

Approximation, Estimation and Optimization Errors for a Deep Neural Network

B. Keene

Department of Mathematics
University of Central Florida
Orlando, FL 32816, USA

keene@ucf.edu

G. Welper

Department of Mathematics
University of Central Florida
Orlando, FL 32816, USA

gerrit.welper@ucf.edu

Reviewed on OpenReview: <https://openreview.net/forum?id=dzND5haNvA>

Abstract

The error of supervised learning is typically split into three components: approximation, estimation and optimization errors. While all three have been extensively studied in the literature, a unified treatment is less frequent, in part because of conflicting assumptions. Current approximation results rely on carefully hand crafted weights or practically unavailable information, which are difficult to achieve by gradient descent. Optimization theory is best understood in over-parametrized regimes with more weights than samples, while classical estimation errors require the opposite regime with more samples than weights. This paper contains two results which bound all three error components simultaneously for (non-convex) training of the second but last layer of deep fully connected networks on the unit sphere. The first uses a regular least squares loss and shows convergence in the under-parametrized regime. The second uses a kernel based loss function and shows convergence in both under and over-parametrized regimes.

1 Introduction

In this paper, we consider supervised learning of fully connected neural networks without bias: For network f_θ with weights θ and normalized training samples (x_i, y_i) on the d -dimensional sphere, we minimize the loss

$$\ell(\theta) = \frac{1}{2N} \sum_{i=1}^N |f_\theta(x_i) - y_i|^2$$

by gradient descent. We also consider alternative losses, which allow more flexibility with regard to the number of samples and network size. The main results provide a complete error analysis including *approximation errors*, *estimation errors* and *optimization errors* (Shalev-Shwartz & Ben-David, 2014).

1.1 Literature Review

Approximation Error: If the data points $y_i = f(x_i)$ are generated by some unknown target function f , how well can the network approximate it, i.e. how large is the error $\inf_\theta \|f_\theta - f\|_{L_2(D)}$ on some domain D , ignoring error contributions from sampling and optimization algorithms? Typical results establish bounds

$$\inf_\theta \|f_\theta - f\|_{L_2(D)} \leq cn(\theta)^{-r}, \quad f \in K, \quad (1)$$

for some rate $r > 0$, with an asymptotic rate $n(\theta)^{-r}$, where $n(\theta)$ is a complexity measure of the network, e.g. width, depth or total number of weights. Any quantifiable rate requires some prior conditions on f , here given by membership in a compact set K , which typically consists of functions with bounded Sobolev, Besov, Barron or other smoothness norms.

First results show universal approximation properties (Cybenko, 1989; Hornik et al., 1989; Barron, 1993; Zhou, 2020; Lu et al., 2017; Hanin & Sellke, 2017) and reproduction of classical approximation rates for targets with Sobolev and Besov regularity (Gribonval et al., 2022; Gühring et al., 2020; Opschoor et al., 2020; Li et al., 2019; Suzuki, 2019). More recent papers provide super-convergence results, where networks outperform classical methods, as well as optimality benchmarks like manifold width, for the price of discontinuous weight assignments (Yarotsky, 2017; 2018; Yarotsky & Zhevnerchuk, 2020; Daubechies et al., 2022; Shen et al., 2019; Lu et al., 2021). Approximation results with smoothness requirements more tailored to neural networks use Barron and related spaces (Bach, 2017; Klusowski & Barron, 2018; Weinan et al., 2022; Li et al., 2020; Siegel & Xu, 2020; 2022a; Bresler & Nagaraj, 2020). Several surveys are given in (Pinkus, 1999; DeVore et al., 2021; Weinan et al., 2020; Berner et al., 2022).

In the above results, the weights are not trained by practical optimizers, but rather hand-picked or sampled from practically unknown distributions. Practical networks do not always achieve the theoretical bounds Adcock & Dexter (2021); Grohs & Voigtlaender (2023). The papers Jentzen & Riekert (2022); Ibragimov et al. (2022); Gentile & Welper (2022); Welper (2024b;a) are closely related to the results in this paper and provide an analysis of approximation errors in combination with gradient descent training. Finally Siegel & Xu (2022b); Siegel et al. (2023); Beck et al. (2022); Herrmann et al. (2022) consider approximation or estimation with alternative optimizers like greedy algorithms.

Generalization and Estimation errors: Practically, one can neither evaluate nor optimize the $L_2(D)$ error directly, and therefore trains the sample or empirical loss, resulting in the empirical risk minimizer $\hat{\theta}$. The resulting expected loss is called the *generalization error* and split into the *approximation error* and a remaining *estimation error*, (Shalev-Shwartz & Ben-David, 2014):

$$\|f_{\hat{\theta}} - f\|_{L_2(D)}^2 = \inf_{\theta} \|f_{\theta} - f\|_{L_2(D)}^2 + \left(\|f_{\hat{\theta}} - f\|_{L_2(D)}^2 - \inf_{\theta} \|f_{\theta} - f\|_{L_2(D)}^2 \right).$$

There are many techniques to bound generalization and estimation errors, for example by establishing uniform bounds between the expected and sample loss

$$\sup_{\theta} \left| \|f_{\theta} - f\|_{L_2(D)}^2 - \frac{1}{2N} \sum_{i=1}^N |f_{\theta}(x_i) - f(x_i)|^2 \right| \lesssim \mathcal{C} + N^{-1/2} \quad (2)$$

for some complexity measure \mathcal{C} of the neural networks like VC-dimension or Rademacher complexity. VC-dimension bounds of the form $\text{VC-dim} \lesssim \tilde{\mathcal{O}}(Ln(\theta))$ for total number of weights $n(\theta)$ and depth L are in Neyshabur et al. (2017); Bartlett et al. (1998); Harvey et al. (2017). Bounds for the Rademacher complexity tend to be independent of the size of the network as e.g. $\mathcal{O}\left(\frac{\sqrt{L} \prod_{\ell=1}^L \|W^{\ell}\|_F}{\sqrt{N}}\right)$ for weight W^{ℓ} in the Frobenius norm from Golowich et al. (2018). Similar bounds are in Neyshabur et al. (2015); Liang et al. (2019); Tu et al. (2020). While the norm $\|W^{\ell}\|_F$ may grow for wide networks with standard scaling, newer results depend on the difference $\|W^{\ell} - W_0^{\ell}\|$ to some reference or initial weights W_0^{ℓ} in Frobenius and other matrix norms, which tends to be small in the over-parametrized limit and leads to some combined generalization and gradient descent convergence results (Cao & Gu, 2020). Other Rademacher complexity bounds rely on smoothness (Weinan et al., 2022; 2019). Further techniques include margin theory (Jakubovitz et al., 2019; Neyshabur et al., 2018; Bartlett et al., 2017) mutual information (Asadi et al., 2018; Steinke & Zakynthinou, 2020) compression (Arora et al., 2018) and Besov regularity (Suzuki, 2019). Empirical observations show that against conventional wisdom neural networks generalize well in over-parametrized regimes, with enough weights to fit random data, (Neyshabur et al., 2017; Zhang et al., 2017; Geiger et al., 2019).

Optimization Error: Both the empirical loss minimizer $\hat{\theta} = \operatorname{argmin} \ell(\theta)$ and the approximation error $\inf_{\theta} \|f_{\theta} - f\|_{L_2(D)}^2$ build on an optimization problem. Since these are non-convex, can we practically compute

the minimizers, or at least a sufficiently good substitute? This question is addressed in the optimization literature for neural networks. The approach in this paper relies on the neural tangent kernel (NTK) coined in Jacot et al. (2018) and introduced simultaneously in Li & Liang (2018); Allen-Zhu et al. (2019); Du et al. (2019b;a). The concept is further developed in Zou et al. (2020); Arora et al. (2019a;b); Su & Yang (2019); Lee et al. (2019); Song & Yang (2019); Zou & Gu (2019); Kawaguchi & Huang (2019); Chizat et al. (2019); Oymak & Soltanolkotabi (2020); Ji & Telgarsky (2020); Nguyen & Mondelli (2020); Bai & Lee (2020); Chen et al. (2021); Song et al. (2021); Lee et al. (2022); Gentile & Welper (2022); Welper (2024b;a). In this paper we use lower bounds for the NTK that originate from Bietti & Mairal (2019); Geifman et al. (2020); Ji et al. (2020); Chen & Xu (2021). The optimization literature contains many other approaches that we only mention briefly: Landscape analysis (Nguyen & Hein, 2017; Ge et al., 2018; Du & Lee, 2018; Soltanolkotabi et al., 2019; Venturi et al., 2019), Wasserstein gradient flow (Soudry & Carmon, 2016; Safran & Shamir, 2018; Chizat & Bach, 2018; Mei et al., 2018; Rotskoff & Vanden-Eijnden, 2018; Sirignano & Spiliopoulos, 2020), as well as several overviews (Weinan et al., 2020; Berner et al., 2022; Roberts et al., 2022).

Generalization and Optimization: The definition of stochastic gradient descent intertwines optimization with sampling and the literature provides positive effects on generalization. This can be analyzed in abstract settings, e.g. the paper Hardt et al. (2016) bounds the generalization error between empirical and expected risk of SGD iterates of general convex and non-convex objective functions. In addition, it bounds the expected loss between SGD and empirical risk minimizers in the convex case.

Other papers are more specific to neural networks, i.e. Cao & Gu (2020); Nitanda & Suzuki (2021) consider a combination of estimation and optimization in over-parametrized regimes. The latter proves generalization errors for stochastic gradient descent by comparing its dynamics with an NTK idealization, yielding minimax optimal SGD convergence rates. In addition, it does not require any lower bounds on the NTK eigenvalues, similar to Welper (2024b) and the current paper. The papers Wang & Ma (2023); Liu et al. (2022); Park et al. (2022); Neu et al. (2021) find improved estimation bounds by explicitly incorporating gradient descent dynamics. Another closely related set of results is Drews & Kohler (2022); Kohler & Krzyzak (2022), which consider all three error contributions and control the optimization error based on the contributions of the final layer.

1.2 New Contributions

This paper contains generalization error bounds of the form

$$\|f_\theta - f\|_{L_2(\mathbb{S}^{d-1})}^2 \lesssim m^{-a} + m^b N^{-c},$$

where m is the width of the fully connected network f_θ , N is the number of samples and $a, b, c \geq 0$ are constants specified more closely in the main Theorems 2.2, 2.4. The weights θ on the left hand side are the output of gradient descent training on two different discrete losses. For simplicity, we confine the data to uniform samples on the unit sphere and train only the non-convex second but last layer. The new contributions of this bound are as follows:

Less Over-Parametrization: While over-parametrization in *total number of weights* is common in practical neural networks, current optimization theory heavily relies on much larger networks with more *width* than samples $m \gg N$. In contrast, successful architectures like AlexNet Krizhevsky et al. (2012) or ResNet He et al. (2016) are much thinner $m < N$ (including channels, width and height).

Our result does not require any relation between m and N . For $b > 0$, this provides small bounds in $m < N$ regimes (including under-parametrization). For a kernel based loss, Theorem 2.4 achieves $b = 0$, resulting in small bounds for any regime $m \lesssim N$ (including under- and over-parametrization). To the best of our knowledge, this is the first result that shows gradient descent convergence for deep neural networks independent of the relation between network size and number of samples.

Unified Analysis: The given bound contains approximation errors (m^{-a}), estimation errors ($m^b N^c$) and optimization errors (θ is trained by GD). These bounds can be found in the literature in isolation, but

rarely in one single theorem because of conflicting assumptions. The optimization literature requires over-parametrization so that all data $y_i = f(x_i)$ can be correctly reproduced by the network $y_i = f_\theta(x_i)$. This directly entails maximally large Rademacher complexity and hence conflicts with common assumptions for estimation error bounds. In the infinite sample limit considered in approximation theory, this also confines the target functions f to neural networks themselves, instead of rich classes of practical relevance. Therefore, the approximation literature ignores the optimization problem in favour of simpler hand-picked weights.

Only a few papers in the literature can balance all required assumptions and provide bounds for approximation, estimation and optimization in one single theorem. The papers Drews & Kohler (2022); Kohler & Krzyzak (2022), use non-standard architectures, which allow a derivation of error bounds from the last layer alone. This is a convex training objective, in contrast to this paper, which analyzes the training contributions from non-convex layers of standard fully connected networks. The papers Cao & Gu (2020); Nitanda & Suzuki (2021) include estimation and optimization errors, but require excessive over-parametrization as is typical in the current literature.

Estimation: The prior work Gentile & Welper (2022); Welper (2024b;a) establishes approximation and optimization error bounds. This paper contributes a corresponding analysis of the estimation errors, based on two alternative approaches. The first follows more traditional lines and bounds the complexity of the network directly. The second relies on smoothness bounds that arise from gradient descent training.

1.3 Overview

We use the following two techniques to show estimation errors.

Network Complexity: The most common complexity measures are VC-dimension and Rademacher complexity, which in turn can be bounded by chaining techniques, i.e. by Dudley’s inequality (Shalev-Shwartz & Ben-David, 2014). In our case, it is convenient to skip the Rademacher complexity and use Dudley’s inequality directly because it has already been used to establish NTK concentration inequalities for the approximation and optimization error bounds (Welper, 2024b).

We minimize the sample loss

$$\ell(\theta) = \frac{1}{2} \sum_{i=1}^N |f_\theta(x_i) - f(x_i)|^2 \quad (3)$$

for N uniformly random normalized samples x_i on the unit sphere \mathbb{S}^{d-1} with gradient descent. We show that as long as the error does not satisfy the approximation and estimation error estimate

$$\|f_\theta - f\|_{L_2(\mathbb{S}^{d-1})}^2 \lesssim m^{-a} + m^b N^{-c},$$

with network width m , the gradient descent error decreases exponentially. The rates a , b and c are specified in Theorem 2.2 below and depend on the Sobolev smoothness of the target function f . Although we optimize the discrete sample loss, we bound the error in the continuous $L_2(\mathbb{S}^{d-1})$ norm and therefore obtain the expected or generalization error. The result is comparable to standard machine learning theory. In particular, the second term requires that the number of samples N is larger than the width m of the network (up to some power).

Smoothness: While requiring more samples N than width m matches common wisdom in machine learning theory, it does not explain the empirical observation that neural networks generalize well in over-parametrized regimes. To establish generalization error bounds in this regime, we rely on a different complexity measure: The approximation and optimization results in Welper (2024b;a) establish that the Sobolev norm of the gradient descent iterates $\|f_{\theta^n} - f\|_{H^s(\mathbb{S}^{d-1})}$ remains uniformly bounded, independent of the size of the network. If $s > 1 + d/2$, Sobolev embedding theorems imply that f_θ is uniformly Lipschitz and therefore the estimation error bound (2) can be proven by uniform laws of large numbers (Vershynin, 2018, Section 8.2).

Unfortunately, the current theory only provides bounds for $s < 1/2$, insufficient for the argument. We may, however, proceed with the *kernel loss*

$$\ell^k(\theta) = \frac{1}{2N} \sum_{i=1}^N \langle k(x_i, \cdot), f_\theta - f \rangle^2, \quad (4)$$

with uniformly random x_i , which probes the residual $f_\theta - f$ with an integral kernel $k(x, y)$, $x, y \in \mathbb{S}^{d-1}$ in the L_2 inner product $\langle \cdot, \cdot \rangle$ and is easier to bound in low regularity settings. Moreover, for common kernels like the heat kernel, Gaussian kernel $e^{-|x-y|^2/\sigma^2}$ or Laplacian kernel $e^{-|x-y|/\sigma}$, this loss converges to the standard mean squared loss (3) for $\sigma \rightarrow 0$ and proper normalization.

Although our interest in this kernel loss is of theoretical nature, to explore new arguments for generalization in over-parametrized regimes, it is similar to variational losses in VPINNs (Kharazmi et al., 2019; 2021), used to solve PDEs with neural networks. In this application, it is common that PDE solutions do not admit continuous point evaluations, and instead one probes the residual $\langle f_\theta - f, v \rangle$ with test functions v from some linear subspace, for which the given kernels would be one example.

The kernel loss also bears a resemblance with randomized smoothing (Cohen et al., 2019): In order to mitigate adversarial attacks on a classifier f_θ for \mathcal{Y} classes, these methods choose the class that is most likely under normal perturbations $\epsilon \sim \mathcal{N}(0, \delta^2)$

$$g(x) = \arg \max_{c \in \mathcal{Y}} \Pr[f_\theta(x + \epsilon) = c].$$

In comparison, for a Gaussian kernel with variance δ^2 , the kernel loss is identical to the mean squares loss of the averaged network

$$g_\theta(x) = \mathbb{E}[f_\theta(x + \epsilon)] = \langle k(x, \cdot), f_\theta \rangle.$$

The second main result shows that for the kernel loss, gradient descent decreases exponentially until it reaches the approximation and estimation error

$$\|f_\theta - f\|_{L_2(\mathbb{S}^{d-1})}^2 \lesssim m^{-a} + N^{-c},$$

again with rates a and c dependent on the Sobolev regularity of f as specified in Theorem 2.4. The two error contributions on the left hand side are decoupled and we achieve the worst case of the approximation error m^{-a} and the sample error N^{-c} . Contrary to the first result and conventional machine learning theory, this allows meaningful generalization errors even in over-parametrized regimes with more samples N than width m .

Estimation and Gradient Descent: To obtain the results, we do not bound the difference between empirical and expected loss (2) directly. Instead, we compare the gradient descent evolution to an idealized method trained on the continuous L_2 loss:

$$\begin{aligned} \theta^n : & \quad \text{trained by gradient descent on loss (3) or (4).} \\ \bar{\theta}^n : & \quad \text{trained by gradient descent on loss } \frac{1}{2} \|f_{\bar{\theta}} - f\|_{L_2(\mathbb{S}^{d-1})}^2. \end{aligned}$$

Convergence of the latter is established in Welper (2024a). From this we prove convergence for the former based on perturbation analysis and sample errors for the respective *gradients* (not the loss as in standard analysis (2)). This approach is reminiscent of Cohen et al. (2002), which analyzes adaptive PDE solvers by comparing them with idealized infinite dimensional ones. The sample errors are established by Dudley's inequality for the sample loss (3) and by matrix Bernstein inequalities for the kernel loss (4).

2 Main Results

This section contains the main results of the paper.

2.1 Setup

The setup is almost identical to Welper (2024b;a), with the major difference that the references train on an idealized continuous $L_2(\mathbb{S}^{d-1})$ loss, whereas we train on practical sample losses.

Notations We denote generic constants by c , which may be different in each occurrence, but do not depend on the width m , number of samples N or input dimension d . Alternatively, we use the shorthand $a \lesssim b$, $a \gtrsim b$, $a \sim b$ to denote $a \leq cb$, $a \geq cb$, $a \lesssim b \lesssim a$, respectively.

For integer s , Sobolev spaces $H^s(\mathbb{S}^{d-1})$ consist of all functions on \mathbb{S}^{d-1} with $L_2(\mathbb{S}^{d-1})$ bounded weak derivatives of order s . For non-integer s , these spaces can be defined by the decay of their expansion

$$\|f\|_{H^\alpha(\mathbb{S}^{d-1})}^2 = \sum_{l=0}^{\infty} \sum_{j=1}^{\nu(l)} \left(1 + l^{1/2}(l+d-2)^{1/2}\right)^{2\alpha} \left|\langle Y_l^j, f \rangle\right|^2 \quad (5)$$

in spherical harmonics

$$Y_\ell^j, \quad \ell = 0, 1, 2, \dots, \quad 1 \leq j \leq \nu(\ell), \quad (6)$$

for suitable numbers $\nu(\ell)$, see e.g. Barceló et al. (2020). We denote the corresponding inner product by $\langle \cdot, \cdot \rangle_{H^s(\mathbb{S}^{d-1})}$.

Network We consider fully connected networks without bias

$$\begin{aligned} f^1(x) &= W^0 x, \\ f^{\ell+1}(x) &= W^\ell m_\ell^{-1/2} \sigma(f^\ell(x)), \quad \ell = 1, \dots, L \end{aligned} \quad (7)$$

of depth L , with normalized inputs in the unit sphere $x \in D := \mathbb{S}^{d-1}$ and standard scaling. We summarize all trainable weights in the parameter θ and abbreviate the network by $f_\theta(x) := f^{L+1}(x)$. To obtain a simple non-convex model problem, we optimize the second but last layer W^{L-1} and initialize all weights randomly

$W^L \in \{-1, +1\}^{1 \times m_L}$	i.i.d. Rademacher	not trained,
$W^{L-1} \in \mathbb{R}^{m_L \times m_{L-1}}$,	i.i.d. $\mathcal{N}(0, 1)$	trained
$W^\ell \in \mathbb{R}^{m_{\ell+1} \times m_\ell}$, $\ell = 1, \dots, L-2$	i.i.d. $\mathcal{N}(0, 1)$	not trained
$W^0 \in \mathbb{R}^{m_1 \times d}$,	i.i.d. $\mathcal{N}(0, 1)$	not trained.

All hidden layers are of comparable size, the input d -dimensional and the output scalar:

$$m := m_{L-1}, \quad 1 = m_{L+1} \leq m_L \sim \dots \sim m_1 \geq d.$$

Activation Functions We require smooth activation functions with no more than linear growth and derivatives bounded as follows:

$$|\sigma(x)| \lesssim |x|, \quad |\sigma^{(i)}(x)| \lesssim 1 \quad i = 1, 2, \quad |\sigma^{(j)}(x)| \leq p(x), \quad j = 3, 4, \quad (8)$$

for some polynomial $p(x)$.

Training All networks are trained by gradient descent

$$\theta^{n+1} = \theta^n - \gamma \nabla_\theta \ell(\theta^n), \quad (9)$$

with learning rate $\gamma > 0$. We use different losses $\ell(\theta)$ for the main results and define them in the respective sections.

Neural Tangent Kernel The main results require coercivity of the neural tangent kernel, which has been shown in Welper (2024b) for ReLU activations based on Bietti & Mairal (2019); Geifman et al. (2020); Chen & Xu (2021), but remains open for smoother activations (8) used in this paper. Since we only train the second but last layer, in our case the *neural tangent kernel* (NTK) is informally defined as

$$\Gamma(x, y) = \lim_{\text{width} \rightarrow \infty} \sum_{r=1}^R \partial_r f_\theta(x) \partial_r f_\theta(y), \quad (10)$$

with partial derivatives $\partial_{W_{ij}^{L-1}}$ abbreviated by a single index ∂_r with $r = 1, \dots, R := m_L m_{L-1}$. The coercivity condition is then stated as

$$\langle f, Hf \rangle_{H^s(\mathbb{S}^{d-1})} \gtrsim \|f\|_{H^{s-\beta}(\mathbb{S}^{d-1})}, \quad Hf := \int_D \Gamma(\cdot, y) f(y) dy \quad (11)$$

for some $\beta > 0$, all $S \in \{0, s\}$, some smoothness level $0 \leq s \leq \frac{\beta}{2}$ and all $f \in H^s(\mathbb{S}^{d-1})$. Again, for networks with ReLU activations, bias and all layers trained this is true with $\beta = d/2$, see Welper (2024b).

In addition, the main results require

$$\Sigma^k(1) \neq 0, \quad k = 1, \dots, L, \quad (12)$$

for the Gaussian process that describes the forward evaluation of the random initial network in the infinite width limit (Jacot et al., 2018). Its correlation matrices $\Sigma(x, y) = \Sigma(x^T y)$ only depend on the angle between $x, y \in \mathbb{S}^{d-1}$ and are, defined by

$$\Sigma^{\ell+1}(x, y) := \mathbb{E}_{u, v \sim \mathcal{N}(0, A)} [\sigma(u), \sigma(v)], \quad A = \begin{bmatrix} \Sigma^\ell(x, x) & \Sigma^\ell(x, y) \\ \Sigma^\ell(y, x) & \Sigma^\ell(y, y) \end{bmatrix}, \quad \Sigma^0(x, y) = x^T y,$$

As for coercivity, this property is known for ReLU activations, where $\Sigma(1) = 1$, see Chen & Xu (2021), and is expected to be a minor technical assumption for smoother activations (8). The condition is directly related to the NTK coercivity and with it left for future work.

2.2 Result I: Pointwise Sampling

For the first result, we use the standard least squares loss

$$\ell(\theta) = \frac{1}{2} \frac{1}{N} \sum_{i=1}^N [f_\theta(x_i) - f(x_i)]^2, \quad (13)$$

with N independent uniform samples $x_i \in \mathbb{S}^{d-1}$. We first collect all major assumptions.

Assumption 2.1. *Assume:*

1. The neural network (7) - (8) is trained by gradient descent (9).
2. The NTK satisfies coercivity (11) for $0 \leq 2s \leq \beta$ and the forward process satisfies (12).
3. All hidden layers are of similar size: $m_0 \sim \dots \sim m_{L-1} =: m$.
4. Smoothness is bounded by $0 < s < 1/2$.
5. Define h and τ as follows and choose learning rate γ and an arbitrary α so that

$$h = c_h m^{-\frac{1}{2} \frac{1}{1+\alpha}}, \quad \tau = h^{2\alpha} m, \quad \gamma \lesssim h \sqrt{m}, \quad 0 \leq \alpha < 1 - s.$$

for some constant c_h that may depend on the initial error $\|f_{\theta^0} - f\|_{L_2(\mathbb{S}^{d-1})}$.

The following result is similar to Welper (2024a, Theorem 2.2), which only considers approximation and optimization errors. While the reference trains on the continuous $L_2(\mathbb{S}^{d-1})$ loss, we train on the discrete sample loss and therefore also include estimation errors.

Theorem 2.2. *Assume we train the sample loss (13), let Assumption 2.1 be satisfied, let $\|f\|_{L_\infty(\mathbb{S}^{d-1})} \lesssim m^{1/2}$, and define*

$$\Delta_{\text{sample}}(m, N) = c_\Delta \frac{m^{3/2}}{N^{1/2}} h^{1-\frac{\alpha s}{2\beta}} \|f_{\theta^0} - f\|_{H^s(\mathbb{S}^{d-1})}^{-1}$$

for some sufficiently large c_Δ . Then with residual $\kappa^n := f_{\theta^n} - f$ and probability at least $1 - cL(e^{-m} + e^{-\tau})$, while the gradient descent error exceeds the final approximation and estimation error

$$\|\kappa^k\|_{L_2(\mathbb{S}^{d-1})}^2 \geq c_a \left(m^{-\frac{1}{2} \frac{\alpha}{1+\alpha}} + \Delta_{\text{sample}}(m, N) \right)^{\frac{s}{\beta}} \|\kappa^0\|_{H^s(\mathbb{S}^{d-1})}^2, \quad k < n, \quad (14)$$

we have

$$\|\kappa^n\|_{L_2(\mathbb{S}^{d-1})}^2 \leq C e^{-\gamma[h^\alpha + \Delta_{\text{sample}}(m, N)]n} \|\kappa^0\|_{L_2(\mathbb{S}^{d-1})}^2, \quad \|\kappa^n\|_{H^s(\mathbb{S}^{d-1})}^2 \leq C \|\kappa^0\|_{H^s(\mathbb{S}^{d-1})}^2.$$

for sufficiently large constants c_a , c and C independent of m , κ^0 and κ^n .

The proof is in Section B.2. The assumptions relate the smoothness, the size of the network and number of samples. The only major assumption is the coercivity (11), (12), which is open for our activations but can be easily inferred from the literature (Bietti & Mairal, 2019; Geifman et al., 2020; Chen & Xu, 2021) for ReLU activations. In the latter case, β depends on the input dimension d and therefore all other bounds are also dimension dependent, although this is not explicit in the stated results. See Welper (2024b) for details, and numerical verification for smoother activations required in the theorem.

The result shows that gradient descent converges exponentially fast until the error is sufficiently small (14) and we have

$$\|\kappa^n\|_{L_2(\mathbb{S}^{d-1})}^2 < c_a \left(m^{-\frac{1}{2} \frac{\alpha}{1+\alpha}} + \Delta_{\text{sample}}(m, N) \right)^{\frac{s}{\beta}} \|\kappa^0\|_{H^s(\mathbb{S}^{d-1})}^2, \quad (15)$$

The first summand, $m^{-\frac{1}{2} \frac{\alpha}{1+\alpha}}$ together with the smoothness $\|\kappa^0\|_{H^s(\mathbb{S}^{d-1})}^2$ provides a typical approximation error bound of the form (1). The second term Δ_{sample} , bounds the sample error and has the typical $N^{-1/2}$ dependence on the number of samples together with some factors of m and h that measure the complexity of the network. In particular, for the overall error to be small, we must have more samples N than width m .

Remark 2.3. *We can simplify the final error bound (15) as follows. First, we allow all constants to depend on the initial value κ^0 and drop all terms that contain it. Next, we distribute the outer exponent s/β to the summands, for the price of a slightly worse constant, and obtain*

$$\|\kappa^n\|_{L_2(\mathbb{S}^{d-1})}^2 \lesssim m^{-\frac{1}{2} \frac{\alpha}{1+\alpha} \frac{s}{\beta}} + \Delta_{\text{sample}}(m, N)^{\frac{s}{\beta}}.$$

Next, we unravel the definitions of Δ_{sample} and h to arrive at

$$\|\kappa^n\|_{L_2(\mathbb{S}^{d-1})}^2 \lesssim m^{-\frac{1}{2} \frac{\alpha}{1+\alpha} \frac{s}{\beta}} + \left(\frac{m^{3/2}}{N^{1/2}} h^{1-\frac{\alpha s}{2\beta}} \right)^{\frac{s}{\beta}} \lesssim m^{-\frac{1}{2} \frac{\alpha}{1+\alpha} \frac{s}{\beta}} + \left(\frac{m^{3/2}}{N^{1/2}} \left[m^{-\frac{1}{2} \frac{1}{1+\alpha}} \right]^{1-\frac{\alpha s}{2\beta}} \right)^{\frac{s}{\beta}}$$

Observing the constraints on s , β and α , we abbreviate the exponents by some numbers $a, b, c \geq 0$ to obtain

$$\|f_\theta - f\|_{L_2(\mathbb{S}^{d-1})}^2 \lesssim m^{-a} + m^b N^{-c}.$$

2.3 Result II: Kernel Sampling

Motivation Generalization errors can be derived from bounds of the form

$$\sup_{\theta} \left| \|f_\theta - f\|_{L_2(\mathbb{S}^{d-1})} - \frac{1}{2N} \sum_{i=1}^N |f_\theta(x_i) - f(x_i)| \right| \leq C + N^{-1/2},$$

which controls the difference between the expected and empirical loss, uniformly in all parameters θ contained in a set Θ of all relevant parameters. The bound on the right hand side depends on some complexity measure \mathcal{C} such as Rademacher complexity. By the supremum in the estimate, the complexity bound usually depends on the size of the hypothesis class $\{f_\theta \mid \theta \in \Theta\}$. In Theorem 2.2, this gives rise to the m dependent $\Delta_{\text{sample}}(m, N)$ and therefore to the requirement of more samples than network size (although we technically apply the argument to the gradient, not the loss).

In order to decouple the size of the network from the number of samples, we use the observation from Welper (2024b) or Theorem 2.2 that in the initial NTK regime the Sobolev norm $H^s(\mathbb{S}^{d-1})$ of the residual $\kappa := f_\theta - f$ does not grow, i.e. $\|\kappa^n\|_{H^s(\mathbb{S}^{d-1})} \leq \|\kappa^0\|_{H^s(\mathbb{S}^{d-1})}$. Hence, we may replace the sample error with a bound of the form

$$\sup_{\|\kappa\|_{H^s(\mathbb{S}^{d-1})} \leq \|\kappa^0\|_{H^s(\mathbb{S}^{d-1})}} \left| \|\kappa\|_{L_2(\mathbb{S}^{d-1})}^2 - \frac{1}{2N} \sum_{i=1}^N |\kappa(x_i)|^2 \right| \leq \mathcal{C} + N^{-1/2}, \quad (16)$$

which is *independent of the network size*. If s is sufficiently large so that $H^s(\mathbb{S}^{d-1})$ is embedded into $L^\infty(\mathbb{S}^{d-1})$ with some margin, this estimate can be shown by a uniform law of large numbers, as e.g. in Vershynin (2018, Chapter 8.2). Such an embedding also ensures a bounded loss function, which is assumed for standard VC dimension and Rademacher complexity bounds of the generalization error. For small s , a favorable right hand side cannot be expected. Indeed, the map $\kappa \rightarrow \|\kappa\|_*^2 := \frac{1}{2N} \sum_{i=1}^N |\kappa(x_i)|^2$ from H^s to \mathbb{R} must be continuous, otherwise one can find a perturbation $\tilde{\kappa}$ so that $\|\kappa\|_{L_2} - \|\tilde{\kappa}\|_{L_2} \leq \|\kappa - \tilde{\kappa}\|_{L_2} \leq \|\kappa - \tilde{\kappa}\|_{H^s}$ is small and $\|\kappa\|_* - \|\tilde{\kappa}\|_*$ is large and as a result the left hand side of (16) must be large for either κ or $\tilde{\kappa}$. By the Sobolev embedding theorem (Adams & Fournier, 2008) the continuity holds for $s \geq d/2$ and fails otherwise. Unfortunately, our results provide low Sobolev regularity $s < 1/2$, insufficient for this embedding.

For a first theoretical exploration of alternative complexity measures and in order to stay compatible with earlier approximation and optimization results, we change the point evaluation to localized integrals

$$\ell^k(\theta) := \frac{1}{2N} \sum_{i=1}^N \langle k(x_i, \cdot), f_\theta - f \rangle^2, \quad (17)$$

with uniformly random x_i and some integral kernel $k: \mathbb{S}^{d-1} \times \mathbb{S}^{d-1} \rightarrow \mathbb{R}$ that is smoothing and allows continuous evaluation for all $s > 0$. Note that many standard kernels, like heat, Gaussian and Laplacian kernels, converge to the Dirac delta for their “width” going to zero. As a result, the loss ℓ^k converges to the regular mean square loss ℓ in (13). See Section 2.4 for more details.

Kernels Before we state the main result, we need some properties of the kernel. First, we assume it is zonal, i.e. that $k(x, y) = k(x^T y)$. As a result, by the Funk-Hecke formula (Atkinson & Han, 2012) the eigenfunctions are spherical harmonics Y_l^j (6) and for the corresponding eigenvalues λ_{lj} , we require

$$\begin{aligned} 1 &\lesssim \lambda_{lj} \lesssim 1, \quad l \leq L, \quad 1 \leq j \leq \nu(l), \\ \lambda_{lj} &\lesssim 1, \quad l > L, \quad 1 \leq j \leq \nu(l), \end{aligned} \quad (18)$$

so that up to a limiting level $L > 0$ the eigenvalues are of unit size and falling thereafter. In addition, the kernels are bounded

$$\sup_{x \in D} \|k(x, \cdot)\|_{L_2(\mathbb{S}^{d-1})} \leq C_k. \quad (19)$$

We defer a more thorough discussion of kernels with the given properties to Section 2.4 and for the time being consider a convolutional kernel on the line instead of zonal kernel on the sphere, to avoid technicalities. In this case, a natural example is the Gaussian kernel $k_G(x - y) = \frac{1}{\sqrt{2\pi}t} e^{-\frac{|x-y|^2}{t^2}}$. It is diagonalized by the Fourier transform with continuous eigenvalues

$$\lambda(\omega) = \hat{k}_G(\omega) = \frac{1}{2\sqrt{\pi}} e^{-\frac{\omega^2 t^2}{4}}$$

for which we obtain (18) by the simple observation

$$\lambda(\omega) \lesssim 1, \quad \omega \in \mathbb{R}, \quad \lambda(\omega) \gtrsim 1, \quad \omega^2 t^2 \leq 1 \Leftrightarrow |\omega| \leq \frac{1}{t},$$

Similar kernels on the sphere are more technical and discussed in Section 2.4.

Result Unlike more traditional complexity measures in machine learning, the smoothness $\|\kappa^n\|_{H^s(\mathbb{S}^{d-1})} \leq \|\kappa^0\|_{H^s(\mathbb{S}^{d-1})}$ is a byproduct of the gradient descent method and *independent of the size of the network*. This yields error bounds with decoupled approximation and sampling error in the following theorem.

Theorem 2.4. *Assume we train the kernel loss (17) with conditions (18), (19) and corresponding constants C_k and L . Let Assumption 2.1 be satisfied and for arbitrary $\tau_N \lesssim N$ define*

$$\Delta_{\text{sample}}(m, N) = c_\Delta \left[C_k^2 \left(\frac{\tau_N}{N} \right)^{1/2} + C_k^{-2} \left(\frac{N}{\tau_N} \right)^{1/2} L^{-s} + L^{-s} \right].$$

for some sufficiently large c_Δ . Then with $\kappa^n := f_{\theta^n} - f$ and probability at least $1 - ce^{-\tau} - 2\tau_N [e^{\tau_N} - \tau_N - 1]^{-1}$ while the gradient descent error exceeds the final approximation and estimation error

$$\|\kappa^k\|_{L_2(\mathbb{S}^{d-1})}^2 \geq c_a \left(m^{-\frac{1}{2} \frac{\alpha}{1+\alpha}} + \Delta_{\text{sample}}(m, N) \right)^{\frac{s}{\beta}} \|\kappa^0\|_{H^s(\mathbb{S}^{d-1})}^2, \quad k < n,$$

we have

$$\|\kappa^n\|_{L_2(\mathbb{S}^{d-1})}^2 \leq C e^{-\gamma[h^\alpha + \Delta_{\text{sample}}(m, N)]n} \|\kappa^0\|_{L_2(\mathbb{S}^{d-1})}^2, \quad \|\kappa^n\|_{H^s(\mathbb{S}^{d-1})}^2 \leq C \|\kappa^0\|_{H^s(\mathbb{S}^{d-1})}^2.$$

for sufficiently large constants c_a , c and C independent of m , κ^0 and κ^n .

As for Theorem 2.2 the error decays exponentially until the final sum of approximation and sample error is reached. Unlike Theorem 2.2, the sample error $\Delta_{\text{sample}}(m, N)$ does not depend on the network size m . Hence, the final error is the worse of the approximation and the sample error and provides meaningful error bounds both in under- and over-parametrized regimes.

If we choose the best possible ratio $\frac{N}{\tau_N} = C_k^4 L^s$ between the number of samples N and the success probability parameter τ_N , given the parameters L and C_k of the kernel, we obtain

$$\Delta_{\text{sample}}(m, N) \leq (1 + c_\Delta) L^{-s/2}, \quad (20)$$

converging to zero for large L corresponding to locally concentrated kernels, see Section 2.4.

The theorem contains the inequality $\|\kappa^n\|_{H^s(\mathbb{S}^{d-1})}^2 \leq C \|\kappa^0\|_{H^s(\mathbb{S}^{d-1})}^2$ and therefore $\kappa^n = f_{\theta^n} - f$ remains bounded in the Sobolev norm. Hence, for the generalization error, we consider the hypothesis class of bounded Sobolev functions, instead of neural networks with bounded weights as in typical Rademacher or margin bounds.

Remark 2.5. *Analogous to Remark 2.3, we can obtain a simplified error bound after training*

$$\|f_\theta - f\|_{L_2(\mathbb{S}^{d-1})}^2 \lesssim m^{-a} + N^{-c}.$$

Here, we use that Δ_{sample} does not depend on m and can be bounded by (20). To eliminate L , we can choose e.g. $\tau_N = \sqrt{N}$ and solve the ratio $\frac{N}{\tau_N} = C_k^4 L^s$ for N . Note, however, that C_k and L are not independent, which is worked out for the heat kernel in (22), below.

2.4 Kernels

In this section, we consider kernels that meet our assumptions (18) and (19).

Heat Kernel We first consider the *heat kernel* $k_t(x, y)$ on the sphere \mathbb{S}^{d-1} , defined as the solution of the heat equation with Dirac delta as initial condition (Zhao & Song, 2018)

$$\partial_t k_t(\cdot, y) - \Delta k_t(\cdot, y) = 0, \quad k_0(\cdot, y) = \delta(\cdot, y),$$

where Δ is the Laplace-Beltrami operator on the sphere and $\delta(x, y)$ the Dirac delta distribution on the sphere. On the flat space \mathbb{R}^d , this kernel is identical to the Gaussian kernel $(2\pi)^{-1/2}t^{-1}e^{-|x-y|^2/t^2}$, while on the sphere they differ. We use the heat kernel because it allows a particularly simple verification of our kernel assumptions.

Indeed, since the Laplace-Beltrami operator's eigenfunctions are spherical harmonics Y_l^j with eigenvalues $-l(l+d-2)$ (Atkinson & Han, 2012), the heat equation has the explicit solution

$$k_t(x, y) = \sum_{l,j} e^{-l(l+d-2)t} Y_l^j(x) \langle Y_l^j, \delta(\cdot, y) \rangle = \sum_{l,j} e^{-l(l+d-2)t} Y_l^j(x) Y_l^j(y) \quad (21)$$

in its eigenbasis. Therefore, the eigenvalues of the kernel are $\lambda_{lj} = e^{-l(l+d-2)t}$ and with $l(l+d-2)t \leq 1 \Leftrightarrow l \lesssim t^{-1/2}$ we have

$$\begin{aligned} e^{-1} &\leq \lambda_{lj} \leq 1, & l &\lesssim t^{-1/2}, \\ \lambda_{lj} &\leq 1, & l &\gtrsim t^{-1/2} \end{aligned}$$

so that the kernel assumption (18) is satisfied with $L = ct^{-1/2}$. Since the eigenvalues decay exponentially, the Sobolev norms of the kernel are bounded. More concretely, Lemma D.1 in the supplementary material shows that

$$\|k_t(\cdot, y)\|_{L_2(\mathbb{S}^{d-1})}^2 \leq C_k^2 =: ct^{-d+3/2}.$$

In conclusion, the heat kernel satisfies all assumptions of Theorem 2.4 and the sample error (20) simplifies to

$$\Delta_{\text{sample}}(m, N) \leq (1 + c_\Delta)t^{s/4}, \quad \frac{N}{\tau_N} \sim t^{-\frac{1}{2}s-4d+6} \quad (22)$$

for the given number of samples. By construction, for $t \rightarrow 0$ the kernel converges to the Dirac delta and therefore the kernel loss (17) converges to the sample loss (13).

Gaussian and Laplace Kernels In order to obtain some further insight into permissible kernels, we consider the Gaussian and Laplacian kernels

$$k_G(x - y) = \frac{1}{\sqrt{2\pi}t} e^{-\frac{|x-y|^2}{t^2}}, \quad k_L(x - y) = \frac{1}{2t} e^{-\frac{|x-y|}{t}}$$

on the real line \mathbb{R} . For $s < 1/2$, these are clearly bounded in $H^s(\mathbb{R})$. Moreover, since these are convolutional kernels, the eigenvalues correspond to the Fourier coefficients, given by

$$\hat{k}_G(\omega) = \frac{1}{2\sqrt{\pi}} e^{-\frac{\omega^2 t^2}{4}}, \quad \hat{k}_L(\omega) = \sqrt{\frac{2}{\pi}} \frac{1}{1 + \omega^2 t^2}.$$

Thus, we easily obtain the analogues of the eigenvalue bounds (18):

$$\begin{aligned} \hat{k}_G(\omega) &\lesssim 1, & \omega &\in \mathbb{R}, & \hat{k}_G(\omega) &\gtrsim 1, & \omega^2 t^2 \leq 1 &\Leftrightarrow |\omega| \leq \frac{1}{t}, \\ \hat{k}_L(\omega) &\lesssim 1, & \omega &\in \mathbb{R}, & \hat{k}_L(\omega) &\gtrsim 1, & \omega^2 t^2 \leq 1 &\Leftrightarrow |\omega| \leq \frac{1}{t}. \end{aligned}$$

Similar results on the sphere are significantly more involved and beyond the scope of this paper. See e.g. Geifman et al. (2020, Appendix C) for an analysis of the Laplace kernel, without the fine grained dependence on t required for our purposes.

2.5 Sketch of Proof

Overview This section contains a short overview over the proofs of Theorems 2.2 and 2.4. We start by introducing a scale of continuous loss functions

$$\ell_S(\theta) := \frac{1}{2} \|f_\theta - f\|_{H^S(D)}^2, \quad S \in \{0, s\}.$$

For $S = 0$, this loss is the generalization or $L_2(D)$ error. For $S = s > 0$, the loss is used to control the Sobolev smoothness of the neural networks later in the proof. For convenience, we abbreviate $\langle \cdot, \cdot \rangle_s := \langle \cdot, \cdot \rangle_{H^s(D)}$ and $\kappa := f_\theta - f$. A simple calculation and the mean value theorem yield that the loss evolves as

$$\ell_S(\theta^{n+1}) - \ell_S(\theta^n) = -\gamma \sum_r \langle \kappa, \partial_r f_{\bar{\theta}} \rangle_S \partial_r \ell(\theta^n),$$

for some $\bar{\theta}$ on the line segment between θ^n and θ^{n+1} . On the left hand side, we use the ℓ_S loss, i.e. the generalization error or the smoothness, because these are the quantities we ultimately want to control in the theorems. On the right hand side, we have the discrete loss ℓ that we actually train on. By a perturbation argument, we replace the discrete loss with the continuous one:

$$\begin{aligned} \ell_S(\theta^{n+1}) - \ell_S(\theta^n) &= -\gamma \sum_r \langle \kappa^n, \partial_r f_{\bar{\theta}} \rangle_S \partial_r \ell_0(\theta^n) - \gamma \sum_r \langle \kappa^n, \partial_r f_{\bar{\theta}} \rangle_S [\partial_r \ell(\theta^n) - \partial_r \ell_0(\theta^n)]. \\ &=: \text{GD}_0 + \text{PERTURBATION} \end{aligned}$$

The perturbation term will be bounded by

$$\text{PERTURBATION} \leq \gamma \Delta_{\text{sample}}(m, N) \|\kappa^k\|_0 \|\kappa^k\|_S,$$

giving rise to the terms $\Delta_{\text{sample}}(m, N)$ in the main theorems.

Convergence of GD_0 We first consider the case that the perturbation term PERTURBATION is zero. This corresponds to training on the continuous $L_2(D)$ loss ℓ_0 directly and has been studied in Welper (2024b;a). The convergence analysis is based on the observation that

$$\begin{aligned} \ell_0(\theta^{n+1}) - \ell_0(\theta^n) &\leq -\gamma \langle \kappa^n, H \kappa^n \rangle_0 + \text{perturbation}, \\ \ell_s(\theta^{n+1}) - \ell_s(\theta^n) &\leq -\gamma \langle \kappa^n, H \kappa^n \rangle_s + \text{perturbation}, \end{aligned} \tag{23}$$

where H is the linear integral operator induced by the neural tangent kernel

$$\Gamma(x, y) = \lim_{\text{width} \rightarrow \infty} \sum_{r=1}^R \partial_r f_{\theta^0}(x) \partial_r f_{\theta^0}(y).$$

In short, the terms $-\gamma \langle \kappa^n, H \kappa^n \rangle$ are a linearization of the gradient $-\gamma \|\nabla_\theta f_\theta\|^2$, usually found in gradient descent analysis. This linearization remains accurate because of the crucial observation that the weights $\theta^0 - \theta^n$ do not move far from their initial during training.

Ignoring the linearization error, the left hand sides of (23) are bilinear forms and therefore allow simple convergence proofs if the eigenvalues of H are lower bounded. While this is true in over-parametrized regimes, for the $L_2(D)$ loss, the NTK is a compact operator and the eigenvalues converge to zero. Therefore, we consider a coupled evolution of the ℓ_0 and ℓ_s losses: The latter ensures that the networks remain uniformly bounded in the Sobolev norms $H^s(D)$. This implies that the residual κ^n is concentrated in low frequencies or equivalently eigenspaces corresponding to large eigenvalues. Therefore, the system is comparable to ones with lower bounded eigenvalues and allows us to prove convergence of the ℓ_0 loss.

Since this theory already contains perturbation terms, the new terms PERTURBATION from sampling the gradient can be added to the theory with only minimal changes, shown in Theorem A.1.

Bounds for PERTURBATION The main new contribution of the paper is to show that PERTURBATION , i.e. the difference between the *gradients* of the continuous loss ℓ_0 and the discrete loss ℓ remain small. The arguments are similar to standard generalization error bounds, which address the difference between $\ell_0 - \ell$ directly.

1. Classical techniques to bound generalization errors include VC-dimension or Rademacher complexity, which, in turn, are sometimes bounded by Dudley's inequality (Shalev-Shwartz & Ben-David, 2014).

In Section B, we use the latter directly to bound PERTURBATION for the least squares loss (13). This is convenient because similar techniques are used in the prior work Welper (2024b) to show concentration of the NTK. The argument is reminiscent of the uniform law of large numbers as shown in Vershynin (2018, Chapter 8.2).

2. For the kernel loss (17), we take a different route in Section C. First, in the limit of infinite samples, heat, Gauss and Laplace kernels do not converge to the $L_2(D)$ -loss. They rather converge to a dual norm $\|f_\theta - f\|_{H'}$ for primal norm $\|\cdot\|_H = \|\cdot\|_{L_2(D)}^2 + t\|\cdot\|_{H^\alpha(D)}^2$, some $\alpha > 0$ and a parameter t for the width of the kernel.

In the NTK convergence analysis, we have already shown that the residual $f_{\theta^n} - f$ is concentrated in low frequencies and as a result for carefully chosen t , the $L_2(D)$ part of the above norms dominate the $H^\alpha(D)$ part. Together with concentration results for the kernels, this time shown by matrix Bernstein inequalities, this implies gradient descent convergence in the $L_2(D)$ norm.

Remark on Generalization Errors and Optimization Many generalization error bounds in the literature involve terms like $N^{-1/2}\|W^\ell - W_0^\ell\|_F$ (or other norms) that include the distance of the weights from their initial, which stays small in NTK gradient descent convergence proofs. Concretely, if we use the bound from the first equation in the proof of Theorem A.1, which underlies all main results, we obtain on the trained layer

$$\|W^{L-1} - W_0^{L-1}\|_F \leq m^{1/2}\|W^{L-1} - W_0^{L-1}\| =: m\|\theta - \theta_0\|_* \leq mh = m^{1-\frac{1}{2}+\frac{1}{1+\alpha}} \leq m^{1/2}$$

for some $\alpha \geq 0$, where $\|\cdot\|_F$ is the Frobenius norm, $\|\cdot\|$ the spectral norm and the middle equality the definition of the $\|\cdot\|_*$ norm. Therefore, the generalization error bounds can be further bounded by $N^{-1/2}\|W^{L-1} - W_0^{L-1}\|_F \leq N^{-1/2}m^{1/2}$. Hence, to obtain small generalization error, we must be in an under-parametrized regime with $m \leq N$. This conflicts with over-parametrization assumptions in most NTK convergence results. While the generalization bounds and NTK convergence may be reconciled with a finer analysis, we rely on the NTK results in Welper (2024b;a), which also work in under-parametrized regimes.

2.6 Extensions

This section contains a brief outlook on two major assumptions.

Training of all Layers In the main Theorems 2.2 and 2.4, we train only the second but last layer, while all other weights remain frozen at their random initialization.

Including the remaining layers is relatively simple in standard NTK theory. The training of each layer adds one matrix H^ℓ to the NTK, so that training select layers $\Lambda \subset \{1, \dots, L\}$ the gradient descent loss reduction in (23) becomes

$$\ell_0(\theta^{n+1}) - \ell_0(\theta^n) \leq -\gamma \sum_{\ell \in \Lambda} \langle \kappa^n, H^\ell \kappa^n \rangle_0 + \text{perturbation}.$$

For finite networks and their infinite width limits, the terms $\langle \kappa^n, H^\ell \kappa^n \rangle_0 \geq 0$ are non-negative so that including more layers in the training gives strictly sharper results.

The argument is more delicate for loss $\ell_s(\theta^n)$ with smoother $s > 0$, that we include in (23) to extend standard NTK theory to under-parametrized regimes. In this case, the terms $\langle \kappa^n, H^\ell \kappa^n \rangle_s$, are expected to be non-negative in the infinite width limit and only approximately so for finite width networks. Therefore, a careful additional perturbation and concentration analysis would be required to include the training of all layers in the main theorems, which is beyond the scope of this paper.

Sample distribution The fully connected networks in (7) do not have a bias and as a result, the input $x = 0$ always yields output $f_\theta(x) = 0$, independent of the network weights. This degenerate case is avoided by restricting the input to the sphere. Introducing a bias would likely avoid this issue and allow more general domains, however, this would require a revision of cited results in Appendix D. Non-uniform distributions

would likely require some changes, in particular, the Sobolev norms in the smoothness bounds would need to be adapted accordingly.

3 Numerical Experiments

In this section, we supplement the theoretical findings with some preliminary numerical convergence rates with respect to network width and number of samples. We consider the following test case:

- *Network Architecture*: Fully connected with bias and ReLU activation.
- *Input data*: Uniformly distributed on the cube $[-1, 1]^d$.
- *Labels*: Gaussian density function of the input samples, i.e. $y_i = e^{-|x_i|^2/2}$.
- *Test Loss*: In order to approach the $L_2(D)$ error, the test loss is computed on a large number of 1000 uniformly sampled x_i with mean squared loss, no matter the loss function used for training.
- *Kernel Loss*: To approximate the kernel loss, we uniformly sample N points $z_i \in [-1, 1]^d$ and then for each i we sample $N_k = 10$ samples from the normal distribution $\mathcal{N}(z_i, 0.01)$ to obtain $x_{ij} \in \mathbb{R}^{N \times N_k}$. Then, we approximate the kernel by Monte Carlo integration

$$\ell^k(\theta) \approx \sum_{i=1}^N \left[\sum_{j=1}^{N_k} f_\theta(x_{ij}) - f(x_{ij}) \right]^2.$$

- *Training*: 20000 gradient descent steps with learning rate 0.05.
- *Repetition*: Since the randomness of data and initialization is important for the theoretical results, all reported losses are an average of 20 trials.

After sufficiently many gradient descent steps, in m and N , Theorems 2.2 and 2.4 guarantee the convergence rates

$$\begin{aligned} \|f_\theta - f\|_{L_2(\mathbb{S}^{d-1})}^2 &\leq m^{-0.0277} + m^{0.1991} N^{-0.0833}, \\ \|f_\theta - f\|_{L_2(\mathbb{S}^{d-1})}^2 &\leq m^{-0.0277} + N^{-0.0033}, \end{aligned}$$


respectively, where we have chosen dimension $d = 3$, $\beta = d/2$, which is true for ReLU activations as described after (11), $s = \min\{1/2, \beta/2\} = 1/2$, which is the maximal allowed value in the theorems and $\alpha = 1 - s$, which is the (excluded) upper limit. In addition, for the kernel bound, we use (22), with $\tau_N = N/2$, to balance the number of samples and the kernel width.

For comparison, we also consider the best possible rates for functions with bounded Sobolev norm $H^s(\mathbb{S}^{d-1})$. The best approximation rate can be computed with manifold width (Lorentz et al., 1996) and the best possible sampling rate is the minimax rate (Yang, 1999; Tsybakov, 2009), given by


$$m^{-\frac{s}{d}}, \quad N^{-\frac{s}{2s+d}},$$

respectively. These are higher than the guaranteed rates in the main theorems. For classical approximation methods, the smoothness s is capped by the polynomial's degree, which would be two for ReLU activations, yielding rate $m^{-2/3}$. The definition of the minimax rate is independent of the neural network and therefore allows $s \rightarrow \infty$ and hence rate $N^{-1/2}$ (up to possibly growing constants).

Figure 1 and Table 1 contain the estimated convergence rates of the test loss for training with mean squared loss (MSE) (13) and kernel loss (17). All plots are in log scale so that slopes along the x or y axes correspond to convergence rates. To compare the rates, note that the MSE loss $\|f_\theta - f\|_{L_2}^2$ reported in the tables differs from the error $\|f_\theta - f\|_{L_2}$ in the theoretical bounds by a square.



(a) MSE Loss, dim=3, depth=5.



(b) Kernel Loss, dim=3, depth=5.

Figure 1: Test loss for training with mean squared loss (13) (left) and kernel loss (17) (right). All axes are log-scaled so that the slope corresponds to convergence rates.

Dimension 3 and Depth 5 – Trained with MSE Loss								
	dof rate				N rate			
m/N	100	150	200	250	100	150	200	250
40	0.331	1.33	1.33	1.19	1.9	1.36	0.743	0.896
48	0.833	0.328	0.628	1.06	1.86	1.13	0.933	1.25
56	1.13	0.969	1.25	0.434	2.2	1.07	1.08	0.684
64	-1.03	2.05	0.967	1.47	1.56	2.08	0.582	0.985

Dimension 3 and Depth 5 – Trained with Kernel Loss								
	dof rate				N rate			
m/N	100	150	200	250	100	150	200	250
40	1.08	0.676	0.748	1.07	1.65	0.939	0.797	0.482
48	0.229	0.517	0.433	1.05	1.23	1.07	0.744	0.987
56	1.69	1.77	1.75	1.11	2.02	1.1	0.734	0.539
64	-0.54	-0.218	0.898	1.19	1.71	1.21	1.25	0.713

Table 1: Estimated convergence rates between neighbouring losses for the given m/N . Left: Rate along the column, i.e. with respect to m . Right: Rate along rows, i.e. with respect to number of samples N . The first table is trained with mean squared loss (MSE) (13) and the second with kernel loss (17).

We first observe that the reported rates have significant variance despite being averaged over 20 separate runs. Individual runs are even more noisy. The majority of the table entries remain below the theoretically optimal rates, but are higher than the rates guaranteed by the theorems. While the theorems seem pessimistic, it is worth noting that the optimal rates utilize the high smoothness of the target function $s \geq 2$, whereas the theorems only allow smoothness up to $s < 1/2$. At the expense of high rates, this does allow much rougher target functions for which the optimal rates would be lower, as well. More detailed experiments are left for future work.

Finally note that despite severe over-parametrization, the loss decreases with respect to N , although with slowing rate. This matches the theoretical results for the kernel loss, which provides bounds in over-parametrized regimes, but does not decrease below the approximation error. More detailed results, including the loss and shallow networks, are contained in Appendix E.

4 Conclusion

The literature contains a large number of papers that bound approximation, estimation and optimization errors of neural networks. Because of conflicting assumptions, these results are usually not compatible. This paper contains a unified analysis that bounds all three error components at once. Unlike the contemporary literature, it does not rely on excessively wide networks or over-parametrization. Instead, we have seen for kernel loss that the generalization error can be bounded by the maximum of approximation and estimation errors, independent of any relation between width m and number of samples N .

The generalization errors of this paper are achieved for a bounded number of gradient descent steps, where the training dynamics is dominated by the neural tangent kernel (NTK). Beyond this bound, the NTK loses its relevance and more nonlinear behaviour is possible. A rigorous analysis is left for future work.

References

- Robert A. Adams and John J F. Fournier. *Sobolev spaces*. Number 140 in Pure and applied mathematics. Acad. Press, Amsterdam, 2. ed., [nachdr.]. edition, 2008. ISBN 9780120441433.
- Ben Adcock and Nick Dexter. The gap between theory and practice in function approximation with deep neural networks. *SIAM Journal on Mathematics of Data Science*, 3(2):624–655, 2021.
- Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 242–252, Long Beach, California, USA, 09–15 Jun 2019. PMLR. Full version available at <https://arxiv.org/abs/1811.03962>.
- Sanjeev Arora, Rong Ge, Behnam Neyshabur, and Yi Zhang. Stronger generalization bounds for deep nets via a compression approach. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 254–263. PMLR, 10–15 Jul 2018.
- Sanjeev Arora, Simon Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 322–332, Long Beach, California, USA, 09–15 Jun 2019a. PMLR.
- Sanjeev Arora, Simon S Du, Wei Hu, Zhiyuan Li, Russ R Salakhutdinov, and Ruosong Wang. On exact computation with an infinitely wide neural net. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019b.
- Amir Asadi, Emmanuel Abbe, and Sergio Verdu. Chaining mutual information and tightening generalization bounds. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- Kendall Atkinson and Weimin Han. *Spherical harmonics and approximations on the unit sphere: an introduction*. Number 2044 in Lecture notes in mathematics. Springer-Verlag, Berlin, 2012. ISBN 9783642259821.
- Francis Bach. Breaking the curse of dimensionality with convex neural networks. *Journal of Machine Learning Research*, 18(19):1–53, 2017.
- Yu Bai and Jason D. Lee. Beyond linearization: On quadratic and higher-order approximation of wide neural networks. In *International Conference on Learning Representations*, 2020.
- Juan Antonio Barceló, Teresa Luque, and Salvador Pérez-Esteva. Characterization of sobolev spaces on the sphere. *Journal of Mathematical Analysis and Applications*, 491(1):124240, 2020. ISSN 0022-247X.

- A. R. Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information Theory*, 39(3):930–945, 1993.
- Peter Bartlett, Vitaly Maierov, and Ron Meir. Almost linear vc dimension bounds for piecewise polynomial networks. In M. Kearns, S. Solla, and D. Cohn (eds.), *Advances in Neural Information Processing Systems*, volume 11. MIT Press, 1998.
- Peter L Bartlett, Dylan J Foster, and Matus J Telgarsky. Spectrally-normalized margin bounds for neural networks. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- Christian Beck, Arnulf Jentzen, and Benno Kuckuck. Full error analysis for the training of deep neural networks. *Infinite Dimensional Analysis, Quantum Probability and Related Topics*, 25(02):2150020, June 2022. ISSN 0219-0257, 1793-6306.
- Julius Berner, Philipp Grohs, Gitta Kutyniok, and Philipp Petersen. The Modern Mathematics of Deep Learning. In Philipp Grohs and Gitta Kutyniok (eds.), *Mathematical Aspects of Deep Learning*, pp. 1–111. Cambridge University Press, 1 edition, 2022. ISBN 9781009025096 9781316516782.
- Alberto Bietti and Julien Mairal. On the inductive bias of neural tangent kernels. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Guy Bresler and Dheeraj Nagaraj. Sharp representation theorems for ReLU networks with precise dependence on depth. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 10697–10706. Curran Associates, Inc., 2020.
- Yuan Cao and Quanquan Gu. Generalization Error Bounds of Gradient Descent for Learning Over-Parameterized Deep ReLU Networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(04):3349–3356, April 2020. ISSN 2374-3468, 2159-5399.
- Lin Chen and Sheng Xu. Deep neural tangent kernel and laplace kernel have the same rkhs. In *International Conference on Learning Representations*, 2021.
- Zixiang Chen, Yuan Cao, Difan Zou, and Quanquan Gu. How much over-parameterization is sufficient to learn deep re{lu} networks? In *International Conference on Learning Representations*, 2021.
- Lénaïc Chizat and Francis Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. <https://arxiv.org/abs/1805.09545>.
- Lénaïc Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- A. Cohen, W. Dahmen, and R. DeVore. Adaptive Wavelet Methods II—Beyond the Elliptic Case. *Foundations of Computational Mathematics*, 2(3):203–202, August 2002. ISSN 16153375.
- Jeremy Cohen, Elan Rosenfeld, and Zico Kolter. Certified adversarial robustness via randomized smoothing. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 1310–1320. PMLR, 09–15 Jun 2019.
- G. Cybenko. Approximation by superpositions of a sigmoidal function. *Math. Control Signal Systems*, 2: 303–314, 1989.
- I. Daubechies, R. DeVore, S. Foucart, B. Hanin, and G. Petrova. Nonlinear Approximation and (Deep) ReLU Networks. *Constructive Approximation*, 55(1):127–172, February 2022. ISSN 0176-4276, 1432-0940.

- Ronald DeVore, Boris Hanin, and Guergana Petrova. Neural network approximation. *Acta Numerica*, 30: 327–444, 2021.
- Selina Drews and Michael Kohler. On the universal consistency of an over-parametrized deep neural network estimate learned by gradient descent, 2022. <https://arxiv.org/abs/2208.14283>.
- Simon Du and Jason Lee. On the power of over-parametrization in neural networks with quadratic activation. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1329–1338. PMLR, 10–15 Jul 2018. <https://arxiv.org/abs/1803.01206>.
- Simon Du, Jason Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 1675–1685, Long Beach, California, USA, 09–15 Jun 2019a. PMLR.
- Simon S. Du, Xiyu Zhai, Barnabas Póczos, and Aarti Singh. Gradient descent provably optimizes over-parameterized neural networks. In *International Conference on Learning Representations*, 2019b.
- Rong Ge, Jason D. Lee, and Tengyu Ma. Learning one-hidden-layer neural networks with landscape design. In *International Conference on Learning Representations*, 2018. <https://arxiv.org/abs/1711.00501>.
- Amnon Geifman, Abhay Yadav, Yoni Kasten, Meirav Galun, David Jacobs, and Basri Ronen. On the similarity between the laplace and neural tangent kernels. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 1451–1461. Curran Associates, Inc., 2020.
- Mario Geiger, Arthur Jacot, Stefano Spigler, Franck Gabriel, Levent Sagun, Stéphane d’Ascoli, Giulio Biroli, Clément Hongler, and Matthieu Wyart. Scaling description of generalization with number of parameters in deep learning. *CoRR*, abs/1901.01608, 2019.
- R. Gentile and G. Welper. Approximation results for gradient descent trained shallow neural networks in 1d, 2022. <https://arxiv.org/abs/2209.08399>.
- Noah Golowich, Alexander Rakhlin, and Ohad Shamir. Size-independent sample complexity of neural networks. In Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet (eds.), *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pp. 297–299. PMLR, 06–09 Jul 2018.
- Rémi Gribonval, Gitta Kutyniok, Morten Nielsen, and Felix Voigtlaender. Approximation Spaces of Deep Neural Networks. *Constructive Approximation*, 55(1):259–367, February 2022. ISSN 0176-4276, 1432-0940.
- Philipp Grohs and Felix Voigtlaender. Proof of the Theory-to-Practice Gap in Deep Learning via Sampling Complexity bounds for Neural Network Approximation Spaces. *Foundations of Computational Mathematics*, July 2023. ISSN 1615-3375, 1615-3383.
- Ingo Gühring, Gitta Kutyniok, and Philipp Petersen. Error bounds for approximations with deep ReLU neural networks in w_s, p norms. *Analysis and Applications*, 18(05):803–859, 2020.
- Boris Hanin and Mark Sellke. Approximating continuous functions by ReLU nets of minimal width. <https://arxiv.org/abs/1710.11278>, 2017.
- Moritz Hardt, Ben Recht, and Yoram Singer. Train faster, generalize better: Stability of stochastic gradient descent. In Maria Florina Balcan and Kilian Q. Weinberger (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 1225–1234, New York, New York, USA, 20–22 Jun 2016. PMLR.
- Nick Harvey, Christopher Liaw, and Abbas Mehrabian. Nearly-tight VC-dimension bounds for piecewise linear neural networks. In Satyen Kale and Ohad Shamir (eds.), *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pp. 1064–1068. PMLR, 07–10 Jul 2017.

- K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, June 2016.
- Lukas Herrmann, Joost A. A. Opschoor, and Christoph Schwab. Constructive deep ReLU neural network approximation. *Journal of Scientific Computing*, 90(2):75, 2022.
- Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989. ISSN 0893-6080.
- Daniel Hsu, Sham Kakade, and Tong Zhang. Tail inequalities for sums of random matrices that depend on the intrinsic dimension. *Electronic Communications in Probability*, 17:1 – 13, 2012.
- Shokhrukh Ibragimov, Arnulf Jentzen, and Adrian Riekert. Convergence to good non-optimal critical points in the training of neural networks: Gradient descent optimization with one random initialization overcomes all bad non-global local minima with high probability, 2022. <https://arxiv.org/abs/2212.13111>.
- Arthur Jacot, Franck Gabriel, and Clement Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- Daniel Jakubovitz, Raja Giryes, and Miguel R. D. Rodrigues. Generalization Error in Deep Learning. In Holger Boche, Giuseppe Caire, Robert Calderbank, Gitta Kutyniok, Rudolf Mathar, and Philipp Petersen (eds.), *Compressed Sensing and Its Applications*, pp. 153–193. Springer International Publishing, Cham, 2019. ISBN 9783319730738 9783319730745.
- Arnulf Jentzen and Adrian Riekert. A proof of convergence for the gradient descent optimization method with random initializations in the training of neural networks with relu activation for piecewise linear target functions. *Journal of Machine Learning Research*, 23(260):1–50, 2022.
- Ziwei Ji and Matus Telgarsky. Polylogarithmic width suffices for gradient descent to achieve arbitrarily small test error with shallow ReLU networks. In *International Conference on Learning Representations*, 2020.
- Ziwei Ji, Matus Telgarsky, and Ruicheng Xian. Neural tangent kernels, transportation mappings, and universal approximation. In *International Conference on Learning Representations*, 2020.
- Kenji Kawaguchi and Jiaoyang Huang. Gradient descent finds global minima for generalizable deep neural networks of practical sizes. In *2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 92–99, 2019.
- E. Kharazmi, Z. Zhang, and G. E. Karniadakis. Variational physics-informed neural networks for solving partial differential equations, 2019. <https://arxiv.org/abs/1912.00873>.
- Ehsan Kharazmi, Zhongqiang Zhang, and George E.M. Karniadakis. hp-vpinns: Variational physics-informed neural networks with domain decomposition. *Computer Methods in Applied Mechanics and Engineering*, 374:113547, 2021. ISSN 0045-7825.
- Jason M. Klusowski and Andrew R. Barron. Approximation by combinations of ReLU and squared ReLU ridge functions with ℓ^1 and ℓ^0 controls. *IEEE Transactions on Information Theory*, 64(12):7649–7656, 2018.
- Michael Kohler and Adam Krzyzak. Analysis of the rate of convergence of an over-parametrized deep neural network estimate learned by gradient descent, 2022. <https://arxiv.org/abs/2210.01443>.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (eds.), *Advances in Neural Information Processing Systems 25*, pp. 1097–1105. Curran Associates, Inc., 2012.
- Jaehoon Lee, Lechao Xiao, Samuel Schoenholz, Yasaman Bahri, Roman Novak, Jascha Sohl-Dickstein, and Jeffrey Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

- Jongmin Lee, Joo Young Choi, Ernest K Ryu, and Albert No. Neural tangent kernel analysis of deep narrow neural networks. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 12282–12351. PMLR, 17–23 Jul 2022.
- Bo Li, Shanshan Tang, and Haijun Yu. Better approximations of high dimensional smooth functions by deep neural networks with rectified power units. *Communications in Computational Physics*, 27(2):379–411, 2019. ISSN 1991-7120.
- Yuanzhi Li and Yingyu Liang. Learning overparameterized neural networks via stochastic gradient descent on structured data. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 31*, pp. 8157–8166. Curran Associates, Inc., 2018.
- Zhong Li, Chao Ma, and Lei Wu. Complexity measures for neural networks with general activation functions using path-based norms, 2020. <https://arxiv.org/abs/2009.06132>.
- Tengyuan Liang, Tomaso Poggio, Alexander Rakhlin, and James Stokes. Fisher-rao metric, geometry, and complexity of neural networks. In Kamalika Chaudhuri and Masashi Sugiyama (eds.), *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pp. 888–896. PMLR, 16–18 Apr 2019.
- Fusheng Liu, Haizhao Yang, Soufiane Hayou, and Qianxiao Li. From optimization dynamics to generalization bounds via Łojasiewicz gradient inequality. *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856.
- George G. Lorentz, Manfred v. Golitschek, and Yuly Makovoz. *Constructive Approximation: Advanced Problems*. Springer-Verlag Berlin Heidelberg, 1996.
- Jianfeng Lu, Zuowei Shen, Haizhao Yang, and Shijun Zhang. Deep network approximation for smooth functions. *SIAM Journal on Mathematical Analysis*, 53(5):5465–5506, 2021.
- Zhou Lu, Hongming Pu, Feicheng Wang, Zhiqiang Hu, and Liwei Wang. The expressive power of neural networks: A view from the width. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 30*, pp. 6231–6239. Curran Associates, Inc., 2017.
- Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33):E7665–E7671, 2018. ISSN 0027-8424. <https://arxiv.org/abs/1804.06561>.
- Stanislav Minsker. On some extensions of bernstein’s inequality for self-adjoint operators. *Statistics & Probability Letters*, 127:111–119, 2017. ISSN 0167-7152.
- Gergely Neu, Gintare Karolina Dziugaite, Mahdi Haghifam, and Daniel M. Roy. Information-theoretic generalization bounds for stochastic gradient descent. In Mikhail Belkin and Samory Kpotufe (eds.), *Proceedings of Thirty Fourth Conference on Learning Theory*, volume 134 of *Proceedings of Machine Learning Research*, pp. 3526–3545. PMLR, 15–19 Aug 2021.
- Behnam Neyshabur, Ryota Tomioka, and Srebro Nathan. Norm-based capacity control in neural networks. In Peter Grünwald, Elad Hazan, and Satyen Kale (eds.), *Proceedings of The 28th Conference on Learning Theory*, volume 40 of *Proceedings of Machine Learning Research*, pp. 1376–1401, Paris, France, 03–06 Jul 2015. PMLR.
- Behnam Neyshabur, Srinadh Bhojanapalli, David Mcallester, and Srebro Nati. Exploring generalization in deep learning. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.

- Behnam Neyshabur, Srinadh Bhojanapalli, and Srebro Nathan. A PAC-bayesian approach to spectrally-normalized margin bounds for neural networks. In *International Conference on Learning Representations*, 2018.
- Quynh Nguyen and Matthias Hein. The loss surface of deep and wide neural networks. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 2603–2612. PMLR, 06–11 Aug 2017.
- Quynh N Nguyen and Marco Mondelli. Global convergence of deep networks with one wide layer followed by pyramidal topology. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 11961–11972. Curran Associates, Inc., 2020.
- Atsushi Nitanda and Taiji Suzuki. Optimal rates for averaged stochastic gradient descent under neural tangent kernel regime. In *International Conference on Learning Representations*, 2021.
- Joost A. A. Opschoor, Philipp C. Petersen, and Christoph Schwab. Deep ReLU networks and high-order finite element methods. *Analysis and Applications*, 18(05):715–770, 2020.
- S. Oymak and M. Soltanolkotabi. Toward moderate overparameterization: Global convergence guarantees for training shallow neural networks. *IEEE Journal on Selected Areas in Information Theory*, 1(1):84–105, 2020.
- Sejun Park, Umut Simsekli, and Murat A Erdogdu. Generalization bounds for stochastic gradient descent via localized ϵ -covers. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022.
- Allan Pinkus. Approximation theory of the mlp model in neural networks. *Acta Numerica*, 8:143–195, 1999.
- Daniel A. Roberts, Sho Yaiday, and Boris Hanin. *The Principles of Deep Learning Theory: An Effective Theory Approach to Understanding Neural Networks*. Cambridge University Press, 2022.
- Grant M. Rotskoff and Eric Vanden-Eijnden. Neural networks as interacting particle systems: Asymptotic convexity of the loss landscape and universal scaling of the approximation error. *CoRR*, abs/1805.00915, 2018. <https://arxiv.org/abs/1805.00915>.
- Itay Safran and Ohad Shamir. Spurious local minima are common in two-layer ReLU neural networks. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4433–4441, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR.
- S. Shalev-Shwartz and S. Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Understanding Machine Learning: From Theory to Algorithms. Cambridge University Press, 2014. ISBN 9781107057135.
- Zuowei Shen, Haizhao Yang, and Shijun Zhang. Nonlinear approximation via compositions. *Neural Networks*, 119:74–84, 2019. ISSN 0893-6080.
- Jonathan W. Siegel and Jinchao Xu. Approximation rates for neural networks with general activation functions. *Neural Networks*, 128:313–321, 2020. ISSN 0893-6080.
- Jonathan W. Siegel and Jinchao Xu. High-order approximation rates for shallow neural networks with cosine and ReLU^k activation functions. *Applied and Computational Harmonic Analysis*, 58:1–26, 2022a. ISSN 1063-5203.
- Jonathan W. Siegel and Jinchao Xu. Optimal convergence rates for the orthogonal greedy algorithm. *IEEE Transactions on Information Theory*, 68(5):3354–3361, 2022b.
- Jonathan W. Siegel, Qingguo Hong, Xianlin Jin, Wenrui Hao, and Jinchao Xu. Greedy training algorithms for neural networks and applications to PDEs. *Journal of Computational Physics*, 484:112084, July 2023. ISSN 00219991.

- Justin Sirignano and Konstantinos Spiliopoulos. Mean field analysis of neural networks: A law of large numbers. *SIAM Journal on Applied Mathematics*, 80(2):725–752, 2020.
- Mahdi Soltanolkotabi, Adel Javanmard, and Jason D. Lee. Theoretical insights into the optimization landscape of over-parameterized shallow neural networks. *IEEE Transactions on Information Theory*, 65(2): 742–769, 2019. <https://arxiv.org/abs/1707.04926>.
- ChaeHwan Song, Ali Ramezani-Kebrya, Thomas Pethick, Armin Eftekhari, and Volkan Cevher. Sub-quadratic overparameterization for shallow neural networks. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 11247–11259. Curran Associates, Inc., 2021.
- Zhao Song and Xin Yang. Quadratic suffices for over-parametrization via matrix chernoff bound, 2019. <https://arxiv.org/abs/1906.03593>.
- Daniel Soudry and Yair Carmon. No bad local minima: Data independent training error guarantees for multilayer neural networks, 2016. <https://arxiv.org/abs/1605.08361>.
- Elias M. Stein and Guido Weiss. *Introduction to Fourier Analysis on Euclidean Spaces (PMS-32)*. Princeton University Press, December 1972. ISBN 9781400883899.
- Thomas Steinke and Lydia Zakyntinou. Reasoning About Generalization via Conditional Mutual Information. In Jacob Abernethy and Shivani Agarwal (eds.), *Proceedings of Thirty Third Conference on Learning Theory*, volume 125 of *Proceedings of Machine Learning Research*, pp. 3437–3452. PMLR, 09–12 Jul 2020.
- Lili Su and Pengkun Yang. On learning over-parameterized neural networks: A functional approximation perspective. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Taiji Suzuki. Adaptivity of deep reLU network for learning in besov and mixed smooth besov spaces: optimal rate and curse of dimensionality. In *International Conference on Learning Representations*, 2019.
- Joel A. Tropp. An introduction to matrix concentration inequalities. *Foundations and Trends® in Machine Learning*, 8(1-2):1–230, 2015. ISSN 1935-8237.
- Alexandre B. Tsybakov. *Introduction to nonparametric estimation*. Springer series in statistics. Springer, New York ; London, 2009. ISBN 9780387790510 9780387790527.
- Zhuozhuo Tu, Fengxiang He, and Dacheng Tao. Understanding generalization in recurrent neural networks. In *International Conference on Learning Representations*, 2020.
- Luca Venturi, Afonso S. Bandeira, and Joan Bruna. Spurious valleys in one-hidden-layer neural network optimization landscapes. *Journal of Machine Learning Research*, 20(133):1–34, 2019. <https://arxiv.org/abs/1802.06384>.
- Roman Vershynin. *High-dimensional probability: an introduction with applications in data science*. Number 47 in Cambridge series in statistical and probabilistic mathematics. Cambridge University Press, Cambridge ; New York, NY, 2018. ISBN 9781108415194.
- Mingze Wang and Chao Ma. Generalization error bounds for deep neural networks trained by sgd, 2023. <https://arxiv.org/abs/2206.03299>.
- E Weinan, Chao Ma, and Lei Wu. *A priori* estimates of the population risk for two-layer neural networks. *Communications in Mathematical Sciences*, 17(5):1407–1425, 2019. ISSN 1945-0796.
- E Weinan, Ma Chao, Wu Lei, and Stephan Wojtowytsch. Towards a mathematical understanding of neural network-based machine learning: What we know and what we don’t. *CSIAM Transactions on Applied Mathematics*, 1(4):561–615, 2020. ISSN 2708-0579.

- E Weinan, Chao Ma, and Lei Wu. The Barron Space and the Flow-Induced Function Spaces for Neural Network Models. *Constructive Approximation*, 55(1):369–406, February 2022. ISSN 0176-4276, 1432-0940.
- G. Welper. Approximation and gradient descent training with neural networks, 2024a. <https://arxiv.org/abs/2405.11696>.
- Gerrit Welper. Approximation results for gradient flow trained neural networks. *Journal of Machine Learning*, 3(2):107–175, 2024b.
- Yuhong Yang. Minimax nonparametric classification .i. rates of convergence. *IEEE Transactions on Information Theory*, 45(7):2271–2284, 1999.
- Dmitry Yarotsky. Error bounds for approximations with deep ReLU networks. *Neural Networks*, 94:103–114, 2017. ISSN 0893-6080.
- Dmitry Yarotsky. Optimal approximation of continuous functions by very deep ReLU networks. In Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet (eds.), *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pp. 639–649. PMLR, 06–09 Jul 2018.
- Dmitry Yarotsky and Anton Zhevnerchuk. The phase diagram of approximation rates for deep neural networks. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 13005–13015. Curran Associates, Inc., 2020.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*, 2017.
- Chenchao Zhao and Jun S. Song. Exact Heat Kernel on a Hypersphere and Its Applications in Kernel SVM. *Frontiers in Applied Mathematics and Statistics*, 4:1, January 2018. ISSN 2297-4687.
- Ding-Xuan Zhou. Universality of deep convolutional neural networks. *Applied and Computational Harmonic Analysis*, 48(2):787–794, 2020. ISSN 1063-5203.
- Difan Zou and Quanquan Gu. An improved analysis of training over-parameterized deep neural networks. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Difan Zou, Yuan Cao, Dongruo Zhou, and Quanquan Gu. Gradient descent optimizes over-parameterized deep ReLU networks. *Machine Learning*, 109(3):467 – 492, 2020.

A Gradient Descent Convergence

A.1 Convergence Result

In this section, we prove an abstracted convergence result that is the foundation for Theorems 2.2 and 2.4. To this end, let $\Theta \subset \mathbb{R}^R$ be a set of admissible weights, and \mathcal{H}^s , $s \in \mathbb{R}$ a set of Hilbert spaces. Then we consider

$$f: \Theta \rightarrow \mathcal{H}^0, \quad \theta \rightarrow f_\theta, \quad (24)$$

$$\ell_s: \mathcal{H}^s \rightarrow \mathbb{R}, \quad \theta \rightarrow \ell_s(\theta) =: \frac{1}{2} \|f_\theta - f\|_s^2, \quad (25)$$

$$\ell: \mathcal{H}^0 \rightarrow \mathbb{R}, \quad \theta \rightarrow \ell(\theta), \quad (26)$$

where ℓ_0 corresponds to the continuous loss, or generalization error, and ℓ to the discrete loss. The Hilbert spaces have norms $\|\cdot\|_s = \|\cdot\|_{\mathcal{H}^s}$ and are related by the interpolation inequality

$$\|\cdot\|_b \lesssim \|\cdot\|_a^{\frac{c-b}{c-a}} \|\cdot\|_c^{\frac{b-a}{c-a}} \quad (27)$$

for all $a, b, c \in \mathbb{R}$. Typically, we choose Sobolev spaces $\mathcal{H}^s = H^s(D)$, the neural network $\theta \rightarrow f_\theta(\cdot) \in L_2(D) = \mathcal{H}^0$, and some discrete loss ℓ , but this is not important throughout this section. The statement of the following result relies on the empirical NTK and NTK

$$H_{\theta, \bar{\theta}} := \sum_{r=1}^R (\partial_r f_\theta)(\partial_r f_{\bar{\theta}})^*, \quad H := \lim_{\text{width} \rightarrow \infty} \sum_{r=1}^R \partial_r f_{\theta^0} \partial_r f_{\theta^0}^*,$$

where f^* is the \mathcal{H}^0 -adjoint of f . Note that the adjoint, contained in the dual space $(\mathcal{H}^0)'$, applied to a function v is $f^*(v) = \langle f, v \rangle$ and as a result, H corresponds to the integral operator induced by the integral kernel in (11).

The following convergence result is almost identical to Welper (2024a, Theorem 3.1), up to a new error term for the difference between the Hilbert space loss ℓ_0 and the discrete sample loss ℓ in (33).

Theorem A.1. *Assume we minimize the loss ℓ of the parametrized function $\theta \rightarrow f_\theta \in \mathcal{H}^0$ in (24), with gradient descent (9) and Hilbert spaces (27). Define the residual $\kappa^k = f_{\theta^k} - f$. Let m be an indicator for the network size that satisfies the inequalities below. With $\alpha > 0$ from (31) below, assume that*

1. *H is coercive for $S = 0$ and $S = s$ and some $\beta > 2s > 0$*

$$\|v\|_{S-\beta}^2 \lesssim \langle v, Hv \rangle_S, \quad v \in \mathcal{H}^{S-\beta}. \quad (28)$$

2. *For some norm $\|\cdot\|_*$, the distance of the weights from their initial value is bounded by*

$$\|\theta^k - \theta^0\|_* \lesssim 1, \quad k = 1, \dots, n \quad \Rightarrow \quad \|\theta^{n+1} - \theta^0\|_* \lesssim \frac{\gamma}{\sqrt{m}} \sum_{k=0}^n \|\kappa^k\|_0. \quad (29)$$

3. *The learning rate γ is sufficiently small so that*

$$\gamma \|\nabla_\theta \ell(\theta^n)\|_* \lesssim c_h m^{-\frac{1}{2} \frac{1}{1+\alpha}} =: h \quad (30)$$

for some constant c_h that may depend on the initial error $\|\kappa^0\|_0$.

4. *For $S = 0$ and $S = s$, initial value θ^0 , any $\bar{\theta}, \tilde{\theta} \in \Theta$ and any $\bar{h} > 0$, the bounds $\|\theta^0 - \bar{\theta}\|_* \leq \bar{h}$ and $\|\theta^0 - \tilde{\theta}\|_* \leq \bar{h}$ imply*

$$\|H_{\bar{\theta}, \theta^0} - H_{\tilde{\theta}, \bar{\theta}}\|_{S,0} \leq c \bar{h}^\alpha, \quad \|H_{\theta^0, \tilde{\theta}} - H_{\bar{\theta}, \bar{\theta}}\|_{S,0} \leq c \bar{h}^\alpha. \quad (31)$$

5. For $S = 0$ and $S = s$, we have

$$\|H - H_{\theta^0, \theta^0}\|_{S,0} \leq c_h m^{-\frac{1}{2} \frac{\alpha}{1+\alpha}} = h^\alpha. \quad (32)$$

6. There is a bound $\Delta_{\text{sample}}(m, N)$ and sufficiently large constant c_A so that

$$\sup_{\substack{\|\theta^0 - \theta\|_* \leq c_A h \\ \|\theta^0 - \bar{\theta}\|_* \leq c_A h}} - \sum_{r=1}^R \langle \kappa^k, \partial_r f_{\bar{\theta}} \rangle_S [\partial_r \ell(\theta) - \partial_r \ell_0(\theta)] \leq \Delta_{\text{sample}}(m, N) \|\kappa^k\|_0 \|\kappa^k\|_S \quad (33)$$

for all gradient descent iterates κ^k with $k \leq n$.

Then, while the gradient descent error exceeds the final approximation and estimation error

$$\|\kappa^k\|_0^2 \geq c_a \left(m^{-\frac{1}{2} \frac{\alpha}{1+\alpha}} + \Delta_{\text{sample}}(m, N) \right)^{\frac{s}{\beta}} \|\kappa^0\|_s^2, \quad k \leq n, \quad (34)$$

we have

$$\|\kappa^{n+1}\|_0^2 \leq C e^{-\gamma[h^\alpha + \Delta_{\text{sample}}(m, N)](n+1)} \|\kappa^0\|_0^2, \quad \|\kappa^{n+1}\|_s^2 \leq C \|\kappa^0\|_s^2$$

for sufficiently large constants c_a , c and C independent of m , κ^0 and κ^{n+1} .

The theorem has a long list of assumptions, which we verify for the proof of the main theorems in Sections B and C below. The coercivity (28) is the only major assumption that remains in the main theorems. The second assumption (29) is used to show that the weights do not move far from their initial and the third (30) provides bounds for the learning rate. The next two assumptions (31) and (32) are major components of NTK analysis and require that the NTK is Hölder continuous and that the empirical NTK concentrates close to the infinite width limit. Finally, assumption (33) bounds the difference between the gradient of the continuous L_2 loss and the discrete loss ℓ . The last assumption is the major concern of this paper, while all others have been established in Welper (2024b;a).

The proof is identical to Welper (2024a, Theorem 3.1), with one difference: The error reduction lemma Welper (2024a, Lemma 3.2) is replaced with Lemma A.2 below. This introduces a new error term from the sample loss. While this extra error term only requires minimal changes in the proof, for the convenience of the reader, we include it in Section A.3.

A.2 Gradient Descent Error Reduction

The first step in our convergence proof establishes an error decay in every gradient descent step. It matches Welper (2024a, Lemma 3.2) up to an additional error term $\Delta_{\text{sample}}(m, J)$ for the difference between the continuous and discrete losses.

Lemma A.2. Assume that (30), (31), (32) and (33) hold. Assume that $\|\theta^0 - \theta^n\|_* \leq h$. Then

$$\ell_S(\theta^{n+1}) - \ell_S(\theta^n) \leq -\gamma \langle \kappa, H\kappa \rangle_S + c\gamma h^\alpha \|\kappa\|_0 \|\kappa\|_S + \gamma \Delta_{\text{sample}}(m, N) \|\kappa\|_0 \|\kappa\|_S.$$

Proof. Applying the mean value theorem to the gradient descent step $\theta^{n+1} = \theta^n - \Delta^n$ with $\Delta^n := \gamma \nabla_{\theta} \ell(\theta^n)$, we obtain

$$\begin{aligned} \ell_S(\theta^{n+1}) - \ell_S(\theta^n) &= \ell_S(\theta^n - \Delta^n) - \ell_S(\theta^n) \\ &= -\ell'_S(\theta^n - \xi \Delta^n) \Delta^n, \end{aligned}$$

for some $\xi \in (0, 1)$. Abbreviating $\kappa = \kappa^n$ and plugging in the derivatives $\ell'_S(\theta)^T = [\langle \kappa, \partial_r f_{\theta} \rangle_S]_{r=1}^R$ and $\Delta^n = \gamma [\partial_r \ell(\theta)]_{r=1}^R$, yields

$$\ell_S(\theta^{n+1}) - \ell_S(\theta^n) = -\gamma \sum_r \langle \kappa, \partial_r f_{\theta^n - \xi \Delta^n} \rangle_S \partial_r \ell(\theta^n).$$

Next, we replace the derivative $\partial_r \ell(\theta)$ of the discrete loss by the corresponding derivative of the continuous loss $\partial_r \ell_0(\theta) = \langle \kappa, \partial_r f_\theta \rangle$

$$\begin{aligned} \ell_S(\theta^{n+1}) - \ell_S(\theta^n) &= -\gamma \sum_r \langle \kappa, \partial_r f_{\theta^n - \xi \Delta^n} \rangle_S [\partial_r \ell_0(\theta^n) + [\partial_r \ell(\theta^n) - \partial_r \ell_0(\theta^n)]] \\ &= -\gamma \sum_r \langle \kappa, \partial_r f_{\theta^n - \xi \Delta^n} \rangle_S \langle \kappa, \partial_r f_{\theta^n} \rangle \\ &\quad - \gamma \sum_r \langle \kappa, \partial_r f_{\theta^n - \xi \Delta^n} \rangle_S [\partial_r \ell(\theta^n) - \partial_r \ell_0(\theta^n)] \\ &=: (I) + (II). \end{aligned}$$

Let us first estimate (I). To this end, we express $\langle v, \kappa \rangle$ by the \mathcal{H}^0 dual $v^* \kappa := \langle v, \kappa \rangle$ and obtain

$$\begin{aligned} (I) &= -\gamma \sum_r \langle \kappa, \partial_r f_{\theta^n - \xi \Delta^n} \rangle_S \partial_r (f_{\theta^n})^*(\kappa) \\ &= -\gamma \left\langle \kappa, \left[\sum_r (\partial_r f_{\theta^n - \xi \Delta^n}) (\partial_r f_{\theta^n})^* \right] \kappa \right\rangle_S, \\ &= -\gamma \langle \kappa, H_{\theta^n - \xi \Delta^n, \theta^n} \kappa \rangle_S. \end{aligned}$$

Adding and subtracting terms to compare $f_{\theta^n - \xi \Delta^n}$ and f_{θ^n} with the initial f_{θ^0} , we obtain

$$\begin{aligned} (I) &= -\gamma \langle \kappa, H_{\theta^0, \theta^0} \kappa \rangle_S \\ &\quad + \gamma \langle \kappa, H_{\theta^0, \theta^0} - H_{\theta^0, \theta^n} \kappa \rangle_S \\ &\quad + \gamma \langle \kappa, H_{\theta^0, \theta^n} - H_{\theta^n - \xi \Delta^n, \theta^n} \kappa \rangle_S. \end{aligned}$$

Assumption (32) implies

$$-\langle \kappa, H_{\theta^0, \theta^0} \kappa \rangle_S = -\langle \kappa, H \kappa \rangle_S + \langle \kappa, H - H_{\theta^0, \theta^0} \kappa \rangle_S \leq -\langle \kappa, H \kappa \rangle_S + h^\alpha \|\kappa\|_0 \|\kappa\|_S$$

and Assumption (31), with $\bar{h} = h$ and $\bar{h} = h + \|\Delta^n\|_*$ implies

$$\begin{aligned} \|H_{\theta^0, \theta^0} - H_{\theta^0, \theta^n}\|_{S,0} &\leq ch^\alpha, \\ \|H_{\theta^0, \theta^n} - H_{\theta^n - \xi \Delta^n, \theta^n}\|_{S,0} &\leq c(h + \|\Delta^n\|_*)^\alpha. \end{aligned}$$

Combining these inequalities, yields

$$\begin{aligned} (I) &\leq -\gamma \langle \kappa, H \kappa \rangle_S + 3c\gamma [h + \|\Delta^n\|_*]^\alpha \|\kappa\|_0 \|\kappa\|_S, \\ &\leq -\gamma \langle \kappa, H \kappa \rangle_S + c\gamma h^\alpha \|\kappa\|_0 \|\kappa\|_S, \end{aligned}$$

where in the last step we have used that $\|\Delta^n\|_* = \gamma \|\nabla_\theta \ell(\theta^n)\|_* \lesssim h$ by assumption (30).

It remains to bound (II). Again, using $\|\Delta^n\|_* \lesssim h$ and the assumptions of this lemma, we have $\|\theta^0 - (\theta^n + \xi \Delta^n)\| \lesssim h$ for all $\xi \in [0, 1]$. Thus, with Assumption (33) and our abbreviation $\kappa = \kappa^n$, we obtain

$$\begin{aligned} (II) &= -\gamma \sum_r \langle \kappa, \partial_r f_{\theta^n - \xi \Delta^n} \rangle_S [\partial_r \ell(\theta^n) - \partial_r \ell_0(\theta^n)] \\ &\leq \gamma \Delta_{\text{sample}}(m, N) \|\kappa\|_0 \|\kappa\|_S. \end{aligned}$$

Combining (I) and (II) yields the lemma. □

A.3 Proof of Theorem A.1

The following proof is identical to Welper (2024a, Theorem 3.1) up to the additional error term $\Delta_{\text{sample}}(m, N)$. We include the proof to trace the minor modification and keep the paper self contained.

Proof of Theorem A.1. The proof is based on the gradient descent error reduction in Lemma A.2. All of its assumptions are given, except the weight distance $\|\theta^n - \theta^0\|_*$, which we include in the induction hypothesis: We assume

$$\begin{aligned} \|\kappa^k\|_0^2 &\lesssim e^{-\gamma[h^\alpha + \Delta]k} \|\kappa^k\|_0^2, \\ h^k &:= \max_{l \leq k} \|\theta^l - \theta^0\|_* \lesssim c_h m^{-\frac{1}{2} \frac{1}{1+\alpha}} =: h \end{aligned}$$

for all $k \leq n$ and $\Delta = \Delta_{\text{sample}}(m, N)$ and prove the case $k = n + 1$ by induction. With the induction hypothesis and assumptions (30), (31), (32) and (33), Lemma A.2 together with coercivity (1) implies

$$\begin{aligned} \|\kappa^{n+1}\|_0^2 - \|\kappa^0\|_0^2 &\leq -\gamma \|\kappa^n\|_{-\beta}^2 + c\gamma h^\alpha \|\kappa^n\|_0^2 + \gamma \Delta \|\kappa^n\|_0^2, \\ \|\kappa^{n+1}\|_s^2 - \|\kappa^0\|_s^2 &\leq -\gamma \|\kappa^n\|_{s-\beta}^2 + c\gamma h^\alpha \|\kappa^n\|_0 \|\kappa^n\|_s + \gamma \Delta \|\kappa^n\|_0 \|\kappa^n\|_s, \end{aligned}$$

or shorter

$$\begin{aligned} \|\kappa^{n+1}\|_0^2 - \|\kappa^0\|_0^2 &\leq -\gamma \|\kappa^n\|_{-\beta}^2 + \gamma [ch^\alpha + \Delta] \|\kappa^n\|_0^2, \\ \|\kappa^{n+1}\|_s^2 - \|\kappa^0\|_s^2 &\leq -\gamma \|\kappa^n\|_{s-\beta}^2 + \gamma [ch^\alpha + \Delta] \|\kappa^n\|_0 \|\kappa^n\|_s. \end{aligned}$$

This is not a closed iteration of the $\|\kappa^n\|_0^2$ and $\|\kappa^n\|_s^2$ residuals because of the $\|\cdot\|_{-\beta}$ and $\|\cdot\|_{s-\beta}$ norms. We eliminate them with the interpolation inequalities 27

$$\begin{aligned} \|\kappa\|_0 &\leq \|\kappa\|_{-\beta}^{\frac{s}{\beta+s}} \|\kappa\|_s^{\frac{\beta}{\beta+s}} &\Rightarrow & -\|\kappa\|_{-\beta}^2 \leq -\|\kappa\|_0^{2+\frac{2\beta}{s}} \|\kappa\|_s^{-\frac{2\beta}{s}}, \\ \|\kappa\|_0 &\leq \|\kappa\|_{s-\beta}^{\frac{s}{\beta}} \|\kappa\|_s^{\frac{\beta-s}{\beta}} &\Rightarrow & -\|\kappa\|_{s-\beta}^2 \leq -\|\kappa\|_0^{\frac{2\beta}{s}} \|\kappa\|_s^{2-\frac{2\beta}{s}} \end{aligned}$$

so that

$$\begin{aligned} \|\kappa^{n+1}\|_0^2 - \|\kappa^0\|_0^2 &\lesssim -\gamma \|\kappa^n\|_0^{2+\frac{2\beta}{s}} \|\kappa^n\|_s^{-\frac{2\beta}{s}} + \gamma [h^\alpha + \Delta] \|\kappa^n\|_0^2, \\ \|\kappa^{n+1}\|_s^2 - \|\kappa^0\|_s^2 &\lesssim -\gamma \|\kappa^n\|_0^{\frac{2\beta}{s}} \|\kappa^n\|_s^{2-\frac{2\beta}{s}} + \gamma [h^\alpha + \Delta] \|\kappa^n\|_0 \|\kappa^n\|_s. \end{aligned}$$

Error bounds for this iteration are given in Lemma D.5 with $n+1 := \|\kappa^{n+1}\|_0^2$, $v^{n+1} := \|\kappa\|_s^2$, $\rho = \beta/s$, $a = c = 1$ and $b = d = h^\alpha + \Delta$. To show the lemma's assumption (50), we use $2s \leq \beta$ so that

$$\left(2 - \frac{s}{\beta}\right) \leq 2 \quad \Leftrightarrow \quad \frac{s}{\beta} \leq \frac{2\frac{s}{\beta}}{2 - \frac{s}{\beta}} \quad \Leftrightarrow \quad \frac{1}{\rho} \leq \frac{2}{2\rho - 1}.$$

Hence, assumption (34) implies

$$u^k = \|\kappa^k\|_0^2 \geq \left(\left(m^{-\frac{1}{2} \frac{1}{1+\alpha}} \right)^\alpha + \Delta \right)^{\frac{s}{\beta}} \|\kappa^0\|_s^2 = (h^\alpha + \Delta)^{\frac{s}{\beta}} \|\kappa^0\|_s^2 \gtrsim \left(2 \frac{b}{a} \right)^{\frac{1}{\rho}} v^0 \gtrsim \left(\frac{d}{c} \right)^{\frac{2}{2\rho-1}} v^0.$$

and Lemma D.5 is applicable. Therefore, we obtain

$$\|\kappa^{n+1}\|_0^2 \lesssim e^{-\gamma[h^\alpha + \Delta](n+1)} \|\kappa^0\|_0^2, \quad \|\kappa^{n+1}\|_s^2 \lesssim \|\kappa^0\|_s^2,$$

which shows the first induction hypothesis.

It remains to bound h^{n+1} to show the second induction hypothesis. To this end note that

$$h^{n+1} = \max_{k \leq n+1} \|\theta^k - \theta^0\|_* \lesssim \frac{\gamma}{\sqrt{m}} \sum_{k=1}^n \|\kappa^k\|_0 \lesssim \frac{\gamma}{\sqrt{m}} \sum_{k=1}^n e^{-\gamma[h^\alpha + \Delta]k} \|\kappa^0\|_0,$$

where in the second step we have used assumption (29) and in the third step the induction hypothesis. We bound the latter sum

$$\sum_{k=1}^n e^{-\gamma[h^\alpha + \Delta]k} \leq \int_0^\infty e^{-\gamma[h^\alpha + \Delta]k} dk = \frac{1}{\gamma[h^\alpha + \Delta]},$$

to conclude that

$$h^{n+1} \leq c \frac{\gamma}{\sqrt{m}} \frac{1}{\gamma[h^\alpha + \Delta]} \|\kappa^0\|_0 \leq c \frac{\gamma}{\sqrt{m}} \frac{1}{\gamma h^\alpha} \|\kappa^0\|_0 \leq \frac{c}{\sqrt{m}} h^{-\alpha} \|\kappa^0\|_0 = h.$$

In the last step we have used that the definition of h implies

$$h = c_h m^{-\frac{1}{2} \frac{1}{1+\alpha}} = \frac{c_h}{\sqrt{m}} m^{\frac{1}{2} \frac{\alpha}{1+\alpha}} = \frac{c}{\sqrt{m}} m^{\frac{1}{2} \frac{\alpha}{1+\alpha}} \|\kappa^0\|_0 = \frac{c}{\sqrt{m}} h^{-\alpha} \|\kappa^0\|_0$$

for constant $c_h = c \|\kappa^0\|_0$ dependent on $\|\kappa^0\|_0$. Together with the first induction hypothesis this concludes the proof. \square

B Sampling

In this section, we prove bounds for the error $\Delta_{\text{sample}}(m, N)$ between continuous and sample loss (Lemma B.4) and then Theorem 2.2. Throughout the section, we use the following regularity conditions on the activation function

$$|\sigma(x)| \lesssim |x|, \quad (35)$$

$$|\sigma(x) - \sigma(\bar{x})| \lesssim |x - \bar{x}| \quad (36)$$

$$|\dot{\sigma}(x)| \lesssim 1. \quad (37)$$

Moreover, for the time being, we assume that the weight matrices are bounded

$$\|W^\ell\| n_\ell^{-1/2} \lesssim 1, \quad \|\bar{W}^\ell\| n_\ell^{-1/2} \lesssim 1, \quad \|x\| \lesssim 1 \forall x \in D. \quad (38)$$

Typically W^ℓ will be the initial weight and \bar{W}^ℓ the weight of a later gradient descent step. Likewise, we denote by $\bar{\theta}$, \bar{f}_θ and \bar{f}^ℓ , etc. the weights, networks and layers based on the perturbed \bar{W}^ℓ . For the main theorems, these bounds follow from properties of random matrices at the initial weights and the observation that weights do not move far from their initial in (29).

Throughout the section, we abbreviate $D = \mathbb{S}^{d-1}$, $L_2 = L_2(D)$, $H^s = H^s(D)$ and $\langle \cdot, \cdot \rangle_{H^s} = \langle \cdot, \cdot \rangle_{H^s(D)}$, when convenient.

B.1 Concentration

In this section, for the sample loss $\ell(\theta) = \frac{1}{2N} \sum_{i=1}^N |f_\theta(x_i) - f(x_i)|^2$, we bound the sample error

$$\sup_{\theta, \bar{\theta}} \left| \sum_{r=1}^R \langle \kappa^k, \partial_r f_{\bar{\theta}} \rangle_{H^s} [\partial_r \ell(\theta) - \partial_r \ell_0(\theta)] \right|.$$

This establishes assumption (33) in the abstract convergence Theorem A.1 and leads to the proof of the first main result.

Let us abbreviate

$$Y_\theta := \sum_{r=1}^R \langle \kappa, \partial_r f_{\bar{\theta}} \rangle_{H^s} [\partial_r \ell(\theta) - \partial_r \ell_0(\theta)],$$

which is a random variable with respect to the sample points x_i implicitly contained in the discrete loss $\ell(\theta)$. Since the ℓ_0 loss is the expectation of the sample loss ℓ , it is easy to see that $\mathbb{E}[Y_\theta] = 0$ and it suffices to show concentration results. These follow from Dudley's inequality, for which we prove that Y_θ has sub-gaussian increments, i.e.

$$\|Y_\theta - Y_\vartheta\|_{\psi_2} \lesssim \|\theta - \vartheta\|_*.$$

Throughout the section, we use Orlicz and weight norms

$$\|X\|_{\psi_2} = \inf \left\{ t > 0 : \mathbb{E} [\exp(X^2/t^2)] \leq 2 \right\},$$

$$\|\theta\|_* := \|W^{L-1}\| m_{L-1}^{-1/2},$$

where we have used that $\theta = W^{L-1}$ because we only train the second but last layer. $\|\cdot\|$ denotes the Euclidean norm for vectors and the induced matrix norm for matrices. As in the introduction, we abbreviate $D := \mathbb{S}^{d-1}$.

Before we prove sub-gaussian increments, we bound several components involved in Y_θ separately. We start with the factor $\langle \kappa, \partial_r f_\theta \rangle_{H^s}$.

Lemma B.1. *Let $0 \leq s < 1$ and assume σ satisfies (35), (36), the weights satisfy (38) and are of equivalent size $m_0 \sim \dots \sim m_{L-1}$. Then*

$$\sum_{r=1}^R \langle \kappa, \partial_r f_\theta \rangle_{H^s}^2 \lesssim \|\kappa\|_{H^s}^2.$$

Proof. By Lemma D.2 cited from Welper (2024b), with ϵ sufficiently small so that $s + \epsilon < 1$ the bilinear form

$$\begin{aligned} B(u, v) &= \int_D \int_D u(x) \left(\sum_{r=1}^R \partial_r f_\theta(x) \partial_r f_\theta(y) \right) v(y) dx dy \\ &= \sum_{r=1}^R \langle u, \partial_r f_\theta \rangle \langle \partial_r f_\theta, v \rangle \end{aligned}$$

for u and v in L_2 is bounded by

$$B(u, v) \lesssim \|u\|_{H^{-s}} \|v\|_{H^{-s}}$$

and can therefore be extended to all $u, v \in H^{-s}$. In this case the L_2 inner product $\langle \cdot, \cdot \rangle$ turns into the $H^{-s} \times H^s$ dual pairing. Denoting by $R: H^s \rightarrow H^{-s}$ the Riesz map, we obtain

$$\begin{aligned} \sum_{r=1}^R \langle \kappa, \partial_r f_\theta \rangle_{H^s}^2 &= \sum_{r=1}^R \langle \kappa, \partial_r f_\theta \rangle_{H^s} \langle \partial_r f_\theta, \kappa \rangle_{H^s} \\ &= \sum_{r=1}^R \langle R\kappa, \partial_r f_\theta \rangle \langle \partial_r f_\theta, R\kappa \rangle = B(R\kappa, R\kappa) \lesssim \|R\kappa\|_{H^{-s}}^2 = \|\kappa\|_{H^s}^2, \end{aligned}$$

which proves the lemma □

Next, we consider the loss $\partial_r \ell(\theta) = \frac{1}{N} \sum_{i=1}^N \kappa(x_i) \partial_r f_\theta(x_i)$ and show sub-gaussian increments for the summands. As usual, we abbreviate $L_\infty(\bar{D}) = L_\infty$.

Lemma B.2. *Let $0 < s < 1$ and assume $\|f\|_{L_\infty} \lesssim m^{1/2}$, σ satisfies (35), (36), (37) and the weights satisfy (38). Let X be a uniform random variable with values in D and set*

$$X_\theta := \kappa(X) \nabla f_\theta(X) - \mathbb{E} [\kappa(X) \nabla f_\theta(X)]$$

Then

$$\| \|X_\theta - X_\vartheta\| \|_{\psi_2} \lesssim m^{1/2} \|\theta - \vartheta\|_*, \quad \| \|X_\theta\| \|_{\psi_2} \lesssim m^{1/2}.$$

Note that $\kappa(X)$ is a scalar and $\nabla f_\theta(X) = \nabla_{W^{L-1}} f_\theta(X)$ is a matrix. With the convention $\partial_{W_{ij}^{L-1}} = \partial_r$ for suitable r , we may also regard $\nabla f_\theta(X)$ as a vector and thus $\|\kappa(X) \nabla f_\theta(X)\|$ as the Euclidean vector norm, which we do in the following.

Proof. We split $X_\theta = Z_\theta - \mathbb{E}[Z_\theta]$ into a random part and an expectation, with

$$Z_\theta := \kappa(X) \nabla f_\theta(X)$$

and estimate both terms separately.

1. *Estimate $\|Z_\theta - Z_\vartheta\|_{\psi_2}$:* First, we upper bound the ψ_2 -norm by the L_∞ -norm and separate the factor κ :

$$\begin{aligned} \|Z_\theta - Z_\vartheta\|_{\psi_2}^2 &\leq \|Z_\theta - Z_\vartheta\|_{L_\infty}^2 = \|\kappa \nabla f_\theta - \kappa \nabla f_\vartheta\|_{L_\infty}^2 \\ &= \sup_{x \in D} \sum_{r=1}^R [\kappa(x) [\partial_r f_\theta(x) - \partial_r f_\vartheta(x)]]^2 \\ &\leq \|\kappa\|_{L_\infty}^2 \|\nabla f_\theta - \nabla f_\vartheta\|_{L_\infty}^2. \end{aligned} \quad (39)$$

The first factor is bounded by

$$\|\kappa\|_{L_\infty} \leq \|f\|_{L_\infty} + \|f_\theta\|_{L_\infty} \lesssim m^{1/2},$$

where we have used that $\kappa = f - f_\theta$ for some weights θ that satisfy (38), that $\|f\|_{L_\infty} \lesssim m^{-1/2}$ by assumption and $\|f_\theta\|_{L_\infty} \lesssim m^{-1/2}$ by Lemma D.3, cited from Welper (2024b).

To bound the second factor, recall that the derivatives ∂_r for $r \in \{1, \dots, R\}$ are shorthand for the derivatives $\partial_{W_{ij}^{L-1}}$ of the second but last layer. A short calculation shows that

$$\partial_{W_{ij}^{L-1}} f_\theta = \underbrace{W_i^L m_L^{-1/2} m_{L-1}^{-1/2}}_{:=w_i} \underbrace{\dot{\sigma}(f_i^L)}_{:=u_i(\theta)} \underbrace{\sigma(f_j^{L-1})}_{:=v_j}, \quad (40)$$

see e.g. Welper (2024b, Proof of Lemma 6.18) for details. The weights W^L of the last layer have only one index because the network is scalar valued. The layer f^{L-1} does not depend on W^{L-1} and therefore v_j does not depend on $\theta = W^{L-1}$. Then, for all $x \in D$, we have

$$\begin{aligned} \|\nabla f_\theta(x) - \nabla f_\vartheta(x)\|^2 &= \sum_{ij} [w_i u_i(\theta) v_j - w_i u_i(\vartheta) v_j]^2 \\ &= \sum_{ij} w_i^2 [u_i(\theta) - u_i(\vartheta)]^2 v_j^2 \\ &= \left[\sum_i w_i^2 [u_i(\theta) - u_i(\vartheta)]^2 \right] \left[\sum_j v_j^2 \right] \\ &\leq \|w\|_{\ell_\infty}^2 \|u(\theta) - u(\vartheta)\|^2 \|v\|^2. \end{aligned}$$

Since $W_i^L = \pm 1$, we have

$$\|w\|_{\ell_\infty} \lesssim m^{-1}$$

and from Lemma D.3 cited from Welper (2024b), together with the Lipschitz continuity of σ and $\dot{\sigma}$, we have

$$\|u(\theta) - u(\vartheta)\| \lesssim m^{1/2} \|\theta - \vartheta\|_*, \quad \|v\| \lesssim m^{1/2}.$$

Thus, we obtain

$$\|\nabla f_\theta(x) - \nabla f_\vartheta(x)\| \lesssim \|\theta - \vartheta\|_*$$

for all $x \in D$ and therefore together with (39) that

$$\|Z_\theta - Z_\vartheta\|_{\psi_2} \lesssim m^{1/2} \|\theta - \vartheta\|_*.$$

2. *Estimate $\|\mathbb{E}[Z_\theta] - \mathbb{E}[Z_\vartheta]\|$:* First, using that X is uniform, we factor out the residual κ :

$$\begin{aligned} \|\mathbb{E}[Z_\theta] - \mathbb{E}[Z_\vartheta]\|^2 &\leq \sum_{r=1}^R \left[\frac{1}{|D|} \int_D \kappa(x) [\partial_r f_\theta(x) - \partial_r f_\vartheta(x)] dx \right]^2 \\ &= \frac{1}{|D|} \|\kappa\|_{L_2}^2 \sum_{r=1}^R \frac{1}{|D|} \int_D |\partial_r f_\theta(x) - \partial_r f_\vartheta(x)|^2 dx \\ &\leq \|\kappa\|_{L_\infty}^2 \|\partial_r f_\theta - \partial_r f_\vartheta\|_{L_\infty}^2 \end{aligned}$$

where in the first inequality we have used Cauchy-Schwarz and in the last one exchanged the order of sum and integral. The right hand side is identical to (39) and bounded the same way.

Combining the estimates for Z_θ and its expectation and using that the ψ_2 -norm of a non-negative constant is upper bounded by the constant itself, we obtain

$$\|X_\theta - X_\vartheta\|_{\psi_2} \leq \|Z_\theta - Z_\vartheta\|_{\psi_2} + \|\mathbb{E}[Z_\theta] - \mathbb{E}[Z_\vartheta]\|_{\psi_2} \lesssim m^{1/2} \|\theta - \vartheta\|_*,$$

which shows the first part of the lemma.

The proof of the second bound $\|X_\theta\|_{\psi_2} \lesssim m^{1/2}$ is identical upon replacing X_ϑ and Z_ϑ with 0 throughout the proof. Using $\dot{\sigma}(x) \lesssim 1$, we obtain $\|u(\theta)\| \lesssim m^{1/2}$ instead of $\|u(\theta) - u(\vartheta)\| \lesssim m^{1/2} \|\theta - \vartheta\|_*$ by Lemma D.3. Omitting the factor $\|\theta - \vartheta\|_*$ in the left hand side throughout the rest of the proof shows the second part of the lemma. \square

The next lemma establishes the sub-gaussian increments of the random variable Y_θ .

Lemma B.3. *Let $0 < s < 1$ and assume $\|f\|_{L_\infty} \lesssim m^{1/2}$, σ satisfies (35), (36), (37) and the weights satisfy (38). Let X be a uniform random variable on D and for $\theta, \bar{\theta} \in \Theta$ set*

$$Y_\theta := \sum_{r=1}^R \langle \kappa, \partial_r f_{\bar{\theta}} \rangle_{H^s} [\partial_r \ell(\theta) - \partial_r \ell_0(\theta)].$$

Then

$$\|Y_\theta - Y_\vartheta\|_{\psi_2} \lesssim \left(\frac{m}{N}\right)^{1/2} \|\kappa\|_{H^s} \|\theta - \vartheta\|_*, \quad \|Y_\theta\|_{\psi_2} \lesssim \left(\frac{m}{N}\right)^{1/2} \|\kappa\|_{H^s}.$$

Proof. Plugging in the definitions of the loss ℓ and its continuum limit ℓ_0 , we have

$$\partial_r \ell(\theta) = \frac{1}{N} \sum_{i=1}^N \kappa(x_i) \partial_r f_\theta(x_i), \quad \partial_r \ell_0(\theta) = \langle \kappa, \partial_r f_\theta \rangle = \mathbb{E}[\kappa \partial_r f_\theta]$$

and therefore

$$Y_\theta = \frac{1}{N} \sum_{i=1}^N \sum_{r=1}^R \underbrace{\langle \kappa, \partial_r f_{\bar{\theta}} \rangle_{H^s}}_{=: u_r} \underbrace{[\kappa(x_i) \partial_r f_\theta(x_i) - \mathbb{E}[\kappa(X) \partial_r f_\theta(X)]]}_{=: (X_\theta^i)_r} = \frac{1}{N} \sum_{i=1}^N u^T X_\theta^i.$$

With Hoeffding's inequality, we estimate the ψ_2 -norm by

$$\|Y_\theta - Y_\vartheta\|_{\psi_2}^2 = \left\| \frac{1}{N} \sum_{i=1}^N u^T (X_\theta^i - X_\vartheta^i) \right\|_{\psi_2}^2 \lesssim \frac{1}{N^2} \sum_{i=1}^N \|u^T (X_\theta^i - X_\vartheta^i)\|_{\psi_2}^2.$$

The argument of the ψ_2 -norm is bounded by $|u^T(X_\theta^i - X_\vartheta^i)| \leq \|u\| \|X_\theta^i - X_\vartheta^i\|$ and thus

$$\begin{aligned} \|Y_\theta - Y_\vartheta\|_{\psi_2}^2 &\leq \frac{1}{N^2} \sum_{i=1}^N \|u\|^2 \|X_\theta^i - X_\vartheta^i\|_{\psi_2}^2 \\ &\lesssim \frac{1}{N^2} \sum_{i=1}^N \|\kappa\|_{H^s}^2 m \|\theta - \vartheta\|_*^2 \\ &\lesssim \frac{m}{N} \|\kappa\|_{H^s}^2 \|\theta - \vartheta\|_*^2, \end{aligned}$$

where the last inequalities follows from

$$\|u\| \lesssim \|\kappa\|_{H^s}, \quad \|\|X_\theta^i - X_\vartheta^i\|\|_{\psi_2} \lesssim m^{1/2} \|\theta - \vartheta\|_*,$$

by Lemmas B.1 and B.2, respectively. Taking the square root completes the proof of the first inequality. The second follows analogously by replacing Y_ϑ and X_ϑ with zero throughout the proof. \square

Lemma B.4. *Let $0 < s < 1$ and assume $\|f\|_{L^\infty} \lesssim m^{1/2}$, σ satisfies (35), (36), (37) and the weights satisfy (38). Let $\Theta_h := \{\theta \in \Theta \mid \|\theta - \theta^0\|_* \leq h\}$ for $h \lesssim 1$ and some initial (trained) weight θ^0 that satisfies (38). Then with probability at least $1 - 2\exp(-m^2 h^2)$*

$$\sup_{\theta \in \Theta_h} \left| \sum_{r=1}^R \langle \kappa, \partial_r f_{\bar{\theta}} \rangle_{H^s} [\partial_r \ell(\theta) - \partial_r \ell_0(\theta)] \right| \lesssim \frac{m^{3/2}}{N^{1/2}} h \|\kappa\|_{H^s}.$$

Proof. We abbreviate

$$Y_\theta := \sum_{r=1}^R \langle \kappa, \partial_r f_{\bar{\theta}} \rangle_{H^s} [\partial_r \ell(\theta) - \partial_r \ell_0(\theta)]$$

and prove the lemma with Dudley's inequality. We have established in Lemma B.3 that Y_θ has sub-gaussian increments

$$\|Y_\theta - Y_\vartheta\|_{\psi_2} \lesssim \left(\frac{m}{N}\right)^{1/2} \|\kappa\|_{H^s} \|\theta - \vartheta\|_*.$$

Next, we estimate the covering numbers in Dudley's inequality. The set of eligible parameters Θ_h is contained in the ball of radius h in \mathbb{R}^R in the $\|\cdot\|_*$ -norm. Hence, for every $\epsilon \geq 0$ there is an ϵ -covering of at most

$$N(\epsilon) \leq \left(\frac{3h}{\epsilon}\right)^R$$

ϵ -balls, see e.g. Lorentz et al. (1996, Chapter 15, Proposition 1.3). Then, using $\log x \leq x - 1 \leq x$,

$$\int_0^\infty \sqrt{\log N(\epsilon)} d\epsilon = R^{1/2} \int_0^h \sqrt{\log \left(\frac{3h}{\epsilon}\right)} d\epsilon \leq R^{1/2} \int_0^h \sqrt{\left(\frac{3h}{\epsilon}\right)} d\epsilon = \frac{1}{2} (3hR)^{1/2} h^{1/2} \lesssim mh,$$

where in the last step we have used that $R^{1/2} \sim m$. Hence, Dudley's inequality (with tail bounds) implies that for all $u \geq 0$

$$\sup_{\theta, \vartheta \in \Theta_h} |Y_\theta - Y_\vartheta| \lesssim \left(\frac{m}{N}\right)^{1/2} \|\kappa\|_{H^s} \left[\int_0^\infty \sqrt{\log N(\epsilon)} d\epsilon + uh \right]$$

holds with probability at least $1 - 2\exp(-u^2)$. Choosing $u = m$ yields

$$\sup_{\theta, \vartheta \in \Theta_h} |Y_\theta - Y_\vartheta| \lesssim \left(\frac{m}{N}\right)^{1/2} mh \|\kappa\|_{H^s},$$

with probability at least $1 - 2\exp(-m^2)$. The sub-gaussian bound $\|Y_\theta\|_{\psi_2} \lesssim \left(\frac{m}{N}\right)^{1/2} \|\kappa\|_{H^S}$ from Lemma B.3 implies

$$|Y_{\theta^0}| \lesssim \left(\frac{m}{N}\right)^{1/2} m h \|\kappa\|_{H^S}.$$

with probability at least $1 - 2\exp(-m^2 h^2)$. Then, the lemma follows from $\sup_{\theta \in \Theta_h} |Y_\theta| \leq \sup_{\theta \in \Theta_h} |Y_\theta - Y_{\theta^0}| + |Y_{\theta^0}|$. \square

B.2 Convergence: Proof of Theorem 2.2

Proof of Theorem 2.2. The theorem follows from Theorem A.1 with $\mathcal{H}^s = H^s$ and $\mathcal{H}^0 = L_2$, for which we have to verify all assumptions. Most of them have been established in the proof of Welper (2024a, Theorem 2.2), which is identical to the one of this paper, except that it uses a continuous L_2 loss instead of a sample loss. This results in the extra assumption (33), which is the only one left to verify.

By Lemma B.4, with probability at least $1 - 2\exp(-m^2 h^2)$ we have

$$\sup_{\theta \in \Theta_h} \left| \sum_{r=1}^R \langle \kappa^k, \partial_r f_{\bar{\theta}} \rangle_{H^S} [\partial_r \ell(\theta) - \partial_r \ell_0(\theta)] \right| \lesssim \frac{m^{3/2}}{N^{1/2}} h \|\kappa^k\|_{H^S},$$

which provides an estimate for $\Delta_{\text{sample}}(m, N)$ up to a missing factor $\|\kappa^k\|_{L_2}$. To insert this factor, we use the lower bound from Assumption (34): For $k \leq n$

$$\|\kappa^k\|_{L_2}^2 \geq c_a (h^\alpha + \Delta_{\text{sample}}(m, N))^{\frac{s}{\beta}} \|\kappa^0\|_{H^S}^2 \geq c_a h^{\alpha \frac{s}{\beta}} \|\kappa^0\|_{H^S}^2,$$

which implies

$$1 \lesssim h^{-\frac{\alpha s}{2\beta}} \|\kappa^0\|_{H^S}^{-1} \|\kappa^k\|_{L_2}.$$

Inserting this into the application of Lemma B.4 above, we obtain

$$\sup_{\theta \in \Theta_h} \left| \sum_{r=1}^R \langle \kappa^k, \partial_r f_{\bar{\theta}} \rangle_{H^S} [\partial_r \ell(\theta) - \partial_r \ell_0(\theta)] \right| \lesssim \left[\frac{m^{3/2}}{N^{1/2}} h^{1-\frac{\alpha s}{2\beta}} \|\kappa^0\|_{H^S}^{-1} \right] \|\kappa^k\|_{L_2} \|\kappa^k\|_{H^S}.$$

Comparing to the definition of $\Delta_{\text{sample}}(m, N)$ in assumption (33), we obtain

$$\Delta_{\text{sample}}(m, N) \lesssim \frac{m^{3/2}}{N^{1/2}} h^{1-\frac{\alpha s}{2\beta}} \|\kappa^0\|_{H^S}^{-1},$$

which shows the assumption.

It remains to consider the success probability. From the parts of the proof shown in Welper (2024a), the bounds fail with probability $cL(e^{-m} + e^{-\tau})$ and from above with probability $e^{-cm^2 h^2} \leq e^{cm}$ because $h = c_h m^{-\frac{1}{2} \frac{1}{1+\alpha}} \geq m^{-1/2}$ for any $\alpha > 0$ from (30).

This completes the proof, together with an index shift $n+1 \rightarrow n$ between the statements of Theorems A.1 and 2.2. \square

C Kernels

In this section, we prove Theorem 2.4 based on bounds for the error $\Delta_{\text{sample}}(m, N)$ between the $L_2(D)$ loss and the kernel loss

$$\ell^k(\theta) := \frac{1}{2N} \sum_{i=1}^N \langle k(x_i, \cdot), f_\theta - f \rangle^2$$

for kernels $k(x, y)$ with $x, y \in D = \mathbb{S}^{d-1}$ and x_i uniformly and independently sampled on the domain D . We denote the corresponding expected loss by

$$\bar{\ell}^k(\theta) := \mathbb{E} [\ell^k(\theta)] := \frac{1}{2|D|} \int_D \langle k(x, \cdot), f_\theta - f \rangle^2 dx.$$

We consider this expectation in Section C.1, which is generally not identical to $\frac{1}{2} \|f_\theta - f\|_{L_2(D)}^2$. Then, we show corresponding concentration inequalities based on matrix Bernstein inequalities in Section C.2. The proof of Theorem 2.4 is in Section C.3.

Throughout this section, we abbreviate $D = \mathbb{S}^{d-1}$, $L_2 = L_2(D)$, $H^s = H^s(D)$ and $\langle \cdot, \cdot \rangle_{H^s} = \langle \cdot, \cdot \rangle_{H^s(D)}$, when convenient.

C.1 Expectation

Define the inner product and norm that give rise to the expected kernel loss as

$$\langle u, v \rangle_k = \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N \langle u, k(x_i, \cdot) \rangle \langle k(x_i, \cdot), v \rangle \right], \quad \|\cdot\|_k^2 := \langle \cdot, \cdot \rangle_k, \quad \bar{\ell}^k(\theta) := \frac{1}{2} \|f_\theta - f\|_k^2. \quad (41)$$

Unfortunately, the norm $\|\cdot\|_k$ is not equivalent to the L_2 norm and therefore the expected loss $\bar{\ell}^k(\theta)$ is not equivalent to the L_2 loss $\ell_0(\theta)$. To bridge this gap, in this section, we construct a modified inner product and corresponding norm and loss

$$\langle u, v \rangle_\sharp \quad \|\cdot\|_\sharp^2 := \langle \cdot, \cdot \rangle_\sharp, \quad \bar{\ell}^\sharp(\theta) := \frac{1}{2} \|f_\theta - f\|_\sharp^2, \quad (42)$$

with the following properties:

1. The two inner products $\langle \cdot, \cdot \rangle_k$ and $\langle \cdot, \cdot \rangle_\sharp$ are close for smooth functions.
2. The norm $\|\cdot\|_\sharp^2$ and loss $\bar{\ell}^\sharp(\theta)$ are equivalent to $\|\cdot\|_{L_2}$ and $\ell_0(\theta)$, respectively.

To this end, we first characterize the inner product $\langle \cdot, \cdot \rangle_k$.

Lemma C.1. *Assume the symmetric kernel $k: D \times D \rightarrow \mathbb{R}$ has L_2 -orthogonal eigenfunctions ψ_j with eigenvalues λ_j . Then for any $u, v \in L_2$*

$$\langle u, v \rangle_k = \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N \langle u, k(\cdot, x_i) \rangle \langle k(x_i, \cdot), v \rangle \right] = \sum_{r=1}^{\infty} \lambda_j^2 \langle u, \psi_r \rangle \langle \psi_r, v \rangle.$$

Proof. We denote the expectation by E . By the law of large numbers, we have

$$E := \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N \langle u, k(\cdot, x_i) \rangle \langle k(x_i, \cdot), v \rangle \right] = \frac{1}{|D|} \int_D \langle u, k(\cdot, x) \rangle \langle k(x, \cdot), v \rangle dx.$$

Plugging in u, v in eigenbasis

$$u = \sum_{s=1}^{\infty} \langle u, \psi_s \rangle \psi_s, \quad v = \sum_{t=1}^{\infty} \langle v, \psi_t \rangle \psi_t,$$

we obtain

$$\begin{aligned} E &= \sum_{s,t=1}^{\infty} \langle u, \psi_s \rangle \langle v, \psi_t \rangle \int_D \langle \psi_s, k(\cdot, x) \rangle \langle k(x, \cdot), \psi_t \rangle dx \\ &= \sum_{s,t=1}^{\infty} \langle u, \psi_s \rangle \langle v, \psi_t \rangle \lambda_s \lambda_t \frac{1}{|D|} \int_D \psi_s(x) \psi_t(x) dx \\ &= \sum_{s=1}^{\infty} \lambda_s^2 \langle u, \psi_s \rangle \langle \psi_s, v \rangle, \end{aligned}$$

where in the second step we have used eigenvalue and vector definition $\langle \psi_s, k(\cdot, x) \rangle = \lambda_s \psi_s(x)$ and in the last step we have used that ψ_s is a orthonormal basis, with normalized squared expectation. \square

In the next lemma, we define the modified inner product $\langle \cdot, \cdot \rangle_\#$ and prove the properties from the introduction of this section.

Lemma C.2. *Assume the symmetric kernel $k: D \times D \rightarrow \mathbb{R}$ has L_2 -orthogonal eigenfunctions ψ_j and eigenvalues λ_j with*

$$\begin{aligned} 1 &\lesssim \lambda_j \lesssim 1, & j \leq J, \\ \lambda_j &\lesssim 1, & j > J. \end{aligned}$$

Define a lower bounded perturbation $\bar{\lambda}_j$ of the eigenvalue λ_j and the corresponding inner product

$$\bar{\lambda}_j := \begin{cases} \lambda_j & j \leq J \\ \max\{\lambda_j, 1\} & j \geq J. \end{cases} \quad \langle u, v \rangle_\# := \sum_{r=1}^{\infty} \bar{\lambda}_j^2 \langle u, \psi_r \rangle \langle \psi_r, v \rangle.$$

For some increasing weights $\mu_j \geq 1$ let

$$\|v\|_{\mathcal{H}^s}^2 := \sum_{j=1}^J \mu_j^{2s} \langle \psi_j, v \rangle^2$$

be the norm of a Hilbert space \mathcal{H}^s . Then for all $u, v \in L_2$ and $s \geq 0$

$$\left| \langle u, v \rangle_k - \langle u, v \rangle_\# \right| = \left| \sum_{j=1}^{\infty} [\lambda_j^2 - \bar{\lambda}_j^2] \langle u, \psi_j \rangle \langle \psi_j, v \rangle \right| \leq \mu_J^{-s} \|u\|_{L_2} \|v\|_{\mathcal{H}^s},$$

and the induced norm $\|\cdot\|_\#^2 := \langle \cdot, \cdot \rangle_\#$ is equivalent to the L_2 norm.

The weights μ_j and \mathcal{H}^s define a smoothness space, which is left generic for now but will be replaced by Sobolev spaces below.

Proof. We have

$$\begin{aligned} \left| \sum_{j=1}^{\infty} [\lambda_j^2 - \bar{\lambda}_j^2] \langle u, \psi_j \rangle \langle \psi_j, v \rangle \right| &= \left| \sum_{j>J} [\lambda_j^2 - \bar{\lambda}_j^2] \langle u, \psi_j \rangle \langle \psi_j, v \rangle \right| \\ &\leq \left[\sum_{j>J} |\lambda_j^2 - \bar{\lambda}_j^2| \langle u, \psi_j \rangle^2 \right]^{1/2} \left[\sum_{j>J} |\lambda_j^2 - \bar{\lambda}_j^2| \langle v, \psi_j \rangle^2 \right]^{1/2} \\ &\leq \left[\sum_{j>J} \langle u, \psi_j \rangle^2 \right]^{1/2} \left[\sum_{j>J} \langle v, \psi_j \rangle^2 \right]^{1/2} \\ &\leq \mu_J^{-s} \|u\|_{L_2} \|v\|_{\mathcal{H}^s}, \end{aligned}$$

where the first equality follows from $\lambda_j = \bar{\lambda}_j$ for $j \leq J$, the second from Cauchy-Schwarz and the third from $0 \leq \bar{\lambda}_j^2 - \lambda_j^2 \leq 1$ by definition of $\bar{\lambda}_j$ for $j > J$. The last inequality follows from

$$\sum_{j>J} \langle u, \psi_j \rangle^2 \leq \|u\|_{L_2}^2$$

and

$$\sum_{j>J} \langle u, \psi_j \rangle^2 = \mu_J^{-2s} \sum_{j>J} \mu_j^{2s} \langle u, \psi_j \rangle^2 \leq \mu_J^{-2s} \sum_{j>J} \mu_j^{2s} \langle u, \psi_j \rangle^2 = \mu_J^{-2s} \|u\|_{\mathcal{H}^s}^2.$$

Finally, by construction, the eigenvalues $\bar{\lambda}_j \sim 1$ are upper and lower bounded by one and therefore the norm $\|\cdot\|_\#$ is equivalent to the L_2 norm. \square

C.2 Concentration

We show concentration for

$$\sum_{r=1}^R \langle \kappa, \partial_r f_{\bar{\theta}} \rangle_{H^s} [\partial_r \ell^k(\theta) - \partial_r \bar{\ell}^k(\theta)]$$

which matches Assumption (33), except for the wrong expected loss $\bar{\ell}^k(\theta)$ instead of $\ell_0(\theta)$, which will be considered in the next section. Throughout this section, we denote the adjoint of u , by u^* , so that $\langle f, u \rangle = f^* u$. For weights $\theta, \bar{\theta} \in \Theta$, we use the abbreviation

$$F_{\theta, \bar{\theta}}^s \kappa := \sum_{r=1}^R \partial_r f_{\theta} \langle \partial_r f_{\bar{\theta}}, \kappa \rangle_{H^s} \quad (43)$$

which shows up repeatedly in the following proofs. We first bound its norm.

Lemma C.3. *Assume that σ satisfies the growth and Lipschitz conditions (35), (36) and may be different in each layer. Assume all weights in the networks, $\theta, \bar{\theta}$ and the domain are bounded (38). Then for $0 \leq S \leq s < 1$*

$$\|F_{\theta, \bar{\theta}}^S \kappa\|_{H^s} \leq \|\kappa\|_{H^s}.$$

Proof. Let $\langle \cdot, \cdot \rangle$ be the $H^{-S} \times H^S$ dual pairing and $R : H^S \rightarrow H^{-S}$ the corresponding Riesz map. Since the H^s -norm is the dual of H^{-s} , we have

$$\begin{aligned} \|F_{\theta, \bar{\theta}}^S \kappa\|_{H^s} &= \left\| \sum_{r=1}^R \partial_r f_{\theta} \langle \partial_r f_{\bar{\theta}}, \kappa \rangle_{H^S} \right\|_{H^s} \\ &= \sup_{\|v\|_{H^{-s}} \leq 1} \left\langle v, \sum_{r=1}^R \partial_r f_{\theta} \langle \partial_r f_{\bar{\theta}}, \kappa \rangle_{H^S} \right\rangle = \sup_{\|v\|_{H^{-s}} \leq 1} \left\langle v, \left(\sum_{r=1}^R \partial_r f_{\theta} \partial_r f_{\bar{\theta}}^* \right) R \kappa \right\rangle \end{aligned}$$

where in the last step we have used the definition of the adjoint functional $f_{\bar{\theta}}^* \in H^{-S}$ so that $\langle f_{\bar{\theta}}, \kappa \rangle_{H^S} = \langle f_{\bar{\theta}}^*, R \kappa \rangle = f_{\bar{\theta}}^*(R \kappa)$. Formally, in this argument $\langle \cdot, \cdot \rangle$ is the $H^S \times H^{-S}$ dual pairing, which reduces to the L_2 inner product in case $R \kappa$ is in L_2 , see the proof of Lemma B.1 for more details. By Lemma D.2 cited from Welper (2024b), with ϵ sufficiently small so that $s + \epsilon < 1$, we obtain

$$\|F_{\theta, \bar{\theta}}^s \kappa\|_{H^s} \lesssim \sup_{\|v\|_{H^{-s}} \leq 1} \|v\|_{H^{-s}} \|R \kappa\|_{H^{-s}} \leq \|R \kappa\|_{H^{-s}} \leq \|R \kappa\|_{H^{-S}} = \|\kappa\|_{H^S} \leq \|\kappa\|_{H^s},$$

where we have used the embeddings $\|\cdot\|_{H^{-s}} \leq \|\cdot\|_{H^{-S}}$ and $\|\cdot\|_{H^S} \leq \|\cdot\|_{H^s}$ because $S \leq s$.

□

The next lemma shows the main concentration result.

Lemma C.4. *Assume that σ satisfies the growth and Lipschitz conditions (35), (36) and may be different in each layer. Assume all weights in the networks, $\theta, \bar{\theta}$ and the domain are bounded (38). Assume the kernel is bounded by $\sup_{x \in D} \|k(x, \cdot)\|_{L_2} \leq C_k$. Then for $0 \leq s < 1$ with probability at least $1 - 2t(e^t - t - 1)^{-1}$ we have*

$$\sum_{r=1}^R \langle \kappa, \partial_r f_{\bar{\theta}} \rangle_{H^s} [\partial_r \ell^k(\theta) - \partial_r \bar{\ell}^k(\theta)] \lesssim C_k^2 \left(\sqrt{\frac{t}{N}} + \frac{t}{N} \right) \|\kappa\|_{L_2} \|\kappa\|_{H^s}.$$

for all $\kappa \in H^s$.

Proof. The lemma states that the gradient of the sample loss is close to its average, with high probability, which we show by the matrix Bernstein inequality. To this end, we first unravel the definition of the loss to

define appropriate random matrices:

$$\begin{aligned}
\sum_{r=1}^R \langle \kappa, \partial_r f_{\bar{\theta}} \rangle_{H^s} \partial_r \ell^k(\theta) &= \sum_{r=1}^R \langle \kappa, \partial_r f_{\bar{\theta}} \rangle_{H^s} \left(\frac{1}{N} \sum_{i=1}^N \langle k(x_i, \cdot), \kappa \rangle \langle k(x_i, \cdot), \partial_r f_{\bar{\theta}} \rangle \right) \\
&= \frac{1}{N} \sum_{i=1}^N \sum_{r=1}^R \langle \kappa, k(x_i, \cdot) \rangle \langle k(x_i, \cdot), \partial_r f_{\bar{\theta}} \rangle \langle \kappa, \partial_r f_{\bar{\theta}} \rangle_{H^s} \\
&= \left\langle \kappa, \frac{1}{N} \sum_{i=1}^N k(x_i, \cdot) k(x_i, \cdot)^* \sum_{r=1}^R \partial_r f_{\bar{\theta}} \langle \kappa, \partial_r f_{\bar{\theta}} \rangle_{H^s} \right\rangle \\
&= \left\langle \kappa, M F_{\theta, \bar{\theta}}^s \kappa \right\rangle
\end{aligned}$$

with $F_{\theta, \bar{\theta}}^s \kappa$ defined in (43) and the integral kernel

$$M := \frac{1}{N} \sum_{i=1}^N k(x_i, \cdot) k(x_i, \cdot)^*.$$

With an analogous computation, we obtain the expectation with respect to the samples x_i :

$$\mathbb{E} \left[\sum_{r=1}^R \langle \kappa, \partial_r f_{\bar{\theta}} \rangle_{H^s} \partial_r \ell^k(\theta) \right] = \left\langle \kappa, \bar{M} F_{\theta, \bar{\theta}}^s \kappa \right\rangle,$$

with

$$\bar{M} := \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N k(x_i, \cdot) k(x_i, \cdot)^* \right] = \int_D k(x, \cdot) k(x, \cdot)^* dx.$$

It follows that

$$\sum_{r=1}^R \langle \kappa, \partial_r f_{\bar{\theta}} \rangle_{H^s} [\partial_r \ell^k(\theta) - \partial_r \bar{\ell}^k(\theta)] = \left\langle \kappa, [M - \bar{M}] F_{\theta, \bar{\theta}}^s \kappa \right\rangle \leq \|\kappa\|_{L_2} \|M - \bar{M}\|_{H^s \rightarrow L_2} \|F_{\theta, \bar{\theta}}^s \kappa\|_{H^s} \quad (44)$$

The last term

$$\|F_{\theta, \bar{\theta}}^s \kappa\|_{H^s} \lesssim \|\kappa\|_{H^s} \quad (45)$$

is bounded by Lemma C.3 so that it remains to bound $\|M - \bar{M}\|_{H^s \rightarrow L_2}$. To this end note that M is a sum of rank one operators, each depending on one independent sample x_i . Hence, concentration follows from dimension free matrix Bernstein inequalities (Hsu et al., 2012; Tropp, 2015; Minsker, 2017). We use a corollary shown in Gentile & Welper (2022) and stated in the supplementary material Lemma D.4 for which we only need to bound

$$\|k(x, \cdot)^*\|_{(H^s)^*} = \|k(x, \cdot)\|_{H^{-s}} \leq \|k(x, \cdot)\|_{L_2} \leq C_k$$

for all $x \in D$, provided by the assumptions of the lemma. With Lemma D.4 this provides

$$\Pr \left[\|M - \bar{M}\|_{H^s \rightarrow L_2} \gtrsim C_k^2 \left(\sqrt{\frac{t}{N}} + \frac{t}{N} \right) \right] \leq 2t (e^t - t - 1)^{-1}.$$

Together with (44) and (45) this proves the lemma. □

C.3 Convergence: Proof of Theorem 2.4

We first bound $\Delta_{\text{sample}}(m, N)$ based on the results in the last two sections.

Lemma C.5. Assume that σ satisfies the growth and Lipschitz conditions (35), (36) and may be different in each layer. Assume all weights in the networks, θ , $\bar{\theta}$ and the domain are bounded (38). Assume the kernel $k: D \times D \rightarrow \mathbb{R}$ is zonal, i.e. $k(x, y) = k(x^T y)$, and has eigenvalues λ_{lj} with

$$\begin{aligned} 1 &\lesssim \lambda_{lj} \lesssim 1, \quad l \leq L, \quad 1 \leq j \leq \nu(l), \\ \lambda_{lj} &\lesssim 1, \quad l > L, \quad 1 \leq j \leq \nu(l) \end{aligned}$$

and index structure and $\nu(l)$ matching spherical harmonics (6) and

$$\sup_{x \in D} \|k(x, \cdot)\|_{L_2} \leq C_k.$$

Then for $0 \leq S \leq s < 1$ with probability at least $1 - 2t[e^t - t - 1]^{-1}$ for all $\kappa \in H^s$ we have

$$\sum_{r=1}^R \langle \kappa, \partial_r f_{\bar{\theta}} \rangle_{H^s} [\partial_r \ell^k(\theta) - \partial_r \bar{\ell}^\sharp(\theta)] \lesssim C_k^2 \left[\sqrt{\frac{t}{N}} + \frac{t}{N} \right] \|\kappa\|_{L_2} \|\kappa\|_{H^s} + L^{-s} \|\kappa\|_{L_2} \|\kappa\|_{H^s},$$

with loss $\bar{\ell}^\sharp(\theta)$ defined in (42).

Note that the conclusion of the lemma contains both S and s . This allows $S = s$, but also the choice $S = 0$ for which the last summand in the right hand side does not provide a meaningful bound if we insist that $S = s = 0$.

Proof. We first split the difference into expectation and concentration components:

$$\begin{aligned} &\sum_{r=1}^R \langle \kappa, \partial_r f_{\bar{\theta}} \rangle_{H^s} [\partial_r \ell^k(\theta) - \partial_r \bar{\ell}^\sharp(\theta)] \\ &= \sum_{r=1}^R \langle \kappa, \partial_r f_{\bar{\theta}} \rangle_{H^s} [\partial_r \ell^k(\theta) - \partial_r \bar{\ell}^k(\theta)] + \sum_{r=1}^R \langle \kappa, \partial_r f_{\bar{\theta}} \rangle_{H^s} [\partial_r \bar{\ell}^k(\theta) - \partial_r \bar{\ell}^\sharp(\theta)] \\ &:= (I) + (II). \end{aligned}$$

The first part (I) is bounded by Lemma C.4:

$$(I) \lesssim C_k^2 \left(\sqrt{\frac{t}{N}} + \frac{t}{N} \right) \|\kappa\|_{L_2} \|\kappa\|_{H^s},$$

with probability at least $1 - 2t(e^t - t - 1)^{-1}$.

To estimate the expectation part (II), first note that the Funk-Hecke formula (Atkinson & Han, 2012) implies that the eigenfunctions of the zonal kernel $k(\cdot, \cdot)$ are spherical harmonics Y_l^j and thus L_2 orthogonal. Hence, we obtain the loss

$$\bar{\ell}^\sharp(\theta) := \frac{1}{2} \|f_\theta - f\|_\sharp^2 := \frac{1}{2} \sum_{l,j} \bar{\lambda}_{lj}^2 \langle Y_l^j, f_\theta - f \rangle^2, \quad \bar{\lambda}_j := \begin{cases} \lambda_j & j \leq J, \\ \max\{\lambda_j, 1\} & j \geq J, \end{cases}$$

by Lemma C.2. The partial derivatives of the modified and expected kernel loss are given by

$$\partial_r \bar{\ell}^\sharp(\theta) = \sum_{l,j} \bar{\lambda}_{lj}^2 \langle \kappa, Y_l^j \rangle \langle Y_l^j, \partial_r f_\theta \rangle,$$

and

$$\partial_r \bar{\ell}^k(\theta) = \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N \langle \kappa, k(x_i, \cdot) \rangle \langle k(x_i, \cdot), \partial_r f_\theta \rangle \right] = \sum_{l,j} \lambda_{lj}^2 \langle \kappa, Y_l^j \rangle \langle Y_l^j, \partial_r f_\theta \rangle,$$

by Lemma C.1, respectively. It follows that

$$\begin{aligned}
(II) &= \sum_{r=1}^R \langle \kappa, \partial_r f_{\bar{\theta}} \rangle_{H^s} \left[\sum_{l,j} [\lambda_{lj}^2 - \bar{\lambda}_{lj}^2] \langle \kappa, Y_l^j \rangle \langle Y_l^j, \partial_r f_{\bar{\theta}} \rangle \right] \\
&= \sum_{l,j} [\lambda_{lj}^2 - \bar{\lambda}_{lj}^2] \langle \kappa, Y_l^j \rangle \left\langle Y_l^j, \partial_r f_{\bar{\theta}} \sum_{r=1}^R \langle \kappa, \partial_r f_{\bar{\theta}} \rangle_{H^s} \right\rangle \\
&= \sum_{l,j} [\lambda_{lj}^2 - \bar{\lambda}_{lj}^2] \langle \kappa, Y_l^j \rangle \langle Y_l^j, F_{\bar{\theta}, \bar{\theta}}^S \kappa \rangle,
\end{aligned}$$

with $F_{\bar{\theta}, \bar{\theta}}^S \kappa$ defined in (43). Bounding the latter with Lemma C.3 and using Lemma C.2, we conclude that

$$(II) \lesssim \mu_L^{-s} \|\kappa\|_{L_2} \|F_{\bar{\theta}, \bar{\theta}}^S \kappa\|_{H^s} \lesssim \mu_L^{-s} \|\kappa\|_{L_2} \|\kappa\|_{H^s},$$

where μ_l are the weights in the definition of the Sobolev space H^s via spherical harmonics (5) and therefore $\mu_l \sim l$. Together with the bounds for (I) this completes the proof. \square

Proof of Theorem 2.4. The theorem follows from Theorem A.1. It shows gradient descent convergence in arbitrary scales of Hilbert spaces, for which the natural choice is $\mathcal{H}^s = H^s$. However, we replace the $L_2 = H^0$ norm with the equivalent $\|\cdot\|_{\sharp}$ norm to utilize our concentration result in Lemma C.5. This choice does not alter any assumptions or conclusions, except coercivity which uses the \mathcal{H}^0 inner product instead of the corresponding norm. We show that coercivity of the L_2 inner product implies coercivity of the $\|\cdot\|_{\sharp}$ inner product as well as the sample assumption (33). All other assumptions are proven in Welper (2024b, Theorem 2.2) for an analogous result without sample error.

1. *Coercivity* (28): Since the neural tangent kernel is zonal, by the Funk-Hecke formula (Atkinson & Han, 2012), it has spherical harmonics as eigenfunctions with some eigenvalues μ_{lj} , so that we have

$$\langle v, Hv \rangle_{\sharp} = \sum_{l,j} \bar{\lambda}_l^2 \mu_{lj} v_{lj}^2 \sim \sum_{l,j} \mu_{lj} v_{lj}^2 = \langle v, Hv \rangle,$$

with $v_{lj} = \langle v, Y_l^j \rangle$ and the $\|\cdot\|_{\sharp}$ norm from (42). Hence, $\|\cdot\|_{\sharp}$ -coercivity is equivalent to the regular L_2 -coercivity.

2. *Assumption* (33): We show that with high probability for all $\theta, \bar{\theta}$ with $\|\cdot\|_{*}$ distance to the initial θ^0 smaller than $\lesssim h$ and $S \in \{0, s\}$ we have the bound

$$\sum_{r=1}^R \langle \kappa, \partial_r f_{\bar{\theta}} \rangle_{H^s} [\partial_r \ell^k(\theta) - \ell_{\sharp}^k(\theta)] \leq \Delta_{\text{sample}}(m, N) \|\kappa\|_{L_2} \|\kappa\|_{H^s} \quad (46)$$

with

$$\Delta_{\text{sample}}(m, N) \lesssim C_k^2 \left(\frac{\tau_N}{N} \right)^{1/2} + C_k^{-2} \left(\frac{N}{\tau_N} \right)^{1/2} L^{-s} + L^{-s}, \quad (47)$$

which provides Assumption (33). To this end, by Lemma C.5 with $t = \tau_N$ and probability at least $1 - 2\tau_N [e^{\tau_N} - \tau_N - 1]^{-1}$ for all $\kappa \in \mathcal{H}^s$ we have

$$\sum_{r=1}^R \langle \kappa, \partial_r f_{\bar{\theta}} \rangle_{H^s} [\partial_r \ell^k(\theta) - \bar{\ell}^k(\theta)] \lesssim C_k^2 \left(\frac{\tau_N}{N} \right)^{1/2} \|\kappa\|_{L_2} \|\kappa\|_{H^s} + L^{-s} \|\kappa\|_{L_2} \|\kappa\|_{H^s}. \quad (48)$$

This yields the claimed bounds (46), (47) for $S = s$. For $S = 0$ the last summand has the wrong norm: $L^{-s}\|\kappa\|_{L_2}\|\kappa\|_{H^s}$ instead of $L^{-s}\|\kappa\|_{L_2}\|\kappa\|_{H^s} = L^{-s}\|\kappa\|_{L_2}^2$. To replace the H^s norm with an L_2 norm note that assumption (34) yields

$$\|\kappa^k\|_{L_2}^2 \geq c_a \left(m^{-\frac{1}{2} \frac{\alpha}{1+\alpha}} + \Delta_{\text{sample}}(m, N) \right)^{\frac{s}{\beta}} \|\kappa^0\|_{H^s}^2 \gtrsim C_k^2 \left(\frac{\tau_N}{N} \right)^{1/2} \|\kappa^0\|_{H^s}^2$$

for $k \leq n$. Moreover, by induction, from Theorem A.1 we have $\|\kappa^k\|_{H^s} \lesssim \|\kappa^0\|_{H^s}$ and therefore arrive at

$$\|\kappa^k\|_{H^s}^2 \lesssim C_k^{-2} \left(\frac{N}{\tau_N} \right)^{1/2} \|\kappa^k\|_{L_2}.$$

Together with (48) this yields the claimed bounds (46) and (47) for the case $S = 0$. Together with the case $S = s$ above, this establishes assumption (33).

This completes the proof, together with an index shift $n+1 \rightarrow n$ between the statements of Theorems A.1 and 2.2. □

D Supplementary Material

D.1 Technical Lemmas

Lemma D.1. *Let $k_t(x, y)$ be the heat kernel defined in (21). Then for all $y \in \mathbb{S}^{d-1}$*

$$\|k_t(\cdot, y)\|_{H^s(\mathbb{S}^{d-1})}^2 \lesssim t^{-s-d+3/2}.$$

Proof. Plugging the definition of the heat kernel (21) into the definition of Sobolev norms (5), we obtain

$$\begin{aligned} \|k_t(\cdot, y)\|_{H^s(\mathbb{S}^{d-1})}^2 &= \sum_{l=0}^{\infty} \sum_{j=1}^{\nu(l)} \left(1 + l^{1/2}(l+d-2)^{1/2} \right)^{2s} \left| e^{-l(l+d-2)t} Y_l^j(y) \right|^2 \\ &\lesssim 1 + \sum_{l=0}^{\infty} l^{2s} e^{-2l^2 t} \sum_{j=1}^{\nu(l)} \left| Y_l^j(y) \right|^2. \end{aligned}$$

Since $|Y_l^j(y)|^2 \lesssim \nu(l)$ and $\nu(l) \lesssim l^{d-2}$, see Stein & Weiss (1972, Chapter 4.2, Corollary 2.9), we obtain

$$\begin{aligned} \|k_t(\cdot, y)\|_{H^s(\mathbb{S}^{d-1})}^2 &\lesssim 1 + \sum_{l=0}^{\infty} l^{2s} l^{2d-4} e^{-2l^2 t} \lesssim 1 + \int_0^{\infty} l^{2s+2d-4} e^{-2l^2 t} dl \\ &= 1 + t^{-s-d+3/2} \int_0^{\infty} x^{2s+2d-4} e^{-2x^2} dx \lesssim t^{-s-d+3/2}, \end{aligned}$$

where we have substituted $x = l\sqrt{t}$ and used that the latter integral is bounded. □

D.2 Results from Gentile & Welper (2022); Welper (2024b;a)

To keep the paper self contained, this section contains several results from Gentile & Welper (2022); Welper (2024b;a).

Lemma D.2. *Assume that σ and $\dot{\sigma}$ satisfy the growth and Lipschitz conditions (35), (36) and may be different in each layer. Assume the weights, perturbed weights and domain are bounded (38) and $m_L \sim m_{L-1} \sim \dots \sim m_1$. Then for $0 < s < 1$ and $m_0 := m_1$*

$$\iint_{D \times D} f(x) \left(\sum_{r=1}^R \partial_r f_{\theta}(x) \partial_r f_{\theta}(y) \right) g(y) dx dy \lesssim \|f\|_{H^{-s}(\mathbb{S}^{d-1})} \|g\|_{H^{-s}(\mathbb{S}^{d-1})}.$$

Proof. This is a direct consequence of Welper (2024b, Lemma 7.16) and Welper (2024b, Lemma 6.7). For ϵ sufficiently small so that $s + \epsilon < 1$, the former shows that

$$\iint_{D \times D} f(x)k(x, y)g(y) dx dy \leq \|f\|_{H^{-s}(\mathbb{S}^{d-1})} \|g\|_{H^{-s}(\mathbb{S}^{d-1})} \|k\|_{C^{0;s+\epsilon, s+\epsilon}(\mathbb{S}^{d-1})},$$

for $k(x, y) = \sum_{r=1}^R \partial_r f_\theta(x) \partial_r f_\theta(y)$ and where $\|\cdot\|_{C^{0;s+\epsilon, s+\epsilon}(\mathbb{S}^{d-1})}$ is a Hölder norm of order $s + \epsilon$ in the two variables x and y . Technically, the reference does not include the case $s = 0$, which follows directly from a sup-norm bound and the fact that the domain is bounded. The second reference Welper (2024b, Lemma 6.7) shows that

$$\|k\|_{C^{0;s+\epsilon, s+\epsilon}} \lesssim 1.$$

where $k(x, y) = \sum_{r=1}^R \partial_r f_\theta(x) \partial_r f_\theta(y)$ is denoted by $\bar{\Gamma}$. Combining the two inequalities yields the result. \square

Lemma D.3 (Welper (2024b, Lemma 6.2)). *Assume that $\|x\| \lesssim 1$.*

1. *Assume that σ satisfies the growth condition (35) and may be different in each layer. Assume the weights are bounded (38). Then*

$$\|f^\ell(x)\| \lesssim m_0^{1/2} \prod_{k=0}^{\ell-1} \|W^k\| m_k^{-1/2}.$$

2. *Assume that σ satisfies the growth and Lipschitz conditions (35) and (36) and may be different in each layer. Assume the weights and perturbed weights are bounded (38). Then*

$$\|f^\ell(x) - \bar{f}^\ell(x)\| \lesssim m_0^{1/2} \sum_{k=0}^{\ell-1} \|W^k - \bar{W}^k\| m_k^{-1/2} \prod_{\substack{j=0 \\ j \neq k}}^{\ell-1} \max\{\|W^j\|, \|\bar{W}^j\|\} m_j^{-1/2}.$$

Lemma D.4 (Gentile & Welper (2022, Corollary 6.4)). *Let $\xi_i, i = 1, \dots, n$ be independent random variables, U, V Hilbert spaces and $X_i = X_i(\xi_i) = v_i(\xi_i)u_i(\xi_i)^* = v_i u_i^*$ be Bochner integrable rank one operators with $v_i \in V$ and $u_i^* \in U^*$. Assume there are $\mu > 0$ and $\nu > 0$ such that for all $i = 1, \dots, n$*

$$\|u^*\|_{U^*} \leq \mu, \quad \|v\|_V \leq \nu,$$

almost surely. Then, for any $t > 0$,

$$\Pr \left[\left\| \frac{1}{n} \sum_{i=1}^n X_i - \mathbb{E}[X_i] \right\| > \sqrt{\frac{8\mu^2\nu^2 t}{n}} + \frac{2\mu\nu t}{3n} \right] \leq 2t(e^t - t - 1)^{-1}.$$

Proof. The only difference to the reference is that we assume $\|u^*\|_{U^*} \leq \mu$ instead of $\|u\|_U \leq \mu$, which is equivalent due to the Riesz representation theorem. \square

Lemma D.5 (Welper (2024a, Lemma 3.3)). *Let $a, b, c, d > 0$ and $\rho > 1/2$. Let u_n and v_n be two sequences that satisfy*


$$\begin{aligned} u_{n+1} - u_n &\leq -\gamma a u_n^{1+\rho} v_n^{-\rho} + \gamma b u_n, \\ v_{n+1} - v_n &\leq -\gamma c u_n^\rho v_n^{1-\rho} + \gamma d \sqrt{u_n v_n}. \end{aligned} \quad (49)$$

Furthermore, assume that


$$u_k \geq \left(\frac{d}{c}\right)^{\frac{2}{2\rho-1}} v_0, \quad u_k \geq \left(2\frac{b}{a}\right)^{\frac{1}{\rho}} v_0, \quad \text{for all } k = 0, \dots, n-1. \quad (50)$$

Then

$$u_n \leq e^{-\gamma b n} u_0, \quad v_n \leq v_0.$$



(a) MSE Loss, dim=3, depth=2.



(b) Kernel Loss, dim=3, depth=2.

Figure 2: Test loss for training with mean squared loss (13) (left) and (17) (right). All axes are log-scaled so that the slope corresponds to convergence rates.

E Extra Numerical Experiments

This appendix contains some extra numerical results for the setup in Section 3.

- Section 3 does not report any losses. They are contained in the extended Table 2.
- This section also contains experiments for shallow networks in three dimensions, shown in Figure 2 and Table 2.
- Figure 3 contains results for the shallow network with MSE loss in 7 dimensions and with larger range for samples and width.

The observations from Section 3 remain unchanged. The deep networks performs slightly better than the shallow ones, but have significantly more weights in total.




Figure 3: Test loss for training with mean squared loss (13) for dimension 7 and depth 2. All axes are log-scaled so that the slope corresponds to convergence rates.

Dimension 3 and Depth 2 – Trained with MSE Loss												
m/N	Test Loss				dof rate				N rate			
	100	150	200	250	100	150	200	250	100	150	200	250
40	0.00043	0.00025	0.000238	0.000212	0.719	1.3	0.659	0.955	0.832	1.34	0.169	0.524
48	0.00033	0.000256	0.000203	0.000165	1.45	-0.119	0.883	1.36	1.08	0.631	0.805	0.917
56	0.000279	0.00019	0.000171	0.000163	1.1	1.92	1.1	0.0938	1.6	0.943	0.367	0.222
64	0.000252	0.000197	0.000153	0.000151	0.768	-0.273	0.859	0.58	1.22	0.601	0.893	0.0546

Dimension 3 and Depth 5 – Trained with MSE Loss												
m/N	Test Loss				dof rate				N rate			
	100	150	200	250	100	150	200	250	100	150	200	250
40	0.000268	0.000155	0.000125	0.000102	0.331	1.33	1.33	1.19	1.9	1.36	0.743	0.896
48	0.00023	0.000146	0.000111	8.44e-05	0.833	0.328	0.628	1.06	1.86	1.13	0.933	1.25
56	0.000194	0.000126	9.19e-05	7.89e-05	1.13	0.969	1.25	0.434	2.2	1.07	1.08	0.684
64	0.000222	9.55e-05	8.08e-05	6.49e-05	-1.03	2.05	0.967	1.47	1.56	2.08	0.582	0.985

Dimension 3 and Depth 2 – Trained with Kernel Loss												
m/N	Test Loss				dof rate				N rate			
	100	150	200	250	100	150	200	250	100	150	200	250
40	0.000317	0.000223	0.000208	0.00019	1.29	1.43	1.47	0.677	1.26	0.871	0.246	0.388
48	0.000275	0.000227	0.000187	0.000162	0.784	-0.11	0.563	0.898	1.2	0.469	0.672	0.662
56	0.000271	0.000199	0.000169	0.00015	0.103	0.858	0.658	0.474	1.16	0.756	0.565	0.536
64	0.000233	0.000173	0.000143	0.000136	1.12	1.05	1.26	0.756	1.49	0.732	0.664	0.232

Dimension 3 and Depth 5 – Trained with Kernel Loss												
m/N	Test Loss				dof rate				N rate			
	100	150	200	250	100	150	200	250	100	150	200	250
40	0.00019	0.00013	0.000103	9.29e-05	1.08	0.676	0.748	1.07	1.65	0.939	0.797	0.482
48	0.000183	0.000118	9.56e-05	7.67e-05	0.229	0.517	0.433	1.05	1.23	1.07	0.744	0.987
56	0.000141	9.01e-05	7.29e-05	6.47e-05	1.69	1.77	1.75	1.11	2.02	1.1	0.734	0.539
64	0.000151	9.27e-05	6.47e-05	5.52e-05	-0.54	-0.218	0.898	1.19	1.71	1.21	1.25	0.713

Table 2: Loss and estimated convergence rates between neighbouring losses for the given m/N . Left: Rate along the column, i.e. with respect to m . Right: Rate along rows, i.e. with respect to number of samples N . The first table is trained with mean squared loss (MSE) (13) and the second with kernel loss (17).