\)](#). We propose to view our results as a qualitative explanation of these observations.

3.2. An error estimate for the boundary penalty method

The treatment of Dirichlet boundary conditions corresponds to a constrained optimization problem, as in standard neural network architectures zero boundary values cannot be directly encoded. We use the boundary penalty method as a way to enforce Dirichlet boundary conditions. For ease of presentation, we discuss our approach for the concrete equation

$$\begin{aligned} -\operatorname{div}(A\nabla u) &= f \quad \text{in } \Omega \\ u &= 0 \quad \text{on } \partial\Omega, \end{aligned} \tag{4}$$

where $A \in L^\infty(\Omega, \mathbb{R}^{d \times d})$ is a symmetric and elliptic coefficient matrix. The weak formulation of this equation gives rise to the bilinear form

$$a: H^1(\Omega) \times H^1(\Omega) \rightarrow \mathbb{R}, \quad a(u, v) = \int_{\Omega} A\nabla u \cdot \nabla v dx$$

and the energy

$$E: H^1(\Omega) \rightarrow \mathbb{R}, \quad E(u) = \frac{1}{2}a(u, u) - f(u)$$

where $f \in H^1(\Omega)^*$. Using the boundary penalty method as an approximation for (4) leads to the bilinear form

$$a_\lambda: H^1(\Omega) \times H^1(\Omega) \rightarrow \mathbb{R}, \quad a_\lambda(u, v) = \int_{\Omega} A\nabla u \nabla v dx + \lambda \int_{\partial\Omega} uv ds$$

for a penalty parameter $\lambda > 0$ and the energy

$$E_\lambda: H^1(\Omega) \rightarrow \mathbb{R}, \quad E_\lambda(u) = \frac{1}{2}a_\lambda(u, u) - f(u).$$

The central error estimation is collected in the following Theorem. Note that we require $H^2(\Omega)$ regularity of the solution to equation (4).

Theorem 9 *Let $\Omega \subset \mathbb{R}^d$ be a bounded domain with $C^{1,1}$ boundary, $f \in L^2(\Omega)$ and assume $A \in C^{0,1}(\Omega, \mathbb{R}^{d \times d})$ is symmetric, uniformly elliptic with ellipticity constant $\alpha > 0$. By $u^* \in H_0^1(\Omega)$ we denote the solution of (4) and by u_λ the minimizer of the penalized energy E_λ over $H^1(\Omega)$. Fix an arbitrary subset $V \subset H^1(\Omega)$ and denote the coercivity constants of a_λ by $\alpha_\lambda > 0$ and set $\delta := E_\lambda(v) - \inf_{\tilde{v} \in V} E_\lambda(\tilde{v})$. Then there is a constant $c > 0$, only depending on A and Ω , such that for every $v \in V$ and $\lambda > 0$ it holds that*

$$\|v - u^*\|_{H^1(\Omega)} \leq \sqrt{\frac{2\delta}{\alpha_\lambda} + \frac{1}{\alpha_\lambda} \inf_{\tilde{v} \in V} \|\tilde{v} - u_\lambda\|_{a_\lambda}^2 + c\lambda^{-1} \|f\|_{L^2(\Omega)}}, \tag{5}$$

where $\|u\|_{a_\lambda}^2 := a_\lambda(u, u)$ is the norm induced by a_λ . Further, we can choose

$$c := c_F c_{reg} \sqrt{\|a_1\|} \|T\|_{\mathcal{L}(H^2(\Omega); \mathcal{H}(\Omega))},$$

where $T: H^2(\Omega) \rightarrow H^1(\Omega)$ maps a function u to the A -harmonic extension of $\partial_A u$, c_F denotes the Friedrich constant (see Proposition 1) and c_{reg} is the operator norm of

$$(-\operatorname{div}(A\nabla \cdot))^{-1}: L^2(\Omega) \rightarrow H^2(\Omega) \cap H_0^1(\Omega).$$

3. here $\partial_A u = \nu \cdot A\nabla u$ and ν denotes the outer normal of Ω

Proof The central idea of the proof consists of the error decomposition

$$\|v - u^*\|_{H^1(\Omega)} \leq \|v - u_\lambda\|_{H^1(\Omega)} + \|u_\lambda - u^*\|_{H^1(\Omega)}.$$

The first error can be treated using Céa's Lemma. Note that a_λ is in fact coercive on $H^1(\Omega)$ which is a consequence of Friedrich's inequality, see Proposition 1. For the second term one uses a Fourier series expansion in a Steklov basis. The latter is useful for weakly A -harmonic functions, hence we investigate the equation satisfied by $v_\lambda := u^* - u_\lambda$. Due to the regularity assumption on Ω and A we have $\operatorname{div}(A\nabla u^*) \in L^2(\Omega)$ and may integrate by parts to obtain for all $\varphi \in H^1(\Omega)$

$$\int_{\Omega} f\varphi dx = - \int_{\Omega} \operatorname{div}(A\nabla u^*)\varphi dx = \int_{\Omega} A\nabla u^* \cdot \nabla \varphi dx - \int_{\partial\Omega} \partial_A u^* \varphi ds. \quad (6)$$

Using the optimality condition of u_λ yields

$$\int_{\Omega} (A\nabla u_\lambda) \cdot \nabla \varphi dx + \lambda \int_{\partial\Omega} u_\lambda \varphi ds = \int_{\Omega} f\varphi dx \quad \forall \varphi \in H^1(\Omega).$$

Subtracting these two equations we obtain that v_λ satisfies

$$\int_{\Omega} (A\nabla v_\lambda) \cdot \nabla \varphi dx + \int_{\partial\Omega} (\lambda v_\lambda - \partial_A u^*)\varphi ds = 0 \quad \forall \varphi \in H^1(\Omega).$$

This implies that v_λ is weakly A -harmonic, i.e.,

$$\int_{\Omega} (A\nabla v_\lambda) \cdot \nabla \varphi dx = 0 \quad \forall \varphi \in H_0^1(\Omega),$$

We claim that there exists a basis $(e_j)_{j \in \mathbb{N}}$ of the space of weakly A -harmonic functions and that v_λ can be written in terms of this basis as

$$v_\lambda = \frac{1}{\lambda} \sum_{j=0}^{\infty} c(\lambda)_j e_j \quad (7)$$

for suitable coefficients $c(\lambda)_j \in \mathbb{R}$. Further, we claim that this Fourier expansion leads to the estimate

$$\|v_\lambda\|_{H^1(\Omega)} \leq \frac{c}{\lambda} \|f\|_{L^2(\Omega)} \quad (8)$$

with c as specified in the statement of the Theorem which finishes the proof. The remaining details are provided in the following Section. ■

We presented the proof in its above form to draw attention to its key elements and to discuss possible limitations and generalizations.

Remark 10 (Limitations) *Our proof requires crucially the $H^2(\Omega)$ regularity of the solution u^* to the Dirichlet problem. This is in contrast with the error estimates for non-essential boundary value problems that do not require additional regularity.*

Remark 11 (Generalizations) *The strategy of the proof of Theorem 9 holds for a broader class of elliptic zero boundary value problems. The essential requirement is that the bilinear form a of the differential operator is coercive on $H_0^1(\Omega)$ and that a_λ is coercive on all of $H^1(\Omega)$. Then, regularity of the solution u^* of the zero boundary value problem is required to identify the equation u^* satisfies when tested with functions in $H^1(\Omega)$ and not only $H_0^1(\Omega)$, see (6).*

Remark 12 (Optimality of the rate λ^{-1}) We demonstrate that the rate⁴

$$\|u_\lambda - u^*\|_{H^1(\Omega)} \lesssim \lambda^{-1}$$

can not in general be improved. To this end we consider the concrete example

$$a_\lambda : H^1(0, 1)^2 \rightarrow \mathbb{R}, \quad (u, v) \mapsto \int_0^1 u'v' dx + \lambda(u(0)v(0) + u(1)v(1)).$$

The minimiser of E_λ with $f \equiv 1$ solves the ODE

$$-u'' = 1 \quad \text{in } (0, 1)$$

with Robin boundary conditions

$$\begin{aligned} -u'(0) + \lambda u(0) &= 0 \\ u'(1) + \lambda u(1) &= 0. \end{aligned}$$

Its solution is given by

$$u_\lambda(x) = -\frac{1}{2}x^2 + \frac{1}{2}x + \frac{1}{2\lambda}.$$

On the other hand the associated Dirichlet problem is solving the same ODE subject to $u(0) = u(1) = 0$ and has the solution

$$u^*(x) = -\frac{1}{2}x^2 + \frac{1}{2}x.$$

Consequently the difference $u_\lambda - u^*$ measured in $H^1(0, 1)$ norm is precisely $\frac{1}{2\lambda}$.

3.2.1. A SOLUTION FORMULA BASED ON STEKLOV EIGENFUNCTIONS

The Steklov theory yields the existence of an orthonormal eigenbasis of the space

$$\mathcal{H}(\Omega) := \{w \in H^1(\Omega) \mid a(w, v) = 0 \text{ for all } v \in H_0^1(\Omega)\}.$$

of weakly a -harmonic functions which we can use for a Fourier expansion of v_λ in order to obtain the desired estimate. For a recent and more general discussion of Steklov theory we refer to [Auchmuty \(2012\)](#). The Steklov eigenvalue problem consists of finding $(\mu, w) \in \mathbb{R} \times H^1(\Omega)$ such that

$$a(w, \varphi) = \mu \int_{\partial\Omega} w\varphi ds \quad \text{for all } \varphi \in H^1(\Omega). \quad (9)$$

We call μ a *Steklov eigenvalue* and w a corresponding *Steklov eigenfunction*.

Lemma 13 (Orthogonal decomposition) We can decompose the space $H^1(\Omega)$ into

$$H^1(\Omega) = \mathcal{H}(\Omega) \oplus_{a_1} H_0^1(\Omega)$$

with the decomposition being a_1 -orthogonal.

4. We write \lesssim and \gtrsim if the inequality \leq or \geq holds up to a constant; if both \lesssim and \gtrsim hold, we write \sim .

Proof By the definition of a_1 and $\mathcal{H}(\Omega)$ the two spaces are a_1 orthogonal. To see that it spans all of $H^1(\Omega)$ let $u \in H^1(\Omega)$ be given. Let u_a be the unique solution of $a(u_a, \cdot) = 0$ in $H_0^1(\Omega)^*$ subject to $\text{tr}(u_a) = \text{tr}(u)$. Then the decomposition is given as $u = u_a + (u - u_a) = u_a + u^*$. ■

Theorem 14 (Steklov spectral theorem) *Let $\Omega \subseteq \mathbb{R}^d$ be open and a be a positive semi-definite bilinear form on $H^1(\Omega)$ and $\mathcal{H}(\Omega) \hookrightarrow L^2(\partial\Omega)$ be compact. Then there exists a non decreasing sequence $(\mu_j)_{j \in \mathbb{N}} \subseteq [0, \infty)$ with $\mu_j \rightarrow \infty$ and a sequence $(e_j)_{j \in \mathbb{N}} \subseteq \mathcal{H}(\Omega)$ such that μ_j is a Steklov eigenvalue with eigenfunction e_j . Further, $(e_j)_{j \in \mathbb{N}}$ is a complete orthonormal system in $\mathcal{H}(\Omega)$ with respect to a_1 .*

Proof This can be derived from the spectral theory for compact operators as for example described in (Dobrowolski, 2010, Section 8.10). In the notation of Dobrowolski (2010), set $X = \mathcal{H}(\Omega)$ with inner product a_1 and let $Y = L^2(\partial\Omega)$ equipped with its natural inner product. Then this yields a divergent sequence $0 < \tilde{\mu}_1 \leq \tilde{\mu}_2 \leq \dots$ growing to ∞ and $(e_j)_{j \in \mathbb{N}} \subseteq \mathcal{H}(\Omega)$ with $a_1(e_i, e_j) = \delta_{ij}$ and

$$a_1(e_j, w) = \tilde{\mu}_j \int_{\partial\Omega} e_j w ds \quad \text{for all } w \in \mathcal{H}(\Omega). \quad (10)$$

For $\varphi \in H^1(\Omega)$, let $\varphi = \varphi_a + \varphi_0$ be the orthogonal decomposition of the preceding lemma and compute

$$a_1(e_j, \varphi) = a_1(e_j, \varphi_0) + a_1(e_j, \varphi_a) = a_1(e_j, \varphi_a) = \tilde{\mu}_j \int_{\partial\Omega} e_j \varphi_a ds = \tilde{\mu}_j \int_{\partial\Omega} e_j \varphi ds.$$

Using the definition of a_1 we obtain

$$a(e_j, \varphi) = (\tilde{\mu}_j - 1) \int_{\partial\Omega} e_j \varphi ds \quad \text{for all } \varphi \in H^1(\Omega).$$

Setting $\mu_j := \tilde{\mu}_j - 1$ and noting that the above equality implies $\mu_j \geq 0$ concludes the proof. ■

As a direct consequence we obtain the following representation formula.

Corollary 15 (Fourier expansion in the Steklov eigenbasis) *Let $w \in \mathcal{H}(\Omega)$. Then we have*

$$w = \sum_{j=0}^{\infty} c_j e_j,$$

where

$$c_j = (1 + \mu_j) \int_{\partial\Omega} w e_j ds. \quad (11)$$

Proof Using that e_j is a Steklov eigenvector, we can compute the Fourier coefficients

$$c_j = a_1(w, e_j) = a(w, e_j) + \int_{\partial\Omega} w e_j ds = (1 + \mu_j) \int_{\partial\Omega} w e_j ds. \quad \blacksquare$$

Lemma 16 (Solution formula) Let $v_\lambda \in H^1(\Omega)$ be the unique solution of

$$a(v_\lambda, \varphi) + \int_{\partial\Omega} (\lambda v_\lambda - \partial_A u^*) \varphi \, ds = 0 \quad \text{for all } \varphi \in H^1(\Omega). \quad (12)$$

Then we have

$$v_\lambda = \frac{1}{\lambda} \sum_{j=0}^{\infty} c(\lambda)_j e_j,$$

where

$$c(\lambda)_j = \frac{1 + \mu_j}{1 + \frac{\mu_j}{\lambda}} \cdot \int_{\partial\Omega} (\partial_A u^*) e_j \, ds.$$

Proof Note that v_λ is weakly harmonic and hence, we can apply the previous corollary to compute the Fourier coefficients of v_λ . Using that v_λ solves (12) and that e_j is a Steklov eigenfunction we compute

$$\int_{\partial\Omega} v_\lambda e_j \, ds = \frac{1}{\lambda} \int_{\partial\Omega} (\partial_A u^*) e_j \, ds - \frac{1}{\lambda} a(v_\lambda, e_j) = \frac{1}{\lambda} \int_{\partial\Omega} (\partial_A u^*) e_j \, ds - \frac{\mu_j}{\lambda} \int_{\partial\Omega} v_\lambda e_j \, ds.$$

Rearranging this yields the following equation which completes the proof

$$\int_{\partial\Omega} v_\lambda e_j \, ds = \frac{1}{\lambda} \cdot \frac{1}{1 + \frac{\mu_j}{\lambda}} \int_{\partial\Omega} (\partial_A u^*) e_j \, ds. \quad \blacksquare$$

3.2.2. COMPLETING THE PROOF OF THEOREM 9

We use the explicit solution formula from Lemma 16 for $v_\lambda = u^* - u_\lambda$ to provide the missing claims in the proof of Theorem 9.

Proof [Completing the Proof of Theorem 9] We have already convinced ourselves that $v_\lambda := u_\lambda - u^*$ indeed solves (12). By the means of Lemma 16 it suffices to bound

$$\left\| \sum_{j=0}^{\infty} c(\lambda)_j e_j \right\|_{H^1(\Omega)}$$

independently of $\lambda > 0$. Let us denote the A -harmonic extension of $\partial_A u^*$ with w , we obtain

$$c(\lambda)_j^2 \leq (1 + \mu_j)^2 \left(\int_{\partial\Omega} (\partial_A u^*) e_j \, ds \right)^2 = a_1(w, e_j)^2.$$

Now we can estimate

$$\begin{aligned} \sum_{j=0}^{\infty} (1 + \mu_j)^2 \left(\int_{\partial\Omega} (\partial_A u^*) e_j \, ds \right)^2 &= \sum_{j=0}^{\infty} a_1(w, e_j)^2 = a_1(w, w) \\ &\leq \|a_1\| \|w\|_{H^1(\Omega)}^2 \\ &\leq \|a_1\| \|u^*\|_{H^2(\Omega)}^2 \|T\|_{\mathcal{L}(H^2(\Omega); H^1(\Omega))}^2 \\ &\leq \|a_1\| c_{reg}^2 \|f\|_{L^2(\Omega)}^2 \|T\|_{\mathcal{L}(H^2(\Omega); H^1(\Omega))}^2, \end{aligned}$$

where $T: H^2 \rightarrow H^1$ is the mapping that assigns a function u the harmonic extension of $\partial_A u$. Consequently, we obtain by Parseval's identity

$$\|v_\lambda\|_{H^1(\Omega)} \leq c_F \|v_\lambda\|_{a_1} = \frac{c_F}{\lambda} \sqrt{\sum_{j=1}^{\infty} c(\lambda)_j^2} \leq \frac{c_F}{\lambda} \sqrt{\|a_1\|} c_{reg} \|f\|_{L^2(\Omega)} \|T\|_{\mathcal{L}(H^2(\Omega); H^1(\Omega))}$$

which finishes the proof. \blacksquare

3.3. Estimates under lower regularity of the right-hand side

If $f \notin L^2(\Omega)$, Theorem 9 cannot be applied as the estimation of the term $\|u_\lambda - u^*\|_{H^1(\Omega)}$ requires that u^* is a member of $H^2(\Omega)$. The next Lemma shows that at the expense of a worse rate and norm, we can still estimate this term for distributional right-hand sides $f \in H^1(\Omega)^*$.

Lemma 17 *Let $\Omega \subset \mathbb{R}^d$ be a bounded domain with $C^{1,1}$ boundary, assume that $A \in C^{0,1}(\Omega, \mathbb{R}^{d \times d})$ is symmetric and uniformly elliptic, $f \in H^1(\Omega)^*$ and let u^* and u_λ be as in Theorem 9. Then it holds*

$$\|u_\lambda - u^*\|_{L^2(\Omega)} \leq c \cdot \|f\|_{H^1(\Omega)^*} \lambda^{-1/2}.$$

Proof We set $v_\lambda = u_\lambda - u^*$ and denote by $w \in H_0^1(\Omega)$ the solution to the equation $-\operatorname{div}(A\nabla w) = v_\lambda$ in $H_0^1(\Omega)^*$. Then, by our assumptions on Ω and A , the function w is a member of $H^2(\Omega)$. This yields upon integration by parts and the fact that v_λ is weakly A -harmonic that it holds

$$\begin{aligned} \|v_\lambda\|_{L^2(\Omega)}^2 &= \int_{\Omega} A\nabla w \cdot \nabla v_\lambda dx - \int_{\partial\Omega} \partial_A w v_\lambda ds = - \int_{\partial\Omega} \partial_A w u_\lambda ds \\ &\leq c \|w\|_{H^2(\Omega)} \|u_\lambda\|_{L^2(\partial\Omega)} \\ &\leq c \|v_\lambda\|_{L^2(\Omega)} \|u_\lambda\|_{L^2(\partial\Omega)}. \end{aligned}$$

It remains to estimate $\|u_\lambda\|_{L^2(\partial\Omega)}$. Note that u_λ satisfies

$$0 = \int_{\Omega} A\nabla u_\lambda \cdot \nabla u_\lambda dx + \lambda \int_{\partial\Omega} u_\lambda^2 ds - f(u_\lambda).$$

We get after rearranging and using Young's inequality

$$\begin{aligned} \|u_\lambda\|_{L^2(\partial\Omega)}^2 &= \frac{2}{\lambda} \left(f(u_\lambda) - \left(\int_{\Omega} A\nabla u_\lambda \cdot \nabla u_\lambda dx + \frac{\lambda}{2} \|u_\lambda\|_{L^2(\partial\Omega)}^2 \right) \right) \\ &\leq \frac{2}{\lambda} \left(\|f\|_{H^1(\Omega)^*} \|u_\lambda\|_{H^1(\Omega)} - \alpha_{\lambda/2} \|u_\lambda\|_{H^1(\Omega)}^2 \right) \\ &\leq \frac{\|f\|_{H^1(\Omega)^*}^2}{2\alpha_{\lambda/2}\lambda} \end{aligned}$$

which completes the proof. \blacksquare

4. Penalization strength and error decay

We have seen that the distance of an ansatz function can be bounded in terms of the optimization error, the approximation power of the ansatz class and the penalization strength. In this section we discuss the trade off of choosing the penalization strength λ too large or too small and discuss the implications of different scalings of λ in dependency of the approximation capabilities of the ansatz classes. We combine our general discussion with Theorem 3 to obtain Theorem 19, however, our discussion can be combined with any result guaranteeing approximation rates of a sequence of ansatz classes.

We consider a sequence $(V_n)_{n \in \mathbb{N}} \subseteq H^1(\Omega)$ of ansatz classes and penalization strengths $\lambda_n \sim n^\sigma$. Further, we denote the minimizers of the energies E_{λ_n} over V_n by $v_n^* \in V_n$. It is our goal to choose $\sigma \in \mathbb{R}$ in such a way that the upper bound of $\|v_n^* - u^*\|_{H^1(\Omega)}$ in (5) decays with the fastest possible rate. Neglecting constants, the bound evaluates to

$$\|v_n^* - u^*\|_{H^1(\Omega)} \lesssim \sqrt{\frac{1}{\alpha_{\lambda_n}} \inf_{v \in V_n} \|v - u_{\lambda_n}\|_{a_{\lambda_n}}^2} + \lambda_n^{-1}. \quad (13)$$

We can assume without loss of generality that $\sigma > 0$ and hence $\lambda_n \geq 1$, because otherwise the upper bound will not decrease to zero. Note that in this case we have $\alpha_{\lambda_n} \geq \alpha_1 > 0$ and hence the we obtain

$$\|v_n^* - u^*\|_{H^1(\Omega)} \lesssim \sqrt{\inf_{v \in V_n} \left(\|\nabla(v - u_{\lambda_n})\|_{L^2(\Omega)}^2 + n^\sigma \|v - u_{\lambda_n}\|_{L^2(\partial\Omega)}^2 \right)} + n^{-\sigma}. \quad (14)$$

Here, the trade off in choosing σ and therefore λ_n too large or small is evident. We discuss the implications of this upper bound in three different scenarios.

Approximation rates with zero boundary values Consider the case where there is an element $v_n \in V_n \cap H_0^1(\Omega)$ such that

$$\|v_n - u^*\|_{H^1(\Omega)} \lesssim n^{-r}.$$

Using the Euler-Lagrange equations $a_\lambda(u_\lambda, \cdot) = f(\cdot)$ and $a(u^*, \cdot) = f(\cdot)$ we can estimate

$$\begin{aligned} \frac{1}{2} \inf_{v \in V_n} \|v - u_\lambda\|_{a_\lambda}^2 &= \inf_{v \in V_n} E_\lambda(v) - E_\lambda(u_\lambda) \leq \inf_{v \in V_n \cap H_0^1(\Omega)} E_\lambda(v) - E(u_\lambda) \\ &\leq \inf_{v \in V_n \cap H_0^1(\Omega)} E(v) - E(u^*) = \frac{1}{2} \inf_{v \in V_n \cap H_0^1(\Omega)} \|v - u^*\|_a^2 \lesssim n^{-2r} \end{aligned}$$

independently of λ . Hence, the estimate (13) yields

$$\|v_n^* - u^*\|_{H^1(\Omega)} \lesssim \sqrt{n^{-2r}} + \lambda_n^{-1} \lesssim n^{-r},$$

whenever $\lambda_n \gtrsim n^r$. Note that in this case, no trade off in λ exists and the approximation rate with zero boundary values can always be achieved up to optimization. However, the curvature α_{λ_n} of E_{λ_n} increases with λ_n . Thus, it seems reasonable to choose λ_n as small as possible, i.e., $\lambda_n \sim n^r$. Approximation rates with zero boundary values have not been established so far for neural networks to the best of our knowledge.

Approximation error of u_λ independent of λ Now, we consider the case, without an approximation rate with exact zero boundary values, but where the sequence $(V_n)_{n \in \mathbb{N}}$ of ansatz classes admits approximation rates in both $H^1(\Omega)$ and $L^2(\partial\Omega)$. More precisely, we assume that there are real numbers $s \geq r > 0$ such that for every (sufficiently big) λ and every $n \in \mathbb{N}$ there is an element $v_n \in V_n$ satisfying

$$\|v_n - u_\lambda\|_{H^1(\Omega)} \leq cn^{-r} \quad \text{and} \quad \|v_n - u_\lambda\|_{L^2(\partial\Omega)} \leq cn^{-s},$$

for some $c > 0$ independent on λ . Then the estimate in (14) yields

$$\|v_n^* - u^*\|_{H^1(\Omega)} \lesssim \sqrt{n^{-2r} + n^{\sigma-2s}} + n^{-\sigma}.$$

The resulting rate of decay of the upper bound of the error is then

$$\rho(\sigma) = \min\left(\frac{1}{2} \min(2r, 2s - \sigma), \sigma\right) = \min\left(r, s - \frac{\sigma}{2}, \sigma\right).$$

which is maximised at $\sigma^* = 2s/3$ with a value of

$$\rho^* = \min\left(\frac{2}{3}s, r\right). \quad (15)$$

In this case, the upper bound does not necessarily decay at the same rate as the approximation error, which decays with rate r . Note that because $H^1(\Omega)$ embeds into $L^2(\partial\Omega)$ we can assume without loss of generality that $s \geq r$.

We made the assumption that the approximation rates of r and s holds with the same constant independently of λ . This is for example the case, if the solutions u_λ are uniformly in λ bounded in $H^s(\Omega)$ for some $s > 1$.

Approximation rates for u^* Now we want to discuss the case, where we weaken the approximation assumption from above, which is uniformly in λ . More precisely, we assume that there is a constant $c > 0$ and elements $v_n \in V_n$ satisfying

$$\|v_n - u^*\|_{H^1(\Omega)} \leq cn^{-r} \quad \text{and} \quad \|v_n - u^*\|_{L^2(\partial\Omega)} \leq cn^{-s}.$$

By (8) (or equivalently Theorem 9 with $V = H^1(\Omega)$ and $v = u_\lambda$) and the triangle inequality we have

$$\|v_n - u_\lambda\|_{H^1(\Omega)} \leq \|v_n - u^*\|_{H^1(\Omega)} + \|u^* - u_\lambda\|_{H^1(\Omega)} \leq cn^{-r} + c'n^{-\sigma} \leq \tilde{c}n^{-\tilde{r}}$$

and similarly

$$\|v_n - u_\lambda\|_{L^2(\partial\Omega)} \leq cn^{-s} + c'n^{-\sigma} \leq \tilde{c}n^{-\tilde{s}},$$

where $\tilde{r} = \min(r, \sigma)$ and $\tilde{s} = \min(s, \sigma)$. Hence, we have reduced this case to the previous case and find that the right hand side of (14) decays at a rate of

$$\begin{aligned} \rho(\sigma) &= \min(\min(r, \sigma), \min(s, \sigma) - \sigma/2, \sigma) = \min(r, \min(s, \sigma) - \sigma/2) \\ &= \min(r, \min(s - \sigma/2, \sigma/2)) = \min(r, s - \sigma/2, \sigma/2). \end{aligned} \quad (16)$$

This function is maximized at $\sigma^* = s$ with a value of $\rho^* = \min(s/2, r)$. Like before, we can without loss of generality assume that $s \geq r$.

Remark 18 Note that in general the decay rate $\rho^* = \min(s/2, r)$ of the upper bound (13) can be smaller than the approximation rate r . This is in contrast to problems with non-essential boundary values for which the error decays proportional to the approximation error by Cea's lemma. We stress that the defect in the decay rate of the right hand side of (13) is not an artefact of our computations but in fact sharp.

Let us now come back to the original problem of the Dirichlet problem (4). For a right hand side $f \in H^r(\Omega)$, standard regularity results yield $u^* \in H^{r+2}(\Omega)$. Theorem 3 provides rates for the approximation in $H^s(\Omega)$ for $s \in [0, 1]$, which lead to the following result.

Theorem 19 (Rates for NN training with boundary penalty) Let $\Omega \subset \mathbb{R}^d$ be a bounded domain with $C^{r+1,1}$ boundary for some $r \in \mathbb{N}$, $f \in H^r(\Omega)$ and assume $A \in C^{r,1}(\Omega, \mathbb{R}^{d \times d})$ is symmetric, uniformly elliptic with ellipticity constant $\alpha > 0$ and denote the solution to (4) by $u^* \in H_0^1(\Omega)$. For every $n \in \mathbb{N}$, there is a ReLU network with parameter space Θ_n of dimension $\mathcal{O}(\log_2^2(n^{(r+2)/d}) \cdot n)$ such that if $\lambda_n \sim n^\sigma$ for $\sigma = \frac{2r+3}{2d}$ one has for any $\rho < \frac{2r+3}{4d}$ that

$$\|u_{\theta_n} - u^*\|_{H^1(\Omega)} \lesssim \sqrt{\delta_n + n^{-2\rho}} + n^{-\rho} \quad \text{for all } \theta_n \in \Theta_n, \quad (17)$$

where $\delta_n = E_{\lambda_n}(u_{\theta_n}) - \inf_{\tilde{\theta} \in \Theta_n} E_{\lambda_n}(u_{\tilde{\theta}})$.

Proof For $\varepsilon > 0$ it holds that $H^{1/2+\varepsilon}(\Omega) \hookrightarrow L^2(\partial\Omega)$. Thus, Theorem 3 guarantees the existence of u_{θ_n} with $\theta_n \in \Theta_n$ and the claimed number of parameters such that $\|u_{\theta_n} - u^*\|_{H^1(\Omega)} \lesssim n^{-(r+1)/d} = n^{-\tilde{r}}$ and

$$\|u_{\theta_n} - u^*\|_{L^2(\partial\Omega)} \lesssim \|u_{\theta_n} - u^*\|_{H^{1/2+\varepsilon}(\Omega)} \lesssim n^{-(r+2-(1/2+\varepsilon))/d} = n^{-\tilde{s}},$$

where $\tilde{s} = (2r + 3 - 2\varepsilon)/(2d)$. By (16), the estimate (17) holds for

$$\rho = \min(\tilde{r}, \tilde{s} - \sigma/2, \sigma/2) = \min\left(\frac{r+1}{d}, \frac{2r+3-4\varepsilon}{4d}, \frac{2r+3}{4d}\right) = \frac{2r+3-4\varepsilon}{4d}. \quad \blacksquare$$

Remark 20 (Adaptation to Smoothness) The discussion from Remark 7 carries over to the case of Dirichlet boundary values and boundary penalties, i.e., the error of the deep Ritz method decays at a rate increasing with the smoothness of the problem. This fact can be especially useful in high spatial dimensions, which is consistent with the empirical findings that the deep Ritz method can be effective in the numerical solution of high dimensional problems E and Yu (2018). Note that also finite element methods can achieve rates increasing with the smoothness of data, however they require the delicate construction of higher order elements.

Remark 21 (Combination with different approximation results) We focus on the ReLU activation in this section, whereas in practice often other architectures and activation functions are used, see E and Yu (2018); Hennigh et al. (2021). However, our results from Section 3 can handle arbitrary function classes and hence reduce the computation of error estimates to the computation of approximation bounds. Therefore, they can be combined with other approximation results for neural networks in Sobolev norm including the works of Gühring and Raslan (2020); Xu (2020); Siegel and Xu (2020a); Hon and Yang (2021); Duan et al. (2021a); De Ryck et al. (2021).

Remark 22 (The boundary penalty method for FEM) *The boundary penalty method has been applied in the context of finite element approximations Babuška (1973) and studied in terms of its convergence rates in Babuška (1973); Shi (1984); Barrett and Elliott (1986). The idea of the finite element approach is analogue to the idea of using neural networks for the approximate solution of variational problems. However, one constructs a nested sequence of finite dimensional vectorspaces $V_h \subseteq H^1(\Omega)$ arising from some triangulation with fineness $h > 0$ and computes the minimiser u_h of the penalized energy E_λ over V_h . Choosing a suitable triangulation and piecewise affine linear elements and setting $\lambda \sim h^{-1}$ one obtains the error estimate*

$$\|u_h - u^*\|_{H^1(\Omega)} \lesssim h,$$

see Shi (1984). At the core of those estimates lies a linear version of C ea’s Lemma, which already incorporates boundary values. However, the proof of this lemma heavily relies on the fact that the class of ansatz functions is linear and that its minimizer solves a linear equation. This is not the case for non linear function classes like neural networks. Therefore, our estimates require a different strategy. However, the optimal rate of convergence for the boundary penalty method with finite elements can be deduced from our results. In fact, one can choose a suitable triangulation and an operator $r_h: H^2(\Omega) \rightarrow V_h$ such that

$$\|r_h u - u\|_{H^1(\Omega)} \lesssim h \|u\|_{H^2(\Omega)},$$

where the approximating functions u_h have zero boundary values as they arise from interpolation. By the general discussion from above for the ansatz classes with approximation rates with exact zero boundary values, choosing $\lambda \gtrsim h^{-1}$ yields an error bound decaying like the approximation error $\|u_h - u^\|_{H^1(\Omega)} \lesssim h$.*

Acknowledgments

The authors want to thank Luca Courte, Patrick Dondl and Stephan Wojtowytsch for their valuable comments. JM acknowledges support by the Evangelisches Studienwerk e.V. (Villigst), the International Max Planck Research School for Mathematics in the Sciences (IMPRS MiS) and the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant number 757983). MZ acknowledges support from BMBF within the e:Med program in the SyMBoD consortium (grant number 01ZX1910C) and the Research Council of Norway (grant number 303362).

References

- Mart n Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. TensorFlow: A System for Large-Scale Machine Learning. In *12th USENIX symposium on operating systems design and implementation (OSDI 16)*, pages 265–283, 2016.
- Raman Arora, Amitabh Basu, Poorya Mianjy, and Anirbit Mukherjee. Understanding Deep Neural Networks with Rectified Linear Units. *arXiv preprint arXiv:1611.01491*, 2016.
- Giles Auchmuty. Bases and comparison results for linear elliptic eigenproblems. *Journal of Mathematical Analysis and Applications*, 390(1):394–406, 2012.

- Ivo Babuška. The finite element method with penalty. *Mathematics of computation*, 27(122):221–228, 1973.
- John W Barrett and Charles M Elliott. Finite element approximation of the Dirichlet problem using the boundary penalty method. *Numerische Mathematik*, 49(4):343–366, 1986.
- Wojciech Marian Czarnecki, Simon Osindero, Max Jaderberg, Grzegorz Świrszcz, and Razvan Pascanu. Sobolev training for neural networks. *arXiv preprint arXiv:1706.04859*, 2017.
- Tim De Ryck, Samuel Lanthaler, and Siddhartha Mishra. On the approximation of functions by tanh neural networks. *Neural Networks*, 143:732–750, 2021.
- Manfred Dobrowolski. *Angewandte Funktionalanalysis: Funktionalanalysis, Sobolev-Räume und elliptische Differentialgleichungen*. Springer-Verlag, 2010.
- Patrick Dondl, Johannes Müller, and Marius Zeinhofer. Uniform Convergence Guarantees for the Deep Ritz Method for Nonlinear Problems. *arXiv preprint arXiv:2111.05637*, 2021.
- Chenguang Duan, Yuling Jiao, Yanming Lai, Xiliang Lu, Qimeng Quan, and Jerry Zhijian Yang. Analysis of deep ritz methods for laplace equations with dirichlet boundary conditions. *arXiv preprint arXiv:2111.02009*, 2021a.
- Chenguang Duan, Yuling Jiao, Yanming Lai, Xiliang Lu, and Zhijian Yang. Convergence rate analysis for deep Ritz method. *arXiv preprint arXiv:2103.13330*, 2021b.
- Weinan E and Bing Yu. The Deep Ritz Method: A Deep Learning-Based Numerical Algorithm for Solving Variational Problems. *Communications in Mathematics and Statistics*, 6(1):1–12, 2018.
- Alexandre Ern and Jean-Luc Guermond. *Theory and practice of finite elements*, volume 159. Springer, 2004.
- Ingo Gühring and Mones Raslan. Approximation Rates for Neural Networks with Encodable Weights in Smoothness Spaces. *arXiv preprint arXiv:2006.16822*, 2020.
- Ingo Gühring, Gitta Kutyniok, and Philipp Petersen. Error bounds for approximations with deep ReLU neural networks in $W^{s,p}$ norms. *arXiv preprint arXiv:1902.07896*, 2019.
- Jiequn Han, Arnulf Jentzen, et al. Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations. *Communications in Mathematics and Statistics*, 5(4):349–380, 2017.
- Jiequn Han, Arnulf Jentzen, and E Weinan. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018.
- Oliver Hennigh, Susheela Narasimhan, Mohammad Amin Nabian, Akshay Subramaniam, Kaushtubh Tangsali, Zhiwei Fang, Max Rietmann, Wonmin Byeon, and Sanjay Choudhry. Nvidia simnet™: An ai-accelerated multi-physics simulation framework. In *International Conference on Computational Science*, pages 447–461. Springer, 2021.

- Sean Hon and Haizhao Yang. Simultaneous neural network approximations in Sobolev spaces. *arXiv preprint arXiv:2109.00161*, 2021.
- Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural networks*, 4(2):251–257, 1991.
- Hyung Ju Hwang, Cheolhyeong Kim, Min Sue Park, and Hwijae Son. The deep minimizing movement scheme. *arXiv preprint arXiv:2109.14851*, 2021.
- Arnulf Jentzen, Diyora Salimova, and Timo Welti. A proof that deep artificial neural networks overcome the curse of dimensionality in the numerical approximation of kolmogorov partial differential equations with constant diffusion and nonlinear drift coefficients. *arXiv preprint arXiv:1809.07321*, 2018.
- Yuling Jiao, Yanming Lai, Yisu Luo, Yang Wang, and Yunfei Yang. Error analysis of deep Ritz methods for elliptic equations. *arXiv preprint arXiv:2107.14478*, 2021.
- Kookjin Lee, Nathaniel A Trask, Ravi G Patel, Mamikon A Gulian, and Eric C Cyr. Partition of unity networks: deep hp-approximation. *arXiv preprint arXiv:2101.11256*, 2021.
- Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020.
- Bertrand Maury. Numerical analysis of a finite element/volume penalty method. *SIAM Journal on Numerical Analysis*, 47(2):1126–1148, 2009.
- Siddhartha Mishra and Roberto Molinaro. Estimates on the generalization error of physics informed neural networks (pinns) for approximating a class of inverse problems for pdes. *arXiv preprint arXiv:2007.01138*, 2020.
- Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019.
- Zhong-Ci Shi. On the convergence rate of the boundary penalty method. *International journal for numerical methods in engineering*, 20(11):2027–2032, 1984.
- Jonathan W Siegel and Jinchao Xu. Approximation rates for neural networks with general activation functions. *Neural Networks*, 128:313–321, 2020a.
- Jonathan W Siegel and Jinchao Xu. High-order approximation rates for neural networks with relu^k activation functions. *arXiv preprint arXiv:2012.07205*, 2020b.
- E Weinan, Chao Ma, and Lei Wu. Barron spaces and the compositional function spaces for neural network models. *arXiv preprint arXiv:1906.08039*, 2019.
- Stephan Wojtowytsch et al. Some observations on partial differential equations in barron and multi-layer spaces. *arXiv preprint arXiv:2012.01484*, 2020.

Jinchao Xu. The finite neuron method and convergence analysis. *arXiv preprint arXiv:2010.01458*, 2020.

Linfeng Zhang, Jiequn Han, Han Wang, Roberto Car, and E Weinan. Deep potential molecular dynamics: a scalable model with the accuracy of quantum mechanics. *Physical review letters*, 120(14):143001, 2018.

Yinhao Zhu, Nicholas Zabaras, Phaedon-Stelios Koutsourelakis, and Paris Perdikaris. Physics-constrained deep learning for high-dimensional surrogate modeling and uncertainty quantification without labeled data. *Journal of Computational Physics*, 394:56–81, 2019.

Appendix A. Universal approximation in Sobolev Topology

Theorem 3 is a restatement of the main result in [Gühring et al. \(2019\)](#). Since this formulation is not present in [Gühring et al. \(2019\)](#), we provide a short proof of it, which explains how it can be deduced from the results in [Gühring et al. \(2019\)](#).

Theorem 3 (Quantitative universal approximation) *Let $\Omega \subseteq \mathbb{R}^d$ be a bounded and open set with Lipschitz regular boundary⁵, let $k \in (1, \infty)$, $p \in [1, \infty]$ and fix an arbitrary function $u \in W^{k,p}(\Omega)$. Then, for every $n \in \mathbb{N}$, there is a ReLU network u_n with $\mathcal{O}(\log_2^2(n^{k/d}) \cdot n)$ many parameters and neurons such that*

$$\|u - u_n\|_{W^{s,p}(\Omega)} \leq c(s) \cdot \|u\|_{W^{k,p}(\Omega)} \cdot n^{-(k-s)/d}$$

for every $s \in [0, 1]$.

Proof The approximation results in [Gühring et al. \(2019\)](#) are stated for functions with the unit cube $[0, 1]^d$ as a domain. However, by scaling and possibly extending functions to the whole of \mathbb{R}^d this implies analogue results for functions defined on bounded Sobolev extension domains Ω .

We examine the proof of [Gühring et al. \(2019\)](#) in order to see that the approximating network architectures do not depend on s . Let us fix a function $u \in W^{k,p}([0, 1]^d)$. In their notation, for $M \in \mathbb{N}$, there are functions $(\phi_m)_{m=1, \dots, M^d}$ and polynomials $(p_m)_{m=1, \dots, M^d}$, such that

$$\left\| u - \sum \phi_m p_m \right\|_{W^{s,p}([0,1]^d)} \lesssim \|u\|_{W^{k,p}([0,1]^d)} M^{-(k-s)}$$

and a ReLU network function u_M with $N \lesssim M^d \log(M^k)$ parameters such that

$$\left\| \sum \phi_m p_m - u_M \right\|_{W^{s,p}([0,1]^d)} \leq c(s) \|u\|_{W^{k,p}(\Omega)} M^{-(k-s)}.$$

This follows from the Lemma C.3, C.4 and Lemma C.6 in [Gühring et al. \(2019\)](#) with $\varepsilon = M^{-k}$. Note that the functions and networks provided by those lemmata do not depend on s , which is evident as the estimates are first shown for $s = 0, 1$ and then generalised through interpolation. Now, by the triangle inequality, we have

$$\|u - u_M\|_{W^{s,p}([0,1]^d)} \leq \tilde{c}(s) \|u\|_{W^{k,p}([0,1]^d)} M^{-(k-s)}.$$

5. or more generally, that Ω is a Sobolev extension domain

Now the claim follows by choosing $M \sim n^{1/d}$. The additional square of the logarithm appears since they are considering networks with skip connections and those are then expressed as networks without skip connections, see also Corollary 4.2 in [Gühring et al. \(2019\)](#). ■