Nonclosedness of Sets of Neural Networks in Sobolev Spaces

Scott Mahan[*][†]   Emily J. King[‡]   Alex Cloninger[*][§]

Abstract

We examine the closedness of sets of realized neural networks of a fixed architecture in Sobolev spaces. For an exactly m-times differentiable activation function ρ, we construct a sequence of neural networks $(\Phi_n)_{n\in\mathbb{N}}$ whose realizations converge in order-$(m-1)$ Sobolev norm to a function that cannot be realized exactly by a neural network. Thus, sets of realized neural networks are not closed in order-$(m-1)$ Sobolev spaces $W^{m-1,p}$ for $p \in [1,\infty]$. We further show that these sets are not closed in $W^{m,p}$ under slightly stronger conditions on the m-th derivative of ρ. For a real analytic activation function, we show that sets of realized neural networks are not closed in $W^{k,p}$ for *any* $k \in \mathbb{N}$. The nonclosedness allows for approximation of non-network target functions with unbounded parameter growth. We partially characterize the rate of parameter growth for most activation functions by showing that a specific sequence of realized neural networks can approximate the activation function's derivative with weights increasing inversely proportional to the L^p approximation error. Finally, we present experimental results showing that networks are capable of closely approximating non-network target functions with increasing parameters via training.

Keywords— Fixed-architecture neural networks, Neural network expressivity, Closedness, Sobolev space

1   Introduction

From an approximation theory perspective, neural networks use observed training data to approximate an unknown target function. Studying topological properties of sets of neural

---
[*]Department of Mathematics, University of California San Diego. 9500 Gilman Drive #0112, La Jolla, CA 92037.

[†]Corresponding author. E-mail: scmahan@ucsd.edu

[‡]Department of Mathematics, Colorado State University. 1874 Campus Delivery, Fort Collins, CO 80523.

[§]Halicioglu Data Science Institute, University of California San Diego. 9500 Gilman Drive #0112, La Jolla, CA 92037.

networks will reveal what kinds of functions can be approximated by neural networks. In particular, closedness of sets of networks is a topological property of interest. If these sets are closed with respect to some norm, then one can construct a sequence of neural networks converging to a target function in that norm if and only if that target function is itself a network. On the other hand, nonclosedness would mean that neural networks can approximate target functions that are not networks themselves. An up-to-date survey on the approximation results for neural networks can be found in (Gühring et al., 2020).

To allow neural networks to approximate a wider class of functions, the number of nodes in the network can be increased. As long as the number of hidden nodes is allowed to grow without bound, Hornik's Universal Approximation Theorem shows that neural networks with only one hidden layer can approximate any p-integrable function to arbitrary accuracy (Hornik, 1991). Other approximation theorems show that neural networks are dense in other function classes, depending on the properties of the activation function, but many of these results allow the depth or width of the network to vary (Cybenko, 1989; Hornik et al., 1989, 1990; Krizhevsky et al., 2012). These results suggest that sets of realized neural networks are not closed in the corresponding function spaces, since not all of these functions can be represented exactly by a neural network. However, these are results about sets of networks of any width.

In practice, the architecture of a neural network is fixed before the learning process begins. Hence, we consider properties of sets of neural networks with a fixed architecture. Related to closedness is the best approximation property in L^p spaces, which holds if every $f \in L^p$ has at least one realized neural network g such that $\|f - g\|_{L^p}$ is minimized over all possible networks with the same architecture. In disproving the best approximation property for sigmoidal neural networks of a fixed size, (Girosi and Poggio, 1990) show that sets of realized neural networks with the sigmoid activation function are not closed in L^p spaces and claim that this should be true for all nonlinear activation functions. However, (Kůrková, 1995) proves that sets of networks with the Heaviside activation function are closed in L^p spaces, and (Kainen et al., 2020) proves a similar result for generalizations of Heaviside-type networks. Moreover, (Kainen et al., 2000) shows that Heaviside perceptron networks do have the best approximation property, but that best approximation maps with these networks are not unique or continuous. More generally, no best approximation maps using fixed-architecture neural networks are continuous (Kainen et al., 1999).

Petersen, Raslan, and Voigtlaender discuss the closedness and other topological properties of sets of realized neural networks (Petersen et al., 2019, 2020). Among other results, they prove that for most commonly used activation functions, these sets are not closed with respect to L^p norms. However, sets of all neural networks with a fixed architecture and uniformly bounded parameters *are* closed, indicating that learning a non-network target function requires the parameters to grow without bound. For example, learning a target function that has fewer derivatives than the activation function requires at least one network parameter to tend to infinity even as the network width increases, as seen in (Mhaskar, 1997). We partially investigate the relationship between L^p approximation error and parameter size, finding that the two are approximately inversely proportional when learning the derivative of the activation function.

In (Petersen et al., 2019) (a shortened conference paper version of (Petersen et al., 2020)), the authors speculate that sets of neural networks may be closed in Sobolev spaces,

where convergence is stronger. However, we show that this is not true, extending their *non-closedness* results to convergence in Sobolev norm under additional smoothness assumptions on the activation function using proof techniques similar to those in (Petersen et al., 2020). Our results apply to the case $p = \infty$ after minor modifications, so sets of realized neural networks are also not closed in $W^{m,\infty}$.

In some cases, such as network compression or distillation, we have data about the derivatives of the target function in addition to the training data. In these instances, one can train a network to learn the target function and its derivatives. This approach, introduced in (Czarnecki et al., 2017) as Sobolev training, often requires less training data and performs better on testing data. Hence, it is natural to consider the theoretical properties of neural networks in Sobolev spaces, as we do in this paper. We also provide some experimental results using Sobolev training with activation functions of varying degrees of smoothness, where we are able to approximate non-network target functions in Sobolev norm. This result indicates that sets of realized neural networks are indeed not closed in Sobolev spaces, but also that Sobolev training does not prevent parameter growth and allows us to approximate functions on the boundary of these sets of realized networks. Our experiments exhibit slow parameter growth relative to a fast decrease in approximation error, but there may be target functions that require much faster parameter growth to approximate.

1.1 Contributions of this Work

Our work considers sets of realized neural networks with a fixed architecture and a fixed activation function. In particular, we study the closedness of these sets of realized neural networks in Sobolev spaces. Our main contributions are:

1) We establish in Theorem 3.1 that for an m-times differentiable activation function that is not $(m+1)$-times differentiable, sets of realized neural networks are *not* closed in order-$(m-1)$ Sobolev spaces $W^{m-1,p}$ for $p \in [1, \infty]$. We prove this result by constructing a sequence of neural networks that converges in Sobolev norm to a target function that is not a neural network, which follows the approach in (Petersen et al., 2020).

2) We extend the nonclosedness result of Theorem 3.1 to $W^{m,p}$ under an additional assumption on the activation function.

3) For real analytic activation functions, Theorem 3.3 shows that sets of realized neural networks are not closed in any order Sobolev spaces.

4) We show analytically in Proposition 3.6 that for most activation functions, the L^p approximation error decays inversely proportional to the growth of network parameters for a given sequence of networks approximating the derivative of the activation function. The relationship between approximation error and weight growth relates to Theorem 1 in (Mhaskar, 1997), although that result holds for L^∞ error of shallow networks with increasing width.

5) We present some experiments in Section 4 demonstrating that neural networks can approximate target functions that require increasingly large parameters. Our exam-

ple achieves a fast decay in approximation error with a relatively slow growth in the network parameters, which may not be the case for other non-network target functions. These results appear to be robust to varying degrees of smoothness of the activation function and to certain classes of target functions.

Our nonclosedness results all indicate that neural networks can be trained to approximate non-network target functions in Sobolev norm. However, we will see that doing so will necessarily cause unbounded growth of network parameters. Thus, the training process may be difficult in practice, or regularization techniques may prevent a network from approximating a non-network target function. In our experiments, we train sequences of networks to approximate non-network target functions in Sobolev norm. The networks are able to closely approximate these target functions, providing further evidence that sets of realized neural networks are not closed in Sobolev spaces. Moreover, the ability to numerically approximate functions on the boundary of these sets of realized neural networks speaks to the expressiveness of neural networks in practice.

1.2 Outline of this Paper

Our work first provides background material on neural networks, then discusses the closedness of realized neural networks in Sobolev spaces, and finally presents some related numerical results. Section 2 lays out the definitions and notation required for the rest of the paper. In Section 3 we begin with our main results that sets of realized neural networks are not closed in Sobolev spaces under reasonable conditions on the activation function. On the other hand, Section 3.1 studies realizations of networks with bounded parameters, and presents a result that these sets of realizations *are* closed in Sobolev spaces. Section 3.2 analyzes the relationship between L^p approximation error and parameter growth for a network learning its activation function's derivative. We provide experimental results in Section 4 that demonstrate the nonclosedness of sets of realized neural networks and show that some classes of non-network target functions can indeed be approximated in Sobolev norm by a sequence of networks with increasing parameters.

2 Notation and Definitions

We first define neural networks. Every network has an architecture which specifies the input dimension, the number of layers, and the number of nodes in each layer. In addition, the network consists of matrix-vector pairs that determine the affine transformation between consecutive layers.

Definition 2.1. *(Petersen et al., 2020) Let $d, L \in \mathbb{N}$. A **neural network** Φ **with input dimension** d **and** L **layers** is a sequence of matrix-vector pairs*

$$\Phi = \big((A_1, b_1), \ldots, (A_L, b_L)\big),$$

*where $N_0 = d$ and $N_1, \ldots, N_L \in \mathbb{N}$, and where each A_ℓ is an $N_\ell \times N_{\ell-1}$ matrix, and $b_\ell \in \mathbb{R}^{N_\ell}$. We call $(d, N_1, \ldots, N_L)$ the **architecture** of Φ. N_L is the **output dimension**. Define $\mathcal{NN}(d, N_1, \ldots, N_L)$ to be the **set of all neural networks** Φ with architecture $(d, N_1, \ldots, N_L)$.*

4

To emphasize the role of the activation function, we distinguish between a neural network and a realized neural network. A realized network is a function defined by alternately applying the affine transformations of the network and the activation function. We also define the set of all realized neural networks with a fixed architecture and the same activation function.

Definition 2.2. *(Petersen et al., 2020) Let Φ be a neural network, $\Omega \subset \mathbb{R}^d$, and $\rho : \mathbb{R} \to \mathbb{R}$. The **realization of Φ with activation function ρ over Ω** is the function $R_\rho^\Omega(\Phi) : \Omega \to \mathbb{R}^{N_L}$ defined by*

$$R_\rho^\Omega(\Phi)(x) = W_L(\rho(W_{L-1}(\cdots \rho(W_1(x)))))$$

*where the affine transformation $W_\ell : \mathbb{R}^{N_{\ell-1}} \to \mathbb{R}^{N_\ell}$ is defined by $W_\ell(x) = A_\ell x + b_\ell$ and ρ is evaluated componentwise. Define R_ρ^Ω to be the **realization map** $\Phi \mapsto R_\rho^\Omega(\Phi)$, and let*

$$\mathcal{RNN}_\rho^\Omega(d, N_1, \ldots, N_L) := R_\rho^\Omega\big(\mathcal{NN}(d, N_1, \ldots, N_L)\big).$$

*We call $\mathcal{RNN}_\rho^\Omega(d, N_1, \ldots, N_L)$ the **set of ρ-realizations of networks with architecture** $(d, N_1, \ldots, N_L)$ **over** Ω.*

We will sometimes need to concatenate networks, which creates a new neural network consisting of the matrix-vector pairs of the first network followed by the pairs of the second network.

Definition 2.3. *(Petersen et al., 2020) Let $\Phi_1 = \big((A_1^1, b_1^1), \ldots, (A_{L_1}^1, b_{L_1}^1)\big)$ and $\Phi_2 = \big((A_1^2, b_1^2), \ldots, (A_{L_2}^2, b_{L_2}^2)\big)$ be two neural networks such that the input dimension of Φ_1 equals the output dimension of Φ_2. Then*

$$\Phi_1 \bullet \Phi_2 := \big((A_1^2, b_1^2), \ldots, (A_{L_2-1}^2, b_{L_2-1}^2), (A_1^1 A_{L_2}^2, A_1^1 b_{L_2}^2 + b_1^1), (A_2^1, b_2^1), \ldots, (A_{L_1}^1, b_{L_1}^1)\big)$$

*defines a neural network with $L_1 + L_2 - 1$ layers. We call $\Phi_1 \bullet \Phi_2$ the **concatenation of Φ_1 and Φ_2**.*

Note that for any activation function $\rho : \mathbb{R} \to \mathbb{R}$ and any $\Omega \subset \mathbb{R}^{d_2}$, we have $R_\rho^\Omega(\Phi_1 \bullet \Phi_2) = R_\rho^{\mathbb{R}^{d_1}}(\Phi_1) \circ R_\rho^\Omega(\Phi_2)$, where d_i is the input dimension of Φ_i. That is, concatenation of neural networks corresponds to function composition of the realizations of those networks.

For a fixed network architecture $(d, N_1, \ldots, N_L)$, we want to consider the closedness of the set $\mathcal{RNN}_\rho^\Omega(d, N_1, \ldots, N_L)$ in Sobolev spaces. We define Sobolev spaces below.

Definition 2.4. *Let $k \in \mathbb{N}$, let $\Omega \subset \mathbb{R}^k$ be measurable with non-empty interior, and let $1 \leq p \leq \infty$. The **Sobolev space** $W^{k,p}(\Omega)$ consists of all functions f on Ω such that for all multi-indices α with $|\alpha| \leq k$, the mixed partial derivative $f^{(\alpha)} := D^\alpha f$ exists in the weak sense and belongs to $L^p(\Omega)$. That is,*

$$W^{k,p}(\Omega) = \{f \in L^p(\Omega) : D^\alpha f \in L^p(\Omega) \text{ for all } |\alpha| \leq k\}.$$

*The number k is the **order** of the Sobolev space. The norm*

$$\|f\|_{W^{k,p}(\Omega)} := \sum_{|\alpha| \leq k} \|D^\alpha f\|_{L^p(\Omega)}$$

makes $W^{k,p}(\Omega)$ a Banach space for any $k \in \mathbb{N}$. Note that $W^{0,p}(\Omega) = L^p(\Omega)$.

3   (Non)closedness in Sobolev Spaces

In (Petersen et al., 2020) it is shown that $\mathcal{RNN}_\rho^{[-B,B]^d}(d, N_1, \ldots, N_{L-1}, 1)$ is *not* closed in $L^p([-B, B]^d)$ for any $p \in (0, \infty)$, under mild assumptions satisfied by most commonly used activation functions (including ReLU, the rectified linear unit). Moreover, these sets of realized neural networks of a fixed architecture are *not* closed in $C([-B, B]^d)$ with respect to the L^∞ norm for most commonly used activation functions. However, sets of ReLU-realizations of two-layer networks *are* closed in $C([-B, B]^d)$. These results are shown for $[-B, B]^d$, but generalize to any compact set $\Omega \subset \mathbb{R}^d$ with non-empty interior.

In this work, we investigate the closedness of sets of realized neural networks in Sobolev spaces. Since convergence in Sobolev norm is stronger than L^p convergence, Petersen, Raslan, and Voigtlaender anticipate that $\mathcal{RNN}_\rho^{[-B,B]^d}(d, N_1, \ldots, N_{L-1}, 1)$ may be closed in Sobolev spaces (Petersen et al., 2019). We prove that this is often not the case. Provided that the activation function ρ is m-times differentiable with bounded derivatives, these sets are not closed in $W^{m-1,p}([-B, B]^d)$ for any $p \in [1, \infty]$.

Theorem 3.1. *Let $m, d \in \mathbb{N}$, $p \in [1, \infty]$, and $B > 0$. Define $\Omega = [-B, B]^d$. Consider a network architecture $(d, N_1, \ldots, N_{L-1}, 1)$ with $L \geq 2$ and $N_{L-1} \geq 2$. Suppose that $\rho \in C^m(\mathbb{R}) \setminus C^{m+1}(\mathbb{R})$ and all derivatives of ρ up to order m are locally p-integrable and bounded on compact sets. Then:*

- *The set $\mathcal{RNN}_\rho^\Omega(d, N_1, \ldots, N_{L-1}, 1)$ is not closed in $W^{m-1,p}(\Omega)$.*

- *If additionally $\rho^{(m)}$ is absolutely continuous and the weak derivative $\rho^{(m+1)}$ exists and is in $L^p(\Omega)$, then $\mathcal{RNN}_\rho^\Omega(d, N_1, \ldots, N_{L-1}, 1)$ is not closed in $W^{m,p}(\Omega)$.*

Proof. See Appendix A.1. Similar to (Petersen et al., 2020), we construct a sequence of networks whose ρ-realizations converge to a target function that is not a ρ-realization of some network. In particular, we show order-$(m - 1)$ (or order-m) Sobolev convergence to a target function that is $(m - 1)$-times but not m-times differentiable, while (Petersen et al., 2020) show L^p convergence to a discontinuous step function. $\qquad\square$

Section 3.3 lists several commonly used activation functions and whether they satisfy the assumptions of Theorem 3.1 for some value of m.

Of course, convergence in order-$(m - 1)$ Sobolev norm is stronger than convergence in lower-order Sobolev norm, so $\mathcal{RNN}_\rho^\Omega(d, N_1, \ldots, N_{L-1}, 1)$ is not closed in lower-order Sobolev spaces either.

Corollary 3.2. *Let $d \in \mathbb{N}$ and $p \in [1, \infty]$. Suppose $\rho \in C^m(\mathbb{R}) \setminus C^{m+1}(\mathbb{R})$ with bounded derivatives up to order m. Then $\mathcal{RNN}_\rho^\Omega(d, N_1, \ldots, N_{L-1}, 1)$ is not closed in $W^{k,p}(\Omega)$ for any $k \in \{0, \ldots, m - 1\}$, where $\Omega = [-B, B]^d$.*

Proof. Note that convergence in $W^{m-1,p}(\Omega)$ implies convergence in $W^{k,p}(\Omega)$ for all $k \in \{0, \ldots, m - 1\}$. So we still have $f_n \to f$ in $W^{k,p}(\Omega)$ in the proof of Theorem 3.1, but $f \notin \mathcal{RNN}_\rho^\Omega(d, N_1, \ldots, N_{L-1}, 1)$. $\qquad\square$

In Theorem 3.1, we show that $\mathcal{RNN}_\rho^\Omega(d, N_1, \ldots, N_{L-1}, 1)$ is not closed in order-$(m-1)$ Sobolev spaces for $\rho \in C^m(\mathbb{R})$. For an analytic, bounded, and non-constant activation function ρ, we extend this result and prove that $\mathcal{RNN}_\rho^\Omega(d, N_1, \ldots, N_{L-1}, 1)$ is not closed in any order Sobolev spaces.

Theorem 3.3. *Let $d \in \mathbb{N}$, $p \in [1, \infty]$, and $B > 0$. Suppose that $\rho : \mathbb{R} \to \mathbb{R}$ is real analytic, bounded, and not constant, and that all derivatives $\rho^{(n)}$ of ρ are bounded. Then for all possible neural network architectures $(d, N_1, \ldots, N_{L-1}, 1)$ with $L \geq 2$ and $N_{L-1} \geq 2$ and all $k \in \mathbb{N}$, the set $\mathcal{RNN}_\rho^{[-B,B]^d}(d, N_1, \ldots, N_{L-1}, 1)$ is not closed in $W^{k,p}([-B, B]^d)$.*

Proof. See Appendix A.2. Using arguments similar to those in (Petersen et al., 2020), we construct a sequence of networks whose ρ-realizations converge in any order Sobolev norm to an unbounded function, which cannot be a ρ-realization of some network since ρ is bounded. Lemma A.3 in the proof is interesting in its own right, as it states that realized neural networks with analytic activation functions can approximate the coordinate projection maps to arbitrary accuracy in Sobolev norm. $\square$

Theorems 3.1 and 3.3 show that $\mathcal{RNN}_\rho^{[-B,B]^d}(d, N_1, \ldots, N_{L-1}, 1)$ is not closed in Sobolev spaces, although the smoothness of the activation function dictates the order of the Sobolev space in question. In particular, for $\rho \in C^m(\mathbb{R}) \setminus C^{m+1}(\mathbb{R})$, Theorem 3.1 shows nonclosedness in order-$(m-1)$ or order-m Sobolev spaces, while Theorem 3.3 gives nonclosedness in all orders of Sobolev spaces for analytic activation functions. Both results apply to the $p = \infty$ case, indicating that sets of realized neural networks are not closed with respect to uniform convergence of a network and its derivatives.

3.1   Closedness of Sets of Networks with Bounded Weights

The nonclosedness of sets of realized neural networks is undesirable if we only want to learn network target functions or if we want to prevent unbounded growth of network parameters. If we desire closedness, then we must modify the set of neural networks under consideration in some way. However, requiring closedness will necessarily constrain the set of target functions that we can approximate.

Modifications to enforce closedness of sets of realized neural networks may include relaxing some assumptions on the activation function or placing restrictions on the network parameters. In this section, we discuss the closedness of sets of realized neural networks whose parameters are all bounded by the same constant. We define a norm on $\mathcal{NN}(d, N_1, \ldots, N_L)$ and sets of realized neural networks with bounded norm.

Definition 3.4. *(Petersen et al., 2020) Let $C > 0$. Define*

$$\mathcal{NN}^C(d, N_1, \ldots, N_L) = \{\Phi \in \mathcal{NN}(d, N_1, \ldots, N_L) : \|\Phi\|_{total} \leq C\},$$

*as a set of neural networks with **uniformly bounded** weights, where*

$$\|\Phi\|_{total} = \max_{\ell=1,\ldots,L} \|A_\ell\|_{max} + \max_{\ell=1,\ldots,L} \|b_\ell\|_{max}$$

and $\|\cdot\|_{max}$ *equals the absolute value of the entry of largest magnitude from a matrix or vector. For* $\Omega \subset \mathbb{R}^d$ *and* $\rho : \mathbb{R} \to \mathbb{R}$*, also define*

$$\mathcal{RNN}_\rho^{\Omega,C}(d, N_1, \ldots, N_L) := R_\rho^\Omega\big(\mathcal{NN}^C(d, N_1, \ldots, N_L)\big)$$

as a set of realized neural networks with uniformly bounded weights and biases.

Petersen, Raslan, and Voigtlaender show that $\mathcal{RNN}_\rho^{[-B,B]^d}(d, N_1, \ldots, N_{L-1}, 1)$ is not closed in $L^p([-B,B]^d)$ or in $C([-B,B]^d)$ with respect to the L^∞ norm. However, sets of realized neural networks with uniformly bounded parameters *are* closed (in fact, compact) in these spaces by Proposition 3.5 in (Petersen et al., 2020).

Since Sobolev convergence is stronger than L^p convergence, $\mathcal{RNN}_\rho^{\Omega,C}(d, N_1, \ldots, N_L)$ is also closed in Sobolev spaces.

Corollary 3.5. *Let* $\Omega \subset \mathbb{R}^d$ *be compact,* $C > 0$*,* $p \in [1, \infty]$*,* $k \in \mathbb{N}$*, and* $\rho : \mathbb{R} \to \mathbb{R}$ *be continuous. Then* $\mathcal{RNN}_\rho^{\Omega,C}(d, N_1, \ldots, N_L)$ *is closed in* $W^{k,p}(\Omega)$*.*

Proof. If $(f_n)_{n \in \mathbb{N}} \subset \mathcal{RNN}_\rho^{\Omega,C}(d, N_1, \ldots, N_L)$ satisfies $\|f_n - f\|_{W^{k,p}(\Omega)} \to 0$ for some f, then $f_n \to f$ in L^p norm. Thus, $f \in \mathcal{RNN}_\rho^{\Omega,C}(d, N_1, \ldots, N_L)$ because this set is closed in $L^p(\Omega)$ by Proposition 3.5 in (Petersen et al., 2020). $\square$

The nonclosedness of $\mathcal{RNN}_\rho^\Omega$ in Sobolev spaces has significant consequences for approximating functions using neural networks. Indeed, it says that for any architecture $S = (d, N_1, \ldots, N_{L-1}, 1)$ with $L \geq 2$ and $N_{L-1} \geq 2$, there is a non-network target function $f \in \overline{\mathcal{RNN}_\rho^\Omega(S)} \setminus \mathcal{RNN}_\rho^\Omega(S)$, where the closure can be taken with respect to Sobolev norm of the appropriate order. Combined with the closedness of the set $\mathcal{RNN}_\rho^{\Omega,C}(S)$ with uniformly bounded weights, this means that if $\|R_\rho(\Omega)(\Phi_n) - f\|_{W^{k,p}(\Omega)} \to 0$ for some sequence of networks $(\Phi_n)_{n \in \mathbb{N}}$ with architecture S, then $\|\Phi_n\|_{total} \to \infty$. This may explain the phenomenon of weights growing without bound that sometimes occurs when training neural networks, and it indicates that using Sobolev training for neural networks can still lead to such a growth of parameters for some target functions.

3.2 Network Parameter Growth Rates

We have seen that convergence of a neural network to a non-network target function requires infinite growth of at least one network weight, but we would like to know how quickly the weights must grow relative to the decreasing approximation error. An interesting area of further research would be to describe this relationship for fixed architectures in full generality based on characteristics of the activation function, target function, and network architecture. We describe the approximation error decay for most activation functions in the case of a certain sequence of networks learning the derivative of the activation function.

Proposition 3.6. *Let* $\rho \in C^m(\mathbb{R})$ *for some* $m \geq 2$*, the softsign function, or ELU (cf. Table 1). Let* $(h_n)_{n=1}^\infty = (R_\rho^\mathbb{R}(\Phi_n))_{n=1}^\infty$ *be the sequence of realized neural networks in* $\mathcal{RNN}(1, 2, 1)$ *from the proof of Theorem 3.1. Further let* $p \in [1, \infty]$ *and let* $\Omega \subset \mathbb{R}$ *be a compact, measurable set with nonempty interior. Then* $\|h_n - \rho'\|_{L^p(\Omega)} \leq C_p/\|\Phi_n\|_{total}$ *for a constant* C_p *depending on* p *but not* n*.*

Proof. See Appendix A.3. □

In other words, Proposition 3.6 shows that the L^p approximation error is approximately inversely proportional to the networks' total norms for a sequence of networks learning the activation function derivative. Note that when the activation function is not analytic, its derivative is a non-network target function.

3.3 Commonly Used Activation Functions

Though some of the assumptions on the activation function required by Theorems 3.1 and 3.3 seem strong, they are satisfied by many commonly used activation functions. Thus, the results of these theorems apply, and sets of neural networks are not closed in various orders of Sobolev spaces for Ω compact and $p \geq 1$. Table 1 summarizes which results apply to several common activation functions.

Name	$\rho(x)$	Smoothness/ Boundedness	$\mathcal{RNN}_\rho^\Omega$ not closed in
Rectified Linear Unit (ReLU)	$\max\{0, x\}$	$C(\mathbb{R})$, abs. cont., $\rho' \in L^p(\Omega)$	$W^{0,p}(\Omega)$ [1]
Exponential Linear Unit (ELU)	$x \cdot \chi_{x \geq 0} + (e^x - 1) \cdot \chi_{x < 0}$	$C^1(\mathbb{R})$, ρ' abs. cont., $\rho'' \in L^p(\Omega)$	$W^{1,p}(\Omega)$
Softsign	$\frac{x}{1+|x|}$	$C^1(\mathbb{R})$, ρ' abs. cont., $\rho'' \in L^p(\Omega)$	$W^{1,p}(\Omega)$
Inverse Square Root Linear Unit ($a > 0$)	$x \cdot \chi_{x \geq 0} + \frac{x}{\sqrt{1+ax^2}} \cdot \chi_{x < 0}$	$C^2(\mathbb{R})$, ρ'' abs. cont., $\rho''' \in L^p(\Omega)$	$W^{2,p}(\Omega)$
Inverse Square Root Unit ($a > 0$)	$\frac{x}{\sqrt{1+ax^2}}$	real analytic, all derivatives bounded	$W^{k,p}(\Omega)$ for all k
Sigmoid	$\frac{1}{1+e^{-x}}$	real analytic, all derivatives bounded [2]	$W^{k,p}(\Omega)$ for all k
tanh	$\frac{e^x - e^{-x}}{e^x + e^{-x}}$	real analytic, all derivatives bounded [2]	$W^{k,p}(\Omega)$ for all k
arctan	$\arctan(x)$	real analytic, all derivatives bounded [3]	$W^{k,p}(\Omega)$ for all k

Table 1: Many activation functions used in practice satisfy some smoothness and boundedness properties so that Theorems 3.1 and 3.3 apply. Thus, $\mathcal{RNN}_\rho^\Omega$ is not closed in various orders of Sobolev spaces for Ω compact and $p \in [1, \infty]$. Some results in the table are found in the following references: [1] (Petersen et al., 2020) [2] (Minai and Williams, 1993) [3] (Adegoke and Layeni, 2010).

The smoothness and boundedness assumptions for the ReLU, exponential linear unit, softsign, and inverse square root linear unit can be checked by hand. The real analyticity of the other activation functions is established by properties from (Krantz and Parks, 2002).

For these activation functions, sets of realized neural networks are not closed in Sobolev spaces of certain orders. Thus, using Sobolev training still allows these networks to learn non-network target functions, although doing so will cause unbounded growth of parameters, which follows from Corollary 3.5.

4 Experimental Results

We now show some experimental results that demonstrate the nonclosedness of sets of realized neural networks in Sobolev spaces and examine the rate of parameter growth. Specifically, we use Sobolev training with an Adam optimizer (Kingma and Ba, 2017) to produce sequences of neural networks that approximate non-network target functions.

For activation functions that are m-times but not $(m+1)$-times differentiable, we know from the proof of Theorem 3.1 that there is a sequence of networks that converges in Sobolev norm to the derivative of the activation function. The derivative is not a realized neural network because it is only $(m-1)$-times differentiable. With this as motivation, each trial of our experiment trains a network to learn a randomly generated $(m-1)$-times differentiable target function in Sobolev norm. We consistently observe a rapidly decreasing approximation error and a fairly steady growth of network weights.
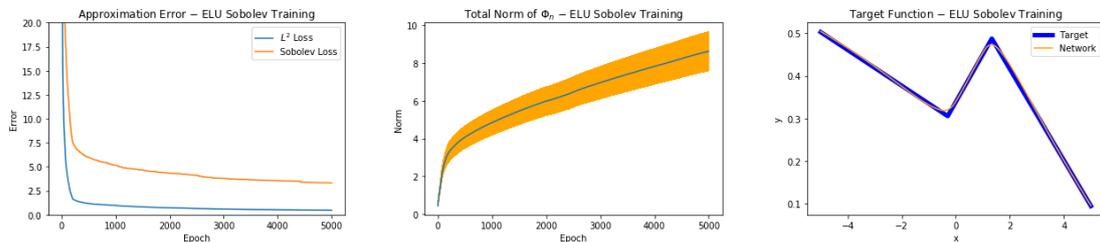
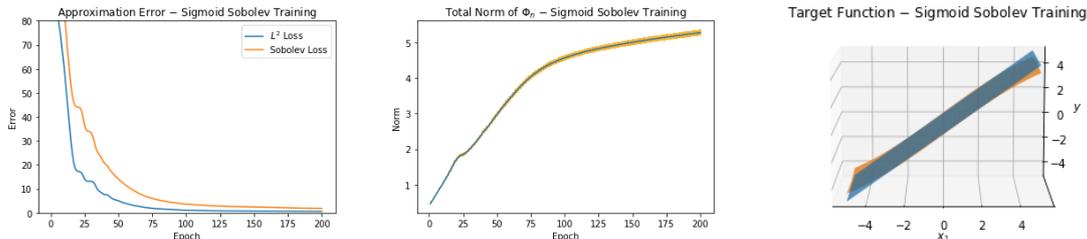

(a) Training error. (b) The network norm. (c) The target function.

Figure 1: For each of 100 target functions f, we train an ELU network Φ in $\mathcal{NN}(1, 10, 1)$ to minimize $\|R_\rho^{[-5,5]}(\Phi) - f\|_{W^{1,2}}$. (a) The best $W^{1,2}$ training loss achieved thus far is plotted at each epoch and averaged over all 100 experiments. (b) The total network norm is averaged over all 100 experiments with 95% confidence bands. (c) An example target function is plotted along with the realized neural network after training.

Figure 1 shows the results of 100 trails of ELU networks learning non-network target functions in Sobolev norm. Since ELU is C^1 but not C^2, we train the networks to learn randomly generated piecewise linear (C^0 but not C^1) functions in order-1 Sobolev space. We see that the L^2 and Sobolev approximation errors decrease rather quickly, while the total network norm increases quickly at first and then at a fairly steady rate.

These results are consistent with Theorem 3.1. We see that networks are able to closely approximate non-network target functions, which is evidence of the nonclosedness of sets of realized neural networks in Sobolev spaces. Appendix B provides similar results for the ISRLU activation function. Since ISRLU is C^2 and not C^3, we train these networks to learn randomly generated piecewise quadratic target functions in order-2 Sobolev norm.

For real analytic activation functions, Lemma A.3 indicates that realized neural networks can approximate coordinate projection maps to arbitrary accuracy. We trained a sigmoid neural network to learn the projection map $P_1 : \mathbb{R}^2 \to \mathbb{R}$ given by $P_1(x_1, x_2) = x_1$ in order-2 Sobolev norm. We reset the network weights and repeat the training process 100 times to assess whether sigmoid networks consistently learn P_1, which is a non-network function.

(a) Training error. (b) The network norm. (c) The target function.

Figure 2: In 100 repetitions, we train a sigmoid network Φ in $\mathcal{NN}(1, 10, 1)$ to minimize $\|R_\rho^{[-5,5]^2}(\Phi) - P_1\|_{W^{2,2}}$. (a) The best $W^{2,2}$ training loss achieved thus far is plotted at each epoch and averaged over all 100 experiments. (b) The total network norm is averaged over all 100 experiments with 95% confidence bands. (c) The target function (blue) is plotted along with the realized neural network (orange) after training.

Figure 2 shows the results of 100 trials of sigmoid networks learning P_1 in order-2 Sobolev norm. Note that P_1 is a non-network target function because it is unbounded, as discussed in the proof of Theorem 3.3. As in Figure 1, we see that the L^2 and Sobolev approximation errors decrease rather quickly, while the total network norm increases at a fairly steady rate.

These results are consistent with Theorem 3.3. We see that networks with analytic activation functions are able to closely approximate non-network target functions, which is further evidence of the nonclosedness of sets of realized neural networks in Sobolev spaces. Our experiments were done with shallow networks, but we expect the same nonclosedness to be demonstrated with deep networks since they are necessarily more expressive.

4.1 Network Parameter Growth Rates

The results in Figures 1 and 2 show that realized neural networks are able to approximate non-network target functions in Sobolev norm with growing parameters, but the relationship between the approximation error and parameter growth is not clear. In this section, we

provide numerical evidence for one case of Proposition 3.6 by showing that L^2 approximation error and the network's total norm are approximately inversely proportional for a softsign network learning the softsign derivative. We analyze the relationship by making a scatter plot of approximation error against the total norm. We do the same for the Sobolev error, even though we do not have a theoretical result for this case.
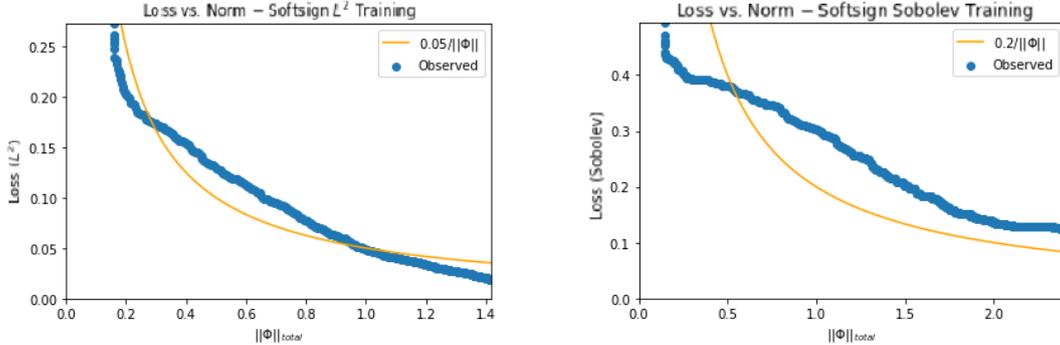
(a) L^2 error vs. network norm.

(b) Sobolev error vs. network norm.

Figure 3: We train a softsign network Φ in $\mathcal{NN}(1, 2, 1)$ to learn the softsign derivative and scatter approximation error against the total norm of Φ. (a) Training error vs. network norm using L^2 training. (b) Training error vs. network norm using order-1 Sobolev training.

Figure 3(a) shows that the relationship between L^2 approximation error and the total norm is approximately inversely proportional, as stated in Proposition 3.6. Note that our networks are trained using an Adam optimizer (Kingma and Ba, 2017) and thus do not follow the sequence $(h_n)_{n=1}^{\infty}$ from Proposition 3.6, but we still observe the expected relationship between approximation error and weights. The relationship for Sobolev error in Figure 3(b) is not as clear. An interesting question of further research is how the relationship between error and norm generalizes to other target functions, sequences of networks, network architectures, and higher-order Sobolev norms.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

12

References

Adegoke, K. and Layeni, O. (2010). The higher derivatives of the inverse tangent function and rapidly convergent BBP-type formulas for pi. *Applied Mathematics E-Notes*, 10:70–75.

Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals, and Systems*, 2(4):303–314.

Czarnecki, W. M., Osindero, S., Jaderberg, M., Swirszcz, G., and Pascanu, R. (2017). Sobolev training for neural networks. In *31st Conference on Neural Information Processing Systems (NIPS 2017)*, pages 1–10.

Folland, G. (1999). *Real Analysis*. Pure and Applied Mathematics. John Wiley & Sons, New York, 2 edition.

Girosi, F. and Poggio, T. (1990). Networks and the best approximation property. *Biological Cybernetics*, 63:169–176.

Gühring, I., Raslan, M., and Kutyniok, G. (2020). Expressivity of deep neural networks. arXiv:2007.04759v1.

Hornik, K. (1991). Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4(2):251–257.

Hornik, K., Stinchcombe, M., and White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366.

Hornik, K., Stinchcombe, M., and White, H. (1990). Universal approximation of an unknown mapping and its derivatives using multilayer feedforward networks. *Neural Networks*, 3:551–560.

Kainen, P. C., Kůrková, V., and Vogt, A. (1999). Approximation by neural networks is not continuous. *Neurocomputing*, 29:47–56.

Kainen, P. C., Kůrková, V., and Vogt, A. (2000). Best approximation by Heaviside perceptron networks. *Neural Networks*, 13:695–697.

Kainen, P. C., Kůrková, V., and Vogt, A. (2020). Approximative compactness of linear combinations of characteristic functions. *Journal of Approximation Theory*, 257:1–17.

Kingma, D. P. and Ba, J. L. (2017). Adam: A method for stochastic optimization. arXiv:1412.6980v9.

Krantz, S. G. and Parks, H. R. (2002). *A Primer of Real Analytic Functions*. Birkhäuser Advanced Texts. Birkhäuser, Boston, MA, 2 edition.

Krizhevsky, A., Sutskever, I., and Hinton, G. (2012). Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, 25(2):1097–1105.

Kůrková, V. (1995). Approximation of functions by perceptron networks with bounded number of hidden units. *Neural Networks*, 8(5):745–750.

Mhaskar, H. N. (1997). On smooth activation functions. In *Mathematics of Neural Networks*, pages 275–279. Springer.

Minai, A. A. and Williams, R. D. (1993). On the derivatives of the sigmoid. *Neural Networks*, 6(6):845–853.

Petersen, P., Raslan, M., and Voigtlaender, F. (2019). Unfavorable structural properties of the set of neural networks with fixed architecture. In *2019 13th International conference on Sampling Theory and Applications (SampTA)*, pages 1–4.

Petersen, P., Raslan, M., and Voigtlaender, F. (2020). Topological properties of the set of functions generated by neural networks of fixed size. *Foundations of Computational Mathematics*, pages 1–70.

A   Proofs of Results from Section 3

A.1   Proof of Theorem 3.1

The proof of Theorem 3.1 uses the following lemma.

Lemma A.1. *Let $d \in \mathbb{N}$, $L \geq 2$, $B > 0$, and $D > 0$. Set $\Omega = [-B, B]^d$. Suppose $\rho \in C^m(\mathbb{R}) \setminus C^{m+1}(\mathbb{R})$ for some $m \in \mathbb{N}$. Then there exists $\Phi \in \mathcal{NN}(d, 1, \ldots, 1)$ with $L - 1$ layers such that $R_\rho^\Omega(\Phi)(\Omega) \supset [-D, D]$.*

Proof of Lemma A.1. Since $\rho \in C^m(\mathbb{R}) \setminus C^{m+1}(\mathbb{R})$ for some $m \in \mathbb{N}$, ρ is not a constant function. Moreover, since $\rho \in C^1(\mathbb{R})$, there exists an open interval $U \subset \mathbb{R}$ such that $0 \notin \rho'(U)$. Hence, $\rho(U)$ is not a single point and must be an interval with non-empty interior. Next, choose $A_1 \in \mathbb{R}^{1 \times d}$ such that

$$A_1 \Omega \supset U.$$

Since U has non-empty interior, we can iteratively choose $A_j \in \mathbb{R}^{1 \times 1}$ such that

$$A_j \rho(A_{j-1} \rho(\cdots \rho(A_1 \Omega) \cdots)) \supset U$$

for $j = 2, \ldots, L - 2$. Again since U has non-empty interior, we can choose $A_{L-1} \in \mathbb{R}^{1 \times 1}$ and $b_{L-1} \in \mathbb{R}$ such that

$$A_{L-1} \rho(A_{L-2} \rho(\cdots \rho(A_1 \Omega) \cdots)) + b_{L-1} \supset [-D, D].$$

Finally, if we let

$$\Phi = \big((A_1, 0), \ldots, (A_{L-2}, 0), (A_{L-1}, b_{L-1})\big)$$

then $R_\rho^\Omega(\Phi)(\Omega) \supset [-D, D]$ by construction.   □

14

We will also use a lemma about the existence of strong L^p derivatives on bounded sets, which is similar to Exercise 8.9 in (Folland, 1999).

Lemma A.2. *Let $p \in [1, \infty]$ and let $K \subset \mathbb{R}$ be compact with $f \in L^p(K)$. If f is absolutely continuous on every bounded interval and its pointwise derivative f' is in $L^p(K)$, then*

$$\lim_{y \to 0} \left\| \frac{f(\cdot + y) - f(\cdot)}{y} - f'(\cdot) \right\|_{L^p(K)} = 0.$$

In this case, we call f' the strong L^p derivative of f on K.

Proof of Lemma A.2. Let $K \subset \mathbb{R}$ be compact with $f \in L^p(K)$. Suppose f is absolutely continuous on every bounded interval and its pointwise derivative f' is in $L^p(K)$. Since f is absolutely continuous on every bounded interval, we can write

$$\frac{f(x+y) - f(x)}{y} - f'(x) = \frac{1}{y} \int_x^{x+y} f'(t)dt - f'(x)$$

$$= \frac{1}{y} \int_0^y f'(x+t)dt - f'(x)$$

$$= \frac{1}{y} \int_0^y [f'(x+t) - f'(x)]dt.$$

For $p < \infty$, Minkowski's inequality for integrals gives

$$\left\| \frac{f(\cdot + y) - f(\cdot)}{y} - f'(\cdot) \right\|_{L^p(K)} = \left[\int_K \left(\frac{1}{y} \int_0^y [f'(x+t) - f'(x)]dt \right)^p dx \right]^{1/p}$$

$$\leq \frac{1}{y} \int_0^y \left(\int_K [f'(x+t) - f'(x)]^p dx \right)^{1/p} dt$$

$$= \frac{1}{y} \int_0^y \left\| f'(\cdot + t) - f'(\cdot) \right\|_{L^p(K)} dt.$$

For $p = \infty$, we can similarly write

$$\left\| \frac{f(\cdot + y) - f(\cdot)}{y} - f'(\cdot) \right\|_{L^\infty(K)} = \sup_{x \in K} \left| \frac{1}{y} \int_0^y [f'(x+t) - f'(x)]dt \right|$$

$$\leq \frac{1}{y} \int_0^y \sup_{x \in K} |f'(x+t) - f'(x)|dt$$

$$= \frac{1}{y} \int_0^y \left\| f'(\cdot + t) - f'(\cdot) \right\|_{L^\infty(K)} dt.$$

(We do not have to consider the essential supremum over K because f is absolutely continuous on K. Moreover, if $y < 0$ then the negative sign can be absorbed by switching the limits of integration, which still allows us to move the supremum inside without flipping the inequality.) Now let $\epsilon > 0$. Since $f' \in L^p(K)$, there exists $\delta > 0$ such that

$$\left\| f'(\cdot + t) - f'(\cdot) \right\|_{L^p(K)} \leq \epsilon$$

15

whenever $|t| < \delta$. If $|y| < \delta$ so we are only integrating over t with $|t| < \delta$, we have

$$\left\| \frac{f(\cdot + y) - f(\cdot)}{y} - f'(\cdot) \right\|_{L^p(K)} \leq \frac{1}{y} \int_0^y \epsilon \, dt = \epsilon.$$

Thus, $\lim_{y \to 0} \left\| y^{-1} \big(f(\cdot + y) - f(\cdot) \big) - f'(\cdot) \right\|_{L^p(K)} = 0$ as desired. $\qquad \square$

Proof of Theorem 3.1. Given $p \in [1, \infty]$ and $\rho \in C^m(\mathbb{R}) \setminus C^{m+1}(\mathbb{R})$ as in the statement of the theorem, we will construct a sequence of functions $(f_n)_{n=1}^\infty$ in $R_\rho^\Omega(d, 1, \ldots, 1, 2, 1)$ whose $W^{m-1,p}(\Omega)$-limit f is not in $C^m(\Omega)$. Since $\rho \in C^m(\mathbb{R})$ implies

$$\mathcal{RNN}_\rho^\Omega(d, 1, \ldots, 1, 2, 1) \subset C^m(\Omega),$$

we have $f \notin \mathcal{RNN}_\rho^\Omega(d, 1, \ldots, 1, 2, 1)$ and hence $\mathcal{RNN}_\rho^\Omega(d, 1, \ldots, 1, 2, 1)$ is not closed in $W^{m-1,p}(\Omega)$. Since $\tilde{N}_\ell \geq N_\ell$ for $\ell = 1, \ldots, L-1$ implies

$$\mathcal{RNN}_\rho^\Omega(d, N_1, \ldots, N_{L-1}, N_L) \subset \mathcal{RNN}_\rho^\Omega(d, \tilde{N}_1, \ldots, \tilde{N}_{L-1}, N_L)$$

by Lemma 2.5 in (Petersen et al., 2020), the nonclosedness of $\mathcal{RNN}_\rho^\Omega(d, N_1, \ldots, N_{L-1}, 1)$ holds for any architecture $(d, N_1, \ldots, N_{L-1}, 1)$ with $L \geq 2$ and $N_{L-1} \geq 2$.

To construct $(f_n)_{n=1}^\infty$, first note that $\rho \notin C^{m+1}(\mathbb{R})$ implies $\rho \notin C^{m+1}([-C, C])$ for some $C > 0$. Let $\Phi \in \mathcal{NN}(d, 1, \ldots, 1)$ have $L-1$ layers such that $J(x) := R_\rho^\Omega(\Phi)(x)$ satisfies $J(\Omega) \supset [-C, C]$ as in Lemma A.1. Define a sequence of neural networks $\Phi_n = \big((A_1^n, b_1^n), (A_2^n, b_2^n) \big) \in \mathcal{NN}(1, 2, 1)$ by

$$A_1^n = \begin{pmatrix} 1 \\ 1 \end{pmatrix} \in \mathbb{R}^{2 \times 1}, \quad b_1^n = \begin{pmatrix} 1/n \\ 0 \end{pmatrix} \in \mathbb{R}^2, \quad A_2^n = \begin{pmatrix} n & -n \end{pmatrix} \in \mathbb{R}^{1 \times 2}, \quad b_2^n = 0 \in \mathbb{R}^1.$$

Next define $h_n : \mathbb{R} \to \mathbb{R}$ by

$$h_n(x) = R_\rho^\mathbb{R}(\Phi_n)(x) = n\rho(x + 1/n) - n\rho(x) = \frac{\rho(x + 1/n) - \rho(x)}{1/n}$$

and let $f_n = h_n \circ J = R_\rho^\Omega(\Phi_n \bullet \Phi)$ so that

$$f_n(x) = \frac{\rho(J(x) + 1/n) - \rho(J(x))}{1/n}$$

for $n \in \mathbb{N}$. Notice that $f_n \in C^m(\mathbb{R}^d)$ for all $n \in \mathbb{N}$, and for any indices $i_1, \ldots, i_l$ (not necessarily distinct) with $l \leq m$, we have

$$\frac{\partial^l}{\partial x_{i_1} \cdots \partial x_{i_l}} f_n(x) = \sum_{\pi \in \Pi} \left(\frac{\rho^{(|\pi|)}(J(x) + 1/n) - \rho^{(|\pi|)}(J(x))}{1/n} \right) \cdot \prod_{B \in \pi} \frac{\partial^{|B|} J(x)}{\prod_{j \in B} \partial x_j}$$

by Faà di Bruno's formula (here the sum is taken over all set partitions π of the set $\{1, \ldots, l\}$ and the product is taken over all blocks B in the partition π). We then have

$$\lim_{n \to \infty} \frac{\partial^l}{\partial x_{i_1} \cdots \partial x_{i_l}} f_n(x) = \sum_{\pi \in \Pi} \left(\rho^{(|\pi|+1)}(J(x)) \right) \cdot \prod_{B \in \pi} \frac{\partial^{|B|} J(x)}{\prod_{j \in B} \partial x_j} = \frac{\partial^l}{\partial x_{i_1} \cdots \partial x_{i_l}} \rho'(J(x)) \quad (1)$$

pointwise for $l < m$. For $l = 0, \ldots, m-1$, $\rho^{(l+1)}$ is continuous and bounded on every bounded interval, so $\rho^{(l)}$ is Lipschitz continuous and therefore absolutely continuous on every bounded interval. Moreover, $\rho^{(l+1)} \in L^p_{\mathrm{loc}}(\mathbb{R})$ by assumption. Hence, by Lemma, A.2, the pointwise derivative of $\rho^{(l)}$ agrees with its strong L^p derivative on $J(\Omega)$. Thus, we have

$$\lim_{n \to \infty} \left\| \frac{\rho^{(l)}(J(\cdot) + 1/n) - \rho^{(l)}(J(\cdot))}{1/n} - \rho^{(l+1)}(J(\cdot)) \right\|_{L^p(\Omega)} = 0$$

for $l = 1, \ldots, m-1$. Since all terms in Equation (1) are bounded on Ω, it follows that

$$\lim_{n \to \infty} \left\| f_n - \rho'(J(\cdot)) \right\|_{W^{m-1,p}(\Omega)} = \lim_{n \to \infty} \sum_{|\alpha| \le m-1} \left\| D^\alpha \big(f_n - \rho'(J(\cdot)) \big) \right\|_{L^p(\Omega)} = 0.$$

That is, the sequence $(f_n)_{n=1}^\infty$ converges in $W^{m-1,p}(\Omega)$ to the function $\Omega \to \mathbb{R}$ given by $x \mapsto f(x) := \rho'(J(x))$. Since $[-C, C] \subset J(\Omega)$, we have $f \notin C^m(\Omega)$ and therefore $f \notin R^\Omega_\rho(d, 1, \ldots, 1, 2, 1)$. Hence, $R^\Omega_\rho(d, 1, \ldots, 1, 2, 1)$ is *not* closed in $W^{m-1,p}(\Omega)$.

If in addition $\rho^{(m)}$ is absolutely continuous and the weak derivative $\rho^{(m+1)}$ exists and is in $L^p_{\mathrm{loc}}(\mathbb{R})$, then Lemma A.2 gives

$$\lim_{n \to \infty} \left\| \frac{\rho^{(m)}(J(\cdot) + 1/n) - \rho^{(m)}(J(\cdot))}{1/n} - \rho^{(m+1)}(J(\cdot)) \right\|_{L^p(\Omega)} = 0$$

and therefore

$$\lim_{n \to \infty} \left\| f_n - \rho'(J(\cdot)) \right\|_{W^{m,p}(\Omega)} = \lim_{n \to \infty} \sum_{|\alpha| \le m} \left\| D^\alpha \big(f_n - \rho'(J(\cdot)) \big) \right\|_{L^p(\Omega)} = 0.$$

Hence, $R^\Omega_\rho(d, 1, \ldots, 1, 2, 1)$ is *not* closed in $W^{m,p}(\Omega)$ under this further assumption on ρ. $\square$

A.2 Proof of Theorem 3.3

The proof of Theorem 3.3 uses an intermediate result which states that neural networks can approximate coordinate projection maps in Sobolev norm under certain conditions on the activation function.

Lemma A.3. *Let $\rho : \mathbb{R} \to \mathbb{R}$ be real analytic, bounded, and not constant. Suppose that all derivatives $\rho^{(n)}$ of ρ are bounded. Then for every $d, L, k \in \mathbb{N}$, $p \in [1, \infty]$, $\epsilon > 0$, $B > 0$, and every $i \in \{1, \ldots, d\}$, one can construct a neural network $\tilde{\Phi}^B_{\epsilon,i,k,p} \in \mathcal{NN}(d, 1, \ldots, 1)$ with L layers such that*

$$\left\| R^{[-B,B]^d}_\rho (\tilde{\Phi}^B_{\epsilon,i,k,p}) - P_i \right\|_{W^{k,p}([-B,B]^d)} \le \epsilon,$$

where $P_i(x) = x_i$.

Proof of Lemma A.3. Without loss of generality, we only consider the case $\epsilon \le 1$. Set $\Omega := [-B, B]^d$ and $\epsilon' := \epsilon / \big((k+1)(2B)^{d/p} L \big)$ (in particular, $\epsilon' = \epsilon/(k+1)L$ for $p = \infty$).

17

Since ρ is in $C^\infty(\mathbb{R})$ and is not constant, there exists $z_0 \in \mathbb{R}$ such that $\rho'(z_0) \neq 0$. For each $C > 0$, define $\tilde{\Phi}_1^C := \big((A_1, b_1), (A_2, b_2)\big) \in \mathcal{NN}(d, 1, 1)$ by

$$A_1 = \begin{pmatrix} \frac{1}{C} & 0 & \cdots & 0 \end{pmatrix} \in \mathbb{R}^{1 \times d}, \quad b_1 = z_0 \in \mathbb{R}^1, \quad A_2 = \frac{C}{\rho'(z_0)} \in \mathbb{R}^{1 \times 1}, \quad b_2 = -\frac{C\rho(z_0)}{\rho'(z_0)} \in \mathbb{R}^1$$

so that

$$R_\rho^\Omega(\tilde{\Phi}_1^C)(x) = \frac{C}{\rho'(z_0)} \cdot \rho\left(\frac{x_1}{C} + z_0\right) - \frac{C\rho(z_0)}{\rho'(z_0)}$$

where $x = (x_1, \ldots, x_d)$. Notice that

$$\lim_{C \to \infty} R_\rho^\Omega(\tilde{\Phi}_1^C)(x) = \lim_{C \to \infty} x_1 \cdot \frac{1}{\rho'(z_0)} \cdot \frac{\rho(z_0 + x_1/C) - \rho(z_0)}{x_1/C} = x_1$$

pointwise. In fact, there exists some $C_0 > 0$ such that $|R_\rho^\Omega(\tilde{\Phi}_1^C)(x) - x_1| \leq \epsilon'$ for all $x \in (-B - L\epsilon, B + L\epsilon)^d$ and all $C \geq C_0$. To see why, notice that by the definition of the derivative there exists $\delta > 0$ such that

$$\left| \frac{\rho(z_0 + t) - \rho(z_0)}{t} - \rho'(z_0) \right| \leq \frac{|\rho'(z_0)| \cdot \epsilon'}{1 + B + L} \tag{2}$$

for all $t \in \mathbb{R}$ with $|t| \leq \delta$. Set $C_0 := (B + L)/\delta$ and let $C \geq C_0$ be arbitrary. Since $\epsilon \leq 1$, every $x \in (-B - L\epsilon, B + L\epsilon)^d$ satisfies $|x_1| \leq B + L$. If we set $t = x_1/C$, then

$$|t| = |x_1|/C \leq (B + L)/C \leq (B + L)/C_0 = \delta.$$

It follows that

$$\begin{aligned} |R_\rho^\Omega(\tilde{\Phi}_1^C)(x) - x_1| &= \left| \frac{C}{\rho'(z_0)} \cdot \rho\left(\frac{x_1}{C} + z_0\right) - \frac{C\rho(z_0)}{\rho'(z_0)} - x_1 \right| \\ &= \left| \frac{C}{\rho'(z_0)} \right| \left| \rho\left(z_0 + \frac{x_1}{C}\right) - \rho(z_0) - \rho'(z_0) \cdot \frac{x_1}{C} \right| \\ &= \left| \frac{C}{\rho'(z_0)} \right| \left| \rho\left(z_0 + t\right) - \rho(z_0) - \rho'(z_0)t \right| \\ &\leq \left| \frac{C}{\rho'(z_0)} \right| \cdot \frac{|\rho'(z_0)| \cdot \epsilon'}{1 + B + L} \cdot |t| \quad &\text{(by (2))} \\ &= \left| \frac{C}{\rho'(z_0)} \right| \cdot \frac{|\rho'(z_0)| \cdot \epsilon'}{1 + B + L} \cdot \left| \frac{x_1}{C} \right| \\ &= \frac{|x_1|}{1 + B + L} \cdot \epsilon' \\ &\leq \epsilon' \end{aligned}$$

for all $x \in (-B - L\epsilon, B + L\epsilon)^d$ and all $C \geq C_0$. We also have

$$\frac{\partial}{\partial x_1} R_\rho^\Omega(\tilde{\Phi}_1^C)(x) = \frac{1}{\rho'(z_0)} \cdot \rho'\left(\frac{x_1}{C} + z_0\right) \xrightarrow{C \to \infty} 1$$

18

pointwise. Note that ρ' is Lipschitz because ρ'' is bounded. Hence, if we define $C_1 := B \cdot \mathrm{Lip}(\rho')/(|\rho'(z_0)| \cdot \hat{\epsilon}_1)$ (where $\hat{\epsilon}_1 > 0$ will be chosen later) then

$$
\begin{aligned}
\left| \frac{\partial}{\partial x_1} R_\rho^\Omega(\tilde{\Phi}_1^C)(x) - 1 \right| &= \left| \frac{1}{\rho'(z_0)} \cdot \rho'\left(\frac{x_1}{C} + z_0 \right) - 1 \right| \\
&= \left| \frac{1}{\rho'(z_0)} \right| \left| \rho'\left(z_0 + \frac{x_1}{C} \right) - \rho'(z_0) \right| \\
&\leq \left| \frac{1}{\rho'(z_0)} \right| \mathrm{Lip}(\rho') \left| \frac{x_1}{C} \right| \\
&\leq \frac{B \cdot \mathrm{Lip}(\rho')}{|\rho'(z_0)|} \cdot \frac{1}{C_1} \\
&= \hat{\epsilon}_1
\end{aligned}
$$

for all $x \in \Omega$ and all $C \geq C_1$. We also have

$$
\frac{\partial^n}{\partial x_1^n} R_\rho^\Omega(\tilde{\Phi}_1^C)(x) = \frac{1}{C^{n-1}} \cdot \frac{1}{\rho'(z_0)} \cdot \rho^{(n)}\left(\frac{x_1}{C} + z_0 \right) \xrightarrow{C \to \infty} 0
$$

pointwise for $n \geq 2$. In fact, if we define $C_n := \max(1, \|\rho^{(n)}\|_\infty/(|\rho'(z_0)| \cdot \hat{\epsilon}_n))$ (where $\hat{\epsilon}_n > 0$ will be chosen later) then

$$
\begin{aligned}
\left| \frac{\partial^n}{\partial x_1^n} R_\rho^\Omega(\tilde{\Phi}_1^C)(x) - 0 \right| &= \left| \frac{1}{C^{n-1}} \cdot \frac{1}{\rho'(z_0)} \cdot \rho^{(n)}\left(\frac{x_1}{C} + z_0 \right) - 0 \right| \\
&= \frac{1}{C^{n-1}} \cdot \frac{1}{|\rho'(z_0)|} \cdot \left| \rho^{(n)}\left(\frac{x_1}{C} + z_0 \right) \right| \\
&\leq \frac{1}{C} \cdot \frac{1}{|\rho'(z_0)|} \cdot \|\rho^{(n)}\|_\infty \\
&\leq \frac{1}{C_n} \cdot \frac{1}{|\rho'(z_0)|} \cdot \|\rho^{(n)}\|_\infty \\
&\leq \hat{\epsilon}_n
\end{aligned}
$$

for all $x \in \Omega$ and all $C \geq C_n$. Set $C^* = \max(C_0, C_1, \ldots, C_k)$. Then

$$
\begin{aligned}
\left| R_\rho^\Omega(\tilde{\Phi}_1^{C^*})(x) - x_1 \right| &\leq \epsilon', \\
\left| \frac{\partial}{\partial x_1} R_\rho^\Omega(\tilde{\Phi}_1^{C^*})(x) - 1 \right| &\leq \hat{\epsilon}_1, \\
\left| \frac{\partial^n}{\partial x_1^n} R_\rho^\Omega(\tilde{\Phi}_1^{C^*})(x) - 0 \right| &\leq \hat{\epsilon}_n, \qquad\qquad n = 2, \ldots, k
\end{aligned}
$$

for all $x \in \Omega$. Furthermore, any partial derivatives involving $x_2, \ldots, x_n$ are identically zero. Next define a network $\tilde{\Phi}_2^{C^*} = \left((A_1', b_1'), (A_2', b_2') \right) \in \mathcal{NN}(1,1,1)$ by

$$
A_1' = \frac{1}{C^*} \in \mathbb{R}^{1 \times 1}, \quad b_1' = z_0 \in \mathbb{R}^1, \quad A_2' = \frac{C^*}{\rho'(z_0)} \in \mathbb{R}^{1 \times 1}, \quad b_2' = -\frac{C^* \rho(z_0)}{\rho'(z_0)} \in \mathbb{R}^1
$$

and define $\tilde{\Phi}^{C^*} = \tilde{\Phi}_2^{C^*} \bullet \cdots \bullet \tilde{\Phi}_2^{C^*} \bullet \tilde{\Phi}_1^{C^*} \in \mathcal{NN}(d, 1, \ldots, 1)$, where we take $L-2$ concatenations to get L layers. Then

$$R_\rho^\Omega(\tilde{\Phi}^{C^*})(x) = (\rho_{C^*} \circ \cdots \circ \rho_{C^*})(x_1)$$

where $\rho_{C^*}(z) = \frac{C^*}{\rho'(z_0)} \cdot \rho\left(\frac{z}{C^*} + z_0\right) - \frac{C^*\rho(z_0)}{\rho'(z_0)}$ is applied L times. Inductively, we have

$$\left|R_\rho^\Omega(\tilde{\Phi}^{C^*})(x) - x_1\right| \le L\epsilon' = \frac{\epsilon}{(k+1)(2B)^{d/p}} \tag{3}$$

by applying ρ_C L times. Working with the derivatives $\frac{\partial^n}{\partial x_1^n}\left(\rho_C \circ \cdots \circ \rho_C\right)(x_1)$ is not as simple because the derivatives of this composition of functions will involve applications of the chain rule and product rule. However, we only apply the chain rule and product rule finitely many times, so Equation (3) and the boundedness of all derivatives of ρ guarantee that $\hat{\epsilon}_1, \ldots, \hat{\epsilon}_k$ can be chosen so that

$$\left|\frac{\partial}{\partial x_1}R_\rho^\Omega(\tilde{\Phi}^{C^*})(x) - 1\right| \le \frac{\epsilon}{(k+1)(2B)^{d/p}}$$

$$\left|\frac{\partial^n}{\partial x_1^n}R_\rho^\Omega(\tilde{\Phi}^{C^*})(x) - 0\right| \le \frac{\epsilon}{(k+1)(2B)^{d/p}} \qquad n = 2, \ldots, k.$$

For $p < \infty$, it follows that

$$\left\|R_\rho^\Omega(\tilde{\Phi}^{C^*}) - P_1\right\|_{W^{k,p}(\Omega)} = \sum_{|\alpha| \le k} \left\|D^\alpha R_\rho^\Omega(\tilde{\Phi}^{C^*}) - D^\alpha P_1\right\|_{L^p(\Omega)}$$

$$= \sum_{n=0}^k \left\|\frac{\partial^n}{\partial x_1^n}R_\rho^\Omega(\tilde{\Phi}^{C^*}) - \frac{\partial^n}{\partial x_1^n}P_1\right\|_{L^p(\Omega)}$$

$$= \sum_{n=0}^k \left(\int_\Omega \left|\frac{\partial^n}{\partial x_1^n}R_\rho^\Omega(\tilde{\Phi}^{C^*})(x) - \frac{\partial^n}{\partial x_1^n}P_1(x)\right|^p dx\right)^{1/p}$$

$$\le \sum_{n=0}^k \left(\int_\Omega \left|\frac{\epsilon}{(k+1)(2B)^{d/p}}\right|^p dx\right)^{1/p}$$

$$= \sum_{n=0}^k \left((2B)^d \cdot \frac{\epsilon^p}{(k+1)^p(2B)^d}\right)^{1/p}$$

$$= \sum_{n=0}^k \frac{\epsilon}{k+1}$$

$$= \epsilon.$$

For $p = \infty$, we have

$$\left\|R_\rho^\Omega(\tilde{\Phi}^{C^*}) - P_1\right\|_{W^{k,\infty}(\Omega)} = \sum_{|\alpha| \le k} \left\|D^\alpha R_\rho^\Omega(\tilde{\Phi}^{C^*}) - D^\alpha P_1\right\|_{L^\infty(\Omega)}$$

$$= \sum_{n=0}^k \left\|\frac{\partial^n}{\partial x_1^n}R_\rho^\Omega(\tilde{\Phi}^{C^*}) - \frac{\partial^n}{\partial x_1^n}P_1\right\|_{L^\infty(\Omega)}$$

$$\leq \sum_{n=0}^{k} \frac{\epsilon}{k+1}$$

$$= \epsilon.$$

For $i = 2, \ldots, n$ we can just permute the columns of A_1 accordingly. $\qquad\square$

We will also need to consider Sobolev convergence of compositions of functions, for which the following lemma will be useful.

Lemma A.4. *Let $k \in \mathbb{N}$ and $p \in [1, \infty]$, and let $\Omega \subset \mathbb{R}^d$ be bounded and compact. Suppose $(g_n)_{n \in \mathbb{N}} \subset C^\infty(\Omega)$ is a sequence of functions such that $\|g_n - h\|_{W^{k,p}} \to 0$ for some $h \in C^\infty(\Omega)$. If $(f_n)_{n \in \mathbb{N}} \subset C^\infty(\Omega)$ is a sequence of functions whose derivatives are all bounded independent of n, then*

$$\lim_{n \to \infty} \|f_n \circ g_n - f_n \circ h\|_{W^{k,p}(\Omega)} = 0.$$

That is, if $g_n \to h$ in Sobolev norm, then $f_n \circ g_n \to f_n \circ h$ in Sobolev norm.

Proof of Lemma A.4. For any $\ell \leq k$, we have

$$\frac{\partial^l}{\partial x_1 \cdots \partial x_l}(f_n \circ g_n)(x) = \sum_{\pi \in \Pi} f_n^{(|\pi|)}(g_n(x)) \cdot \prod_{B \in \pi} \frac{\partial^{|B|} g_n(x)}{\prod_{j \in B} \partial x_j}$$

and

$$\frac{\partial^l}{\partial x_1 \cdots \partial x_l}(f_n \circ h)(x) = \sum_{\pi \in \Pi} f_n^{(|\pi|)}(h(x)) \cdot \prod_{B \in \pi} \frac{\partial^{|B|} h(x)}{\prod_{j \in B} \partial x_j}$$

by Faà di Bruno's formula, where the sum is taken over all set partitions π of the set $\{1, \ldots, l\}$ and the product is taken over all blocks B in the partition π. Subtracting these two expressions, we get a sum over $\pi \in \Pi$ of terms of the form

$$f_n^{(|\pi|)}(g_n(x)) \cdot \prod_{B \in \pi} \frac{\partial^{|B|} g_n(x)}{\prod_{j \in B} \partial x_j} - f_n^{(|\pi|)}(h(x)) \cdot \prod_{B \in \pi} \frac{\partial^{|B|} h(x)}{\prod_{j \in B} \partial x_j}$$

$$= f_n^{(|\pi|)}(g_n(x)) \cdot \prod_{B \in \pi} \frac{\partial^{|B|} g_n(x)}{\prod_{j \in B} \partial x_j} - f_n^{(|\pi|)}(g_n(x)) \cdot \prod_{B \in \pi} \frac{\partial^{|B|} h(x)}{\prod_{j \in B} \partial x_j} \qquad (4)$$

$$+ f_n^{(|\pi|)}(g_n(x)) \cdot \prod_{B \in \pi} \frac{\partial^{|B|} h(x)}{\prod_{j \in B} \partial x_j} - f_n^{(|\pi|)}(h(x)) \cdot \prod_{B \in \pi} \frac{\partial^{|B|} h(x)}{\prod_{j \in B} \partial x_j} \qquad (5)$$

(note that we added and subtracted a cross term). For the first term (4), we have

$$\left\| f_n^{(|\pi|)}(g_n(\cdot)) \prod_{B \in \pi} \frac{\partial^{|B|} g_n(\cdot)}{\prod_{j \in B} \partial x_j} - f_n^{(|\pi|)}(g_n(\cdot)) \prod_{B \in \pi} \frac{\partial^{|B|} h(\cdot)}{\prod_{j \in B} \partial x_j} \right\|_{L^p(\Omega)}$$

$$\leq C \left\| \prod_{B \in \pi} \frac{\partial^{|B|} g_n(\cdot)}{\prod_{j \in B} \partial x_j} - \prod_{B \in \pi} \frac{\partial^{|B|} h(\cdot)}{\prod_{j \in B} \partial x_j} \right\|_{L^p(\Omega)}$$

21

for some constant C since $f_n^{(|\pi|)}$ is bounded independent of n. Note that for each factor in the product over $B \in \pi$, we have

$$\lim_{n \to \infty} \left\| \frac{\partial^{|B|} g_n(\cdot)}{\prod_{j \in B} \partial x_j} - \frac{\partial^{|B|} h(\cdot)}{\prod_{j \in B} \partial x_j} \right\|_{L^p(\Omega)} = 0$$

since $\lim_{n \to \infty} \|g_n - h\|_{W^{k,p}(\Omega)} = 0$ and $|B| \le \ell \le k$. It follows that the entire product converges to 0 in L^p norm as $n \to \infty$. Indeed, for $\phi_1, \phi_2, \psi_1, \psi_2$ bounded and $p < \infty$,

$$\|\phi_1 \phi_2 - \psi_1 \psi_2\|_{L^p(\Omega)}^p = \int_\Omega |\phi_1(x)\phi_2(x) - \psi_1(x)\psi_2(x)|^p dx$$

$$= \int_\Omega |\phi_1(x)\phi_2(x) - \psi_1(x)\phi_2(x) + \psi_1(x)\phi_2(x) - \psi_1(x)\psi_2(x)|^p dx$$

$$\le 2^{p-1} \left(\int_\Omega |\phi_1(x)\phi_2(x) - \psi_1(x)\phi_2(x)|^p dx + \int_\Omega |\psi_1(x)\phi_2(x) - \psi_1(x)\psi_2(x)|^p dx \right)$$

$$\le 2^{p-1} C_{\psi_1,\phi_2} \left(\|\phi_1 - \psi_1\|_{L^p(\Omega)}^p + \|\phi_2 - \psi_2\|_{L^p(\Omega)}^p \right)$$

by Minkowski's inequality, where C_{ψ_1,ϕ_2} bounds both $|\psi_1|^p$ and $|\phi_2|^p$ on Ω. For $p = \infty$,

$$\|\phi_1 \phi_2 - \psi_1 \psi_2\|_{L^\infty(\Omega)} = \sup_{x \in \Omega} |\phi_1(x)\phi_2(x) - \psi_1(x)\psi_2(x)|$$

$$= \sup_{x \in \Omega} |\phi_1(x)\phi_2(x) - \psi_1(x)\phi_2(x) + \psi_1(x)\phi_2(x) - \psi_1(x)\psi_2(x)|$$

$$\le \sup_{x \in \Omega} \left(|\phi_2(x)||\phi_1(x) - \psi_1(x)| + |\psi_1(x)||\phi_2(x) - \psi_2(x)| \right)$$

$$\le C_{\psi_1,\phi_2} \left(\sup_{x \in \Omega} |\phi_1(x) - \psi_1(x)| + \sup_{x \in \Omega} |\phi_2(x) - \psi_2(x)| \right)$$

$$= C_{\psi_1,\phi_2} \left(\|\phi_1 - \psi_1\|_{L^\infty(\Omega)} + \|\phi_2 - \psi_2\|_{L^\infty(\Omega)} \right)$$

where C_{ψ_1,ϕ_2} bounds both $|\psi_1|$ and $|\phi_2|$ on Ω. Thus, if $\phi_{1,n} \to \psi_1$ and $\phi_{2,n} \to \psi_2$ in L^p norm as $n \to \infty$ with C_{ψ_1,ϕ_2} independent of n, then $\phi_{1,n}\phi_{2,n} \to \psi_1 \psi_2$ in L^p norm. For the second term (5), we have

$$\left\| f_n^{(|\pi|)}(g_n(\cdot)) \prod_{B \in \pi} \frac{\partial^{|B|} h(\cdot)}{\prod_{j \in B} \partial x_j} - f_n^{(|\pi|)}(h(\cdot)) \prod_{B \in \pi} \frac{\partial^{|B|} h(\cdot)}{\prod_{j \in B} \partial x_j} \right\|_{L^p(\Omega)}$$

$$\le C_h \left\| f_n^{(|\pi|)}(g_n(\cdot)) - f_n^{(|\pi|)}(h(\cdot)) \right\|_{L^p(\Omega)}$$

for some constant C_h since all derivatives of h are bounded because $h \in C^\infty(\Omega)$ with Ω compact. Since all derivatives of f_n are bounded independent of n, $f_n^{(|\pi|)}$ is Lipschitz with constant M independent of n. Thus, for $p < \infty$, we have

$$\left\| f_n^{(|\pi|)}(g_n(\cdot)) - f_n^{(|\pi|)}(h(\cdot)) \right\|_{L^p(\Omega)}^p = \int_\Omega |f_n^{(|\pi|)}(g_n(x)) - f_n^{(|\pi|)}(h(x))|^p dx$$

$$\le M^p \int_\Omega |g_n(x) - h(x)|^p dx$$

22

$$= M^p \|g_n - h\|_{L^p(\Omega)}^p.$$

For $p = \infty$, we have

$$\left\| f_n^{(|\pi|)}(g_n(\cdot)) - f_n^{(|\pi|)}(h(\cdot)) \right\|_{L^\infty(\Omega)} = \sup_{x \in \Omega} \left| f_n^{(|\pi|)}(g_n(x)) - f_n^{(|\pi|)}(h(x)) \right|$$

$$\leq M \sup_{x \in \Omega} |g_n(x) - h(x)|$$

$$= M \|g_n - h\|_{L^\infty(\Omega)}.$$

In either case, the second term converges to 0 as $n \to \infty$. It follows that

$$\lim_{n \to \infty} \left\| D^\alpha (f_n \circ g_n - f_n \circ h) \right\|_{L^p(\Omega)} = 0$$

for any multi-index α with $|\alpha| \leq k$, so $\lim_{n \to \infty} \|f_n \circ g_n - f_n \circ h\|_{W^{k,p}(\Omega)} = 0$. $\qquad\square$

Proof of Theorem 3.3. Set $\Omega := [-B, B]^d$. Since $\rho \in C^\infty(\mathbb{R})$ is not constant, there exists $z_0 \in \mathbb{R}$ such that $\rho'(z_0) \neq 0$. For each $n \in N$, define $\Phi_n^1 := ((A_1^n, b_1^n), (A_2^n, b_2^n)) \in \mathcal{NN}(1, 2, 1)$ by

$$A_1^n = \begin{pmatrix} 1 \\ 1/n \end{pmatrix} \in \mathbb{R}^{2 \times 1}, \quad b_1^n = \begin{pmatrix} 0 \\ z_0 \end{pmatrix} \in \mathbb{R}^2, \quad A_2^n = \begin{pmatrix} 1 & n \end{pmatrix} \in \mathbb{R}^{1 \times 2}, \quad b_2^n = -n\rho(z_0) \in \mathbb{R}^1$$

so that

$$R_\rho^{\mathbb{R}}(\Phi_n^1)(x) = \rho(x) + n\rho(x/n + z_0) - n\rho(z_0) = \rho(x) + x \cdot \frac{\rho(z_0 + x/n) - \rho(z_0)}{x/n}$$

for all $x \in \mathbb{R}$. By Lemma A.3, there exists a sequence of neural networks $(\Phi_n^2)_{n \in \mathbb{N}} \subset \mathcal{NN}(d, 1, \ldots, 1)$ with $L - 1$ layers such that

$$\left\| R_\rho^\Omega(\Phi_n^2) - P_1 \right\|_{W^{k,p}(\Omega)} \leq \frac{1}{n} \tag{6}$$

for each $n \in \mathbb{N}$. Define $\Phi_n = \Phi_n^1 \bullet \Phi_n^2 \in \mathcal{NN}(d, 1, \ldots, 1, 2, 1)$. Consider the map $F : \mathbb{R}^d \to \mathbb{R}$ defined by $F(x) = \rho(x_1) + \rho'(z_0)x_1$. We have

$$\left\| R_\rho^\Omega(\Phi_n) - F \right\|_{W^{k,p}((-B,B)^d)} = \left\| R_\rho^\Omega(\Phi_n) - R_\rho^{\mathbb{R}}(\Phi_n^1) \circ P_1 + R_\rho^{\mathbb{R}}(\Phi_n^1) \circ P_1 - F \right\|_{W^{k,p}(\Omega)}$$

$$\leq \left\| R_\rho^{\mathbb{R}}(\Phi_n^1) \circ R_\rho^\Omega(\Phi_n^2) - R_\rho^{\mathbb{R}}(\Phi_n^1) \circ P_1 \right\|_{W^{k,p}(\Omega)}$$

$$+ \left\| R_\rho^{\mathbb{R}}(\Phi_n^1) \circ P_1 - F \right\|_{W^{k,p}(\Omega)}.$$

Since $\frac{d^\ell}{dx^\ell} R_\rho^{\mathbb{R}}(\Phi_n^1)(x) = \rho^{(\ell)}(x) + \rho^{(\ell)}(x/n + z_0)/n^{\ell-1}$, each derivative of $R_\rho^{\mathbb{R}}(\Phi_n^1)$ is bounded independent of n. Combined with equation (6) and the boundedness of all derivatives of ρ, this implies that

$$\left\| R_\rho^{\mathbb{R}}(\Phi_n^1) \circ R_\rho^\Omega(\Phi_n^2) - R_\rho^{\mathbb{R}}(\Phi_n^1) \circ P_1 \right\|_{W^{k,p}(\Omega)} \xrightarrow{n \to \infty} 0$$

23

by Lemma A.4. Next note that

$$R_\rho^{\mathbb{R}}(\Phi_n^1) \circ P_1(x) = \rho(x_1) + x_1 \cdot \frac{\rho(z_0 + x_1/n) - \rho(z_0)}{x_1/n} \xrightarrow{n \to \infty} \rho(x_1) + \rho'(z_0)x_1 = F(x)$$

pointwise. An argument similar to the one in the proof of Theorem 3.1 shows that

$$\left\| R_\rho^{\mathbb{R}}(\Phi_n^1) \circ P_1 - F \right\|_{W^{k,p}(\Omega)} \xrightarrow{n \to \infty} 0$$

since the pointwise derivative agrees with the strong L^p derivative. In total, we have

$$\left\| R_\rho^\Omega(\Phi_n) - F \right\|_{W^{k,p}(\Omega)} \xrightarrow{n \to \infty} 0.$$

However, $F|_\Omega \notin \mathcal{RNN}(d, 1, \ldots, 1, 2, 1)$. To see why suppose we had $F|_\Omega = R_\rho^\Omega(\Psi)$ for some neural network Ψ. Since ρ is analytic, F and $R_\rho^{\mathbb{R}^d}(\Psi)$ are both analytic and coincide on $\Omega = [-B, B]^d$. Hence, we must have $F = R_\rho^{\mathbb{R}^d}(\Psi)$ on all of $\mathbb{R}^d$. Note that $F(x) = \rho(x_1) + \rho'(z_0)x_1$ is unbounded because ρ is bounded and $\rho'(z_0) \neq 0$. However, $R_\rho^{\mathbb{R}^d}(\Psi)$ is bounded because ρ is. This contradicts $F = R_\rho^{\mathbb{R}^d}(\Psi)$, so $F|_\Omega \neq R_\rho^\Omega(\Psi)$. Since

$$\left\| R_\rho^\Omega(\Phi_n) - F \right\|_{W^{k,p}((-B,B)^d)} \xrightarrow{n \to \infty} 0$$

with $R_\rho^\Omega(\Phi_n) \in \mathcal{RNN}(d, 1, \ldots, 1, 2, 1)$ but $F|_\Omega \notin \mathcal{RNN}(d, 1, \ldots, 1, 2, 1)$, we have shown that $\mathcal{RNN}(d, 1, \ldots, 1, 2, 1)$ is not closed. $\qquad\square$

A.3 Proof of Proposition 3.6

Let $\rho : \mathbb{R} \to \mathbb{R}$ be arbitrary, and define $(h_n)_{n=1}^\infty = (R_\rho^{\mathbb{R}}(\Phi_n))_{n=1}^\infty$ to be the realized networks in $\mathcal{RNN}(1, 2, 1)$ as in the proof of Theorem 3.1 (even if ρ does not satisfy the hypotheses of the theorem). Then

$$\|\Phi_n\|_{total} = \max\{1, n\} + \max\{0, 1/n\} = n + (1/n) < 2n.$$

Thus, our goal will be to prove for compact, measurable $\Omega \subset \mathbb{R}$ with non-empty interior and various ρ that there exists C_p, dependent on ρ, p, and possibly Ω but independent of n such that

$$\left\| h_n - \rho' \right\|_{L^p(\Omega)}^p \leq C_p n^{-p}.$$

For $p = \infty$, we will show directly that

$$\left\| h_n - \rho' \right\|_{L^\infty(\Omega)} \leq C n^{-1}$$

for some C independent of n. For some of the ρ, the bound will be over all of $\mathbb{R}$ rather than over Ω.

Let $\rho \in C^m(\mathbb{R})$ for some $m \geq 2$. Since Ω is compact, $\|\rho''\|_{L^\infty(\widetilde{\Omega})}$ is finite, where $\widetilde{\Omega} = \{x : x \text{ or } x - (1/n) \in \Omega\}$. Let $x \in \Omega$. It follows from a first order Taylor approximation that there exists $\xi \in (x, x + (1/n))$ such that

$$\left| h_n(x) - \rho'(x) \right| = \left| \frac{\rho(x + (1/n)) - \rho(x)}{1/n} - \rho'(x) \right| = \frac{1}{2n}|\rho''(\xi)| \leq \frac{1}{2n}\|\rho''\|_{L^\infty(\widetilde{\Omega})}.$$

Thus, for $p < \infty$, we have

$$\left\|h_n - \rho'\right\|_{L^p(\Omega)}^p \leq \left(\frac{1}{2n}\left\|\rho''\right\|_{L^\infty(\widetilde{\Omega})}\right)^p |\Omega| =: C_p n^{-p}$$

where $|\Omega|$ is the measure of Ω and C_p is independent of n. For $p = \infty$, we have

$$\left\|h_n - \rho'\right\|_{L^\infty(\Omega)} \leq \frac{1}{2n}\left\|\rho''\right\|_{L^\infty(\widetilde{\Omega})} =: Cn^{-1}$$

where C is independent of n.

Now let ρ be the softsign function $\rho(x) = x/(1 + |x|)$ from Table 1. Then $\rho'(x) = 1/(1 + |x|)^2$. It follows that

$$\left|h_n(x) - \rho'(x)\right| = \left|\frac{\rho(x + (1/n)) - \rho(x)}{1/n} - \rho'(x)\right|$$

$$= \left|n\left(\frac{x + (1/n)}{1 + |x + (1/n)|} - \frac{x}{1 + |x|}\right) - \frac{1}{(1 + |x|)^2}\right|$$

$$= \begin{cases} \frac{1}{n(1+x+(1/n))(1+x)^2} & x \geq 0 \\ \frac{|2nx^3 + 2(1-n)x^2 - 4x - (1/n)|}{(1+x+(1/n))(1-x)^2} & -1/n \leq x < 0 \\ \frac{1}{n(1-x-(1/n))(1-x)^2} & x < -1/n. \end{cases}$$

For $p < \infty$, we have

$$\left\|h_n - \rho'\right\|_{L^p(\Omega)}^p$$

$$= \int_0^\infty \left(\frac{1}{n(1 + x + (1/n))(1 + x)^2}\right)^p dx + \int_{-1/n}^0 \left(\frac{|2nx^3 + 2(1 - n)x^2 - 4x - (1/n)|}{(1 + x + (1/n))(1 - x)^2}\right)^p dx$$

$$+ \int_{-\infty}^{-1/n} \left(\frac{1}{n(1 - x - (1/n))(1 - x)^2}\right)^p dx$$

$$=: A + B + C.$$

We begin by bounding $n^p A$:

$$n^p A = \int_0^\infty \frac{1}{(1 + x + (1/n))^p (1 + x)^{2p}} dx < \int_0^\infty \frac{1}{(1 + x)^{3p}} dx = \frac{1}{3p - 1}. \tag{7}$$

Similarly,

$$n^p C = \int_{-\infty}^{-1/n} \frac{1}{(1 - x - (1/n))^p (1 - x)^{2p}} dx < \int_{-\infty}^{-1/n} \frac{1}{(1 - x - (1/n))^{3p}} dx = \frac{1}{3p - 1}. \tag{8}$$

We now seek to bound B. Let $q(x) = 2nx^3 + 2(1 - n)x^2 - 4x - (1/n)$ so that $|q(x)|^p$ is the numerator of the integrand of B. We would like to bound $|q(x)|$ over $[-1/n, 0]$. Note that $q(0) = -1/n$ and $q(-1/n) = 1/n$. Also,

$$q'(x) = 6nx^2 + 4(1 - n)x - 4,$$

25

and thus $q'(x)$ is zero when

$$x = \frac{n - 1 \pm \sqrt{1 + 4n + n^2}}{3n}.$$

Since $(n - 1 + \sqrt{1 + 4n + n^2})/(3n) > 2/3 > 0$, we only care about the case $x_0 = (n - 1 - \sqrt{1 + 4n + n^2})/(3n)$. We claim that $q(x_0) > 0$ and thus an upper bound on $q(x_0)$ is an upper bound on $|q(x_0)|$. To that end, we note that

$$\frac{-1}{n} = \frac{n - 1 - \sqrt{4 + 4n + n^2}}{3n} \leq x_0 \leq \frac{n - 1 - \sqrt{1 + 2n + n^2}}{3n} = \frac{-2}{3n},$$

where $q(-1/n) = 1/n$ and $q(-2/(3n)) = 2/(9n^2) + 7/(9n)$. Since q'' is negative over the interval $[-1/n, 0)$, this implies that q is positive over the interval $[-1/n, -2/(3n)]$, which includes x_0. So we now bound $q(x_0)$:

$$\begin{aligned}
27n^2 q(x_0) &= -4n^3 - 24n^2 - 3n + 4 + (4n^2 + 16n + 4)\sqrt{n^2 + 4n + 1} \\
&\leq -4n^3 - 24n^2 - 3n + 4 + (4n^2 + 16n + 4)\sqrt{n^2 + 4n + 4} \\
&= 32n + 14.
\end{aligned}$$

As

$$\frac{64}{27n} > \frac{32n}{27n^2}\left(1 + \frac{14}{32n}\right) = \frac{32n + 14}{27n^2} > \frac{1}{n},$$

it is an upper bound for $|q(x)|$ over $[-1/n, 0]$. We now bound B:

$$\begin{aligned}
\left(\frac{27n}{64}\right)^p B &\leq \int_{-1/n}^0 \frac{1}{((1 + x + (1/n))(1 - x)^2)^p} dx \\
&\leq \int_{-1/n}^{-1/(2n)} \frac{1}{(1 + x + (1/n))^{3p}} dx + \int_{-1/(2n)}^0 \frac{1}{(1 - x)^{3p}} dx \\
&= 2 \int_1^{1 + 1/(2n)} \frac{1}{x^{3p}} dx = \frac{2}{3p - 1}\left(1 - \left(1 + \frac{1}{2n}\right)^{1 - 3p}\right) \\
&\leq \frac{2}{3p - 1}.
\end{aligned} \tag{9}$$

Putting together (7), (8), and (9), we obtain for any measurable $\Omega \subseteq \mathbb{R}$

$$\left\|h_n - \rho'\right\|_{L^p(\Omega)}^p \leq \left\|h_n - \rho'\right\|_{L^p(\mathbb{R})}^p \leq \frac{1}{(3p - 1)}\left(\frac{1}{n^p} + \frac{1}{n^p} + 2\left(\frac{64}{27n}\right)^p\right) =: C_p n^{-p},$$

where C_p does not depend on n. For $p = \infty$, we can use the piecewise expression for $|h_n(x) - \rho'(x)|$ and the bound for $|q(x)|$ to see that

$$|h_n(x) - \rho'(x)| \leq \begin{cases} \frac{1}{n} & x \geq 0 \\ \frac{64}{27n} & -1/n \leq x < 0 \\ \frac{1}{n} & x < -1/n. \end{cases}$$

26

Thus, we have
$$\left\|h_n - \rho'\right\|_{L^\infty(\Omega)} \le \frac{64}{27n} =: Cn^{-1}$$
where C is independent of n.

Finally, let ρ be the ELU function $\rho(x) = x\chi_{x\ge 0} + (e^x - 1)\chi_{x<0}$ from Table 1. Then $\rho'(x) = \chi_{x\ge 0} + e^x\chi_{x<0}$ and for all $x \ne 0$, $\rho''(x) = e^x\chi_{x<0}$.

We first compute for $x \ge 0$:
$$\left|h_n(x) - \rho'(x)\right| = \left|\frac{\rho(x + (1/n)) - \rho(x)}{1/n} - \rho'(x)\right| = \left|\frac{x + (1/n) - x}{1/n} - 1\right| = 0. \qquad (10)$$

For $x < -1/n$, we may use Taylor approximation to prove the existence of a $\xi \in (x, x + (1/n))$ such that
$$\left|h_n(x) - \rho'(x)\right| = \left|\frac{\rho(x + (1/n)) - \rho(x)}{1/n} - \rho'(x)\right| = \frac{1}{2n}\left|\rho''(\xi)\right| = \frac{1}{2n}e^\xi \le \frac{e^{1/n}}{2n}e^x,$$

where for $p < \infty$,
$$\left\|\frac{e^{1/n}}{2n}e^\cdot\right\|_{L^p((-\infty,-1/n])}^p = \int_{-\infty}^{-1/n}\left(\frac{e^{1/n}}{2n}e^x\right)^p dx = \left(\frac{e^{1/n}}{2n}\right)^p \int_{-\infty}^{-1/n} e^{px}dx = \left(\frac{e^{1/n}}{2n}\right)^p \frac{1}{p}e^{-p/n}$$
$$= \frac{1}{2^p p}\frac{1}{n^p}. \qquad (11)$$

What remains is to analyze $|h_n(x) - \rho'(x)|$ for $-1/n \le x < 0$. To this end, we compute
$$\left|h_n(x) - \rho'(x)\right| = \left|\frac{\rho(x + (1/n)) - \rho(x)}{1/n} - \rho'(x)\right| = \left|\frac{(x + (1/n)) - (e^x - 1)}{1/n} - e^x\right|$$
$$= |nx + 1 - ne^x + n - e^x| = |(n+1)(1 - e^x) + nx|$$
$$=: |g(x)|.$$

We will use Taylor's theorem over $[-1/n, 0)$ both to show that $g(x)$ is non-negative and to give an upper bound. To this end, for each $x \in [-1/n, 0)$, there exists an $\xi \in (x, 0)$ such that
$$g(x) = (n+1)\left(-x - \frac{e^\xi}{2}x^2\right) + nx = -x - (n+1)\frac{e^\xi}{2}x^2 = -x\left(1 + (n+1)\frac{e^\xi}{2}x\right).$$

Since $x < 0$, $-x > 0$ and $1 + (n+1)e^\xi x/2 < 1$. Thus $-x$ is an upper bound of g over $[-1/n, 0)$. All that is left to show is $1 + (n+1)e^\xi x/2 > 0$. As $e^\xi < e^0 = 1$,
$$0 < 1 + \frac{n+1}{2} \cdot \frac{-1}{n} < 1 + \frac{(n+1)e^\xi}{2}x,$$

as desired. This yields
$$\|g\|_{L^p([-1/n,0)}^p = \int_{-1/n}^0 |g(x)|^p dx \le \int_{-1/n}^0 |x|^p dx = \frac{1}{(p+1)n^{p+1}} < \frac{1}{(p+1)n^p}. \qquad (12)$$

Putting (10), (11), and (12) together, we obtain

$$\left\| h_n - \rho' \right\|_{L^p(\Omega)}^p < \frac{1}{(p+1)n^p} + \frac{1}{2^p p}\frac{1}{n^p} =: C_p n^{-p},$$

where C_p does not depend on n. For $p = \infty$, we have seen that $|h_n(x) - \rho'(x)| \leq \frac{1}{n}$ for all x, and therefore $\|h_n - \rho'\|_{L^\infty(\Omega)} \leq n^{-1}$ as desired.

B More Experimental Results

In this section, we include experimental results using the ISRLU activation function. Since ISRLU is C^2 but not C^3, we train the networks to learn randomly generated piecewise quadratic (C^1 but not C^2) functions in order-2 Sobolev space.
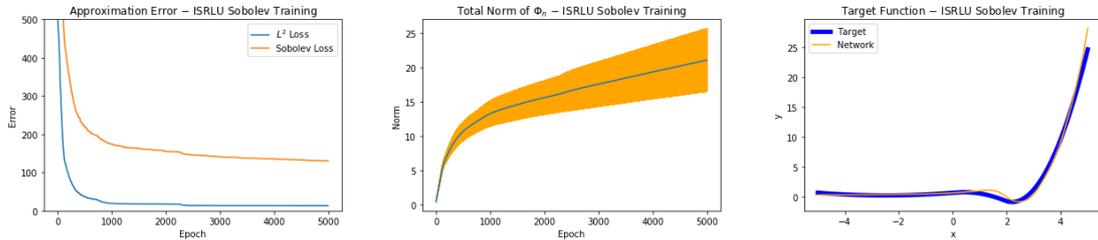
(a) Training error.   (b) The network norm.   (c) The target function.

Figure 4: For each of 100 target functions f, we train an ISRLU network Φ in $\mathcal{NN}(1, 10, 1)$ to minimize $\|R_\rho^{[-5,5]}(\Phi) - f\|_{W^{2,2}}$. (a) The best $W^{2,2}$ training loss achieved thus far is plotted at each epoch and averaged over all 100 experiments. (b) The total network norm is averaged over all 100 experiments with 95% confidence bands. (c) An example target function is plotted along with the realized neural network after training.

Figure 4 shows the results of 100 trails of ISRLU networks learning non-network target functions in Sobolev norm. We see that the L^2 and Sobolev approximation errors decrease rather quickly, though the Sobolev error is still decreasing, indicating that the networks are still learning the target functions. As with the ELU experiments, the total network norm increases quickly at first and then at a fairly steady rate. In this case, however, the 95% confidence bands are wider because the norm grew fairly large in a few cases and did not grow much in some other cases.

These results are again consistent with Theorem 3.1. We see that networks are able to closely approximate non-network target functions, which is evidence of the nonclosedness of sets of realized neural networks in Sobolev spaces. These experiments demonstrate nonclosedness for a C^2 activation function in order-2 Sobolev space. Combining this with the ELU and sigmoid results, we see that these numerical observations of nonclosedness and parameter growth hold for varying degrees of smoothness of the activation function.