

Training Two-Layer ReLU Networks with Gradient Descent is Inconsistent

David Holzmüller

DAVID.HOLZMUELLER@MATHEMATIK.UNI-STUTTGART.DE

University of Stuttgart

Faculty of Mathematics and Physics

Institute for Stochastics and Applications

Ingo Steinwart

INGO.STEINWART@MATHEMATIK.UNI-STUTTGART.DE

University of Stuttgart

Faculty of Mathematics and Physics

Institute for Stochastics and Applications

Editor: Yoshua Bengio

Abstract

We prove that two-layer (Leaky)ReLU networks initialized by e.g. the widely used method proposed by He et al. (2015) and trained using gradient descent on a least-squares loss are not universally consistent. Specifically, we describe a large class of one-dimensional data-generating distributions for which, with high probability, gradient descent only finds a bad local minimum of the optimization landscape, since it is unable to move the biases far away from their initialization at zero. It turns out that in these cases, the found network essentially performs linear regression even if the target function is non-linear. We further provide numerical evidence that this happens in practical situations, for some multi-dimensional distributions and that stochastic gradient descent exhibits similar behavior. We also provide empirical results on how the choice of initialization and optimizer can influence this behavior.

Keywords: Neural networks, consistency, gradient descent, initialization, neural tangent kernel

1. Introduction

In recent years, neural networks (NNs) have achieved various success stories in areas such as image classification and natural language processing. For this reason, NNs are commonly viewed as one of the state-of-the-art machine learning algorithms. Unlike for other machine learning algorithms, however, our theoretical understanding of their learning behavior, e.g. in terms of a-priori learning guarantees, is still rather limited.

We consider the classical setting of statistical learning theory, where a random data set $D = ((x_1, y_1), \dots, (x_n, y_n))$ consists of n i.i.d. pairs (x_j, y_j) sampled from an unknown probability distribution P^{data} on $\mathbb{R}^d \times \mathbb{R}$. In our case, we consider the empirical, respectively

population least-squares risk of a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, that is, the quantities

$$R_D(f) := \frac{1}{2n} \sum_{j=1}^n (y_j - f(x_j))^2 \quad \text{and} \quad R_{P^{\text{data}}}(f) := \frac{1}{2} \mathbb{E}_{(x,y) \sim P^{\text{data}}}(y - f(x))^2 . \quad (1)$$

The Bayes risk is the lowest possible risk, which might be nonzero due to noise in the y component, that is

$$R_{P^{\text{data}}}^* := \inf_{f: \mathbb{R}^d \rightarrow \mathbb{R}} R_{P^{\text{data}}}(f) .$$

It is well known that this infimum is achieved by the conditional mean function, i.e., by $f_{P^{\text{data}}}^*(x) := \mathbb{E}_{P^{\text{data}}}(Y|X = x)$. Now, a learning method, i.e., a method that assigns a function f_D to each data set D , is called consistent for a distribution P^{data} , if its population risk converges in probability to the Bayes risk as the number n of samples goes to infinity. Or, to phrase it more formally, if for each $\varepsilon > 0$, the probability of observing a “bad” data set D with

$$R_{P^{\text{data}}}(f_D) - R_{P^{\text{data}}}^* \geq \varepsilon$$

converges to zero for $n \rightarrow \infty$. The quantity $R_{P^{\text{data}}}(f_D) - R_{P^{\text{data}}}^*$ is also called *excess risk*. Arguably the simplest a-priori guarantee for a learning method is the notion of universal consistency, which requires the learning method to be consistent for all bounded, or more generally, for all distributions P^{data} with $R_{P^{\text{data}}}(0) < \infty$. For simple learning methods such as histogram rules, kernel regression, and k -nearest neighbor rules, universal consistency has been long known, see for example the books by Devroye et al. (1996) and Györfi et al. (2002) for a variety of universally consistent learning methods. Similarly, kernel-based methods including support vector machines are universally consistent, see for example the book by Steinwart and Christmann (2008).

1.1 Contribution

We prove that training under-parameterized ReLU or Leaky ReLU networks with one hidden layer using gradient descent (GD) on a least-squares loss does not yield an universally consistent estimator if e.g. the common initialization method by He et al. (2015) is used. To this end, we specify a large family of one-dimensional data-generating distributions P^{data} , for which we prove that the probability of getting stuck in a bad local minimum converges to one as the width of the network and the number of data points simultaneously go to infinity in a controlled manner. We further show that these one-dimensional distributions, when embedded into higher-dimensional spaces, also provide examples of inconsistency for NNs with multi-dimensional input. Moreover, we prove that there also exist a multitude of data sets for which over-parameterized versions of such NNs cannot be properly trained by gradient descent.

We experimentally investigate the effect of deviating from our theoretically investigated assumptions by changing bias initialization, optimizer, network parameterization, and the dimensionality of the input distribution. We observe similar shortcomings of the zero bias

initialization used in He et al. (2015) even with stochastic gradient descent and moderate-dimensional input distributions. In addition, our experimental results give rise to practical recommendations on ways to mitigate the investigated problems of zero bias initialization.

This paper is an improved version of the first author’s master’s thesis (Holzmüller, 2019).

1.2 Related Work

Despite our result, there do exist some consistency results for certain classes of NNs, see e.g. (White, 1990) for regression, (Faragó and Lugosi, 1993) for classification, as well as Chapter 30 in the book by Devroye et al. (1996) and Chapter 16 in the book by Györfi et al. (2002). However, all these results consider a training algorithm that finds a global minimum, as well as an *under-parameterized regime*, in which the number m of hidden neurons grows more slowly than the sample size. This shows that consistency does not require m to grow fast, although it does require m to converge to infinity such that the neural networks are able to approximate any target function $f_{P_{\text{data}}}^*$ arbitrarily well in the limit. Unfortunately, finding a global minimum for small network sizes can be NP-hard, see e.g. the classical paper by Blum and Rivest (1989) as well as Boob et al. (2020), who establish similar results for certain ReLU-networks. Moreover, Safran and Shamir (2018) have empirically shown that the probability of finding a bad local minimum in certain two-layer ReLU NNs with gradient descent can increase with increasing numbers of samples and neurons. Finally, coming from another angle, Lee (2000) established a consistency result for Bayesian inference over NN parameters, which comes with a huge computational overhead. Consequently, it remains possible that the consistency results mentioned above only apply to computationally infeasible NN training algorithms. On the other hand, a consistent NN training algorithm does not necessarily need to find a global optimum, and hence it also remains an open question whether practical NN algorithms such as variants of (stochastic) GD are consistent in the under-parameterized regime when our assumptions are not satisfied. In this respect we finally note that the described theoretical gap resulting from the difficulty of solving the optimization problem is somewhat specific to neural networks: For example, beginning with Yao et al. (2007) optimizing kernel methods with gradient descent and early stopping has been well understood in the sense that e.g. both finite sample guarantees and learning rates, i.e., convergence rates for the excess risk, have been established.

Based on the empirical observations by Zhang et al. (2017), it has been more recently shown that finding a global minimum for *over-parameterized* NNs, i.e., for NNs whose number of neurons (significantly) exceeds the number of samples, is easier. For example, Arora et al. (2018) present a poly-time algorithm for finding a global minimum for NNs with one hidden layer and Mücke and Steinwart (2019) present a quadratic time training algorithm for NNs with two hidden layers. However, both algorithms are not based on (stochastic) GD. By imposing rather mild assumptions on the data set, Du et al. (2019a,b), Allen-Zhu et al. (2019), and Zou et al. (2018) show that (stochastic) GD also reaches a global optimum with high probability. Their analysis is based on the so-called neural tangent kernel (NTK) approach proposed by Jacot et al. (2018). Many subsequent works have established bounds requiring a smaller amount of overparameterization, see e.g. the papers by Zou

and Gu (2019); Ji and Telgarsky (2019); Song and Yang (2019); Ge et al. (2019); Chen et al. (2021). While these papers address the optimization problem of NNs, generalization guarantees such as consistency are either not discussed, or are only established under special assumptions on the data-generating distribution. In fact, Mücke and Steinwart (2019) show that over-parameterized NNs of sufficient size always have both global minima with good generalization performance in the sense of consistency and global minima with very bad generalization performance. However, their global minima are more interesting from a theoretical than a practical point of view. Finally, Zhang et al. (2017) and Belkin et al. (2018) discuss in detail why common techniques from statistical learning theory cannot work for the analysis of over-parameterized NNs.

There are some partial results regarding consistency of over-parameterized networks. Arora et al. (2019a) present a generalization bound for certain ReLU-NNs with one hidden layer that involves a novel complexity measure in terms of a Gram matrix of a neural tangent kernel. However, their results are only stated for a special class of *noise-free* data-generating distributions. Cao and Gu (2020) investigate a similar scenario for deeper networks. Belkin et al. (2018) have argued that generalization properties of unregularized kernel methods should be studied, which has become even more interesting since Arora et al. (2019b) have shown that regression with sufficiently wide neural networks essentially behaves like unregularized kernel regression with the NTK. For the Laplace kernel, Rakhlin and Zhai (2019) have shown inconsistency of unregularized kernel regression for a realistic class of data-generating distributions *with label noise*. On the contrary, Liang et al. (2020) show that for NTK-like kernels, the excess risk can converge to zero, if one considers a sequence of “uniformly easy” learning problems for which the dimension increases with the number of data points. Clearly, their setting is different from classical statistical learning questions such as universal consistency, and although label noise is permitted, very strong assumptions are imposed on the optimal regression function given by P^{data} . In summary, it can thus be hypothesized that very wide (over-parameterized) unregularized neural networks are not universally consistent, i.e., their excess risk does not converge to zero with increasing number of samples, although the excess risk may become very small if the input-dimension is high.

Similar to us, Williams et al. (2019) analyze two-layer ReLU networks with one-dimensional input. They define certain quantities δ_i that are based on the network weights and identify a “kernel” regime for $\delta_i \rightarrow -\infty$ and an “adaptive” regime for $\delta_i \rightarrow \infty$. They analyze the NN behavior in these idealized regimes and argue heuristically, why NNs that are close to these regimes should yield solutions similar to the idealized solutions. In our case, the NN behaves similar to their “kernel” regime with a finite-rank kernel, although a simple statistical analysis reveals that $\delta_i = -O(1)$ for most i and $\delta_i > 0$ for some i .

Finally, Steinwart (2019) observed experimentally that on one-dimensional data sets, two-layer ReLU NNs with the initialization method of He et al. (2015) often get stuck in bad local minima, and this was the initial motivation for our theoretical analysis. However, the existence of local minima in the loss landscape of neural networks has been investigated before in various settings with various results, see e.g. the papers by Sontag and Sussmann (1989); Gori and Tesi (1992); Soudry and Carmon (2016); Yun et al. (2019); Fukumizu and Amari (2000); Safran and Shamir (2018); He et al. (2019). Except for some experiments in

Safran and Shamir (2018), however, the probability of reaching a spurious local minimum is not investigated in these works. In this respect it is interesting to note that Yun et al. (2019) investigate a type of spurious local minimum that is similar to the ones that are found by NNs in our scenario. Moreover, He et al. (2019) prove the existence of similar local minima in a fairly general setting that includes deeper architectures and other loss functions.

The rest of this work is organized as follows: In Section 2, we define our considered NN architecture and training process. In Section 3, we then present some intuition behind our result. Our main theorem is then presented in Section 4, with applications to the over-parameterized case in Section 5 and the under-parameterized (inconsistent) case in Section 6. We outline the ideas behind the proof of our main theorem in Section 7 and discuss the relation to NTK-based analyses in Section 8. We support our theory with experimental evidence in Section 9, provide further experiments on ways to resolve the problems discussed in this paper in Section 10, and discuss further research questions in Section 11. Most proofs are deferred to the appendix.

2. Neural Network Architecture and Training

In this section we present the considered network architecture, the initialization, and the training process. For simplicity, we will mostly focus on one-dimensional inputs, but show in Remark 6 that our results for one-dimensional inputs can be easily generalized to multi-dimensional inputs.

Definition 1. Let $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ be the Leaky ReLU function with fixed parameter $\alpha \in \mathbb{R} \setminus \{\pm 1\}$, that is

$$\varphi(x) := \begin{cases} x & , x \geq 0 \\ \alpha x & , x \leq 0 . \end{cases}$$

We consider two-layer single-input single-output neural networks with $m \in \mathbb{N}$ hidden neurons. Such a neural network defines a function $f_W : \mathbb{R} \rightarrow \mathbb{R}$ via

$$f_W(x) := c + \sum_{i=1}^m w_i \varphi(a_i x + b_i) , \tag{2}$$

where $W = (\mathbf{a}, \mathbf{b}, c, \mathbf{w}) \in \mathbb{R}^{3m+1}$ with $\mathbf{a}, \mathbf{b}, \mathbf{w} \in \mathbb{R}^m$ and $c \in \mathbb{R}$. Note that φ is piecewise linear with a kink (non-differentiable point) at 0, and therefore the function f_W is piecewise affine linear with potential kinks at $-b_i/a_i$.

Assumption 2 (Initialization). The components of the initial parameter vector W_0 of (2) are initialized independently with distributions

$$b_{i,0} = 0, \quad c_0 = 0, \quad a_{i,0} \sim Z_a, \quad w_{i,0} \sim \frac{1}{\sqrt{m}} Z_w ,$$

where Z_a, Z_w are \mathbb{R} -valued random variables satisfying:

- (Q1) Z_a and Z_w have symmetric and bounded probability density functions $p_a, p_w : \mathbb{R} \rightarrow [0, B_Z^{\text{wa}}]$, where B_Z^{wa} is a suitable constant.
(Q2) $\mathbb{E}|Z_a|^p < \infty$ and $\mathbb{E}|Z_w|^p < \infty$ for all $p \in (0, \infty)$.

We denote the distribution of W_0 by P_m^{init} .

Importantly, Assumption 2 is satisfied, for example, by the initialization method of He et al. (2015), where $Z_a, Z_w \sim \mathcal{N}(0, 2)$. Our assumption is also satisfied for e.g. uniformly distributed Z_a and Z_w . Moreover, recall that the assumed zero bias initialization is the default in both `Tensorflow` and `Keras`. Finally, note that Z_a and Z_w do not necessarily need to have distributions that belong to the same family of distributions.

Assumption 3 (Data distribution I). *Let P^{data} be a distribution on $\mathbb{R} \times \mathbb{R}$ for which all moments are finite, that is*

$$\int |x|^p + |y|^p dP^{\text{data}}(x, y) < \infty$$

for all $p \in (0, \infty)$. Note that this is in particular satisfied if P^{data} is bounded, that is, if there exist bounded (measurable) subsets $X, Y \subset \mathbb{R}$ with $P^{\text{data}}(X \times Y) = 1$. Finally, P_X^{data} denotes the marginal distribution of P^{data} with respect to the x -component.

Definition 4. For a weight vector $W \in \mathbb{R}^{3m+1}$, a data set $D = ((x_1, y_1), \dots, (x_n, y_n)) \in (\mathbb{R} \times \mathbb{R})^n$, and a bounded distribution P on $\mathbb{R} \times \mathbb{R}$, we define the empirical, respectively population (least-squares) loss of the function $f_W : \mathbb{R} \rightarrow \mathbb{R}$ in (2) by

$$L_D(W) := R_D(f_W) \quad \text{and} \quad L_P(W) := R_P(f_W),$$

where the risks $R_D(f_W)$ and $R_P(f_W)$ have already been introduced in (1).

Assumption 5 (Training). For $m, n \in \mathbb{N}$, step size $h > 0$, a distribution P_m^{init} as in Assumption 2, and a data set $D \in (\mathbb{R} \times \mathbb{R})^n$, training the neural network (2) produces a random sequence of functions $(f_{W_k})_{k \in \mathbb{N}_0}$ via initializing $W_0 \sim P_m^{\text{init}}$ and applying GD, i.e., for all $k \geq 0$ we have

$$W_{k+1} := W_k - h \nabla L_D(W_k).$$

Note that for a fixed data set D the randomness of the training sequence $(f_{W_k})_{k \in \mathbb{N}_0}$ only comes from the initialization. When we investigate inconsistency, however, we also need to view the data set D as a random variable.

Remark 6 (Multi-dimensional input). We show in Section H that NNs with multi-dimensional input behave similarly in the following sense: For a fixed $\mathbf{z} \in \mathbb{R}^d$ with $\|\mathbf{z}\|_2 = 1$ and D as above, consider the data set $\tilde{D} := ((x_1 \mathbf{z}, y_1), \dots, (x_n \mathbf{z}, y_n)) \in (\mathbb{R}^d \times \mathbb{R})^n$, that is, we embed D into \mathbb{R}^d by mapping all x_j to a point on the line $\text{Span}\{\mathbf{z}\}$ without changing the mutual distances between the samples. Similarly, a distribution P^{data} on $\mathbb{R} \times \mathbb{R}$ can be mapped to $\mathbb{R}^d \times \mathbb{R}$ by considering the image measure \tilde{P}^{data} of P^{data} under the mapping $(x, y) \mapsto (x\mathbf{z}, y)$. Note that the marginal distribution $\tilde{P}_X^{\text{data}}$ of such an image measure lives on a one-dimensional

(linear) manifold in \mathbb{R}^d , and if its conditional mean function is smooth, conventional wisdom thus suggests that it is particularly easy to learn from such a probability measure \tilde{P}^{data} . In this respect note that the distributions P^{data} on which inconsistency occurs, see Assumption 16 for a detailed description, include distributions with arbitrarily smooth conditional mean function and by the transformation described above, the same applies to \tilde{P}^{data} .

Now consider a neural network with multivariate input $\mathbf{x} \in \mathbb{R}^d$ defined by

$$f_{\tilde{W}}(\mathbf{x}) = \tilde{c} + \sum_{i=1}^m \tilde{w}_i \varphi \left(\tilde{b}_i + \sum_{j=1}^d \tilde{a}_{ij} x_j \right),$$

where $\tilde{W} = (\tilde{\mathbf{a}}, \tilde{\mathbf{b}}, \tilde{c}, \tilde{\mathbf{w}}) \in \mathbb{R}^{(d+2)n+1}$ are the parameters of the NN. Again, we assume that the NN is trained using gradient on the same loss function and initialized via

$$\tilde{b}_{i,0} = 0, \quad \tilde{c}_0 = 0, \quad \tilde{a}_{i,j,0} \sim Z_a, \quad \tilde{w}_{i,0} \sim \frac{1}{\sqrt{m}} Z_w,$$

where Z_a and Z_w satisfy the conditions in Assumption 2. Again, this is satisfied for example by the initialization method of He et al. (2015).

We show in Section H that there exists a random $W_0 \in \mathbb{R}^{3n+1}$ satisfying Assumption 2 such that the GD iterates

$$\begin{aligned} \tilde{W}_{k+1} &= \tilde{W}_k - h \nabla L_{\tilde{D}}(\tilde{W}_k), \\ W_{k+1} &= W_k - h \nabla L_D(W_k) \end{aligned}$$

satisfy

$$f_{\tilde{W}_k}(x\mathbf{z}) = f_{W_k}(x).$$

for all $x \in \mathbb{R}, k \in \mathbb{N}_0$. In other words, if we have a data set where all x_j lie on a line $\text{Span}\{\mathbf{z}\} \subseteq \mathbb{R}^d$, then an NN on the d -dimensional input space behaves on this line like an NN with one-dimensional input space.

3. A Closer Look at the Training Behavior

In this section we illustrate why the combination of zero bias initialization and GD training may produce poorly predicting NNs. In the subsequent sections we then rigorously prove that the illustrated training behavior does occur with high probability.

Our first, rather basic observation is that zero bias initialization $b_{i,0} = 0$ and $c_0 = 0$ as in Assumption 2 places all kinks $-b_{i,0}/a_{i,0}$ of the initial f_{W_0} at zero. Consequently, f_{W_0} is linear on both $(-\infty, 0]$ and $[0, \infty)$. In contrast, the function to be learned is typically nonlinear on these two sets, and finding a suitable NN approximation during training may thus require to substantially move at least some of the kinks. To illustrate this statement, consider Figure 1, in which a data set that requires a nonlinear predictor is depicted. It is obvious from Figure 1 that any reasonable NN approximation requires at least a few kinks in both $[-1.5, -0.5]$ and $[0.5, 1.5]$. The training algorithm thus needs to move a few kinks from




Figure 1: A data set with $n = 300$ samples (black crosses), a close-to-optimal NN predictor (green line, kinks marked as circles) with $m = 16$, and functions f_{W_k} for $k = \lceil 1.1^l \rceil - 1, l \in \{0, \dots, 120\}, m = 16, h = 0.002$, trained on the visualized data set. The colors of f_{W_k} transition from blue to red to yellow, i.e., the lower blue function is f_{W_0} and the upper red line is reached at an intermediate stage before the NN converges to the prominent yellow line.

0 into these two areas. Unfortunately, such a behavior can in general not be guaranteed. For example, Figure 2 illustrates that on the data set D of Figure 1, GD does not move the kinks outside the interval $[-0.2, 0.2]$ and in particular, no kink is moved across a sample, since no x -component of a sample of D falls inside $[-0.2, 0.2]$. As a consequence, the corresponding NN predictors f_{W_k} , for $k \geq 1$, remain affine linear on the left part D_{-1} and the right part D_1 of the data set D , as Figure 1 illustrates. This problem has been observed experimentally by Steinwart (2019), who suggests to use a data-dependent initialization method that places kinks randomly in between the data points.

A closer look at Figure 1 further indicates that f_{W_k} approaches the optimal affine linear regression lines for the two data sets D_{-1} and D_1 with increasing k . As a consequence, GD gets stuck in a bad local minimum of the loss surface. Note that the existence of structurally similar bad local minima has already been shown in Yun et al. (2019), but there it remained an open question whether GD can avoid such bad minima. In the following sections, we rigorously show that under some assumptions on D or P , on the step size h , and on n and m , the predictors f_{W_k} produced by GD remain affine linear on the negative and positive parts of D with high probability. Consequently, GD does not escape from the corresponding bad local minima in such situations.

Figure 2 also shows that the kinks initially move very fast but then slow down. Empirically, this slowdown is related to the loss, whose evolution relative to the reached optimum in our example is also shown in Figure 2. In our theoretical analysis, we find that the convergence of the neural network and the movement of the kinks are related to a four-dimensional linear iteration equation. Here, the system matrix has two eigenvalues of order $-\Theta(m)$ leading to a fast convergence and two eigenvalues of order $-\Theta(1)$ leading to a slow convergence. In




Figure 2: This figure shows the evolution of several quantities during training in Figure 1, where the two plots show a different number of epochs. The thick red line shows $L_D(W_k) - \inf_{k' \in \mathbb{N}_0} L_D(W_{k'})$, where the latter infimum is the minimal loss that can be reached by fitting D_1 and D_{-1} with linear regression similar to the yellow line in Figure 1. The thin black lines show the 16 NN kinks in Figure 1 (one can see only 14 kinks since there are two almost identical pairs of paths). The right y -axis in this plot corresponds to the x -positions of the kinks in Figure 1.

the example of this section, the initialization of this system is close to the fast eigenvectors, which leads to the two-step decay of the loss in Figure 2.

Our theory suggests the following intuitive explanation of this two-step decay: Since the weights in the second layer have a lower standard deviation ($O(1/\sqrt{m})$) than the weights in the first layer, updates to second-layer weights have a stronger effect and the gradient with respect to second-layer weights is larger. Since the first-layer biases are initially zero, the faster learning of second-layer weights can only adjust the slopes of the affine network parts. Adjusting the first-layer biases and the single second-layer bias happens at a substantially slower speed. In cases where almost no bias adjustment is needed for learning the optimal affine regression lines, learning in the second layer is so fast that kinks do not succeed in moving far. The relative training speeds of the layers can be adjusted, for example, using the NTK parameterization (Jacot et al., 2018), which we further discuss in Section 10.

A necessary condition for consistency is that the NN is able to approximate the target function $f_{P^{\text{data}}}^*$ arbitrarily well. In general, this requires the width $m = m_n$ of the NN, which can depend on the number of samples n , to converge to infinity:

$$\lim_{n \rightarrow \infty} m_n = \infty .$$

Note that wide networks are more likely to have neurons with very small $|a_i|$, for which the kinks $-b_i/a_i$ can move fast. Because of this, we find that in order to prove the failure case described in this section, the intercept of the optimal affine regression lines needs to be closer to zero the more neurons are used. When considering the setting of consistency, where the variance of this intercept decreases with the size of the randomly sampled data set, our negative results only apply to the under-parameterized regime $m_n < o(n)$. As explained in Section 1.2, this is also the regime where universal consistency results are known for NNs with perfect optimization. For suitable fixed data sets, our negative results also extend to the over-parameterized regime $m_n > n$.

The insights above suggest that making the first layer learn faster via a different initialization / parameterization and initializing the kinks randomly should improve the training behavior of neural networks. Indeed, Du et al. (2019b) consider such a setting and show convergence to a global optimum with high probability. However, their result is not sufficient for universal consistency since only the training error is analyzed. In Section 10, we provide further experimental results on the benefits of such modifications.

4. Main Result

Before we can state the main result of our work, we need to introduce some more notions.

Definition 7. For given data set $D = ((x_1, y_1), \dots, (x_n, y_n)) \in (\mathbb{R} \times \mathbb{R})^n$, the subsequence of samples (x_j, y_j) with $x_j > 0$ is denoted by D_1 . Analogously, D_{-1} denotes the samples with $x_j < 0$. Moreover, for a distribution P^{data} , we define the measures P_1 and P_{-1} via

$$\begin{aligned} P_1(E) &:= P^{\text{data}}(E \cap ((0, \infty) \times \mathbb{R})) \\ P_{-1}(E) &:= P^{\text{data}}(E \cap ((-\infty, 0) \times \mathbb{R})) . \end{aligned}$$

Throughout this work, we require D_1 and D_{-1} to be non-empty and P_1 and P_{-1} to be non-trivial. However, it is also possible to prove analogous results for the cases where D or P^{data} are solely concentrated on $(-\infty, 0)$ or $(0, \infty)$, cf. Remark D.4.

With the help of $D_{\pm 1}$ and $P_{\pm 1}$ we can now define the “half-sided linear regression optima”, which will be crucial for our analysis.

Definition 8. For $P := P^{\text{data}}$ and D as in Definition 4, $x, y \in \mathbb{R}$, and $\sigma \in \{\pm 1\}$, we define

$$\begin{aligned} \mathbf{M}_x &:= \begin{pmatrix} x^2 & x \\ x & 1 \end{pmatrix}, & \hat{\mathbf{u}}_{(x,y)}^0 &:= \begin{pmatrix} xy \\ y \end{pmatrix}, \\ \mathbf{M}_{D,\sigma} &:= \frac{1}{n} \sum_{(x,y) \in D_\sigma} \mathbf{M}_x, & \hat{\mathbf{u}}_{D,\sigma}^0 &:= \frac{1}{n} \sum_{(x,y) \in D_\sigma} \hat{\mathbf{u}}_{(x,y)}^0, \end{aligned}$$




Figure 3: Optimal affine regression lines for D_1 (orange, upper line) and D_{-1} (blue, lower line) for an example data set D . The data points are shown in black. The slope and intercept of the optimal affine regression line for D_σ are given by $p_{D,\sigma}^{\text{opt}}$ and $q_{D,\sigma}^{\text{opt}}$, respectively.

$$\mathbf{M}_{P,\sigma} := \int \mathbf{M}_x \, dP_\sigma(x, y), \quad \hat{\mathbf{u}}_{P,\sigma}^0 := \int \hat{\mathbf{u}}_{(x,y)}^0 \, dP_\sigma(x, y).$$

Moreover, if the 2×2 matrix $\mathbf{M}_{D,\sigma}$ is invertible, see Remark 9 for a simple characterization of this situation, we write

$$\mathbf{v}_{D,\sigma}^{\text{opt}} := \begin{pmatrix} p_{D,\sigma}^{\text{opt}} \\ q_{D,\sigma}^{\text{opt}} \end{pmatrix} := \mathbf{M}_{D,\sigma}^{-1} \hat{\mathbf{u}}_{D,\sigma}^0.$$

Note that $p_{D,\sigma}^{\text{opt}}$ and $q_{D,\sigma}^{\text{opt}}$ are the slope and intercept of the optimal linear regression line for D_σ , see Remark D.6 for a proof and Figure 3 for an illustration.

In the sequel, we are mostly interested in the maximal absolute slopes and intercepts as well as the distance of D to 0, i.e.,

$$\begin{aligned} \psi_{D,p} &:= \max \left\{ \left| p_{D,1}^{\text{opt}} \right|, \left| p_{D,-1}^{\text{opt}} \right| \right\}, \\ \psi_{D,q} &:= \max \left\{ \left| q_{D,1}^{\text{opt}} \right|, \left| q_{D,-1}^{\text{opt}} \right| \right\}, \\ \underline{x}_D &:= \min \{ |x_1|, \dots, |x_n| \}. \end{aligned}$$

If $\mathbf{M}_{P,\sigma}$ is invertible, the quantities $\mathbf{v}_{P,\sigma}^{\text{opt}}$, $q_{P,\sigma}^{\text{opt}}$, $p_{P,\sigma}^{\text{opt}}$, $\psi_{P,p}$ and $\psi_{P,q}$ are defined and interpreted analogously. Finally, the smallest and the largest eigenvalue of a symmetric matrix A are denoted by $\lambda_{\min}(A)$ and $\lambda_{\max}(A)$. Since the matrix $\mathbf{M}_{D,\sigma}$ is positive semi-definite, we note that it is invertible if and only if $\lambda_{\min}(\mathbf{M}_{D,\sigma}) > 0$.

The following remark, whose proof can be found in Remark D.5, shows that the matrices $\mathbf{M}_{D,\pm 1}$ are invertible for all interesting data sets D . Since the eigenvalues of $\mathbf{M}_{D,\pm 1}$ will play

an important role in our main result, the following remark also provides simple estimates for the eigenvalues of $\mathbf{M}_{D,\pm 1}$.

Remark 9. *Given a $\sigma \in \{\pm 1\}$, the matrix $\mathbf{M}_{D,\sigma}$ is invertible, if and only if the data set D_σ contains at least two samples (x_i, y_i) and (x_j, y_j) with $x_i \neq x_j$. Moreover, if D_σ contains n_σ samples and $D_{X,\sigma}$ are the x components of the samples in D_σ , we have*

$$\begin{aligned}\lambda_{\min}(\mathbf{M}_{D,\sigma}) &\geq \frac{n_\sigma}{n} \cdot \frac{\text{Var } D_{X,\sigma}}{\text{Var } D_{X,\sigma} + (\mathbb{E}D_{X,\sigma})^2 + 1} , \\ \lambda_{\max}(\mathbf{M}_{D,\sigma}) &\leq \frac{n_\sigma}{n} (\text{Var } D_{X,\sigma} + (\mathbb{E}D_{X,\sigma})^2 + 1) ,\end{aligned}$$

and these bounds are off by a factor of at most 2. Therefore, as long as $n_\sigma/n = \Theta(1)$, $\mathbb{E}D_{X,\sigma} = \Theta(1)$ and $\text{Var } D_{X,\sigma} = \Theta(1)$, the eigenvalues of $\mathbf{M}_{D,\sigma}$ are also $\Theta(1)$. Especially, Proposition 21 shows that this holds with high probability if D is sampled from a suitable distribution P^{data} . For example, it holds for the distribution used in Figure 1.

We can now state our main theorem, which is proven at the end of Section G. In Corollary 12 and Theorem 22 we then apply this main theorem to over-parameterized and under-parameterized NNs, respectively. Note that in the formulation of Theorem 10, the constants $C_P, C_{\text{lr}}, C_{\text{weights}}$ can depend on all previous constants and especially on ε , which prohibits removing the term m^ε by taking the limit $\varepsilon \rightarrow 0$.

Theorem 10. *Let $\varepsilon > 0$ and let $\gamma_\psi, \gamma_{\text{data}}, \gamma_P \geq 0$ with $\gamma_\psi + \gamma_{\text{data}} + \gamma_P < 1/2$. Then, for all given constants $K_{\text{data}}, K_\psi, K_M > 0$, there exist constants $C_P, C_{\text{lr}}, C_{\text{weights}} > 0$ such that the following statement holds:*

For all neural networks (2) with width $m \in \mathbb{N}$ that are initialized and trained according to Assumption 5 with step size h satisfying

$$0 < h \leq C_{\text{lr}} m^{-1} \tag{3}$$

and all data sets D satisfying:

- (a) $\underline{x}_D \geq K_{\text{data}} m^{-\gamma_{\text{data}}}$
- (b) $K_M^{-1} \leq \lambda_{\min}(\mathbf{M}_{D,\pm 1}) \leq \lambda_{\max}(\mathbf{M}_{D,\pm 1}) \leq K_M$
- (c) $\psi_{D,p} \leq K_\psi$
- (d) $\psi_{D,q} \leq K_\psi m^{\gamma_\psi - 1}$

the random training sequence $(f_{W_k})_{k \in \mathbb{N}_0}$ has the following properties with probability not less than $1 - C_P m^{-\gamma_P}$:

- (i) *For all $i = 1, \dots, m$ and $k \geq 0$ it holds $|a_{i,k} - a_{i,0}| \leq C_{\text{weights}} m^{\gamma_\psi + \varepsilon - 3/2}$.*
- (ii) *For all $i = 1, \dots, m$ and $k \geq 0$ it holds $|b_{i,k} - b_{i,0}| \leq C_{\text{weights}} m^{\gamma_\psi + \varepsilon - 3/2}$.*
- (iii) *For all $i = 1, \dots, m$ and $k \geq 0$ it holds $|w_{i,k} - w_{i,0}| \leq C_{\text{weights}} m^{\gamma_\psi + \varepsilon - 1}$.*
- (iv) *For all $k \geq 0$ it holds $|c_k - c_0| \leq C_{\text{weights}} m^{\gamma_\psi + \varepsilon - 1}$.*
- (v) *For all $k \geq 0$, the neural network function f_{W_k} is affine linear on the intervals $(-\infty, -K_{\text{data}} m^{-\gamma_{\text{data}}}]$ and $[K_{\text{data}} m^{-\gamma_{\text{data}}}, \infty)$.*

To fully appreciate Theorem 10 a couple remarks are necessary: Conditions (a) and (b) are conditions on the x -parts $D_{X,\sigma}$ of the data sets $D_{\pm 1}$, whereas (c) and (d) impose conditions

on the optimal half-sided affine linear regression estimators. In a nutshell, (c) only requires bounded slopes of these linear estimates, whereas (d) demands the intercepts to become closer to 0 for increasing network sizes. Moreover, note that the exponents $\gamma_\psi, \gamma_{\text{data}}, \gamma_P$ describe a balance between data set characteristics (a) and (d) on the one-hand and the accuracy of the guarantees on the other hand. For example, in the extreme case, in which we are only interested in data sets that are bounded away from the origin and whose intercepts vanish, we may choose $\gamma_{\text{data}} = \gamma_\psi = 0$ and $\gamma_P = 1/2 - \varepsilon$ for all sufficiently small $\varepsilon > 0$. Ignoring the nuisance term ε , the statements (i) – (v) then hold with probability $\geq 1 - C_P m^{-1/2}$, the guarantees (i) and (ii) for the parameters of the hidden layer are of the order $O(m^{-3/2})$, and guarantees (iii) and (iv) for the parameters of the output neuron are of the order $O(m^{-1})$. In particular, the parameters “typically” move less during training the larger the network size is, and in fact, the probability of this event also increases with the network size. Moreover, the guarantees for the hidden layer are stronger in m suggesting that the parameters in the hidden layer move less during training than those of the output neuron. The guarantee (v) simply confirms the illustrative example in Figure 1.

In a nutshell, Theorem 10 thus states that if D is inside certain bounds and the step size h is sufficiently small, the NN function f_{W_k} remains affine on D_1 and D_{-1} with high probability and the weights W_k do not change much. This behavior is independent of when the iteration is stopped.

Remark 11. *The data-independent condition (3) involves a conservative constant C_{lr} . In Proposition F.4 and the proof of Theorem 10, we show that (3) can be replaced by $h \leq \lambda_{\max}(\mathbf{H})^{-1}$, where $\mathbf{H} \in \mathbb{R}^{4 \times 4}$ is a symmetric positive (semi-)definite matrix, which is defined in Definition F.1 and which can be computed from D and W_0 .*

5. Over-parameterization

In this case section we investigate the consequences of Theorem 10 for neural networks of arbitrary width m . We begin with the following corollary that describes these consequences for a fixed data set D of size $n \geq 4$, where $n \geq 4$ is necessary for $\text{Var } D_{X,\pm 1} > 0$ as defined in Remark 9. Note that in particular it applies to the over-parameterized case $m > n$.

Corollary 12. *Let D be fixed data set such that $\underline{x}_D > 0$, $\text{Var } D_{X,\pm 1} > 0$ and $\psi_{D,q} = 0$. Then for all $\varepsilon > 0$ and $0 < \gamma_P < 1/2$, there exist constants $C_P, C_{\text{lr}}, C_{\text{weights}} > 0$ such that for all widths $m \geq 1$ and all step sizes*

$$0 < h \leq C_{\text{lr}} m^{-1} ,$$

the random sequence $(f_{W_k})_{k \in \mathbb{N}_0}$ obtained by initializing and training according to Assumption 5 satisfies (i) – (v) in Theorem 10 for $K_{\text{data}} = \underline{x}_D$ and $\gamma_{\text{data}} = \gamma_\psi = 0$ with probability not less than $1 - C_P m^{-\gamma_P}$.

Proof By Remark 9 we know that there is a constant $K_M > 0$ such that (b) of Theorem 10 is satisfied. By choosing $K_\psi := \psi_{D,p} + 1$, where 1 is added to ensure $K_\psi > 0$, we then see that D also satisfies the remaining assumptions of Theorem 10. ■

Remark 13. In Corollary 12, it is also possible to choose $\gamma_{\text{data}} > 0$, since then conclusion (v) of Theorem 10 tells us that the kinks move less with increasing m . The price for this is that γ_P , and therefore the probability of the kinks moving less, decreases accordingly.

Remark 14. Lemma J.1 shows that for one-dimensional data sets D with $x_j \neq 0$ for all j , the assumptions on the data set in Corollary 12 can always be satisfied by suitably adding three more points to D . By Corollary 12, these additional points can play the role of adversarial training samples in the sense that they provably hinder the neural network training algorithm to converge to a good predictor if the best possible predictor is not half-sided affine linear, see again Figure 1. Moreover, by Corollary 12 the success probability of such an attack increases with increasing network size m .

Let us finally present a particularly simple data set to which Corollary 12 applies and for which the best possible predictor is not half-sided affine linear.

Example 1. *The data set*

$$D := ((-3, -1), (-2, 2), (-1, -1), (1, 1), (2, -2), (3, 1))$$

satisfies $\underline{x}_D = 1$ and $\text{Var } D_{X,\pm 1} > 0$. Moreover, $\hat{\mathbf{u}}_{D,\pm 1}^0 = \mathbf{0}$ implies $\psi_{D,p} = \psi_{D,q} = 0$. In particular, the best possible empirical risk a half-sided affine linear predictor can achieve is 1. Now assume that $m \geq 6$. Then one can easily construct a network f_W of the form (2) with $R_D(f_W) = 0$. On the other hand, Corollary 12 shows that with high probability f_{W_k} is affine linear on $(-\infty, -1]$ and $[1, \infty)$ for all $k \geq 1$, and hence we have $R_D(f_{W_k}) \geq 1$. Consequently, training with gradient descent does not come even close to a global optimum.

6. Inconsistency

In this section we present the inconsistency result for initializing and training a neural network according to Assumption 5. To this end, we describe a class of distributions that produce data sets satisfying the assumptions of Theorem 10 with high probability, and for which the conditional mean function is not half-sided affine linear.

To describe the failure of gradient descent in this case, we need to recall the classical notion of consistency from statistical learning theory, see e.g. Steinwart and Christmann (2008, Chapter 6). Since the predictors obtained by our training scheme described in Assumption 5 are probabilistic in both the data set D and the initialization W_0 , the following definition of consistency has been adapted in this respect.

Definition 15. *A learning method, i.e., a method that produces for each data set D a potentially random predictor f_D , is called consistent for a bounded distribution P^{data} , if for all $\varepsilon > 0$, the probability of sampling a data set D with n i.i.d. samples $(x_j, y_j) \sim P^{\text{data}}$ satisfying*

$$R_{P^{\text{data}}}(f_D) \geq R_{P^{\text{data}}}^* + \varepsilon$$

converges to 0 as $n \rightarrow \infty$, where $R_{P^{\text{data}}}^$ is the optimal risk, i.e., $R_{P^{\text{data}}}^* := \inf_{f: \mathbb{R} \rightarrow \mathbb{R}} R_{P^{\text{data}}}(f)$. The learning method is universally consistent if this holds for all such P^{data} .*

Roughly speaking, universally consistent learning methods are guaranteed to produce close-to-optimal predictors for $n \rightarrow \infty$, *independently* of the data generating distribution P^{data} . Universal consistency is therefore widely accepted as a minimal requirement for statistically sound learning methods, see e.g. the books by Devroye et al. (1996) and Györfi et al. (2002).

Next we will show that neural networks initialized and trained as in Assumption 5 are not universally consistent in a strong sense. Namely, we show that inconsistency occurs for all distributions satisfying the following assumption. For its formulation we denote the set of all $f : \mathbb{R} \rightarrow \mathbb{R}$ that are affine linear on both $(-\infty, 0)$ and $(0, \infty)$ by $\mathcal{F}_{\text{hsal}}$.

Assumption 16 (Data distribution II). *The distribution P^{data} satisfies Assumption 3 and the following four conditions:*

(P1) *For all $\sigma \in \{\pm 1\}$, the matrix $\mathbf{M}_{P^{\text{data}}, \sigma}$ is invertible.*

(P2) *There exists an $\eta \in (4, \infty]$ such that we have*

$$\begin{aligned} P_X^{\text{data}}([-x, x]) &= O(x^\eta) \quad \text{for } x \searrow 0, & (\text{if } \eta < \infty) \\ P_X^{\text{data}}((-\delta, \delta)) &= 0 \quad \text{for some } \delta > 0. & (\text{if } \eta = \infty) \end{aligned}$$

(P3) *The intercepts of Definition 8 satisfy $\psi_{P^{\text{data}}, q} = 0$.*

(P4) *We have $\inf_{f \in \mathcal{F}_{\text{hsal}}} R_{P^{\text{data}}}(f) > R_{P^{\text{data}}}^*$.*

The next remarks show, for example, that P^{data} satisfies (P1) and (P2) if the distribution P_X^{data} of the x component has a density that is sufficiently small around 0. They further show that we can always enforce (P3) by suitably modifying P^{data} and that (P4) means that the target function to be learned is not half-sided affine linear in the sense of $\mathcal{F}_{\text{hsal}}$. As an example, we note that the data set in Figure 1 has been sampled from a distribution satisfying Assumption 16.

Remark 17. *Assumption (P1) is satisfied if $P_\sigma^{\text{data}}((\mathbb{R} \setminus \{x\}) \times \mathbb{R}) > 0$ for all $x \in \mathbb{R}$. Indeed, the kernel of the matrix \mathbf{M}_x is $\text{Span}\{(1, -x)^\top\}$, and hence there is, for any vector $0 \neq \mathbf{v} \in \mathbb{R}^2$, at most one $x \in \mathbb{R}$ with $\mathbf{v}^\top \mathbf{M}_x \mathbf{v} = 0$. This gives*

$$\mathbf{v}^\top \mathbf{M}_{P^{\text{data}}, \sigma} \mathbf{v} = \int \mathbf{v}^\top \mathbf{M}_x \mathbf{v} dP_\sigma^{\text{data}}(x, y) > 0.$$

In particular, (P1) is satisfied if, for example, P_X^{data} has a density that does not completely vanish on $(-\infty, 0)$ or $(0, \infty)$.

Remark 18. *Assumption (P2) holds, for example, if P_X^{data} has a density p with*

$$\begin{aligned} p(x) &= O(|x|^{\eta-1}) \quad \text{for } x \rightarrow 0, & (\text{if } \eta < \infty) \\ p(x) &= 0 \quad \text{for } x \in (-\delta, \delta) \text{ for some } \delta > 0. & (\text{if } \eta = \infty) \end{aligned}$$

Verifying this claim is a straight-forward exercise.

Remark 19. *Let P^{data} be a distribution satisfying Assumption 3, (P1), (P2), and (P4), and let us fix a pair of random variables $(X, Y) \sim P^{\text{data}}$. Then the distribution \tilde{P}^{data} of the vertically shifted random variables*

$$(X, Y - q_{P^{\text{data}}, \text{sgn}(X)}^{\text{opt}})$$




Figure 4: Support of the distribution P^{data} from Example 2 for $\lambda = 1/4$. More specifically, if $(x, y) \sim P^{\text{data}}$, then (x, y) is uniformly distributed on the blue intervals. Note that this distribution is symmetric for simplicity, but symmetry is not necessary for the conclusions of Example 2.

satisfies Assumption 16. This assertion immediately follows from the meaning of the intercepts $p_{P^{\text{data}}, -1}^{\text{opt}}$ and $p_{P^{\text{data}}, 1}^{\text{opt}}$ and the definition of the maximal absolute intercept $\psi_{P^{\text{data}}, q}$.

Remark 20. Recall from e.g. Steinwart and Christmann (2008, Example 2.6) that the risk $R_{P^{\text{data}}}$ is minimized by the conditional expectation $f^*(x) := \mathbb{E}_{P^{\text{data}}}(Y|x)$ and that the excess risk of a predictor $f : \mathbb{R} \rightarrow \mathbb{R}$ is

$$R_{P^{\text{data}}}(f) - R_{P^{\text{data}}}^* = \frac{1}{2} \int |f(x) - f^*(x)|^2 dP_X^{\text{data}}(x).$$

Assumption (P4) thus states that the least squares target function f^* cannot be approximated by half-sided affine linear functions in the sense of $\mathcal{F}_{\text{hsal}}$.

Assumption 16 can be satisfied even by very simple distributions:

Example 2. Let $\lambda > 0$. Consider the distribution P^{data} shown in Figure 4, which is given by $P^{\text{data}}(y = 1) = P^{\text{data}}(y = -1) = \frac{1}{2}$ and

$$P^{\text{data}}(x|y = \sigma) = \mathcal{U}([\sigma\lambda, 2\sigma\lambda] \cup [-3\sigma\lambda, -5\sigma\lambda] \cup [6\sigma\lambda, 7\sigma\lambda]),$$

for $\sigma \in \{\pm 1\}$. In other words, given $y = \sigma$, y is uniformly distributed across four intervals. The intervals are also visualized in Figure 4. By construction, we have $\text{Var}(y) = 1$ and by a suitable choice of λ , we can even achieve that $\text{Var}(x) = 1$. The distribution P^{data} can be interpreted as a simple binary classification problem with classes ± 1 without label noise where the classes are separated by a positive margin. For many classical learning algorithms, the classification error on such a problem behaves asymptotically like $O(n^{-1})$, or even like $O(0)$. For example, a simple calculation shows that histogram rules with fixed width smaller than λ achieve zero classification error with high probability even for moderately sized data sets.

It is not hard to show that P^{data} satisfies Assumption 16 with $\eta = \infty$. For example, we can compute $\hat{\mathbf{u}}_{P^{\text{data}}, \pm 1}^0 = \mathbf{0}$, which implies $\psi_{P^{\text{data}}, p} = \psi_{P^{\text{data}}, q} = 0$. When employing the quadratic loss function used in this paper for this binary classification problem, Theorem 22 will show that the training setup in Section A with high probability produces a NN with affine

linear predictions on the left and right halves P_{-1}^{data} and P_{+1}^{data} of P^{data} . However, such affine linear predictions must have a classification error of at least 25% since at least one of the three intervals on each side of the y axis must be entirely misclassified. As mentioned in Definition 7, using a modification of the arguments in this paper, it is also possible to show the same result if only the part of P^{data} with $x > 0$ or only the part with $x < 0$ is considered.

The next proposition, whose proof is delegated to Proposition E.5, helps us to show that the assumptions of Theorem 10 are satisfied with high probability for data sets D sampled from any distribution P^{data} satisfying Assumption 16. In its formulation we use the convention $\infty \cdot 0 := \infty$.

Proposition 21. *Let P^{data} satisfy Assumption 3 and (P1) – (P3) from Assumption 16 and let $\varepsilon > 0, K_{\text{data}} > 0, m \geq 1$, and $\gamma_{\text{data}}, \gamma' \geq 0$. In the case $\eta = \infty$, we further assume that K_{data} is chosen such that it satisfies $P^{\text{data}}((-K_{\text{data}}, K_{\text{data}}) \times \mathbb{R}) = 0$. Finally, let D be a data set with $n \geq 1$ i.i.d. samples $(x_j, y_j) \sim P^{\text{data}}$. Then with probability $1 - O(n^{-\gamma'} + nm^{-\eta\gamma_{\text{data}}})$ the following statements simultaneously hold:*

- (D1) $\underline{x}_D \geq K_{\text{data}} m^{-\gamma_{\text{data}}}$,
 (D2) For all $\sigma \in \{\pm 1\}$ we have

$$\begin{aligned} \frac{1}{2} \lambda_{\min}(\mathbf{M}_{P^{\text{data}}, \sigma}) &\leq \lambda_{\min}(\mathbf{M}_{D, \sigma}) \\ \lambda_{\max}(\mathbf{M}_{D, \sigma}) &\leq 2 \lambda_{\max}(\mathbf{M}_{P^{\text{data}}, \sigma}) , \end{aligned}$$

- (D3) $\|\mathbf{v}_D^{\text{opt}} - \mathbf{v}_{P^{\text{data}}}^{\text{opt}}\|_{\infty} \leq n^{\varepsilon-1/2}$.

By combining Proposition 21 with Theorem 10, we obtain the following theorem, which is the key result for proving our inconsistency result. Due to the imposed constraint $m \leq K_{\text{param}} n^{\frac{1}{2-2\gamma_{\psi}}} < o(n)$, it is situated in the under-parameterized regime.

Theorem 22. *Let P^{data} satisfy Assumption 16, let $\varepsilon > 0$, and let $\gamma_{\psi}, \gamma_{\text{data}}, \gamma_P \geq 0$ with $\gamma_{\psi} + \gamma_{\text{data}} + \gamma_P < 1/2$. Then, for given constants $K_{\text{param}}, K_{\text{data}} > 0$, where K_{data} needs to satisfy $P^{\text{data}}((-K_{\text{data}}, K_{\text{data}}) \times \mathbb{R}) = 0$ in the case $\eta = \infty$, there exist constants $C_P, C_{\text{lr}}, C_{\text{weights}} > 0$ such that the following statement holds:*

For all data sets D with $n \geq 1$ i.i.d. samples $(x_j, y_j) \sim P^{\text{data}}$ and all neural networks (2) with width $m \geq 1$ that are initialized and trained according to Assumption 5 with step size h , the random sequence $(f_{W_k})_{k \in \mathbb{N}_0}$ satisfies the conclusions (i) – (v) of Theorem 10 with probability not less than $1 - C_P(m^{-\gamma_P} + nm^{-\eta\gamma_{\text{data}}})$ provided that h and m satisfy the constraints:

$$\begin{aligned} 0 < h &\leq C_{\text{lr}} m^{-1} \\ m &\leq K_{\text{param}} n^{\frac{1}{2-2\gamma_{\psi}}} . \end{aligned}$$

Proof Let $\tilde{\varepsilon} > 0$ be small enough such that $\gamma'_{\psi} := 2\tilde{\varepsilon} + \gamma_{\psi}$ satisfies $\gamma'_{\psi} + \gamma_{\text{data}} + \gamma_P < 1/2$. Under the assumptions above, there is a constant $K_M > 0$ such that for any data set D satisfying (D1) – (D3) of Proposition 21, we have the following three estimates:

$$\underline{x}_D \geq K_{\text{data}} m^{-\gamma_{\text{data}}} ,$$

$$\frac{1}{K_M} \leq \lambda_{\min}(\mathbf{M}_{D,\pm 1}) \leq \lambda_{\max}(\mathbf{M}_{D,\pm 1}) \leq K_M,$$

$$\|\mathbf{v}_D^{\text{opt}} - \mathbf{v}_{P^{\text{data}}}^{\text{opt}}\|_{\infty} \leq n^{(\tilde{\varepsilon}-1/2)} \leq O(m^{(2-2\gamma_{\psi})(\tilde{\varepsilon}-1/2)}) \leq O(m^{2\tilde{\varepsilon}+\gamma_{\psi}-1}) = O(m^{\gamma'_{\psi}-1}).$$

It follows that

$$\psi_{D,q} \leq \psi_{P^{\text{data}},q} + \|\mathbf{v}_D^{\text{opt}} - \mathbf{v}_{P^{\text{data}}}^{\text{opt}}\|_{\infty} \stackrel{\text{(P3)}}{\leq} O(m^{\gamma'_{\psi}-1})$$

$$\psi_{D,p} \leq \psi_{P^{\text{data}},p} + \|\mathbf{v}_D^{\text{opt}} - \mathbf{v}_{P^{\text{data}}}^{\text{opt}}\|_{\infty} \leq \psi_{P^{\text{data}},p} + O(m^{\gamma'_{\psi}-1}) \leq O(1).$$

Then, by Theorem 10, there exists $C_{\text{lr}} > 0$ such that the conclusions (i) – (v) of Theorem 10 hold with probability $\geq 1 - O(m^{-\gamma_P})$ if $0 < h \leq C_{\text{lr}} m^{-1}$. Let us now fix a $\gamma' \geq \gamma_P$. Because of (P1) and (P2), Proposition 21 then shows that the conditions (D1) – (D3) hold with probability

$$1 - O(n^{-\gamma'} + nm^{-\eta\gamma_{\text{data}}}).$$

Moreover, $\gamma' \geq \gamma_P$ implies $n^{-\gamma'} \leq O(m^{-(2-2\gamma_{\psi})\gamma'}) \leq O(m^{-\gamma'}) \leq O(m^{-\gamma_P})$. By the union bound, the conclusions therefore hold with the specified probabilities. \blacksquare

Remark 23. *A simple calculation shows that the condition $\eta > 4$ imposed in Assumption 16 is necessary and sufficient for the existence of a sequence $(m_n)_{n \geq 1}$ that satisfies the following two conditions:*

$$nm_n^{-\eta\gamma_{\text{data}}} \rightarrow 0 \quad \text{and} \quad m_n \leq K_{\text{param}} n^{\frac{1}{2-2\gamma_{\psi}}}.$$

Note that the first condition ensures that the probability considered in Theorem 22 converges to 1, while the second condition is a prerequisite in Theorem 22. Finally, note that such a sequence necessarily satisfies $m_n \rightarrow \infty$, and hence the approximation error of the considered networks converge to zero by the universal approximation theorem. In other words, the considered neural networks can represent predictors $f_{m_n} : \mathbb{R} \rightarrow \mathbb{R}$ such that $R_{P^{\text{data}}}(f_{m_n}) \rightarrow R_{P^{\text{data}}}^$ for $n \rightarrow \infty$, but gradient descent together with the considered initialization scheme is not able to find such predictors as the following three corollaries, whose proofs are given at the end of Section I, show.*

The first inconsistency result applies to all distributions P^{data} satisfying Assumption 16 for $\eta = \infty$, where we recall that $\eta = \infty$ simply means that P_X^{data} has no mass in a small vicinity around 0.

Corollary 24. *Consider the learning method that, given a data set $D \in (\mathbb{R} \times \mathbb{R})^n$, chooses a function $f_D = f_{W_k}$, where k can be arbitrarily chosen and f_{W_k} is found according to Assumption 5 for network size m_n and step size h_n . If m_n and h_n satisfy*

$$m_n \leq O(n^{1-\varepsilon}), \quad \lim_{n \rightarrow \infty} m_n = \infty, \quad \text{and} \quad h_n < o(m_n^{-1}),$$

then this learning method is not consistent for every distribution P^{data} satisfying Assumption 16 for $\eta = \infty$.

Note that the first condition $m_n \leq O(n^{1-\varepsilon})$ in Corollary 24 means that we are in the under-parameterized regime of neural networks. However, since we may consider $\varepsilon \rightarrow 0$ we can get at least arbitrarily close to the limiting regime, in which the number of parameters grows linearly in n .

Our next inconsistency result applies to distributions satisfying Assumption 16 for some $4 < \eta < \infty$, i.e., for distributions that only have a small positive mass around 0.

Corollary 25. *Let $\eta \in (4, \infty)$ and consider the neural network learning method of Corollary 24, but with $m_n \leq O(n^{1-\varepsilon})$ replaced by $m_n = \Theta(n^\gamma)$ for some $\gamma \in (\frac{2}{\eta}, 1 - \frac{2}{\eta})$. Then this learning method is not consistent for every distribution P^{data} satisfying Assumption 16 for the chosen η .*

Note that the price for extending the inconsistency from $\eta = \infty$ to $\eta < \infty$ is that we can no longer get arbitrarily close to the limiting regime discussed above. In fact, for $\eta \rightarrow 4$, the constraint on the network size m_n becomes stronger. This is consistent with the observations made in Remark 23.

Our final result shows that neural networks acting on higher dimensional data are also not consistent provided that they are initialized and trained in analogy to Assumption 5.

Corollary 26. *Consider a learning method as in Corollary 24 but for d -dimensional data sets $D \in (\mathbb{R}^d \times \mathbb{R})^n$ and with the initialization and training adaptations discussed in Remark 6. Then this learning method is not universally consistent.*

Let us finally set our results in a broader context of learning guarantees for other learning methods derived in statistical learning theory. Here, besides the worst-case approach leading to the notion of universal consistency, refined results that, for example, establish learning rates for restricted classes of distributions exists. For such rates, one is typically interested in minmax optimal rates, that is, in rates that, independent of the algorithm, cannot be improved for the considered class of distributions. In the regression context such restricted classes are almost always described by smoothness assumptions on the target function to be learned. Typical and rather classical results in this direction show that many learning algorithms learn faster the smoother this target function is. This corresponds to the intuition that learning should be easier if, on average, close-by input points lead to close-by labels. Now, the classes of distributions considered in this section contain examples with arbitrarily high smoothness, meaning that at least some traditional learning methods including gradient descent with early stopping in suitable reproducing kernel Hilbert spaces can learn these examples with up to the rate $\mathcal{O}(n^{-1})$, see Yao et al. (2007). In contrast, our results show that the considered neural networks do not learn at all in the case $d = 1$. In higher dimensions, the well-known curse of dimensionality prevents fast learning rates for moderately smooth target functions unless additional assumptions are imposed. One such assumption is the so-called manifold assumption, which, roughly speaking, assumes that the input data lies on a manifold with low intrinsic dimension $\rho < d$ and it is widely believed that at least some real-world data sets satisfy this assumption approximately. Theoretical results, see e.g. Hamm and Steinwart (2021) and the references therein, show that in this case actually ρ plays the role of the dimension in the resulting rates, or to phrase

it differently, the smaller the intrinsic dimension is the faster some learning algorithms can learn.

Both the smoothness and the manifold assumption were postulated to better describe common features of real-world distributions, and our results show that the considered neural networks do not learn at all for some of these distributions, at least if $\rho = 1$. While the case $\rho = 1$ may seem unrealistic, our experiments in Section 10 indicate that postulating lower bounds of the form $\rho \geq 2$ or even $\rho \geq 4$ may not solve the problem for (stochastic) gradient descent with zero bias initialization.

7. Proof Idea

Here, we want to give an overview over the proof of Theorem 10. We omit technical terms with exponent ε for simplicity and choose a different order than in the appendix.

As explained in Section 3, we want to show that the kinks $-b_{i,k}/a_{i,k}$ do not move much during training. To this end, we show that $|a_{i,0}| \geq \Omega(m^{-1-\gamma_P})$ with probability $1 - O(m^{-\gamma_P})$. Hence, if $|a_{i,k} - a_{i,0}| < o(m^{-1-\gamma_P})$, then we still have $|a_{i,k}| \geq \Omega(m^{-1-\gamma_P})$. Moreover, if $|b_{i,k}| = |b_{i,k} - b_{i,0}| \leq O(m^{-1-\gamma_P - \gamma_{\text{data}}})$, then $|b_{i,k}/a_{i,k}| \leq O(m^{-\gamma_{\text{data}}})$ and the main conclusion (v) of Theorem 10 follows.

Note that the (Leaky)ReLU φ satisfies $\varphi(a_i x + b_i) = \varphi'(\text{sgn}(a_i x + b_i)) \cdot (a_i x + b_i)$. We then investigate GD on a modified loss function $L_{D,\tau}$, where we replace $\varphi'(\text{sgn}(a_i x + b_i))$ by $\varphi'(\text{sgn}(a_{i,0} x + b_{i,0}))$. If GD on this modified loss does not move the kinks much, then $\text{sgn}(a_i x + b_i)$ remains constant and $\nabla L_{D,\tau}(W_k) = \nabla L_D(W_k)$. On both loss functions GD will thus yield the same result. Such a strategy has also been used by Li and Liang (2018).

Next, we will explain how to bound the change in $b_{i,k}$, the situation for $a_{i,k}$ is analogous. One can show that with $\sigma := \text{sgn}(a_{i,0})$, we have $b_{i,l+1} = b_{i,l} + h w_{i,l} s_{\sigma,l}$ for a quantity $s_{\sigma,l}$ defined in Definition C.1. We can then derive

$$\begin{aligned} |b_{i,k} - b_{i,0}| &\leq h \sum_{l=0}^{k-1} |w_{i,l} s_{\sigma,l}| \leq \left(\sup_{0 \leq l < k} |w_{i,l}| \right) \cdot h \sum_{l=0}^{k-1} |s_{\sigma,l}| \\ &\leq \left(|w_{i,0}| + \sup_{0 \leq l < k} |w_{i,l} - w_{i,0}| \right) \cdot h \sum_{l=0}^{k-1} |s_{\sigma,l}|. \end{aligned} \tag{4}$$

Given a bound on $h \sum_{l=0}^{k-1} |s_{\sigma,l}|$, bounding $|b_{i,k} - b_{i,0}|$ in this way requires bounding $|w_{i,l} - w_{i,0}|$. Bounding the latter with a similar argument would require bounding $|b_{i,l'} - b_{i,0}|$ and so on. While one can proceed by proving bounds using induction, we resolve the problem in a different but similar fashion in Proposition G.2 which does not require guessing an induction hypothesis.

As mentioned before, the neural network functions f_{W_k} are piecewise affine. Hence there are affine functions $f_{W_{k,1}}(x) = p_{1,k}x + q_{1,k}$ and $f_{W_{k,-1}}(x) = p_{-1,k}x + q_{-1,k}$ such that $f_{W_k}(x) = f_{W_{k,1}}(x)$ for sufficiently large $x > 0$ and $f_{W_k}(x) = f_{W_{k,-1}}(x)$ for sufficiently small

$x < 0$. A central quantity in our proof is the vector

$$\bar{\mathbf{v}}_k := \begin{pmatrix} p_{1,k} - p_{D,1}^{\text{opt}} \\ p_{-1,k} - p_{D,-1}^{\text{opt}} \\ q_{1,k} - q_{D,1}^{\text{opt}} \\ q_{-1,k} - q_{D,-1}^{\text{opt}} \end{pmatrix}$$

containing the difference of the slope and intercept parameters to their affine regression optimum. We show in Section C that $s_{\sigma,l}$ is a linear combination of components of $\bar{\mathbf{v}}_l$. Thus, any bound on $h \sum_{l=0}^{k-1} \|\bar{\mathbf{v}}_l\|$ directly yields a bound on $h \sum_{l=0}^{k-1} |s_{\sigma,l}|$. We also show that

$$\bar{\mathbf{v}}_{k+1} = (\mathbf{I}_4 - h\mathbf{A}_k\mathbf{M}_D)\bar{\mathbf{v}}_k, \quad (5)$$

where \mathbf{I}_4 denotes the four-dimensional identity matrix, $\mathbf{A}_k \in \mathbb{R}^{4 \times 4}$ depends on W_k , and $\mathbf{M}_D \in \mathbb{R}^{4 \times 4}$ is assembled using $\mathbf{M}_{D,1}$ and $\mathbf{M}_{D,-1}$. Under the hypothesis that W_k is close to W_0 , we have $\mathbf{A}_k \approx \mathbf{A}^{\text{ref}}$, where $\mathbf{A}^{\text{ref}} \in \mathbb{R}^{4 \times 4}$ is a suitable matrix only depending on W_0 .

We first consider a reference system $\bar{\mathbf{v}}_{k+1} = (\mathbf{I}_4 - h\mathbf{A}^{\text{ref}}\mathbf{M}_D)\bar{\mathbf{v}}_k$. It can be shown that \mathbf{A}^{ref} and \mathbf{M}_D are symmetric and positive definite (s.p.d.) with probability one. By a change of basis, we obtain the s.p.d. matrix $\mathbf{H} := \mathbf{M}_D^{1/2}\mathbf{A}^{\text{ref}}\mathbf{M}_D^{1/2} = \mathbf{M}_D^{1/2}(\mathbf{A}^{\text{ref}}\mathbf{M}_D)\mathbf{M}_D^{-1/2}$. Hence, $\mathbf{A}^{\text{ref}}\mathbf{M}_D$ has positive real eigenvalues $\lambda_1, \dots, \lambda_4$ with eigenvectors $\mathbf{v}_1, \dots, \mathbf{v}_4$. If $\bar{\mathbf{v}}_0 = \sum_i C_i \mathbf{v}_i$, then

$$\bar{\mathbf{v}}_k = \sum_{i=1}^4 (1 - h\lambda_i)^k C_i \mathbf{v}_i$$

and for $0 < h \leq (\max_i \lambda_i)^{-1}$, i.e., $1 - h\lambda_i \in [0, 1)$, we thus find

$$h \sum_{l=0}^{\infty} \|\bar{\mathbf{v}}_l\| \leq \sum_{i=1}^4 h \sum_{l=0}^{\infty} (1 - h\lambda_i)^l |C_i| \|\mathbf{v}_i\| = \sum_{i=1}^4 \lambda_i^{-1} |C_i| \|\mathbf{v}_i\|.$$

The idea is now to show that, with high probability, $\lambda_1, \lambda_2 = \Theta(m)$ and $\lambda_3, \lambda_4 = \Theta(1)$, while $\|\mathbf{v}_i\| \leq O(1)$, $|C_1|, |C_2| \leq O(1)$ and $|C_3|, |C_4| \leq O(m^{\gamma_\psi - 1})$ in order to obtain the bound

$$h \sum_{l=0}^{\infty} \|\bar{\mathbf{v}}_l\| \leq O(m^{\gamma_\psi - 1}).$$

Indeed, we show in Proposition F.4 that $\text{Span}\{\mathbf{v}_1, \mathbf{v}_2\} \approx \text{Span}\{\mathbf{e}_1, \mathbf{e}_2\}$ with the first two standard unit vectors $\mathbf{e}_1, \mathbf{e}_2 \in \mathbb{R}^4$. With $q_{\sigma,0} = 0$ and $\psi_{D,q} \leq O(m^{\gamma_\psi - 1})$, we also show that $\bar{\mathbf{v}}_0$ is close to $\text{Span}\{\mathbf{e}_1, \mathbf{e}_2\}$ and therefore $|C_3|, |C_4| \leq O(m^{\gamma_\psi - 1})$.

In Proposition G.6, we perform an induction showing that W_k is close to W_0 and that the solution of $\bar{\mathbf{v}}_{k+1} = (\mathbf{I}_4 - h\mathbf{A}_k\mathbf{M}_D)\bar{\mathbf{v}}_k$ behaves similar to the solution of the reference system $\bar{\mathbf{v}}_{k+1} = (\mathbf{I}_4 - h\mathbf{A}^{\text{ref}}\mathbf{M}_D)\bar{\mathbf{v}}_k$. Inserting the result into Eq. (4) yields the asymptotics

$$|b_{i,k} - b_{i,0}| = (O(m^{-1/2}) + o(m^{-1/2}))O(m^{\gamma_\psi - 1})$$

$$= O(m^{\gamma_\psi - 3/2}) .$$

By our assumption $\gamma_\psi + \gamma_{\text{data}} + \gamma_P < 1/2$, we have $m^{\gamma_\psi - 3/2} < o(m^{-1 - \gamma_{\text{data}} - \gamma_P})$ and as outlined in the beginning of this section, all kinks only move by $O(m^{-\gamma_{\text{data}}})$.

The idea of deriving a system as in Eq. (5) and using induction to prove that W does not change much over time has already been used by e.g. Du et al. (2019b). The main novelties in this part of our proof are:

- In our scenario, we are able to use a four-dimensional system instead of a n -dimensional system. We explain the relation between these systems in Section 8.
- We find different eigenvalue asymptotics and results for the alignment of the corresponding eigenvectors depending on the data set, which requires more sophisticated arguments to exploit.
- We prove different bounds on the change of weights in different layers, which is a consequence of using a more standard parameterization of the NN. We also prove strong bounds on certain “second-moment” weight statistics, cf. Remark G.4.

8. Relation to Neural Tangent Kernels

In this section, we illustrate that a part of our approach essentially consists of factoring and analyzing the singular NTK matrix associated with our neural network. To this end, consider the continuous gradient flow dynamics

$$\frac{d}{dt}W(t) = -\nabla L_D(W(t)) .$$

As shown e.g. by Du et al. (2019b), the vector

$$\bar{\mathbf{f}}(t) := \begin{pmatrix} f_{W(t)}(x_1) - y_1 \\ \vdots \\ f_{W(t)}(x_n) - y_n \end{pmatrix}$$

then satisfies the differential equation

$$\dot{\bar{\mathbf{f}}}(t) = -\frac{1}{n}\mathbf{K}(t)\bar{\mathbf{f}}(t) ,$$

where $\mathbf{K}(t) \in \mathbb{R}^{n \times n}$ is the empirical NTK matrix defined by

$$[\mathbf{K}(t)]_{ij} = \left\langle \frac{\partial f_{W(t)}(x_i)}{\partial W}, \frac{\partial f_{W(t)}(x_j)}{\partial W} \right\rangle .$$

Assuming without loss of generality that $x_1, \dots, x_{n'} > 0$ and $x_{n'+1}, \dots, x_n < 0$ for suitable n' , define

$$\mathbf{X} := \begin{pmatrix} x_1 & 0 & 1 & 0 \\ \vdots & \vdots & \vdots & \vdots \\ x_{n'} & 0 & 1 & 0 \\ 0 & x_{n'+1} & 0 & 1 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & x_n & 0 & 1 \end{pmatrix} .$$

We show in Proposition K.1 that in our scenario with no kink reaching a data point (setting $h = 0$),

$$\begin{aligned} \mathbf{K}(t) &= \mathbf{X} \mathbf{A}(t) \mathbf{X}^\top \\ \mathbf{M}_D &= \frac{1}{n} \mathbf{X}^\top \mathbf{X} \\ \bar{\mathbf{v}} &= (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \bar{\mathbf{f}}. \end{aligned}$$

Therefore $\frac{1}{n} \mathbf{K}(t)$ has the same non-zero eigenvalues and the same rank as $\frac{1}{n} \mathbf{A}(t) \mathbf{X}^\top \mathbf{X} = \mathbf{A}(t) \mathbf{M}_D$, i.e., it has rank four, two $\Theta(m)$ eigenvalues and two $\Theta(1)$ eigenvalues.¹ It follows that

$$\begin{aligned} \dot{\bar{\mathbf{v}}} &= (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \dot{\bar{\mathbf{f}}} \\ &= -\frac{1}{n} \mathbf{A} \mathbf{X}^\top \bar{\mathbf{f}} = -\frac{1}{n} \mathbf{A} (\mathbf{X}^\top \mathbf{X}) (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \bar{\mathbf{f}} \\ &= -\frac{1}{n} \mathbf{A} \mathbf{X}^\top \mathbf{X} \bar{\mathbf{v}} = -\mathbf{A} \mathbf{M}_D \bar{\mathbf{v}}, \end{aligned}$$

which is the continuous-time analog of our update equation $\bar{\mathbf{v}}_{k+1} = (\mathbf{I}_4 - h \mathbf{A}_k \mathbf{M}_D) \bar{\mathbf{v}}_k$. In conclusion, after removing its null space, the rescaled kernel matrix $\frac{1}{n} \mathbf{K}$ is related to $\mathbf{A} \mathbf{M}_D$ via a change of basis. Working with $\mathbf{A} \mathbf{M}_D$ is more comfortable in our case, because $\mathbf{A} \mathbf{M}_D$ is invertible and it facilitates bounding the eigenvalues and eigenvectors.

9. Inconsistency Experiments

In this section, we present empirical evidence from some Monte Carlo experiments that:

- (1) The failure of kinks to move sufficiently far to reach data points can occur for realistic network sizes with high probability.
- (2) Stochastic gradient descent (SGD) can exhibit a similar behavior when combined with an early stopping rule.
- (3) A similar failure also occurs on multi-dimensional data sets that do not lie in a one-dimensional subspace as in Remark 6.

Data for the figures in this section can be reproduced using the code at

`github.com/dholzmueLLer/nn_inconsistency`,

which is archived at <https://doi.org/10.18419/darus-2978>.

9.1 One-dimensional Distribution

For (1) and (2), we use the following experimental setup: We compute each estimated probability using 10^4 Monte Carlo trials. We choose P^{data} as the uniform distribution on

1. In general, if v is an eigenvector of AB with eigenvalue $\lambda \neq 0$, then $Bv \neq 0$ is an eigenvector of BA with eigenvalue λ .




Figure 5: Monte-Carlo estimates (10^4 repetitions) of the probability of a kink crossing a sample for different values of m , $n = m^2$, and different optimization and termination strategies. The data set is described in the text and ES stands for early stopping. The line $y = x^{-1/2}$ is added to better illustrate the asymptotic behavior of the probabilities in m .

the data set D from Example 1 and sample a data set D' of size $n = m^2$ from P^{data} . We initialize the weights independently as

$$a_{i,0} \sim \mathcal{N}(0, 2), \quad b_{i,0} = 0, \quad c_0 = 0, \quad w_{i,0} \sim \mathcal{N}(0, 2/m).$$

We then use either gradient descent or stochastic gradient descent with batch size 16 in order to train the network and check whether a kink $-b_{i,k}/a_{i,k}$ leaves the interval $(-1, 1)$, in which case we can stop training.² For some experiments, we also stop training if an early stopping (ES) criterion is satisfied.³ For GD, we can also stop, if the techniques used in Proposition F.4, Proposition G.2, and Proposition G.6 guarantee that no kink will ever leave the interval $(-1, 1)$.

For the step size h , unless specified otherwise, we use our maximal upper bound $h = \lambda_{\max}(\mathbf{H})^{-1}$ as mentioned in Remark 11 in order to reduce the number of iterations needed. We observe experimentally that $\lambda_{\max}(\mathbf{H})^{-1} \approx 0.4m^{-1}$ for our choice of P^{data} .

Figure 5 shows how the probabilities behave as m , and thus $n = m^2$, increases. In our scenario, we can apply Theorem 22 with $\gamma_\psi = \gamma_{\text{data}} = 0$ and obtain that the probabilities should behave like $O(m^{-\gamma_P})$ for all $\gamma_P < 1/2$. In Figure 5, this behavior can be observed

2. Ignoring the unlikely case $a_{i,k} = 0$, f_{W_k} remains affine on $(-\infty, -1]$ and $[1, \infty)$ iff no kink leaves $[-1, 1]$.

3. We use early stopping as implemented in `Keras` (Chollet and others, 2015) with `patience = 10` and `min_delta = 10-8`: Every 1000 epochs (GD) or 1000 batches (SGD), we monitor the loss on an independently drawn validation set of size n . Whenever $L_{\text{val}} < L_{\text{ref}} - 10^{-8}$, where L_{val} is the validation loss, we set $L_{\text{ref}} := L_{\text{val}}$. Training is stopped, when L_{ref} did not decrease within the last ten checks.

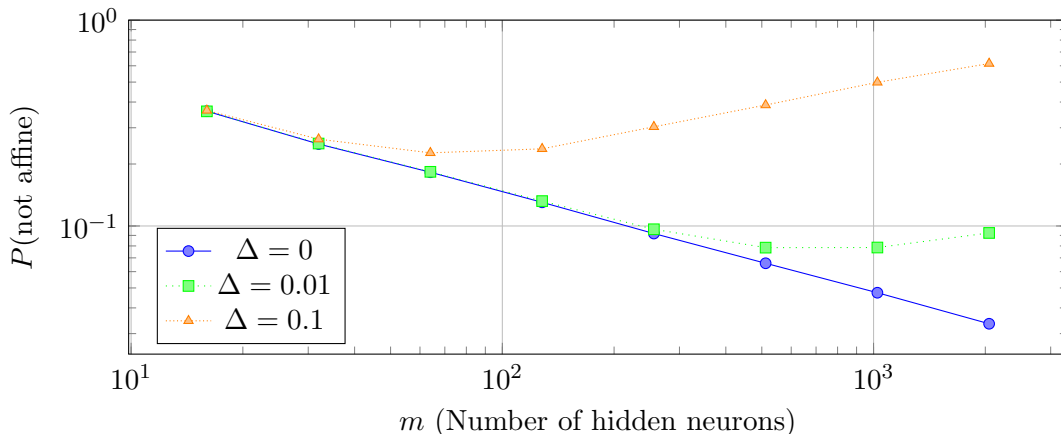


Figure 6: Monte-Carlo estimates (10^4 trials) of the probability of a kink crossing a sample for different values of m . We use $n = m^2$ and GD without early stopping. Here, the data distribution P^{data} of Figure 5 is shifted upwards in y -direction by Δ . As described in the text, this change means that the condition $\psi_{D,q} \leq K_\psi m^{\gamma_\psi - 1}$ in Theorem 10 is eventually violated for increasing m .

even for small m and also for SGD.⁴ In this respect recall that it is shown in Figure 2 that kinks may also move by significant amounts in later stages of the optimization. It is therefore not surprising that compared to considering f_{W_k} for all $k \in \mathbb{N}_0$, using early stopping can significantly reduce the probability that a kink leaves $(-1, 1)$ during training. We see this especially for the small step size $h = 0.01m^{-1}$ in Figure 5. Finally, if we shift P^{data} upwards by adding some $\Delta \in \mathbb{R}$ to all y values, we have $\psi_{P^{\text{data}},q} = |\Delta|$ and assumption (P3) from Assumption 16 is violated. We can see in Figure 6 that this changes the asymptotic behavior, but for small m and $|\Delta|$, the probabilities are still similarly low. Note that this behavior is due to $\psi_{D,q} \leq O(m^{-1} + |\Delta|)$, where the $|\Delta|$ term dominates once we have entered the regime $m \gg 1/|\Delta|$, as it can be seen in Figure 6.

9.2 Multi-dimensional Distribution

In order to show (3), i.e., that NNs can also perform poorly on multi-dimensional data sets, we choose the following experimental setup: We consider the uniform distribution P^{data} on the data set from Figure 7, which consists of the 33 points (x_{ij}, y_{ij}) defined by

$$x_{ij} = (1 + 0.1 \cdot i) \begin{pmatrix} \cos(2\pi j/11) \\ \sin(2\pi j/11) \end{pmatrix}, \quad y_{ij} = 3|i| - 2, \quad j \in \{0, \dots, 10\}, \quad i \in \{-1, 0, 1\}.$$

As above, we sample $n = m^2$ data points from P^{data} , but for $m \in \{128, 256, 512, \dots, 8192\}$ with 1000 Monte Carlo runs on 10^6 epochs to ensure that the NNs have sufficiently converged. In the 1D case, we picked the step size $h = \lambda_{\max}(\mathbf{H})^{-1} = \lambda_{\max}(\mathbf{A}^{\text{ref}} \mathbf{M}_D) =$

4. Without early stopping, SGD might still be able to move kinks far enough after a large number of iterations due to noisy gradients.




Figure 7: Illustration of P^{data} of the two-dimensional data sets used for the experiments in Figure 8. The middle blue points have a y value of -2 and the inner and outer orange points have a y value of 1 .

$n\lambda_{\max}(\mathbf{K}_0)^{-1}$, where the last equality follows from Section 8 for the kernel matrix \mathbf{K}_0 at initialization. Since we have not defined \mathbf{H} for multi-dimensional data sets, we choose $h = n\lambda_{\max}(\mathbf{K}_0)^{-1}$ in this case. As in the previous experiments this approach again results in different values for each considered NN instance. For the initialization, we again follow He et al. (2015) and use

$$a_{i,j,0} \sim \mathcal{N}(0, 1), \quad b_{i,0} = 0, \quad c_0 = 0, \quad w_{i,0} \sim \mathcal{N}(0, 2/m) .$$

Figure 8 shows that for this multi-dimensional distribution, the probability of a kink reaching a data point still decreases with increasing width of the network, although the probabilities are higher than in the one-dimensional setting above. Moreover, we also monitor the probability that the mean squared error (MSE) drops below 0.999999 times the best MSE achievable by functions that are affine on each of the 11 “arms” $\{x_{j,-1}, x_{j,0}, x_{j,1}\}$. We observe that even when a kink reaches a data point, the MSE often does not drop below this threshold.

10. Exploring Possible Improvements

The inconsistency of the considered ReLU networks prompts the question on how they can be improved. In this section, we experimentally compare two-layer ReLU networks with He initialization and (stochastic) gradient descent to other two-layer ReLU Networks. When the input dimension d is larger than one, we use SGD instead of GD in order to deal with larger data sets efficiently. Note that in our experiments in Section 9, SGD exhibits similar behavior to GD.

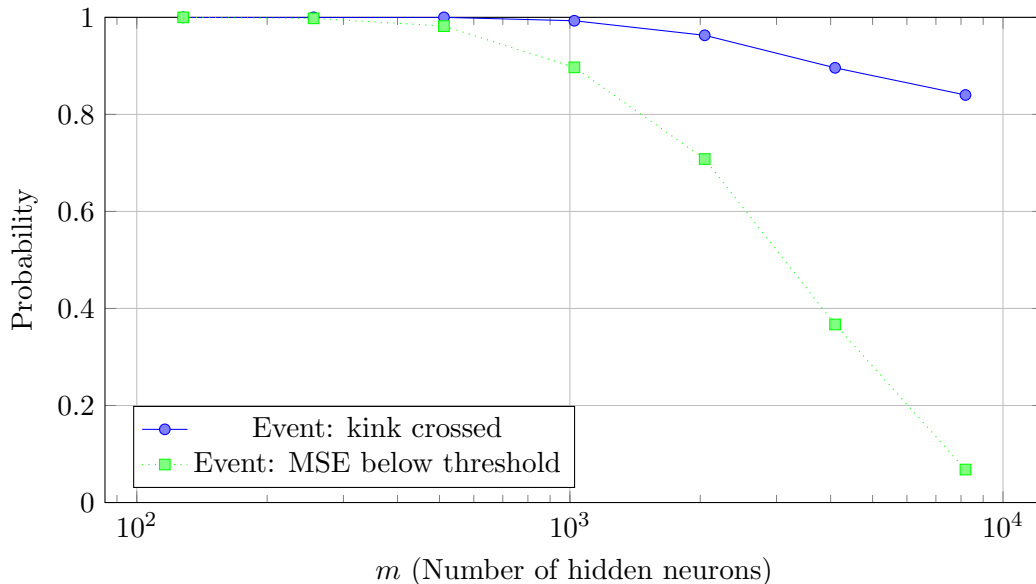


Figure 8: Monte-Carlo estimates (1000 trials) of the probability of a kink crossing a sample / the MSE passing below 0.999999 times the best possible “affine” MSE for different values of m , $n = m^2$ using GD over 10^6 epochs. Here, a uniform distribution on the data set from Figure 7 is used.

For possible improvements of the considered two-layer ReLU networks, a wide variety of options is available, such as

- (1) changing the initialization / parameterization of the weights,
- (2) changing the optimizer,
- (3) using a different bias initialization method,
- (4) making the network deeper or wider, and
- (5) using a different activation function.

For all of these options, it is currently unknown whether the resulting learning method will be universally consistent. For this reason, we investigate options (1) – (3) using numerical experiments,⁵ keeping a two-layer ReLU NN structure: For input $\mathbf{x} \in \mathbb{R}^d$, let

$$f_W(\mathbf{x}) := c + \alpha_2 \sum_{i=1}^m w_i \varphi(\alpha_1 \langle \mathbf{a}_i, \mathbf{x} \rangle + b_i)$$

with parameters $c, w_i, b_i \in \mathbb{R}$ and $\mathbf{a}_i \in \mathbb{R}^d$ and constants $\alpha_1, \alpha_2 \in \mathbb{R}$. In our experiments, we set $m = 512$ and we always initialize the second-layer bias c to zero since it is not followed by a ReLU activation φ . We consider the following options:

5. Our results can be reproduced using the code at github.com/dholzmueLLer/nn_inconsistency, which is archived at <https://doi.org/10.18419/darus-2978>.

10.1 Parameterization

We consider two parameterizations with corresponding weight initializations:

- Standard parameterization (**He**): Following He et al. (2015), we set $\alpha_1 = \alpha_2 = 1$ and initialize $\mathbf{a}_i \sim \mathcal{N}(0, \frac{2}{d}\mathbf{I}_d)$ and $w_i \sim \mathcal{N}(0, \frac{2}{m})$ independently. This parameterization is used in our theoretical analysis.
- NTK parameterization (**NTK**): Similar to Jacot et al. (2018), we set $\alpha_1 = \sqrt{\frac{2}{d}}, \alpha_2 = \sqrt{\frac{2}{m}}$ and initialize $\mathbf{a}_i \sim \mathcal{N}(0, \mathbf{I}_d)$ and $w_i \sim \mathcal{N}(0, 1)$ independently. For gradient-based optimization, this parameterization usually increases the relative training speed of the first layer compared to the second layer.⁶ Compared to Jacot et al. (2018), we include the factor $\sqrt{2}$ in α_1 and α_2 such that the initial network function is distributed in the same way as for the standard parameterization.

10.2 Optimization

We consider the following optimizers, using hyperparameter optimization for the learning rate (see Training procedure below):

- Stochastic Gradient Descent (**SGD**): The non-stochastic version of this optimizer is used in our theoretical analysis and in the experiments for $d = 1$.
- SGD with Momentum (**SGDM**): We set the momentum hyperparameter to 0.9, which corresponds to the default momentum used for Adam.
- Adam (**Adam**): A very popular adaptive optimizer proposed by Kingma and Ba (2015). Except for the learning rate, we use the standard hyperparameters $\beta_1 = 0.9$, $\beta_2 = 0.999$ and $\epsilon = 10^{-8}$.

10.3 Bias Initialization

For the initialization of the first-layer biases, we consider multiple choices:

- Zero initialization (**Zero**): This is the initialization $b_i = 0$ that we have used throughout this paper.
- PyTorch default (**PyTorch**): The default initialization in PyTorch initializes $b_i \sim \mathcal{U}[-1/\sqrt{d}, 1/\sqrt{d}]$.
- Uniform initialization ($U(1)$): We initialize $b_i \sim \mathcal{U}[-\sqrt{6}, \sqrt{6}]$. For a normalized input distribution and for both parameterizations, we have $\mathbb{E}_{\mathbf{x} \sim P_X} \mathbb{E}_{\mathbf{a}_i} (\alpha_1 \langle \mathbf{a}_i, \mathbf{x} \rangle)^2 = 2 = \mathbb{E} b_i^2$.
- Positive uniform initialization ($U_+(1)$): We initialize $b_i \sim \mathcal{U}[0, \sqrt{6}]$.

6. When going from standard parameterization to NTK parameterization, the magnitude of GD updates of $\alpha_1 \mathbf{a}_i$ is multiplied by $\alpha_1^2 = \frac{2}{d}$, while the magnitude of updates of $\alpha_2 w_i$ is multiplied by $\alpha_2^2 = \frac{2}{m}$. In the regime $m \gg d$ considered here, this means that the relative training speed of the first layer to the second layer is multiplied by $m/d \gg 1$. After a corresponding increase of the learning rate, the first layer will train faster than for the standard parameterization. Since we especially want the biases to train faster, we do not employ the option to multiply the biases by a (small) constant to reduce the training speed of the biases to a level similar to the standard parameterization (Jacot et al., 2018).

- Negative uniform initialization ($U_-(1)$): We initialize $b_i \sim \mathcal{U}[-\sqrt{6}, 0]$.
- Negative uniform kink initialization ($U_{k-}(1)$): We initialize $b_i = \alpha_1 \|\mathbf{a}_i\|_2 \beta_i$, where $\beta_i \sim \mathcal{U}[-\sqrt{3}, 0]$ is independent of \mathbf{a}_i . This again ensures $\mathbb{E}b_i^2 = (\mathbb{E}(\alpha_1 \|\mathbf{a}_i\|)^2)(\mathbb{E}\beta_i^2) = 2 \cdot 1 = 2$. Here, $|\beta_i|$ is the distance of the hyperplane $\{\mathbf{x} \in \mathbb{R}^d \mid \alpha_1 \langle \mathbf{a}_i, \mathbf{x} \rangle + b_i = 0\}$ to zero. This initialization method is inspired by the initialization methods considered by Steinwart (2019).

10.4 Training Procedure

We choose a training set with $256 \cdot d$ samples, use a batch size of 256 and train for $8192/d$ epochs such that we always use 8192 iterations. Hence, the optimization only uses stochastic minibatches for $d \geq 2$. Every $64/d$ epochs, we compute errors on the validation and test sets, which contain 1024 samples each. Test errors are reported for those epochs with the best validation errors. We run each hyperparameter configuration for 100 times to obtain more accurate results. This is repeated for learning rates on a grid $\{\eta_0 \cdot 2^{k/2}, k \in \{-12, -11, \dots, 10\}\}$, where η_0 is a sensible default learning rate depending on the parameterization and optimization. For the learning rate η on the grid where the 100 repetitions have the lowest RMSE on average, another grid search is performed on the finer grid $\{\eta \cdot 2^{k/8}, k \in \{-3, -2, \dots, 3\}\}$ to find the finally used best learning rate. Using the same optimized learning rate for all 100 repetitions resembles a practitioner using a well-tuned default value for the learning rate. *For training, we normalize the training data such that y and each component of \mathbf{x} has zero mean and unit variance.*

10.5 One-Dimensional Example

Figure 9 shows typical results of various training configurations on a one-dimensional data set similar to Figure 1. In the remainder of this section, we consider a different set of data-generating distributions in order to study the influence of the input dimension on the performance of training configurations:

10.6 Data Distribution

For $d \in \{1, 2, 4, 8, 16, 32, 64\}$, we randomly draw samples (\mathbf{x}, y) as

$$\begin{aligned} \mathbf{x} &= u \frac{\tilde{\mathbf{x}}}{\|\tilde{\mathbf{x}}\|}, & \tilde{\mathbf{x}} &\sim \mathcal{N}(0, \mathbf{I}_d), & u &\sim \mathcal{U}[0, 1] \\ y &= \cos(2\pi\|\mathbf{x}\|) = \cos(2\pi u). \end{aligned} \tag{6}$$

Here, $\tilde{\mathbf{x}}$ and u are independent random variables. This makes the distribution of \mathbf{x} rotationally invariant, with its radius being uniformly distributed on $[0, 1]$. The definition of y implies that no label noise is used. We chose y such that it only depends on the radial component of the data and hence cannot be fit well with zero biases. We denote the distribution of (\mathbf{x}, y) by P_d^{data} .




Figure 9: Median-performing NNs for various configurations on a different data set that was drawn randomly from the distribution $P_{\text{ex}}^{\text{data}}$ specified in Section L. The training set is shown as black crosses and is identical to Figure 1 except that in Figure 1, $q_{D,\pm 1}^{\text{opt}} = 0$ was ensured by shifting both halves of the data set as in Remark 19. Training setup and learning rate optimization are identical to the other results in Section 10 for $d = 1$. For each configuration, eleven NNs were trained using the best found learning rate without early stopping. After 8192 epochs, the NN function realizing the median of the eleven validation RMSEs was plotted. It is apparent that the combination He/SGD/Zero performs badly although not all assumptions of our theoretical results are satisfied: Here, the x and y values of the data set are normalized before training, which violates (P2) and (P3) from Assumption 16 as well as the assumption that the (x_i, y_i) are i.i.d. Moreover, the learning rate is not constrained to a small regime during hyperparameter optimization.

10.7 Discussion

For each combination of optimizer and parameterization, relative comparisons of the bias initialization methods are shown in Figure 10. A direct comparison of optimizers and parameterizations for zero bias initialization is shown in Figure 11. Detailed results are provided in Section L. We observe the following trends in the experimental results:

- **Bias initialization matters less for NTK parameterization.** This is likely because the first layer trains faster in NTK parameterization than in standard parameterization.
- **NTK parameterization mostly outperforms standard parameterization.** A similar trend has been observed by Arora et al. (2019c). Counterexamples to this trend occur mostly for good bias initializations and small dimension $d \in \{1, 2\}$. This can also be observed in Section L and also in Figure 9. We conjecture that in this case, the faster training speed of the first layer when using NTK parameterization is detrimental because it allows the kinks to move away from their initially well-distributed locations.
- **Zero bias initialization performs badly for small input dimensions.** This is in agreement with our theoretical observations even though the data-generating

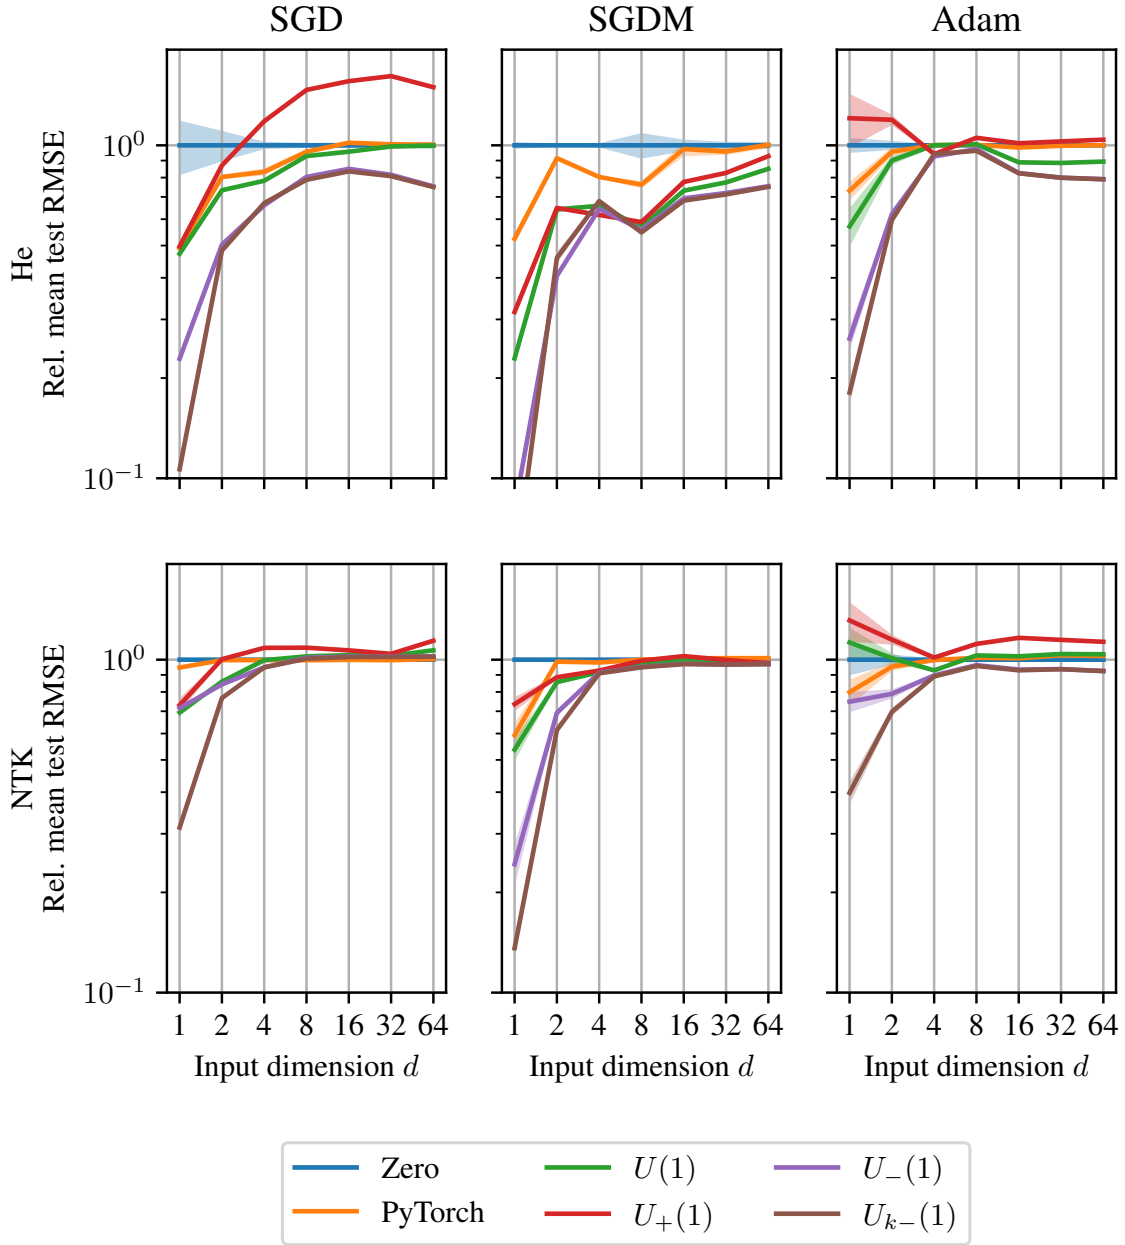


Figure 10: Relative Mean Test RMSEs of trained NNs using the data-generating distribution from Eq. (6) for varying input dimension d , bias initialization method, optimizer and parameterization. The Mean Test RMSEs are plotted relative to the corresponding NNs with zero bias initialization on a logarithmic scale.




Figure 11: Relative Mean Test RMSEs of trained NNs with zero-bias initialization using the data-generating distribution from Eq. (6) for varying input dimension d , optimizer and parameterization. The Mean Test RMSEs are plotted relative to the combination of He parameterization/initialization and SGD.

distribution is more strongly concentrated around zero than (P2) in Assumption 16 allows.

- Negative bias initialization helps.** In our experiments, negative uniform bias initialization $U_-(1)$ performs better than uniform bias initialization $U(1)$, which in turn performs better than positive uniform bias initialization $U_+(1)$. Note that the two neurons $\varphi(\langle \mathbf{a}, \mathbf{x} \rangle + b)$ and $\varphi(\langle -\mathbf{a}, \mathbf{x} \rangle + (-b))$ have the same “kink hyperplane” $\{\mathbf{x} \in \mathbb{R}^d \mid \langle \mathbf{a}, \mathbf{x} \rangle + b = 0\}$ but are nonzero on opposite sides of the hyperplane. Since our distribution of the weights \mathbf{a} is invariant under negation, the bias initializations $U_-(1)$, $U(1)$ and $U_+(1)$ generate the same distribution of kink hyperplanes. However, negative biases cause less data points to fall in the active region of the ReLU activation, creating a sparser activation pattern. This appears to be helpful, potentially because it eases optimization. A potential downside of negative bias initialization is that too small biases can create dead neurons, i.e., neurons that are always zero. Hence, the variance of the biases at initialization should not be too large.
- Initializing the kinks directly helps for small input dimensions.** For $d \in \{1, 2\}$, initializing the biases according to $U_{k-}(1)$ instead of $U_-(1)$, i.e., drawing the kinks uniformly instead of the biases, is mostly beneficial. For larger input dimensions, the differences between $U_{k-}(1)$ and $U_-(1)$ are very small. This is plausible: If $\xi \sim$

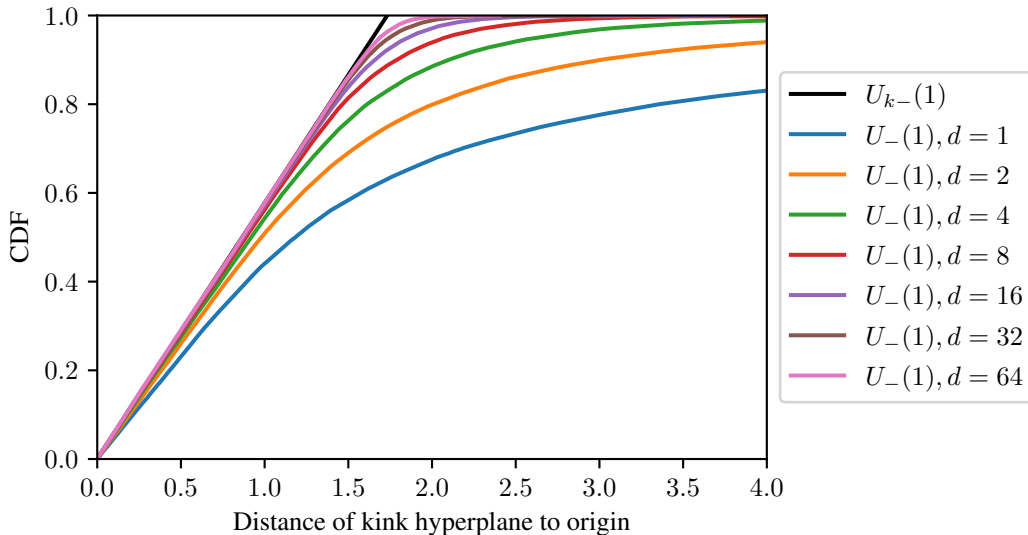


Figure 12: CDFs of the distribution of kink distances to the origin given by $\frac{|b_i|}{|\alpha_1| \|\mathbf{a}_i\|}$ for different bias initializations and input dimensions d . The CDF is independent of d when using $U_{k-}(1)$.

$\mathcal{N}(0, 1)$, then

$$\text{Var}(\alpha_1^2 \|\mathbf{a}_i\|_2^2) = \sum_{k=1}^d \text{Var}(\alpha_1^2 a_{i,k}^2) = \sum_{k=1}^d \text{Var}\left(\frac{2}{d} \xi^2\right) = \frac{4}{d} \text{Var}(\xi^2) = \frac{8}{d}.$$

Therefore, the factor $\alpha_1 \|\mathbf{a}_i\|_2$ becomes almost constant for high input dimensions d and the distributions $U_-(1)$ and $U_{k-}(1)$ become very similar. This is shown in Figure 12, which also shows that for small d and bias initialization $U_-(1)$, the distribution of distances of kink hyperplanes from the origin has heavy tails and can therefore create many dead neurons, which might explain the superior performance of $U_{k-}(1)$.

- **Adam is beneficial for small input dimensions.** In low dimensions, we observe that Adam performs better than SGDM and SGD. In higher dimensions, where the discrepancy between train and test errors becomes much larger, no clear picture emerges.

10.8 Limitations

We have demonstrated that negative kink-based bias initializations can improve the test errors of NNs on a family of synthetic data sets. While encouraging, this result still leaves unanswered questions for future research. First of all, the practical use of such bias initialization methods can only be answered by a large-scale empirical study. Steinwart (2019) performed such a large-scale study on tabular data sets and found that data-dependent

bias initialization methods can, on average, outperform zero initialization. As shown in Section L, $U_{k-}(1)$ usually yields better results on the synthetic data sets considered here than one of the initialization methods proposed by Steinwart (2019). Second, while our synthetic data distributions allow us to vary the input dimension in our experiments, distinguishing the effects of varying dimension from varying amounts of data and varying number of minibatches per epoch would require significantly more experiments. Third, a large-scale study including real-world data sets would also be necessary to determine the best kink distribution. In our experiments, we were not able to find a clearly best (data-dependent) kink distribution. For example, placing the kinks on random data points is not always better than $U_{k-}(1)$, cf. Section L. We also do not study the effects of bias initialization methods in deeper layers of deeper NNs. The inputs of deeper layers can exhibit different characteristics than normalized input data in the first layer and might hence require different bias initialization methods. For example, a good data-dependent bias initialization might help to counteract sample variance decay (Luther and Seung, 2019) in deep NNs and the resulting dying ReLU problem.

11. Conclusion

We have proven that NNs can fail with high probability in the over-parameterized regime for certain data sets and in the under-parameterized regime even for sampled data sets. In these cases, the NN converges to a local “linear regression” optimum. In particular, our analysis reveals that the difference of the NN to this linear regression optimum consists of a fast-decaying and a slow-decaying component and that for certain data sets, the slow-decaying component is already small at initialization. In essence, the reason is that learning is done mainly by the last layer, which contains m weights but only one bias and therefore, the bias is learned more slowly than the weights. Especially, in our case, the NN operates in a “lazy regime” (Chizat et al., 2019), where the NN stays close to a “linearized version” around its initialization.

Using slightly different assumptions⁷, Du et al. (2019b) show convergence of over-parameterized NNs to a global minimum. One of their assumptions is that no two data points x_i, x_j should be parallel. This assumption is violated in our scenario, where *all* x_j are parallel.⁸ Ironically, the case where samples lie on a low-dimensional submanifold of the input space is often considered a strength of deep learning methods, but at least in the extreme case of non-curved one-dimensional submanifolds our results show that neural networks with one hidden layer and zero bias initialization can be inconsistent. In this respect we also like to note that for a data set with parallel x_j , we can satisfy the assumption of Du et al. (2019b) by considering the transformed covariate samples $\tilde{x}_j := (x_j, 1)$. In fact, this transformation essentially corresponds to considering the original data set but initializing

7. For example, they use no biases \mathbf{b}, c and they use the NTK parameterization. The NTK parameterization is essentially equivalent to using a smaller learning rate for the second layer, which might lead to a different optimization result than the usual parameterization.

8. While we use biases, which can be emulated in NNs without biases by transforming the inputs to $\tilde{x}_j := (x_j, 1)$, our biases are initialized differently from the weights, and hence such a transformation is not equivalent to the NNs considered here.

the hidden biases by $b_i \sim \mathcal{N}(0, 1)$ instead of by $b_i = 0$. Note that this ensures that the kinks are now distributed randomly over \mathbb{R} after initialization. Consequently, given enough over-parameterization, there are already enough kinks available such that an NN could achieve zero training loss without moving them at all.

The training of NNs can be viewed as a race between linear and nonlinear dynamics trying to fit the data. In both our setting and the setting of Du et al. (2019b), linear dynamics reach a critical point of the optimization landscape before nonlinear dynamics can contribute substantially. From a theoretical viewpoint, a major contribution of this paper is to show that this can happen with high probability even in a setting where the critical point is a spurious local minimum and nonlinear dynamics would be necessary to find a global optimum. While our result is a worst-case result in the sense that it only considers a very specific class of data-generating distributions, our experiments demonstrate that the considered training configuration performs badly beyond the assumptions of our theoretical analysis.

It should be noted that our result does not exclude the possibility of proving worst-case generalization guarantees, such as universal consistency, for NNs. On the contrary, it might assist the search for such guarantees by showing that a necessary assumption of such a guarantee would be to rule out our considered training configuration. While the training configuration of Du et al. (2019b) allows to find a global optimum and therefore provides an interesting candidate setting for worst-case guarantees, it is yet unknown whether this leads to a universally consistent learning method. Moreover, Du et al. (2019b) require the number of hidden neurons to grow polynomially with the number of training samples, which can quickly become infeasible even for moderately sized data sets. It is still an open question whether kinks move suitably in an intermediate case, e.g. if they are initialized randomly but the NN is not strongly over-parameterized. It also remains an open question for future research whether or not our results can be extended to deep NNs, other loss functions, or other optimization methods.

While our analysis focuses on the role of biases during training, as we have already discussed above, biases can be interpreted as weights by considering inputs of the form $(x, 1)$. Hence, certain weight initialization methods such as sparse initialization methods might suffer from similar problems as zero bias initialization, as they can lead to similar degeneracies in the distribution of kink hyperplanes. Our insights might therefore help to explain why such weight initialization techniques are not frequently used. Moreover, it might be interesting to consider implications for NNs that are forced to be sparse, either by pruning after training or by imposing a sparse structure during training.

From a practical point of view, we provide experimental results on the effect of different alterations to our training setup on synthetic data sets in different dimensions. These data sets are generated using cosine-type functions of the norm of the input, such that the learning of biases is crucial to achieve good results. For high-dimensional learning problems, where neural networks have been employed with great success, the radial behavior of the target function may not be very interesting. For example, in image classification, changing the norm of the input can be interpreted as changing the illumination of the image, which should not change the desired classification outcome. On the other hand, our experimental

results are potentially relevant to the training of NNs on tabular data, a domain where other learning methods are still competitive with NNs (Shwartz-Ziv and Armon, 2022; Kadra et al., 2021; Gorishniy et al., 2021). Here, the input dimensions can be much smaller and the prediction result can strongly depend on the norm of the input. This can be intuitively seen for tasks such as the prediction of house prices from features such as floor area, or the task of predicting the runtime of a program from features such as the size of its input. Similarly, the learning of surrogate models for parameterized simulations may involve non-trivial radial behavior of the target function in low dimensions. Indeed, for tabular data, an extensive experimental evaluation by Steinwart (2019) has revealed that the results of zero bias initialization can be improved considerably by certain random bias initializations, which is in agreement with our experimental results on synthetic data sets. We hope that our analysis motivates further research to improve bias initializations and the distributions of ReLU kinks in general.

Acknowledgments

Funded by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy - EXC 2075 - 390740016. The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting David Holzmüller.

Appendix A. Notation and Matrix Algebra

In this section, we will introduce some notation that is used throughout the appendix. We will also list some results about matrices, especially involving matrix norms and eigenvalues, cf. e.g. Bhatia (2013) as well as Golub and Van Loan (1989). The rest of the appendix is structured as follows: In Section B, we will define a modification of the NN with fixed activation pattern and prove some elementary results. We will then investigate gradient descent dynamics for derived quantities like $\bar{\mathbf{v}}_k$ in Section C and discuss them in Section D. In Section E, we prove stochastic properties of the initialization W_0 and the data set D . We then investigate a simplified “linearized” case in Section F and show in Section G that the true behavior is close to the linearized case, which concludes the proof of our main theorem. In Section H, we show how NNs with one-dimensional input and NNs with multi-dimensional input are related. Our inconsistency corollaries are proved in Section I, a miscellaneous statement in Section J and relations of our quantities to NTK-related quantities in Section K. Finally, detailed results accompanying the experiments in Section 10 are given in Section L.

Definition A.1. We denote the sign of a real number $x \in \mathbb{R}$ by

$$\text{sgn}(x) = \begin{cases} 1 & , \text{ if } x > 0 \\ 0 & , \text{ if } x = 0 \\ -1 & , \text{ if } x < 0 . \end{cases}$$

For a set S , we denote its indicator function by $\mathbb{1}_S$, i.e.,

$$\mathbb{1}_S(x) = \begin{cases} 1 & , \text{ if } x \in S \\ 0 & , \text{ otherwise.} \end{cases}$$

Definition A.2 (Asymptotic notation). We use standard asymptotic notation $f < o(g)$, $f \leq O(g)$, $f = \Theta(g)$, $f \geq \Omega(g)$ (we do not need $f > \omega(g)$). The constant in such an asymptotic (in)equality should not depend on

- the number m of hidden neurons,
- the number n of data points,
- the step size $h > 0$,
- step count variables such as $k, l, l' \in \mathbb{N}_0$,
- the initialization W_0 ,
- the data set D ,

as long as these variables satisfy the imposed assumptions.⁹ For example, we could write $K_{\text{data}} m^{-\gamma_{\text{data}}} = \Theta(m^{-\gamma_{\text{data}}})$ for $K_{\text{data}} > 0$, but not $nm^{-\gamma_{\text{data}}} \leq O(m^{-\gamma_{\text{data}}})$.

Definition A.3. Let $n, m \geq 1$, let $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{m \times m}$ and $\mathbf{C} \in \mathbb{R}^{n \times m}$. (We sometimes use m, n to denote arbitrary vector space dimensions instead of numbers of hidden neurons and data points.)

9. Sometimes, we also need to assume that m, n is sufficiently large to be able to write $f \leq O(g)$ even if f is infinite or undefined for small m, n .

- (1) We write $\mathbf{A} \succ 0$ iff \mathbf{A} is symmetric and positive definite and $\mathbf{A} \succeq 0$ iff \mathbf{A} is symmetric and positive semidefinite. We define \preceq and \prec analogously.
- (2) A symmetric matrix \mathbf{A} has an orthogonal eigendecomposition $\mathbf{A} = \mathbf{U}\mathbf{D}\mathbf{U}^\top$ with $\mathbf{U} \in \mathbb{R}^{m \times m}$ orthogonal and $\mathbf{D} \in \mathbb{R}^{m \times m}$ diagonal such that \mathbf{D} contains the (real) eigenvalues of \mathbf{A} . We denote the set of eigenvalues of \mathbf{A} by $\text{eig}(\mathbf{A})$ and define

$$\begin{aligned}\lambda_{\max}(\mathbf{A}) &:= \max \text{eig}(\mathbf{A}) \\ \lambda_{\min}(\mathbf{A}) &:= \min \text{eig}(\mathbf{A}) .\end{aligned}$$

The matrix \mathbf{A} is invertible iff $0 \notin \text{eig}(\mathbf{A})$ and we have $\mathbf{A} \succeq 0$ iff $\text{eig}(\mathbf{A}) \subseteq [0, \infty)$. In the latter case, we can define the (symmetric) square root of \mathbf{A} as $\mathbf{A}^{1/2} := \mathbf{U}\mathbf{D}^{1/2}\mathbf{U}^\top$, where $\mathbf{D}^{1/2}$ contains the square roots of the entries of \mathbf{D} . Similarly, $\mathbf{A}^{-1} = \mathbf{U}\mathbf{D}^{-1}\mathbf{U}^\top$, which yields

$$\begin{aligned}\lambda_{\max}(\mathbf{A}^{1/2}) &= \lambda_{\max}(\mathbf{A})^{1/2}, & \lambda_{\min}(\mathbf{A}^{1/2}) &= \lambda_{\min}(\mathbf{A})^{1/2}, \\ \lambda_{\max}(\mathbf{A}^{-1}) &= \lambda_{\min}(\mathbf{A})^{-1}, & \lambda_{\min}(\mathbf{A}^{-1}) &= \lambda_{\max}(\mathbf{A})^{-1} .\end{aligned}$$

- (3) As matrix norms, we use the Frobenius norm as well as the induced 2- and ∞ -norms:

$$\begin{aligned}\|\mathbf{C}\|_F &= \left(\sum_{i,j} C_{i,j}^2 \right)^{1/2} \\ \|\mathbf{C}\|_2 &= \sup_{x \neq 0} \frac{\|\mathbf{C}x\|_2}{\|x\|_2} \\ \|\mathbf{C}\|_\infty &= \sup_{x \neq 0} \frac{\|\mathbf{C}x\|_\infty}{\|x\|_\infty} = \max_i \sum_j |C_{ij}| .\end{aligned}$$

If $\mathbf{C} \succeq 0$, then $\|\mathbf{C}\|_2 = \lambda_{\max}(\mathbf{C})$.

These matrix norms satisfy the following inequalities (cf. e.g. Section 2.3 in Golub and Van Loan (1989)):

$$\begin{aligned}\|\mathbf{C}\|_2 &\leq \|\mathbf{C}\|_F \leq \sqrt{m} \|\mathbf{C}\|_2 \\ \frac{1}{\sqrt{m}} \|\mathbf{C}\|_\infty &\leq \|\mathbf{C}\|_2 \leq \sqrt{n} \|\mathbf{C}\|_\infty .\end{aligned}$$

Moreover, if \mathbf{C}' is a subblock of \mathbf{C} , then $\|\mathbf{C}'\|_p \leq \|\mathbf{C}\|_p$ for $p \in \{2, F, \infty\}$.

- (4) We define the condition number of a matrix $\mathbf{A} \succ 0$ by

$$\text{cond}(\mathbf{A}) := \|\mathbf{A}\|_2 \cdot \|\mathbf{A}^{-1}\|_2 = \lambda_{\max}(\mathbf{A}) \lambda_{\max}(\mathbf{A}^{-1}) = \frac{\lambda_{\max}(\mathbf{A})}{\lambda_{\min}(\mathbf{A})} .$$

- (5) We occasionally use element-wise operations on matrices. For example, $|\mathbf{A}|$ is the matrix containing as entries the absolute values of the entries of \mathbf{A} and $\text{sup}_s \mathbf{A}(s)$ consists of the element-wise suprema. Also, $\mathbf{A} \leq \mathbf{B}$ means that $\mathbf{A}_{ij} \leq \mathbf{B}_{ij}$ for all i, j .

There are some more facts about matrices that we will use during some proofs. We show some typical arguments here:

- We will use the fact that for symmetric \mathbf{A} ,

$$\lambda_{\max}(\mathbf{A}) = \sup_{\|\mathbf{v}\|_2=1} \mathbf{v}^\top \mathbf{A} \mathbf{v}, \quad \lambda_{\min}(\mathbf{A}) = \inf_{\|\mathbf{v}\|_2=1} \mathbf{v}^\top \mathbf{A} \mathbf{v},$$

which is a special case of the Courant-Fischer-Weyl min-max principle (e.g. Corollary III.1.2 in Bhatia (2013)). This shows $\mathbf{A} \succeq 0 \Leftrightarrow \lambda_{\min}(\mathbf{A}) \geq 0$. For $\mathbf{A}, \mathbf{B} \succeq 0$, we can use such an argument to show that

$$\begin{aligned} \lambda_{\max}(\mathbf{A} + \mathbf{B}) &\leq \lambda_{\max}(\mathbf{A}) + \lambda_{\max}(\mathbf{B}) \\ \lambda_{\min}(\mathbf{A} + \mathbf{B}) &\geq \lambda_{\min}(\mathbf{A}) + \lambda_{\min}(\mathbf{B}). \end{aligned}$$

- If

$$\mathbf{M} = \begin{pmatrix} \mathbf{M}_{11} & \mathbf{M}_{12} \\ \mathbf{M}_{12}^\top & \mathbf{M}_{22} \end{pmatrix} \succeq 0,$$

we know that

$$\mathbf{x}^\top \mathbf{M}_{11} \mathbf{x} = \begin{pmatrix} \mathbf{x} \\ 0 \end{pmatrix}^\top \mathbf{M} \begin{pmatrix} \mathbf{x} \\ 0 \end{pmatrix} \geq \lambda_{\min}(\mathbf{M}) \|\mathbf{x}\|_2^2 \geq 0,$$

hence $\mathbf{M}_{11} \succeq 0$ with $\lambda_{\min}(\mathbf{M}_{11}) \geq \lambda_{\min}(\mathbf{M})$. Similarly, $\lambda_{\max}(\mathbf{M}_{11}) \leq \lambda_{\max}(\mathbf{M})$ and analogous identities hold for \mathbf{M}_{22} . We also have

$$\text{eig} \begin{pmatrix} \mathbf{M}_1 & \\ & \mathbf{M}_2 \end{pmatrix} = \text{eig}(\mathbf{M}_1) \cup \text{eig}(\mathbf{M}_2)$$

and therefore

$$\begin{pmatrix} \mathbf{M}_1 & \\ & \mathbf{M}_2 \end{pmatrix} \succ 0 \text{ iff } \mathbf{M}_1, \mathbf{M}_2 \succ 0.$$

Appendix B. Gradient Descent with Fixed Activation Pattern

In this section, we construct a modified loss function $L_{D,\tau}$ which fixes the activation pattern of the neurons to its state at initialization. We also show that $L_{D,\tau}(W) = L_D(W)$ for $W \approx W_0$ and introduce a shorter notation for gradient descent updates.

Definition B.1 (Fixed activation pattern). *Define*

$$\begin{aligned} \tau_i &:= \text{sgn}(a_{i,0}) \\ I &:= \{1, \dots, m\} \\ J &:= \{1, \dots, n\} \\ I_\sigma &:= \{i \in I \mid \tau_i = \sigma\} \\ J_\sigma &:= \{j \in J \mid \text{sgn}(x_j) = \sigma\} \\ f_{W,\tau,\sigma}(x) &:= c + \sum_{i \in I} w_i \varphi'(\sigma \tau_i) \cdot (a_i x + b_i) \\ L_{D,\tau}(W) &:= \frac{1}{2n} \sum_{j \in J} (y_j - f_{W,\tau,\text{sgn}(x_j)}(x_j))^2. \end{aligned}$$

The previous definition is motivated by the following lemma:

Lemma B.2. For $\underline{x} > 0$ and $W_0 \in \mathbb{R}^{3m+1}$, consider the open set

$$\mathcal{S}_{W_0}(\underline{x}) := \{W \in \mathbb{R}^{3m+1} \mid \forall i \in I : |b_i| < (|a_{i,0}| - |a_{i,0} - a_i|)\underline{x}\} .$$

The functions $f_{W,\tau,\sigma}$ are affine and for all $W \in \mathcal{S}_{W_0}(\underline{x})$ and all $x \in \mathbb{R}$ with $|x| \geq \underline{x}$, we have

$$f_W(x) = f_{W,\tau,\text{sgn}(x)}(x) .$$

If $\underline{x}_D > 0$, we have

$$L_D(W) = L_{D,\tau}(W), \quad \nabla L_D(W) = \nabla L_{D,\tau}(W)$$

for all $W \in \mathcal{S}_{W_0}(\underline{x}_D)$.

Proof Trivially, $f_{W,\tau,\sigma}$ is affine. Now, let $W \in \mathcal{S}_{W_0}(\underline{x})$, let $i \in I$ and let $|x| \geq \underline{x}$. We then obtain that $|a_{i,0}| - |a_{i,0} - a_i| > 0$ and therefore

$$|b_i| < (|a_{i,0}| - |a_{i,0} - a_i|)|x| .$$

Since $b_{i,0} = 0$, we have

$$|(a_i x + b_i) - (a_{i,0} x + b_{i,0})| \leq |a_i - a_{i,0}| \cdot |x| + |b_i| < |a_{i,0}| \cdot |x| = |a_{i,0} x + b_{i,0}| .$$

This shows $\text{sgn}(a_i x + b_i) = \text{sgn}(a_{i,0} x + b_{i,0})$, where

$$\text{sgn}(a_{i,0} x + b_{i,0}) = \text{sgn}(a_{i,0} x) = \text{sgn}(a_{i,0}) \text{sgn}(x) = \tau_i \text{sgn}(x) .$$

Due to our special choice of φ , we have

$$\begin{aligned} \varphi(a_i x + b_i) &= \varphi'(a_i x + b_i) \cdot (a_i x + b_i) = \varphi'(\text{sgn}(a_i x + b_i)) \cdot (a_i x + b_i) \\ &= \varphi'(\tau_i \text{sgn}(x)) \cdot (a_i x + b_i) , \end{aligned}$$

which yields

$$f_W(x) = f_{W,\tau,\text{sgn}(x)}(x) .$$

Since all data points x_j satisfy $|x_j| \geq \underline{x}_D$ by definition of \underline{x}_D , we find that

$$L_{D,\tau}(W) = \frac{1}{2n} \sum_{j \in J} (y_j - f_{W,\tau,\text{sgn}(x)}(x_j))^2 = \frac{1}{2n} \sum_{j \in J} (y_j - f_W(x_j))^2 = L_D(W) .$$

In addition, because $L_{D,\tau}$ and L_D are equal on the open set $\mathcal{S}_{W_0}(\underline{x}_D)$, their derivatives must also be equal on $\mathcal{S}_{W_0}(\underline{x}_D)$. ■

We can now define gradient descent iterates with respect to the “linearized” loss function $L_{D,\tau}$.

Definition B.3 (Gradient descent). *Given the random initial vector W_0 , its activation pattern τ and a step size $h > 0$, we recursively define*

$$W_{k+1} := W_k - h \nabla L_{D,\tau}(W_k) .$$

Moreover, we write $W_k = (\mathbf{a}_{\cdot,k}, \mathbf{b}_{\cdot,k}, c_k, \mathbf{w}_{\cdot,k})$ and we may implicitly omit the index k when deriving identities that hold for each $k \in \mathbb{N}_0$. For any derived quantity $\xi := g(W)$, define

$$\delta \xi := \delta g(W) := g(W - h \nabla L_{D,\tau}(W)) - g(W)$$

such that

$$\xi_{k+1} = g(W_{k+1}) = g(W_k) + (g(W_{k+1}) - g(W_k)) = \xi_k + \delta \xi_k$$

and hence

$$\delta g(W) = g(W + \delta W) - g(W) .$$

We can now write iteration rules differently: Instead of

$$W_{k+1} = W_k - h \nabla L_{D,\tau}(W_k) ,$$

we will use the more convenient notation

$$\delta W = -h \nabla L_{D,\tau}(W)$$

which suppresses the iteration index k and reads more like the negative gradient flow ODE

$$\dot{W} = -h \nabla L_{D,\tau}(W) .$$

The following lemma introduces some convenient rules for using δ .

Lemma B.4 (Differential calculus for δ). *Let $g : \mathbb{R}^{3m+1} \rightarrow \mathbb{R}^N$ for some $m, N \geq 1$.*

(a) *If g is linear, then $\delta g(W) = g(\delta W) = -hg(\nabla L_{D,\tau}(W))$.*

(b) *If g is constant, then $\delta g = 0$.*

(c) *If $g_1, g_2 : \mathbb{R}^{3m+1} \rightarrow \mathbb{R}$ are linear, then*

$$\delta(g_1 \cdot g_2) = (\delta g_1) \cdot g_2 + g_1 \cdot (\delta g_2) + (\delta g_1) \cdot (\delta g_2) .$$

(d) *If $g_1, g_2 : \mathbb{R}^{3m+1} \rightarrow \mathbb{R}^N$, then*

$$\delta(g_1 + g_2) = \delta g_1 + \delta g_2 .$$

(e) *If $g_2 : \mathbb{R}^{3m+1} \rightarrow \mathbb{R}^N, g_1 : \mathbb{R}^N \rightarrow \mathbb{R}^{N'}$ and g_1 is linear, then*

$$\delta(g_1 \circ g_2) = g_1 \circ (\delta g_2) .$$

(f) *If $g_1, \dots, g_N : \mathbb{R}^{3m+1} \rightarrow \mathbb{R}$, then*

$$\delta \begin{pmatrix} g_1 \\ \vdots \\ g_N \end{pmatrix} = \begin{pmatrix} \delta g_1 \\ \vdots \\ \delta g_N \end{pmatrix} .$$

Proof

(a) If g is linear, then

$$\delta g(W) = g(W + \delta W) - g(W) = g(\delta W) = g(-h\nabla L_{D,\tau}(W)) = -hg(\nabla L_{D,\tau}(W)) .$$

(b) Trivial.

(c) In this case,

$$\begin{aligned} \delta g(W) &= g(W + \delta W) - g(W) \\ &= g_1(W)g_2(\delta W) + g_1(\delta W)g_2(W) \\ &\quad + g_1(\delta W)g_2(\delta W) \\ &\stackrel{(a)}{=} \delta g_1(W)g_2(W) + g_1(W)\delta g_2(W) + \delta g_1(W)\delta g_2(W) . \end{aligned}$$

(d) We have

$$\begin{aligned} \delta(g_1 + g_2)(W) &= (g_1 + g_2)(W + \delta W) - (g_1 + g_2)(W) \\ &= (g_1(W + \delta W) - g_1(W)) + (g_2(W + \delta W) - g_2(W)) \\ &= \delta g_1(W) + \delta g_2(W) . \end{aligned}$$

(e) For $W \in \mathbb{R}^{3m+1}$,

$$\begin{aligned} \delta(g_1 \circ g_2)(W) &= g_1(g_2(W + \delta W)) - g_1(g_2(W)) = g_1(g_2(W + \delta W) - g_2(W)) \\ &= g_1(\delta g_2(W)) . \end{aligned}$$

(f) This follows from

$$\begin{pmatrix} g_1 \\ \vdots \\ g_N \end{pmatrix} (W + \delta W) - \begin{pmatrix} g_1 \\ \vdots \\ g_N \end{pmatrix} (W) = \begin{pmatrix} g_1(W + \delta W) - g_1(W) \\ \vdots \\ g_N(W + \delta W) - g_N(W) \end{pmatrix} . \quad \blacksquare$$

Appendix C. Reformulation of Gradient Descent

In this section, we will derive equations that describe how different aspects of the neural network behave during gradient descent. A summary and interpretation of the derived equations is presented in Section D.

Definition C.1 (Derived quantities).

- (a) For $\sigma \in \{\pm 1\}$, we write $\Sigma_{\sigma,a^2} := \sum_{i \in I_\sigma} a_i^2$, $\Sigma_{\sigma,wa} := \sum_{i \in I_\sigma} w_i a_i$ and so on.
 (b) The matrix $\mathbf{M}_\sigma := \mathbf{M}_{D,\sigma}$ from Definition 8 helps in relating different interesting quantities. For $\mathbf{M}_\sigma \succ 0$, let

$$\begin{aligned} \hat{\mathbf{v}}_\sigma &:= \begin{pmatrix} \hat{p}_\sigma \\ \hat{q}_\sigma \end{pmatrix} := \begin{pmatrix} \Sigma_{\sigma,wa} \\ \Sigma_{\sigma,wb} \end{pmatrix} & \mathbf{v}_\sigma &:= \begin{pmatrix} p_\sigma \\ q_\sigma \end{pmatrix} := \begin{pmatrix} \hat{p}_\sigma + \alpha \hat{p}_{-\sigma} \\ c + \hat{q}_\sigma + \alpha \hat{q}_{-\sigma} \end{pmatrix} \\ \hat{\mathbf{u}}_\sigma &:= \begin{pmatrix} \hat{r}_\sigma \\ \hat{s}_\sigma \end{pmatrix} := \begin{pmatrix} -\frac{1}{n} \sum_{j \in J_\sigma} (f_{W,\tau,\sigma}(x_j) - y_j) x_j \\ -\frac{1}{n} \sum_{j \in J_\sigma} (f_{W,\tau,\sigma}(x_j) - y_j) \end{pmatrix} & \mathbf{u}_\sigma &:= \begin{pmatrix} r_\sigma \\ s_\sigma \end{pmatrix} := \begin{pmatrix} \hat{r}_\sigma + \alpha \hat{r}_{-\sigma} \\ \hat{s}_\sigma + \alpha \hat{s}_{-\sigma} \end{pmatrix} \end{aligned}$$

and

$$\bar{\mathbf{v}}_\sigma := \mathbf{v}_\sigma - \mathbf{v}_\sigma^{\text{opt}} .$$

We will show in Lemma C.4 that $\hat{\mathbf{u}}_\sigma = -\mathbf{M}_\sigma \bar{\mathbf{v}}_\sigma$. The \mathbf{u} -vectors are interesting since their components can be used to simplify δW . As we will see in Lemma C.4, \mathbf{v}_σ is interesting since $f_{W,\tau,\sigma}(x) = p_\sigma x + q_\sigma$ for $x \in \mathbb{R}$. The notation of the different variants is motivated as follows: Expressions with a hat such as $\hat{\mathbf{v}}_\sigma$ and $\hat{\mathbf{u}}_\sigma$ only sum over one sign σ while hat-less expressions include both $\sigma = 1$ and $\sigma = -1$.

We will also use the matrices

$$\mathbf{G}_\sigma^{\text{w}} := \begin{pmatrix} \Sigma_{\sigma,w^2} & 0 \\ 0 & \Sigma_{\sigma,w^2} \end{pmatrix}, \quad \mathbf{G}_\sigma^{\text{ab}} := \begin{pmatrix} \Sigma_{\sigma,a^2} & \Sigma_{\sigma,ab} \\ \Sigma_{\sigma,ab} & \Sigma_{\sigma,b^2} \end{pmatrix}, \quad \mathbf{G}_\sigma^{\text{wab}} := (r_\sigma \Sigma_{\sigma,wa} + s_\sigma \Sigma_{\sigma,wb}) \mathbf{I}_2 ,$$

where \mathbf{I}_2 is the 2×2 identity matrix.

- (c) For any two vectors $\mathbf{z}_1, \mathbf{z}_{-1} \in \mathbb{R}^2$ defined in step (c) and any two matrices $\mathbf{F}_1, \mathbf{F}_{-1} \in \mathbb{R}^{2 \times 2}$ defined in step (b), we define

$$\tilde{\mathbf{z}} := \begin{pmatrix} \mathbf{z}_1 \\ \mathbf{z}_{-1} \end{pmatrix} \in \mathbb{R}^4, \quad \tilde{\mathbf{F}} := \begin{pmatrix} \mathbf{F}_1 & \\ & \mathbf{F}_{-1} \end{pmatrix} \in \mathbb{R}^{4 \times 4} .$$

For example, this means that

$$\tilde{\mathbf{u}} = \begin{pmatrix} \mathbf{u}_1 \\ \mathbf{u}_{-1} \end{pmatrix} = \begin{pmatrix} r_1 \\ s_1 \\ r_{-1} \\ s_{-1} \end{pmatrix} .$$

In addition, we define new matrices

$$\tilde{\mathbf{C}} := \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \end{pmatrix}, \quad \tilde{\mathbf{B}} := \begin{pmatrix} 1 & 0 & \alpha & 0 \\ 0 & 1 & 0 & \alpha \\ \alpha & 0 & 1 & 0 \\ 0 & \alpha & 0 & 1 \end{pmatrix} = \begin{pmatrix} \mathbf{I}_2 & \alpha \mathbf{I}_2 \\ \alpha \mathbf{I}_2 & \mathbf{I}_2 \end{pmatrix},$$

$$\tilde{\mathbf{A}} := \tilde{\mathbf{B}}(\tilde{\mathbf{G}}^{\text{w}} + \tilde{\mathbf{G}}^{\text{ab}} + h\tilde{\mathbf{G}}^{\text{wab}})\tilde{\mathbf{B}} + \tilde{\mathbf{C}} .$$

We will prove in Proposition C.5 that $\delta \tilde{\mathbf{v}} = h\tilde{\mathbf{A}}\tilde{\mathbf{u}} = -h\tilde{\mathbf{A}}\tilde{\mathbf{M}}\tilde{\mathbf{v}}$.

- (d) We want to perform a change of basis using the permutation matrix

$$\tilde{\mathbf{P}} := \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} .$$

which satisfies $\tilde{\mathbf{P}} = \tilde{\mathbf{P}}^\top = \tilde{\mathbf{P}}^{-1}$: For any vector $\tilde{\mathbf{z}} \in \mathbb{R}^4$ and any matrix $\tilde{\mathbf{F}} \in \mathbb{R}^{4 \times 4}$ defined in step (d), we define

$$\mathbf{z} := \tilde{\mathbf{P}}\tilde{\mathbf{z}}, \quad \mathbf{F} := \tilde{\mathbf{P}}\tilde{\mathbf{F}}\tilde{\mathbf{P}}^{-1} = \tilde{\mathbf{P}}\tilde{\mathbf{F}}\tilde{\mathbf{P}} .$$

For example, this yields

$$\mathbf{u} = \tilde{\mathbf{P}}\tilde{\mathbf{u}} = \begin{pmatrix} r_1 \\ r_{-1} \\ s_1 \\ s_{-1} \end{pmatrix}, \quad \mathbf{C} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{pmatrix}, \quad \mathbf{B} = \begin{pmatrix} 1 & \alpha & & \\ \alpha & 1 & & \\ & & 1 & \alpha \\ & & \alpha & 1 \end{pmatrix} =: \begin{pmatrix} \hat{\mathbf{B}} & \\ & \hat{\mathbf{B}} \end{pmatrix}.$$

We see that this change of basis by $\tilde{\mathbf{P}}$ makes the matrices $\tilde{\mathbf{B}}$ and $\tilde{\mathbf{C}}$ block-diagonal while it destroys the block-diagonal structure of $\tilde{\mathbf{G}}^{\text{ab}}$ and $\tilde{\mathbf{M}}$. We will see in Lemma F.3 that \mathbf{G}^{ab} is still block-diagonal at initialization. We will use the tilde quantities as an intermediate step to derive equations for the non-tilde quantities, since the latter will be more suitable for us to analyze eigenvectors and eigenvalues.

Elementary arguments show that

$$(\mathbf{M}_1 \succ 0 \text{ and } \mathbf{M}_{-1} \succ 0) \Leftrightarrow \tilde{\mathbf{M}} \succ 0 \Leftrightarrow \mathbf{M} = \tilde{\mathbf{P}}\tilde{\mathbf{M}}\tilde{\mathbf{P}}^\top \succ 0.$$

Therefore, we need to require $\mathbf{M} \succ 0$ so that \mathbf{v}^{opt} and $\bar{\mathbf{v}}$ can be defined.

- (e) Many of the quantities above depend on the data set D , which we may highlight later by indexing them with D . For example, we may write \mathbf{u}_D instead of \mathbf{u} .
- (f) Finally, let

$$\boldsymbol{\theta}_i := \begin{pmatrix} a_i \\ b_i \\ w_i \end{pmatrix}, \quad \boldsymbol{\Sigma}_\sigma := \sum_{i \in I_\sigma} \boldsymbol{\theta}_i \boldsymbol{\theta}_i^\top = \begin{pmatrix} \Sigma_{\sigma,a^2} & \Sigma_{\sigma,ab} & \Sigma_{\sigma,wa} \\ \Sigma_{\sigma,ab} & \Sigma_{\sigma,b^2} & \Sigma_{\sigma,wb} \\ \Sigma_{\sigma,wa} & \Sigma_{\sigma,wb} & \Sigma_{\sigma,w^2} \end{pmatrix}, \quad \mathbf{Q}_\sigma := \begin{pmatrix} 0 & 0 & r_\sigma \\ 0 & 0 & s_\sigma \\ r_\sigma & s_\sigma & 0 \end{pmatrix}.$$

These quantities will be analyzed in the next proposition.

Proposition C.2. For $i \in I_\sigma, \sigma \in \{\pm 1\}$, we have

$$\begin{aligned} \delta \boldsymbol{\theta}_i &= h \mathbf{Q}_\sigma \boldsymbol{\theta}_i \\ \delta c &= h(\hat{s}_1 + \hat{s}_{-1}) \\ \delta \boldsymbol{\Sigma}_\sigma &= h \mathbf{Q}_\sigma \boldsymbol{\Sigma}_\sigma + h \boldsymbol{\Sigma}_\sigma \mathbf{Q}_\sigma + h^2 \mathbf{Q}_\sigma \boldsymbol{\Sigma}_\sigma \mathbf{Q}_\sigma \end{aligned}$$

and the latter identity can also be written as

$$\boldsymbol{\Sigma}_{\sigma,k+1} = (\mathbf{I}_3 + h \mathbf{Q}_{\sigma,k}) \boldsymbol{\Sigma}_{\sigma,k} (\mathbf{I}_3 + h \mathbf{Q}_{\sigma,k}).$$

Proof The first two equations can also be written as

$$\begin{aligned} \delta a_i &= h r_\sigma w_i \\ \delta b_i &= h s_\sigma w_i \\ \delta w_i &= h r_\sigma a_i + h s_\sigma b_i \\ \delta c &= h(\hat{s}_1 + \hat{s}_{-1}). \end{aligned}$$

We will prove the first of these equations, the other ones follow similarly. Set $g(W) := a_i$. With Lemma B.4 (a), we obtain

$$\delta a_i = \delta g(W) = -h g(\nabla L_{D,\tau}(W)) = -h \frac{\partial L_{D,\tau}}{\partial a_i}(W)$$

$$\begin{aligned}
 &= -h \frac{1}{n} \sum_{j \in J} (f_{W, \tau, \text{sgn}(x_j)}(x_j) - y_j) \varphi'(\tau_i \cdot \text{sgn}(x_j)) w_i x_j \\
 &= -h \frac{1}{n} \left(\sum_{j \in J_\sigma} (f_{W, \tau, \sigma}(x_j) - y_j) w_i x_j + \alpha \sum_{j \in J_{-\sigma}} (f_{W, \tau, -\sigma}(x_j) - y_j) w_i x_j \right) \\
 &= h(\hat{r}_\sigma + \alpha \hat{r}_{-\sigma}) w_i = h r_\sigma w_i .
 \end{aligned}$$

Now for Σ_σ : Since $\mathbf{Q}_\sigma = \mathbf{Q}_\sigma^\top$, we have

$$\begin{aligned}
 \Sigma_{\sigma, k+1} &= \sum_{i \in I_\sigma} \theta_{i, k+1} \theta_{i, k+1}^\top = \sum_{i \in I_\sigma} (\mathbf{I}_3 + h \mathbf{Q}_{\sigma, k}) \theta_{i, k} \theta_{i, k}^\top (\mathbf{I}_3 + h \mathbf{Q}_{\sigma, k})^\top \\
 &= (\mathbf{I}_3 + h \mathbf{Q}_{\sigma, k}) \left(\sum_{i \in I_\sigma} \theta_{i, k} \theta_{i, k}^\top \right) (\mathbf{I}_3 + h \mathbf{Q}_{\sigma, k})^\top = (\mathbf{I}_3 + h \mathbf{Q}_{\sigma, k}) \Sigma_{\sigma, k} (\mathbf{I}_3 + h \mathbf{Q}_{\sigma, k}) ,
 \end{aligned}$$

which means that

$$\delta \Sigma_k = \Sigma_{k+1} - \Sigma_k = h \mathbf{Q}_{\sigma, k} \Sigma_{\sigma, k} + h \Sigma_{\sigma, k} \mathbf{Q}_{\sigma, k} + h^2 \mathbf{Q}_{\sigma, k} \Sigma_{\sigma, k} \mathbf{Q}_{\sigma, k} . \quad \blacksquare$$

Remark C.3. The term $h^2 \mathbf{Q}_\sigma \Sigma_\sigma \mathbf{Q}_\sigma$ in Proposition C.2 corresponds to the term $\delta g_1 \cdot \delta g_2$ in the “product rule” for δ (Lemma B.4 (c)). It vanishes when using negative gradient flow. In our case, it does not affect the qualitative behavior of gradient descent.

The following lemma shows relations between several quantities from Definition C.1.

Lemma C.4. Let $\mathbf{M} \succ 0$. For $\sigma \in \{\pm 1\}$ and $x \in \mathbb{R}$, we have

$$\begin{aligned}
 f_{W, \tau, \sigma}(x) &= p_\sigma x + q_\sigma \\
 \hat{\mathbf{u}}_\sigma &= -\mathbf{M}_\sigma \bar{\mathbf{v}}_\sigma .
 \end{aligned}$$

Moreover,

$$\tilde{\mathbf{u}} = \tilde{\mathbf{B}} \tilde{\mathbf{u}}, \quad \tilde{\hat{\mathbf{u}}} = -\tilde{\mathbf{M}} \tilde{\mathbf{v}}, \quad \tilde{\mathbf{v}} = \tilde{\mathbf{B}} \tilde{\mathbf{v}} + \begin{pmatrix} 0 \\ c \\ 0 \\ c \end{pmatrix} .$$

Proof For $x \in \mathbb{R}$,

$$\begin{aligned}
 f_{W, \tau, \sigma}(x) &= c + \sum_{i \in I} w_i \varphi'(\tau_i \sigma) (a_i x + b_i) \\
 &= c + \sum_{i \in I_\sigma} (w_i a_i x + w_i b_i) + \alpha \sum_{i \in I_{-\sigma}} (w_i a_i x + w_i b_i) = p_\sigma x + q_\sigma .
 \end{aligned}$$

Therefore, using Definition 8,

$$\hat{\mathbf{u}}_\sigma = \begin{pmatrix} \hat{r}_\sigma \\ \hat{s}_\sigma \end{pmatrix} = -\frac{1}{n} \begin{pmatrix} \sum_{j \in J_\sigma} (f_{W, \tau, \sigma}(x_j) - y_j) x_j \\ \sum_{j \in J_\sigma} (f_{W, \tau, \sigma}(x_j) - y_j) \end{pmatrix}$$

$$\begin{aligned}
 &= -\frac{1}{n} \left(\frac{\sum_{j \in J_\sigma} (p_\sigma x_j + q_\sigma - y_j) x_j}{\sum_{j \in J_\sigma} (p_\sigma x_j + q_\sigma - y_j)} \right) \\
 &= -\frac{1}{n} \left(\frac{p_\sigma \sum_{j \in J_\sigma} x_j^2 + q_\sigma \sum_{j \in J_\sigma} x_j - \sum_{j \in J_\sigma} x_j y_j}{p_\sigma \sum_{j \in J_\sigma} x_j + q_\sigma \sum_{j \in J_\sigma} 1 - \sum_{j \in J_\sigma} y_j} \right) \\
 &= -\mathbf{M}_\sigma \begin{pmatrix} p_\sigma \\ q_\sigma \end{pmatrix} + \frac{1}{n} \sum_{j \in J_\sigma} \hat{\mathbf{u}}_{(x_j, y_j)}^0 = -\mathbf{M}_\sigma \mathbf{v}_\sigma + \hat{\mathbf{u}}_\sigma^0 = -\mathbf{M}_\sigma (\mathbf{v}_\sigma - \mathbf{v}_\sigma^{\text{opt}}) = -\mathbf{M}_\sigma \bar{\mathbf{v}}_\sigma .
 \end{aligned}$$

We now obtain

$$\begin{aligned}
 \tilde{\mathbf{u}} &= \begin{pmatrix} r_1 \\ s_1 \\ r_{-1} \\ s_{-1} \end{pmatrix} = \begin{pmatrix} \hat{r}_1 + \alpha \hat{r}_{-1} \\ \hat{s}_1 + \alpha \hat{s}_{-1} \\ \hat{r}_{-1} + \alpha \hat{r}_1 \\ \hat{s}_{-1} + \alpha \hat{s}_1 \end{pmatrix} = \tilde{\mathbf{B}} \tilde{\mathbf{u}} \\
 \tilde{\hat{\mathbf{u}}} &= \begin{pmatrix} \hat{\mathbf{u}}_1 \\ \hat{\mathbf{u}}_{-1} \end{pmatrix} = \begin{pmatrix} \mathbf{M}_1 & \\ & \mathbf{M}_{-1} \end{pmatrix} \begin{pmatrix} \bar{\mathbf{v}}_1 \\ \bar{\mathbf{v}}_{-1} \end{pmatrix} = \tilde{\mathbf{M}} \tilde{\mathbf{v}} \\
 \tilde{\mathbf{v}} &= \begin{pmatrix} p_1 \\ q_1 \\ p_{-1} \\ q_{-1} \end{pmatrix} = \begin{pmatrix} \hat{p}_1 + \alpha \hat{p}_{-1} \\ c + \hat{q}_1 + \alpha \hat{q}_{-1} \\ \hat{p}_{-1} + \alpha \hat{p}_1 \\ c + \hat{q}_{-1} + \alpha \hat{q}_1 \end{pmatrix} = \tilde{\mathbf{B}} \tilde{\mathbf{v}} + \begin{pmatrix} 0 \\ c \\ 0 \\ c \end{pmatrix} .
 \end{aligned}$$

This enables us to compute another iteration equation:

Proposition C.5. *Let $\mathbf{M} \succ 0$. Then,*

$$\begin{aligned}
 \delta \tilde{\mathbf{v}} &= -h \tilde{\mathbf{A}} \tilde{\mathbf{M}} \tilde{\mathbf{v}} = -h (\tilde{\mathbf{B}} (\tilde{\mathbf{G}}^w + \tilde{\mathbf{G}}^{\text{ab}} + h \tilde{\mathbf{G}}^{\text{wab}}) \tilde{\mathbf{B}} + \tilde{\mathbf{C}}) \tilde{\mathbf{M}} \tilde{\mathbf{v}} \\
 \delta \bar{\mathbf{v}} &= -h \mathbf{A} \mathbf{M} \bar{\mathbf{v}} = -h (\mathbf{B} (\mathbf{G}^w + \mathbf{G}^{\text{ab}} + h \mathbf{G}^{\text{wab}}) \mathbf{B} + \mathbf{C}) \mathbf{M} \bar{\mathbf{v}} .
 \end{aligned}$$

Hence,

$$\bar{\mathbf{v}}_{k+1} = \bar{\mathbf{v}}_k + \delta \bar{\mathbf{v}}_k = (\mathbf{I}_4 - h \mathbf{A}_k \mathbf{M}) \bar{\mathbf{v}}_k .$$

Proof Consider

$$\hat{\mathbf{v}}_\sigma = \begin{pmatrix} \hat{p}_\sigma \\ \hat{q}_\sigma \end{pmatrix} = \begin{pmatrix} \Sigma_{\sigma, wa} \\ \Sigma_{\sigma, wb} \end{pmatrix} = (\mathbf{I}_2 \quad 0) \Sigma_\sigma \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} .$$

Using Proposition C.2 and Lemma B.4, we obtain

$$\begin{aligned}
 \delta \hat{\mathbf{v}}_\sigma &= (\mathbf{I}_2 \quad 0) (h \mathbf{Q}_\sigma \Sigma_\sigma + h \Sigma_\sigma \mathbf{Q}_\sigma + h^2 \mathbf{Q}_\sigma \Sigma_\sigma \mathbf{Q}_\sigma) \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \\
 &= h \begin{pmatrix} 0 & 0 & r_\sigma \\ 0 & 0 & s_\sigma \end{pmatrix} \begin{pmatrix} \Sigma_{\sigma, wa} \\ \Sigma_{\sigma, wb} \\ \Sigma_{\sigma, w^2} \end{pmatrix} + h \begin{pmatrix} \Sigma_{\sigma, a^2} & \Sigma_{\sigma, ab} & \Sigma_{\sigma, wa} \\ \Sigma_{\sigma, ab} & \Sigma_{\sigma, b^2} & \Sigma_{\sigma, wb} \end{pmatrix} \begin{pmatrix} r_\sigma \\ s_\sigma \\ 0 \end{pmatrix}
 \end{aligned}$$

$$\begin{aligned}
 & + h^2 \begin{pmatrix} 0 & 0 & r_\sigma \\ 0 & 0 & s_\sigma \end{pmatrix} \begin{pmatrix} \Sigma_{\sigma,a^2} & \Sigma_{\sigma,ab} & \Sigma_{\sigma,wa} \\ \Sigma_{\sigma,ab} & \Sigma_{\sigma,b^2} & \Sigma_{\sigma,wb} \\ \Sigma_{\sigma,wa} & \Sigma_{\sigma,wb} & \Sigma_{\sigma,w^2} \end{pmatrix} \begin{pmatrix} r_\sigma \\ s_\sigma \\ 0 \end{pmatrix} \\
 & = h \begin{pmatrix} \Sigma_{\sigma,w^2} & 0 \\ 0 & \Sigma_{\sigma,w^2} \end{pmatrix} \mathbf{u}_\sigma + h \begin{pmatrix} \Sigma_{\sigma,a^2} & \Sigma_{\sigma,ab} \\ \Sigma_{\sigma,ab} & \Sigma_{\sigma,b^2} \end{pmatrix} \mathbf{u}_\sigma + h^2 (r_\sigma \Sigma_{\sigma,wa} + s_\sigma \Sigma_{\sigma,wb}) \mathbf{u}_\sigma \\
 & = h \left(\mathbf{G}_\sigma^w + \mathbf{G}_\sigma^{ab} + h \mathbf{G}_\sigma^{wab} \right) \mathbf{u}_\sigma .
 \end{aligned}$$

Therefore, $\delta \tilde{\mathbf{v}} = h \left(\tilde{\mathbf{G}}^w + \tilde{\mathbf{G}}^{ab} + h \tilde{\mathbf{G}}^{wab} \right) \tilde{\mathbf{u}}$. Also,

$$\delta \begin{pmatrix} 0 \\ c \\ 0 \\ c \end{pmatrix} = h \begin{pmatrix} 0 \\ \hat{s}_1 + \hat{s}_{-1} \\ 0 \\ \hat{s}_1 + \hat{s}_{-1} \end{pmatrix} = h \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \end{pmatrix} \begin{pmatrix} \hat{r}_1 \\ \hat{s}_1 \\ \hat{r}_{-1} \\ \hat{s}_{-1} \end{pmatrix} = h \tilde{\mathbf{C}} \tilde{\mathbf{u}} .$$

We can now use the identities from Lemma C.4 and the fact that $\tilde{\mathbf{v}} - \tilde{\tilde{\mathbf{v}}} = \tilde{\mathbf{v}}^{\text{opt}}$ is constant to compute

$$\begin{aligned}
 \delta \tilde{\tilde{\mathbf{v}}} & = \delta \tilde{\mathbf{v}} = \tilde{\mathbf{B}} \delta \tilde{\mathbf{v}} + \delta \begin{pmatrix} 0 \\ c \\ 0 \\ c \end{pmatrix} = \tilde{\mathbf{B}} h \left(\tilde{\mathbf{G}}^w + \tilde{\mathbf{G}}^{ab} + h \tilde{\mathbf{G}}^{wab} \right) \tilde{\mathbf{u}} + h \tilde{\mathbf{C}} \tilde{\mathbf{u}} \\
 & = h \left(\tilde{\mathbf{B}} \left(\tilde{\mathbf{G}}^w + \tilde{\mathbf{G}}^{ab} + h \tilde{\mathbf{G}}^{wab} \right) \tilde{\mathbf{B}} + \tilde{\mathbf{C}} \right) \tilde{\mathbf{u}} = h \tilde{\mathbf{A}} \tilde{\mathbf{u}} = -h \tilde{\mathbf{A}} \tilde{\mathbf{M}} \tilde{\mathbf{v}} .
 \end{aligned}$$

Since $\tilde{\mathbf{P}}^2 = \mathbf{I}_4$, it follows that

$$\delta \bar{\mathbf{v}} = \delta (\tilde{\mathbf{P}} \tilde{\mathbf{v}}) = \tilde{\mathbf{P}} \delta \tilde{\mathbf{v}} = -h \tilde{\mathbf{P}} \tilde{\mathbf{A}} \tilde{\mathbf{M}} \tilde{\mathbf{v}} = -h \tilde{\mathbf{P}} \tilde{\mathbf{A}} \tilde{\mathbf{P}} \tilde{\mathbf{P}} \tilde{\mathbf{M}} \tilde{\mathbf{P}} \tilde{\mathbf{v}} = -h \tilde{\mathbf{A}} \tilde{\mathbf{M}} \bar{\mathbf{v}}$$

and

$$\begin{aligned}
 \mathbf{A} & = \tilde{\mathbf{P}} \tilde{\mathbf{A}} \tilde{\mathbf{P}} = \tilde{\mathbf{P}} \left(\tilde{\mathbf{B}} \tilde{\mathbf{P}} \tilde{\mathbf{P}} \left(\tilde{\mathbf{G}}^w + \tilde{\mathbf{G}}^{ab} + h \tilde{\mathbf{G}}^{wab} \right) \tilde{\mathbf{P}} \tilde{\mathbf{P}} \tilde{\mathbf{B}} + \tilde{\mathbf{C}} \right) \tilde{\mathbf{P}} \\
 & = \mathbf{B} \left(\mathbf{G}^w + \mathbf{G}^{ab} + h \mathbf{G}^{wab} \right) \mathbf{B} + \mathbf{C} . \quad \blacksquare
 \end{aligned}$$

Appendix D. Comments

In this section, we provide some remarks on the interpretation of the gradient descent equations derived in Section C.

Remark D.1 (System decomposition). *We have so far derived different “systems”, i.e., results on how quantities evolve during gradient descent. These systems and their dependencies are depicted in Figure D.1. In particular, we see that the systems for Σ_1, Σ_{-1} and $\bar{\mathbf{v}}$ together yield a 22-dimensional system that does not depend on any other quantities. This 22-dimensional system describes some central properties of the neural network parameters W although its dimension does not depend on m . These properties include:*

- Slope p_σ and intercept q_σ for both signs $\sigma \in \{\pm 1\}$.

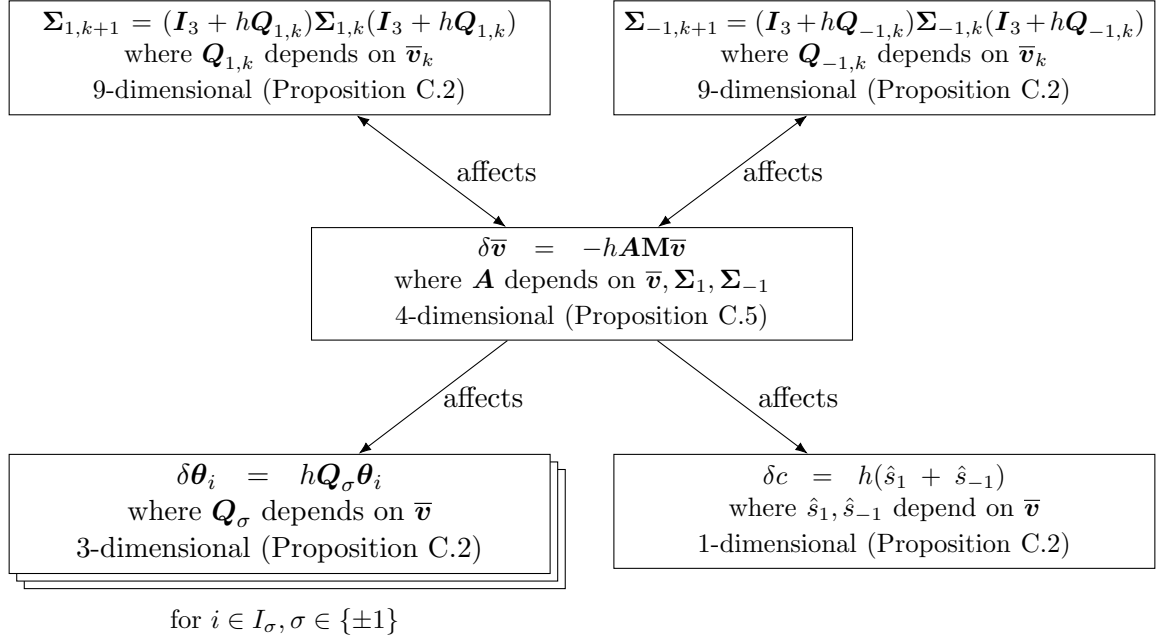


Figure D.1: Decomposition into different systems that can be used to analyze the behavior of gradient descent.

- The loss $L_{D,\tau}(W)$, which can be computed from p_σ and q_σ .

While this system has a dimension independent of m , the probability distribution over its initialization may well depend on m . If its evolution is known, the evolution $(W_k)_{k \in \mathbb{N}_0}$ can be determined by solving m independent three-dimensional systems and the one-dimensional system $\delta c = h(\hat{s}_1 + \hat{s}_{-1})$. Here, we will proceed along similar lines: We will first analyze the behavior of the 22-dimensional system and then apply our results to the three-dimensional systems.

In fact, the 22-dimensional system can be reduced to a 14-dimensional system: The matrices Σ_σ are always symmetric and thus effectively 6-dimensional, which reduces the dimension from 22 to 16. Moreover, we always have

$$\begin{pmatrix} p_1 \\ p_{-1} \end{pmatrix} = \begin{pmatrix} 1 & \alpha \\ \alpha & 1 \end{pmatrix} \begin{pmatrix} \Sigma_{1,wa} \\ \Sigma_{-1,wa} \end{pmatrix}.$$

However, removing these redundancies is not beneficial for our analysis.

Remark D.2. The components of the equation $\delta \bar{\mathbf{v}} = -h\mathbf{A}\mathbf{M}\bar{\mathbf{v}}$ in Proposition C.5 can be interpreted as follows: Recall that

$$\mathbf{G}_\sigma^w = \begin{pmatrix} \Sigma_{\sigma,w^2} & 0 \\ 0 & \Sigma_{\sigma,w^2} \end{pmatrix}, \quad \mathbf{G}_\sigma^{\text{ab}} = \begin{pmatrix} \Sigma_{\sigma,a^2} & \Sigma_{\sigma,ab} \\ \Sigma_{\sigma,ab} & \Sigma_{\sigma,b^2} \end{pmatrix}, \quad \mathbf{G}_\sigma^{\text{wab}} = (r_\sigma \Sigma_{\sigma,wa} + s_\sigma \Sigma_{\sigma,wb}) \mathbf{I}_2,$$

$$\tilde{\mathbf{B}} = \begin{pmatrix} \mathbf{I}_2 & \alpha \mathbf{I}_2 \\ \alpha \mathbf{I}_2 & \mathbf{I}_2 \end{pmatrix}, \quad \tilde{\mathbf{C}} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \end{pmatrix},$$

$$\tilde{\mathbf{A}} = \tilde{\mathbf{B}} \begin{pmatrix} \mathbf{G}_1^w + \mathbf{G}_1^{\text{ab}} + h\mathbf{G}_1^{\text{wab}} & \\ & \mathbf{G}_{-1}^w + \mathbf{G}_{-1}^{\text{ab}} + h\mathbf{G}_{-1}^{\text{wab}} \end{pmatrix} \tilde{\mathbf{B}} + \tilde{\mathbf{C}}.$$

- The matrix $\mathbf{G}_\sigma^w \succeq 0$ describes the improvement of \bar{v}_σ by updating the weights $(a_i)_{i \in I_\sigma}$ and $(b_i)_{i \in I_\sigma}$. The larger $|w_i|$, the larger the gradients $\frac{\partial L_{D,\tau}}{\partial a_i}$, $\frac{\partial L_{D,\tau}}{\partial b_i}$ and the more effect does a change in a_i, b_i have on the overall function $f_{W,\tau,\sigma}$.
- The matrix $\mathbf{G}_\sigma^{\text{ab}}$ is also positive semidefinite since $\text{tr}(\mathbf{G}_\sigma^{\text{ab}}) \geq 0$ and $\det(\mathbf{G}_\sigma^{\text{ab}}) = \Sigma_{\sigma,a^2} \Sigma_{\sigma,b^2} - \Sigma_{\sigma,ab}^2 \geq 0$ due to Cauchy-Schwarz. It describes the improvement of \bar{v}_σ by updating the weights $(w_i)_{i \in I_\sigma}$. Larger values of $|a_i|, |b_i|$ mean stronger effects of changing w_i . If the vectors $(a_i)_{i \in I_\sigma}$ and $(b_i)_{i \in I_\sigma}$ are linearly dependent (perfectly correlated), e.g. at initialization because of $b_{i,0} = 0$, then $\tilde{\mathbf{G}}_\sigma^{\text{ab}}$ only has rank one and changing the w_i cannot independently update both components of \bar{v}_σ . Recall that the components of \bar{v}_σ are the differences of the slope and intercept of $f_{W,\tau,\sigma}$ to the optimal linear regression slope and intercept, respectively.
- The matrix $\tilde{\mathbf{B}}$ causes an interaction between both signs $\sigma \in \{\pm 1\}$ if the leaky parameter α is nonzero. If it is zero, the hidden neurons are only active for one sign σ and do only interact indirectly via the bias c .
- The matrix $\tilde{\mathbf{C}}$ describes the improvement of v_σ by updating the bias c . It is not block-diagonal since c is active for both signs $\sigma \in \{\pm 1\}$. However, $\tilde{\mathbf{C}}$ only has rank one since changing c can only change q_1 and q_{-1} by the same amount. $\tilde{\mathbf{C}}$ is positive semidefinite since it is symmetric and it has eigenvectors $\mathbf{e}_1, \mathbf{e}_3, (0, 1, 0, -1)$ to the eigenvalue 0 and $(0, 1, 0, 1)$ to the eigenvalue 2.
- The matrix $\tilde{\mathbf{G}}^{\text{wab}}$ represents parts of the error that (discrete) gradient descent makes when trying to approximate negative gradient flow. It arises from the additional term $\delta g_1 \cdot \delta g_2$ in the product rule for δ (Lemma B.4 (c)) and does not need to be positive semidefinite. If h is too large, the matrix $\tilde{\mathbf{A}}$ might therefore not be positive semidefinite.

Remark D.3 (Discretization error). We have already seen that the systems for Σ_σ and \bar{v} are affected by terms that arise from the term $\delta g_1 \cdot \delta g_2$ in the “discrete product rule” of Lemma B.4 (c). We will see that in our scenario (with small enough step size), these “disturbances” are small enough to not influence the qualitative behavior of gradient descent. There is also an invariant that holds when using negative gradient flow but breaks down when using gradient descent: In the former case, $a_i^2 + b_i^2 - w_i^2$ remains constant during the optimization for each $i \in I$. An analogous identity for linear networks has been observed by Saxe et al. (2014).

Remark D.4 (Alternative systems). In some special cases, the approach presented here only works if we modify the systems. For example, the assumption $\mathbf{M} \succ 0$ is not satisfied if the data set is contained in $(0, \infty)$ since this implies $\mathbf{M}_{-1} = 0$. In this case, the system $\delta \bar{v} = -h\mathbf{A}\mathbf{M}\bar{v}$ could be reduced to a two-dimensional system since p_{-1} and q_{-1} are irrelevant for the loss. We will also see that the argument here does not work for $|\alpha| = 1$ since this renders the matrix \mathbf{B} singular. The case $\alpha = 1$ corresponds to a linear activation function

$\varphi(x) = x$, which implies $p_1 = p_{-1}$ and $q_1 = q_{-1}$. Similarly, the case $\alpha = -1$ corresponds to $\varphi(x) = |x|$, which implies $p_1 = -p_{-1}$. In both cases, the dimension of \mathbf{v} could be reduced.

Remark D.5 (Calculations for Remark 9). *To verify the assertions made in Remark 9, we first note that $\mathbf{M}_{D,\sigma}$ is given by*

$$\mathbf{M}_{D,\sigma} := \frac{n_\sigma}{n} \begin{pmatrix} \alpha & \beta \\ \beta & 1 \end{pmatrix}, \quad (7)$$

where

$$\begin{aligned} \alpha &:= \frac{1}{n_\sigma} \sum_{(x,y) \in D_\sigma} x^2 = \text{Var } D_{X,\sigma} + (\mathbb{E} D_{X,\sigma})^2, \\ \beta &:= \frac{1}{n_\sigma} \sum_{(x,y) \in D_\sigma} x = \mathbb{E} D_{X,\sigma}. \end{aligned}$$

Consequently, we have

$$\det(\mathbf{M}_{D,\sigma}) = \left(\frac{n_\sigma}{n}\right)^2 (\alpha - \beta^2) = \left(\frac{n_\sigma}{n}\right)^2 \text{Var } D_{X,\sigma},$$

and since $\text{Var } D_{X,\sigma} \neq 0$ if and only if $D_{X,\sigma}$ contains at least two distinct samples, it follows that $\mathbf{M}_{D,\sigma}$ is invertible if and only if D_σ contains at least two samples with different x values.

Now assume that $\text{Var } D_{X,\sigma} > 0$. Since $\det(\mathbf{M}_{D,\sigma})$ and $\text{tr}(\mathbf{M}_{D,\sigma})$ are positive, the minimum and maximum eigenvalues $\lambda_{\min}, \lambda_{\max}$ of $\mathbf{M}_{D,\sigma}$ are both positive. Thus,

$$\begin{aligned} \lambda_{\max} &\leq \lambda_{\min} + \lambda_{\max} = \text{tr}(\mathbf{M}_{D,\sigma}) = \frac{n_\sigma}{n}(\alpha + 1), \\ \lambda_{\min} &= \frac{\lambda_{\min}\lambda_{\max}}{\lambda_{\max}} \geq \frac{\lambda_{\min}\lambda_{\max}}{\lambda_{\min} + \lambda_{\max}} = \frac{\det(\mathbf{M}_{D,\sigma})}{\text{tr}(\mathbf{M}_{D,\sigma})} = \frac{\left(\frac{n_\sigma}{n}\right)^2 (\alpha - \beta^2)}{\left(\frac{n_\sigma}{n}\right) (\alpha + 1)} = \frac{n_\sigma}{n} \cdot \frac{\text{Var } D_{X,\sigma}}{\alpha + 1}. \end{aligned}$$

Since $\frac{1}{2}(\lambda_{\min} + \lambda_{\max}) \leq \lambda_{\max} \leq \lambda_{\min} + \lambda_{\max}$, it is easy to see that the bounds are off by a factor of at most two.

Remark D.6 (Affine regression optimum). *Here, we review some well-known properties of performing affine least-squares regression on a data set D . This analysis also applies to D_1 and D_{-1} . Consider*

$$\mathbf{X} := \begin{pmatrix} x_1 & 1 \\ \vdots & \vdots \\ x_n & 1 \end{pmatrix}, \quad \mathbf{y} := \begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix}, \quad \mathbf{v} := \begin{pmatrix} p \\ q \end{pmatrix}.$$

We always have $\mathbf{X}^\top \mathbf{X} \succeq 0$. Assume that \mathbf{X} has full column rank such that $\mathbf{X}^\top \mathbf{X} \succ 0$. The least-squares risk of an affine function $f_{\mathbf{v}}(x) = px + q$ is $R_D(f_{\mathbf{v}})$ with

$$2nR_D(f_{\mathbf{v}}) = \sum_{j=1}^n (y_j - f_{\mathbf{v}}(x_j))^2 = \|\mathbf{y} - \mathbf{X}\mathbf{v}\|_2^2 = (\mathbf{y} - \mathbf{X}\mathbf{v})^\top (\mathbf{y} - \mathbf{X}\mathbf{v})$$

$$\begin{aligned}
 &= \left(\mathbf{v} - (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y} \right)^\top \mathbf{X}^\top \mathbf{X} \left(\mathbf{v} - (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y} \right) \\
 &\quad + \left(\mathbf{y}^\top \mathbf{y} - \mathbf{y}^\top \mathbf{X} (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y} \right), \tag{8}
 \end{aligned}$$

where we performed a completion of the square. Therefore, the optimal affine function has parameters

$$\begin{aligned}
 \mathbf{v}^{\text{opt}} &= (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y} = \begin{pmatrix} \sum_j x_j^2 & \sum_j x_j \\ \sum_j x_j & \sum_j 1 \end{pmatrix}^{-1} \begin{pmatrix} \sum_j x_j y_j \\ \sum_j y_j \end{pmatrix} \\
 &= \left(\frac{1}{n} \begin{pmatrix} \sum_j x_j^2 & \sum_j x_j \\ \sum_j x_j & \sum_j 1 \end{pmatrix} \right)^{-1} \left(\frac{1}{n} \begin{pmatrix} \sum_j x_j y_j \\ \sum_j y_j \end{pmatrix} \right).
 \end{aligned}$$

Note that applying Eq. (8) to D_1 and D_{-1} after some rearrangement yields

$$L_{D,\tau}(W) = \frac{1}{2} (\mathbf{v} - \mathbf{v}_D^{\text{opt}})^\top \mathbf{M}_D (\mathbf{v} - \mathbf{v}_D^{\text{opt}}) + \text{const} = \frac{1}{2} \bar{\mathbf{v}}^\top \mathbf{M}_D \bar{\mathbf{v}} + \text{const},$$

where the constant term is the optimal achievable loss by affine regression on D_1 and D_{-1} .

Appendix E. Stochastic Proofs

In this section, we show that W_0 and D likely have certain properties. The results are formulated in Proposition E.3 and Proposition E.5, respectively. In order to obtain these results, we employ concentration inequalities. Besides Markov's inequality, we use Hoeffding's inequality:

Lemma E.1 (Hoeffding's inequality, e.g. Lemma A.3 in Györfi et al. (2002)). *Let (Ω, \mathcal{F}, P) be a probability space, $a < b, n \geq 1$ and $X_1, \dots, X_n : \Omega \rightarrow [a, b]$ be independent random variables. Then, for $\tau \geq 0$, we have*

$$P \left(\left| \frac{1}{n} \sum_{i=1}^n (X_i - \mathbb{E} X_i) \right| \geq (b-a) \sqrt{\frac{\tau}{2n}} \right) \leq 2e^{-\tau}.$$

Using Markov and Hoeffding, we can prove an asymptotic concentration result. The intuition behind this result is that for random variables X_1, \dots, X_n with mean zero and finite variance, the value $n^{-1/2}(X_1 + \dots + X_n)$ asymptotically has a Gaussian distribution by the central limit theorem. The tail of the Gaussian distribution decreases stronger than any inverse polynomial: If Φ is the CDF of a Gaussian distribution, then $\Phi(\beta n^\varepsilon) \leq O(n^{-\gamma})$ for all $\beta, \varepsilon, \gamma > 0$, where the constant in $O(n^{-\gamma})$ depends on $\beta, \varepsilon, \gamma$. However, the central limit theorem does not tell us how close the CDF of $n^{-1/2}(X_1 + \dots + X_n)$ is to Φ , so we use Markov's and Hoeffding's inequalities instead.

Lemma E.2. *Let Q be a probability distribution on \mathbb{R} with $\mu_p := \int |x|^p dQ(x) < \infty$ for all $p \in (1, \infty)$. For $n \in \mathbb{N}$, let $(\Omega_n, \mathcal{F}_n, P_n)$ be probability spaces with independent Q -distributed random variables $X_{n1}, X_{n2}, \dots, X_{nn} : \Omega_n \rightarrow \mathbb{R}$. Then, the random variables $S_n := \frac{1}{n} \sum_{i=1}^n X_{ni}$ satisfy*

$$P_n \left(|S_n - \mathbb{E} S_n| \geq \beta n^{\varepsilon-1/2} \right) \leq O(n^{-\gamma})$$

for all $\beta, \varepsilon, \gamma > 0$, where the constant in $O(n^{-\gamma})$ may depend on $\beta, \varepsilon, \gamma$ (cf. Definition A.2).

Proof Let $\beta, \varepsilon, \gamma > 0$ be fixed. For $n \in \mathbb{N}$ and $b > 0$ to be determined later, define $B := \{\max_{1 \leq i \leq n} |X_{ni}| \leq b\}$. Then, for all $p > 0$,

$$P_n(B^c) \leq \sum_{i=1}^n P_n(|X_{ni}| \geq b) \leq n P_n(|X_{n1}|^p \geq b^p) \stackrel{\text{Markov}}{\leq} n \frac{\mathbb{E}_{P_n} |X_{n1}|^p}{b^p} = n \frac{\mu_p}{b^p}. \quad (9)$$

Since $S_n = S_n \mathbf{1}_B + S_n \mathbf{1}_{B^c}$, we can now bound

$$\begin{aligned} P_n(|S_n - \mathbb{E}S_n| \geq \beta n^{\varepsilon-1/2}) &\leq \underbrace{P_n(|S_n \mathbf{1}_B - \mathbb{E}(S_n \mathbf{1}_B)| \geq \beta n^{\varepsilon-1/2}/2)}_{\text{I}} \\ &\quad + \underbrace{P_n(|S_n \mathbf{1}_{B^c} - \mathbb{E}(S_n \mathbf{1}_{B^c})| \geq \beta n^{\varepsilon-1/2}/2)}_{\text{II}}. \end{aligned}$$

With $\tau := \gamma \log n$ and $b := \beta n^\varepsilon \sqrt{\frac{1}{8\gamma \log n}}$, we have

$$(b - (-b)) \sqrt{\frac{\tau}{2n}} = 2\beta n^\varepsilon \sqrt{\frac{1}{8\gamma \log n}} \cdot \sqrt{\frac{\gamma \log n}{2n}} = \beta n^{\varepsilon-1/2}/2$$

and hence, Hoeffding (Lemma E.1) applied to $X_i := X_{ni} \mathbf{1}_{|X_{ni}| \leq b}$ yields

$$\text{I} \leq 2e^{-\tau} = 2n^{-\gamma}.$$

Moreover, we have

$$\begin{aligned} |\mathbb{E}_{P_n}(S_n \mathbf{1}_{B^c})| &\leq \|S_n \mathbf{1}_{B^c}\|_{\mathcal{L}_1(P_n)} \stackrel{\text{Hölder}}{\leq} \|S_n\|_{\mathcal{L}_2(P_n)} \|\mathbf{1}_{B^c}\|_{\mathcal{L}_2(P_n)} \\ &\leq \left(\frac{1}{n} \sum_{i=1}^n \|X_{ni}\|_{\mathcal{L}_2(P_n)} \right) \|\mathbf{1}_{B^c}\|_{\mathcal{L}_2(P_n)} = \sqrt{\mu_2} \sqrt{P_n(B^c)} \\ &\stackrel{(9)}{\leq} \sqrt{\mu_2} \sqrt{n \frac{\mu_p}{b^p}} = \sqrt{\frac{\mu_2 \mu_p}{\beta^p}} (8\gamma \log n)^{p/4} n^{(1-\varepsilon p)/2}. \end{aligned}$$

If we choose $p \geq 2/\varepsilon$, we have $(1 - \varepsilon p)/2 \leq -1/2 < \varepsilon - 1/2$ and hence $|\mathbb{E}(S_n \mathbf{1}_{B^c})| < \beta n^{\varepsilon-1/2}/2$ for n large enough. Now, let n be sufficiently large. For $\omega \in B$, we have $S_n(\omega) \mathbf{1}_{B^c}(\omega) = 0$ and hence $|S_n(\omega) \mathbf{1}_{B^c}(\omega) - \mathbb{E}(S_n \mathbf{1}_{B^c})| < \beta n^{\varepsilon-1/2}/2$. Thus,

$$\text{II} \leq P(B^c) \leq n \frac{\mu_p}{b^p} = \frac{\mu_p}{\beta^p} \cdot (8\gamma \log n)^{p/2} n^{1-\varepsilon p}.$$

If we choose $p > (1 + \gamma)/\varepsilon$, then $1 - \varepsilon p < -\gamma$ and hence $\text{II} \leq O(n^{-\gamma})$. ■

Now, we can prove that certain properties of the initialization W_0 hold with high probability. We will see that in all properties except (W4), the tails of the probability distributions decrease so quickly that only the parameter γ_P in (W4) is relevant for the rate of convergence.

Proposition E.3. *Let $\varepsilon, \gamma_P > 0$ and let W_0 be distributed as in Assumption 2. Then, the properties*

- (W1) $b_{i,0} = c_0 = 0$,
 (W2) $\max_i |w_{i,0}| \leq m^{-1/2+\varepsilon}$,
 (W3) $\max_i |a_{i,0}| \leq m^\varepsilon$,
 (W4) $\min_i |a_{i,0}| \geq m^{-(1+\gamma_P)}$,
 (W5) $\Sigma_{\sigma,a^2,0} \in [m \text{Var}(Z_a)/4, m \text{Var}(Z_a)]$ for all $\sigma \in \{\pm 1\}$,
 (W6) $\Sigma_{\sigma,w^2,0} \in [\text{Var}(Z_w)/4, \text{Var}(Z_w)]$ for all $\sigma \in \{\pm 1\}$,
 (W7) $|\Sigma_{\sigma,wa,0}| \leq m^\varepsilon$ for all $\sigma \in \{\pm 1\}$.

are satisfied with probability $\geq 1 - O(m^{-\gamma_P})$, where the constant in $O(m^{-\gamma_P})$ may depend on ε and γ_P (cf. Definition A.2).

Proof We will show the statement for each of the properties (W1) – (W7) individually, the rest follows by the union bound. Let Z_a, Z_w be the random variables from Assumption 2.

By property (Q2) in Assumption 2, $\mathbb{E}|Z_a|^p, \mathbb{E}|Z_w|^p < \infty$ for all $p \in (0, \infty)$. It can be shown (using the Minkowski and Hölder inequalities) that all other random variables used below satisfy the same property, which we will use in order to apply Lemma E.2.

(W1) True by Assumption 2.

(W2) By the Markov inequality, for $p > 0$,

$$\begin{aligned} P_m^{\text{init}} \left(|w_{i,0}| \geq m^{-1/2+\varepsilon} \right) &= P_m^{\text{init}} \left(|w_{i,0}|^p \geq m^{(-1/2+\varepsilon)p} \right) \\ &\leq \frac{\mathbb{E}|w_{i,0}|^p}{m^{(-1/2+\varepsilon)p}} = \frac{\mathbb{E}|Z_w|^p}{(\sqrt{m})^p m^{(-1/2+\varepsilon)p}} = \mathbb{E}(|Z_w|^p) m^{-\varepsilon p}. \end{aligned}$$

By choosing $p = (1 + \gamma_P)/\varepsilon$, we can use the union bound to conclude

$$P_m^{\text{init}} \left(\max_i |w_{i,0}| \geq m^{-1/2+\varepsilon} \right) \leq m \cdot \mathbb{E}(|Z_w|^p) m^{-\varepsilon p} \leq O(m^{1-\varepsilon p}) = O(m^{-\gamma_P}).$$

(W3) Similar to (W2).

(W4) By property (Q1) of Assumption 2, Z_a has a probability density p_a that is bounded by B_Z^{wa} . Thus, for all $\delta \geq 0$, we obtain

$$P_m^{\text{init}}(|a_{i,0}| \leq \delta) = P(|Z_a| \leq \delta) = \int_{-\delta}^{\delta} p_a(x) dx \leq 2\delta \cdot B_Z^{\text{wa}}.$$

Therefore,

$$\begin{aligned} P_m^{\text{init}} \left(\min_i |a_{i,0}| \leq m^{-(1+\gamma_P)} \right) &\leq \sum_{i=1}^m P_m^{\text{init}} \left(|a_{i,0}| \leq m^{-(1+\gamma_P)} \right) \\ &\leq m \cdot 2m^{-(1+\gamma_P)} \cdot B_Z^{\text{wa}} \leq O(m^{-\gamma_P}). \end{aligned}$$

(W5) For the next three properties, we need some preparation. Let

$$\begin{aligned} A_{\sigma,i} &:= \mathbb{1}_{(0,\infty)}(\sigma a_{i,0}) a_{i,0} \\ W_{\sigma,i} &:= \mathbb{1}_{(0,\infty)}(\sigma a_{i,0}) w_{i,0}. \end{aligned}$$

Note that the indicator function is applied to $\sigma a_{i,0}$ in both definitions. Then, $\Sigma_{\sigma,a^2,0} = \sum_{i \in I_\sigma} a_{i,0}^2 = \sum_{i=1}^m A_{\sigma,i}^2$ and similarly for $\Sigma_{\sigma,w^2,0}$ and $\Sigma_{\sigma,wa,0}$. We obtain

$$\mathbb{E}_{P_m^{\text{init}}} A_{\sigma,i}^2 = \int A_{\sigma,i}^2 dP_m^{\text{init}} = \int (\mathbb{1}_{(0,\infty)}(\sigma a_{i,0}) a_{i,0})^2 dP_m^{\text{init}}$$

$$\begin{aligned}
 &= \int_{\{\sigma a_{i,0} > 0\}} a_{i,0}^2 dP_m^{\text{init}} = \int_{\sigma(0,\infty)} x^2 p_a(x) dx \\
 &\stackrel{\text{(Q1)}}{=} \frac{1}{2} \int_{\mathbb{R}} x^2 p_a(x) dx = \frac{\mathbb{E}(Z_a^2)}{2} \stackrel{\text{(Q1)}}{=} \frac{\text{Var}(Z_a)}{2} . \\
 \mathbb{E}_{P_m^{\text{init}}} W_{\sigma,i}^2 &= \mathbb{E}_{P_m^{\text{init}}} \left((\mathbb{1}_{(0,\infty)}(\sigma a_{i,0}))^2 w_{i,0}^2 \right) \\
 &\stackrel{\text{indep.}}{=} \left(\mathbb{E}_{P_m^{\text{init}}} (\mathbb{1}_{(0,\infty)}(\sigma a_{i,0}))^2 \right) \cdot \left(\mathbb{E}_{P_m^{\text{init}}} w_{i,0}^2 \right) \\
 &= P_m^{\text{init}}(\sigma a_{i,0} > 0) \cdot \mathbb{E} \left(m^{-1/2} Z_w \right)^2 \stackrel{\text{(Q1)}}{=} \frac{1}{2} \cdot \frac{\text{Var}(Z_w)}{m} . \\
 \mathbb{E}_{P_m^{\text{init}}} W_{\sigma,i} A_{\sigma,i} &= \mathbb{E}_{P_m^{\text{init}}} (\mathbb{1}_{(0,\infty)}(\sigma a_{i,0}) w_{i,0} a_{i,0}) \\
 &\stackrel{\text{indep.}}{=} \underbrace{\left(\mathbb{E}_{P_m^{\text{init}}} w_{i,0} \right)}_{\stackrel{\text{(Q1)}}{=} 0} \cdot \left(\mathbb{E}_{P_m^{\text{init}}} \mathbb{1}_{(0,\infty)}(\sigma a_{i,0}) a_{i,0} \right) = 0 .
 \end{aligned}$$

Now, define

$$S_m := \frac{\Sigma_{\sigma,a^2,0}}{m} = \frac{1}{m} \sum_{i=1}^m A_{\sigma,i}^2 ,$$

which is an average of m i.i.d. variables that are p -integrable for every $p > 0$. Then, $\mathbb{E}_{P_m^{\text{init}}} S_m = \mathbb{E}_{P_m^{\text{init}}} A_{\sigma,1}^2 = \text{Var}(Z_a)/2$ and Lemma E.2 with $\varepsilon = 1/2, \beta = \text{Var}(Z_a)/4$ yields:

$$P_m^{\text{init}} \left(\left| S_m - \frac{\text{Var}(Z_a)}{2} \right| \geq \frac{\text{Var}(Z_a)}{4} \right) \leq O(m^{-\gamma_P}) .$$

Hence,

$$\begin{aligned}
 P_m^{\text{init}}(\Sigma_{\sigma,a^2,0} \notin [m \text{Var}(Z_a)/4, m \text{Var}(Z_a)]) &= P_m^{\text{init}}(S_m \notin [\text{Var}(Z_a)/4, \text{Var}(Z_a)]) \\
 &\leq P_m^{\text{init}}(S_m \notin [\text{Var}(Z_a)/4, 3 \text{Var}(Z_a)/4]) \\
 &\leq O(m^{-\gamma_P}) .
 \end{aligned}$$

(W6) An analogous argument yields $P_m^{\text{init}}(\Sigma_{\sigma,w^2,0} \notin [\text{Var}(Z_w)/4, \text{Var}(Z_w)]) \leq O(m^{-\gamma_P})$.

(W7) Let $S_m := \frac{1}{m} \sum_{i=1}^m A_{\sigma,i} \cdot \sqrt{m} W_{\sigma,i} = \Sigma_{\sigma,wa,0}/\sqrt{m}$. Then, $\mathbb{E}_{P_m^{\text{init}}} S_m = 0$ and thus

$$P_m^{\text{init}}(|\Sigma_{\sigma,wa,0}| \geq m^\varepsilon) = P_m^{\text{init}}(|S_m| \geq m^{\varepsilon-1/2}) \stackrel{\text{Lemma E.2}}{\leq} O(m^{-\gamma_P}) . \quad \blacksquare$$

Now, we want to investigate stochastic properties of the data set. In order to show that \mathbf{M}_D^{-1} is likely close to $\mathbf{M}_{P^{\text{data}}}^{-1}$ (both are defined in Definition 8), we need the following lemma, which is similar for example to Theorem 2.3.4 in Golub and Van Loan (1989):

Lemma E.4. *Let $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{m \times m}$ and let $\|\cdot\|$ be a matrix norm on $\mathbb{R}^{m \times m}$. If \mathbf{A} is invertible and $\|\mathbf{A}^{-1}\| \|\mathbf{A} - \mathbf{B}\| < 1$, then \mathbf{B} is invertible with*

$$\|\mathbf{B}^{-1} - \mathbf{A}^{-1}\| \leq \|\mathbf{A}^{-1}\| \|\mathbf{A} - \mathbf{B}\| \|\mathbf{B}^{-1}\|, \quad \|\mathbf{B}^{-1}\| \leq \frac{\|\mathbf{A}^{-1}\|}{1 - \|\mathbf{A}^{-1}\| \|\mathbf{A} - \mathbf{B}\|} .$$

Proof We have $\mathbf{B} = \mathbf{A}(\mathbf{I} - \mathbf{A}^{-1}(\mathbf{A} - \mathbf{B}))$ and since $\|\mathbf{A}^{-1}(\mathbf{A} - \mathbf{B})\| \leq \|\mathbf{A}^{-1}\| \|\mathbf{A} - \mathbf{B}\| < 1$, the Neumann series implies that

$$(\mathbf{I} - \mathbf{A}^{-1}(\mathbf{A} - \mathbf{B}))^{-1} = \sum_{k=0}^{\infty} (\mathbf{A}^{-1}(\mathbf{A} - \mathbf{B}))^k .$$

Hence \mathbf{B} is invertible with $\mathbf{B}^{-1} - \mathbf{A}^{-1} = \mathbf{A}^{-1}(\mathbf{A} - \mathbf{B})\mathbf{B}^{-1}$ and

$$\mathbf{B}^{-1} = (\mathbf{I} - \mathbf{A}^{-1}(\mathbf{A} - \mathbf{B}))^{-1} \mathbf{A}^{-1} = \sum_{k=0}^{\infty} (\mathbf{A}^{-1}(\mathbf{A} - \mathbf{B}))^k \mathbf{A}^{-1} ,$$

which yields both bounds using the submultiplicativity of $\|\cdot\|$. \blacksquare

Now, we can show that for large n , a sampled data set D likely has characteristics that are close to P^{data} . We use the convention $\infty \cdot 0 := \infty$.

Proposition E.5. *Let P^{data} satisfy Assumption 16, let $\varepsilon, K_{\text{data}} > 0$, $m \geq 1$ and $\gamma_{\text{data}}, \gamma' \geq 0$. If $\eta = \infty$, we further assume that K_{data} satisfies $P^{\text{data}}((-K_{\text{data}}, K_{\text{data}}) \times \mathbb{R}) = 0$. Finally, let D be a data set with n data points (x_j, y_j) sampled independently from P^{data} . Then with probability $1 - O(n^{-\gamma'} + nm^{-\eta_{\text{data}}})$ the following hold:*

- (D1) $\|\mathbf{v}_D^{\text{opt}} - \mathbf{v}_{P^{\text{data}}}^{\text{opt}}\|_{\infty} \leq n^{\varepsilon-1/2}$,
- (D2) For $\sigma \in \{\pm 1\}$, $\frac{1}{2}\lambda_{\min}(\mathbf{M}_{P^{\text{data}}, \sigma}) \leq \lambda_{\min}(\mathbf{M}_{D, \sigma})$ and $\lambda_{\max}(\mathbf{M}_{D, \sigma}) \leq 2\lambda_{\max}(\mathbf{M}_{P^{\text{data}}, \sigma})$,
- (D3) $\underline{x}_D \geq K_{\text{data}} m^{-\gamma_{\text{data}}}$.

Proof We use the shorthand $P := P^{\text{data}}$. Again, we bound the probabilities for each property separately.

(D1) For $\sigma \in \{\pm 1\}$, define

$$S_n := (\mathbf{M}_{D, \sigma})_{11} = \frac{1}{n} \sum_{j=1}^n \mathbf{1}_{(0, \infty)}(\sigma x_j) x_j^2 .$$

Then,

$$\mathbb{E}_{P^n} S_n = \frac{1}{n} \sum_{j=1}^n \mathbb{E}_{D \sim P^n} (\mathbf{1}_{(0, \infty)}(\sigma x_j) x_j^2) = \mathbb{E}_{(x, y) \sim P} \mathbf{1}_{(0, \infty)}(\sigma x) x^2 = (\mathbf{M}_{P, \sigma})_{11} .$$

Because P is bounded, it has finite moments and we can apply Lemma E.2: For all $\beta > 0$,

$$P^n \left(|(\mathbf{M}_{D, \sigma})_{11} - (\mathbf{M}_{P, \sigma})_{11}| \geq \beta n^{(\varepsilon-1)/2} \right) \leq O(n^{-\gamma}) .$$

We can get similar bounds for other entries of $\mathbf{M}_{D, \sigma}$ and $\hat{\mathbf{u}}_{D, \sigma}^0$. Since

$$\mathbf{M}_D - \mathbf{M}_P = \tilde{\mathbf{P}} \begin{pmatrix} \mathbf{M}_{D,1} - \mathbf{M}_{P,1} & 0 \\ 0 & \mathbf{M}_{D,-1} - \mathbf{M}_{P,-1} \end{pmatrix} \tilde{\mathbf{P}}, \quad \hat{\mathbf{u}}_D^0 = \tilde{\mathbf{P}} \begin{pmatrix} \hat{\mathbf{u}}_{D,1}^0 \\ \hat{\mathbf{u}}_{D,-1}^0 \end{pmatrix},$$

and $\|\tilde{\mathbf{P}}\|_\infty = 1$, the union bound implies that the following properties hold with probability $1 - O(n^{-\gamma'})$:

$$\begin{aligned} \|\mathbf{M}_{D,\pm 1} - \mathbf{M}_{P,\pm 1}\|_\infty &\leq 2\beta n^{(\varepsilon-1)/2}, \quad \|\mathbf{M}_D - \mathbf{M}_P\|_\infty \leq 2\beta n^{(\varepsilon-1)/2}, \\ \|\hat{\mathbf{u}}_D^0 - \hat{\mathbf{u}}_P^0\|_\infty &\leq \beta n^{(\varepsilon-1)/2}. \end{aligned} \quad (10)$$

Now assume that (10) holds. Set $\mathbf{A} := \mathbf{M}_P, \mathbf{B} := \mathbf{M}_D, \mathbf{a} := \hat{\mathbf{u}}_P^0, \mathbf{b} := \hat{\mathbf{u}}_D^0$. By condition (P1), \mathbf{A} is invertible. Without loss of generality, we can assume $\varepsilon < 1/2$. Then, for n large enough,

$$\|\mathbf{A}^{-1}\|_\infty \|\mathbf{A} - \mathbf{B}\|_\infty \leq \|\mathbf{A}^{-1}\|_\infty 2\beta n^{(\varepsilon-1)/2} \leq \frac{1}{2}.$$

Hence, Lemma E.4 implies that $\mathbf{B} = \mathbf{M}_D$ is invertible with $\|\mathbf{B}^{-1}\|_\infty \leq 2\|\mathbf{A}^{-1}\|_\infty$ and

$$\begin{aligned} \|\mathbf{v}_D^{\text{opt}} - \mathbf{v}_P^{\text{opt}}\|_\infty &= \|\mathbf{B}^{-1}\mathbf{b} - \mathbf{A}^{-1}\mathbf{a}\|_\infty \\ &\leq \|\mathbf{B}^{-1}\|_\infty \|\mathbf{b} - \mathbf{a}\|_\infty + \|\mathbf{B}^{-1} - \mathbf{A}^{-1}\|_\infty \|\mathbf{a}\|_\infty \\ &\leq \|\mathbf{B}^{-1}\|_\infty \|\mathbf{b} - \mathbf{a}\|_\infty + \|\mathbf{A}^{-1}\|_\infty \|\mathbf{A} - \mathbf{B}\|_\infty \|\mathbf{B}^{-1}\|_\infty \|\mathbf{a}\|_\infty \\ &\leq 2\|\mathbf{A}^{-1}\|_\infty (\|\mathbf{b} - \mathbf{a}\|_\infty + \|\mathbf{A}^{-1}\|_\infty \|\mathbf{a}\|_\infty \|\mathbf{B} - \mathbf{A}\|_\infty) \\ &\stackrel{(10)}{\leq} 4\|\mathbf{A}^{-1}\|_\infty (1 + \|\mathbf{A}^{-1}\|_\infty \|\mathbf{a}\|_\infty) \beta n^{(\varepsilon-1)/2}. \end{aligned}$$

We can choose $\beta > 0$ such that $4\|\mathbf{A}^{-1}\|_\infty (1 + \|\mathbf{A}^{-1}\|_\infty \|\mathbf{a}\|_\infty) \beta \leq 1$. Therefore,

$$\|\mathbf{v}_{P,\sigma}^{\text{opt}} - \mathbf{v}_{D,\sigma}^{\text{opt}}\|_\infty \leq n^{(\varepsilon-1)/2}$$

with probability $1 - O(n^{-\gamma'})$.

(D2) For $\sigma \in \{\pm 1\}$ and each $\mathbf{v} \in \mathbb{R}^2$, we have

$$|\mathbf{v}^\top \mathbf{M}_{D,\sigma} \mathbf{v} - \mathbf{v}^\top \mathbf{M}_{P,\sigma} \mathbf{v}| \leq \|\mathbf{v}\|_2 \|\mathbf{M}_{D,\sigma} - \mathbf{M}_{P,\sigma}\|_2 \|\mathbf{v}\|_2 \leq \sqrt{2} \|\mathbf{M}_{D,\sigma} - \mathbf{M}_{P,\sigma}\|_\infty \|\mathbf{v}\|_2^2$$

since $\|\cdot\|_2 \leq \sqrt{2}\|\cdot\|_\infty$ on $\mathbb{R}^{2 \times 2}$ as mentioned in Definition A.3. If we choose $\beta > 0$ small enough such that (10) implies $\sqrt{2}\|\mathbf{M}_{D,\sigma} - \mathbf{M}_{P,\sigma}\|_\infty \leq \lambda_{\min}(\mathbf{M}_{P,\sigma})/2$, it follows that

$$\begin{aligned} \lambda_{\min}(\mathbf{M}_{D,\sigma}) &= \inf_{\|\mathbf{v}\|_2=1} \mathbf{v}^\top \mathbf{M}_{D,\sigma} \mathbf{v} \geq \inf_{\|\mathbf{v}\|_2=1} \mathbf{v}^\top \mathbf{M}_{P,\sigma} \mathbf{v} - |\mathbf{v}^\top \mathbf{M}_{P,\sigma} \mathbf{v} - \mathbf{v}^\top \mathbf{M}_{D,\sigma} \mathbf{v}| \\ &\geq \lambda_{\min}(\mathbf{M}_{P,\sigma}) - \sqrt{2} \|\mathbf{M}_{D,\sigma} - \mathbf{M}_{P,\sigma}\|_\infty \geq \lambda_{\min}(\mathbf{M}_{P,\sigma})/2. \end{aligned}$$

Since (10) holds with probability $1 - O(n^{-\gamma'})$, we have $\lambda_{\min}(\mathbf{M}_{D,\sigma}) \geq \lambda_{\min}(\mathbf{M}_{P,\sigma})/2$ with probability $1 - O(n^{-\gamma})$. The probability for $\lambda_{\max}(\mathbf{M}_{D,\sigma}) \leq 2\lambda_{\max}(\mathbf{M}_{P,\sigma})$ can be bounded similarly.

(D3) In the case $\eta = \infty$ and $P_X((-K_{\text{data}}, K_{\text{data}})) = 0$, this obviously holds with probability one since $K_{\text{data}} m^{-\gamma_{\text{data}}} \leq K_{\text{data}} \leq \underline{x}_D$ almost surely. Otherwise, using property (P2) from Assumption 16 and the union bound yields

$$\begin{aligned} P^n(\underline{x}_D < K_{\text{data}} m^{-\gamma_{\text{data}}}) &\leq \sum_{j=1}^n P^n(|x_j| < K_{\text{data}} m^{-\gamma_{\text{data}}}) \\ &\stackrel{(P2)}{\leq} n \cdot O((K_{\text{data}} m^{-\gamma_{\text{data}}})^\eta) = O(nm^{-\eta\gamma_{\text{data}}}). \quad \blacksquare \end{aligned}$$

Appendix F. Reference Dynamics

In this section, we define a matrix $\mathbf{A}^{\text{ref}} \approx \mathbf{A}_0$ and study the asymptotic behavior of $h \sum_{k=0}^{\infty} \|\bar{\mathbf{v}}_k\|$ when $\bar{\mathbf{v}}_k$ satisfies a reference system

$$\bar{\mathbf{v}}_{k+1} = \bar{\mathbf{v}}_k - h\mathbf{A}^{\text{ref}}\mathbf{M}_D\bar{\mathbf{v}}_k,$$

instead of the actual dynamics $\bar{\mathbf{v}}_{k+1} = \bar{\mathbf{v}}_k - h\mathbf{A}_k\mathbf{M}_D\bar{\mathbf{v}}_k$. The solution of the reference system is simply $\bar{\mathbf{v}}_k = (\mathbf{I}_4 - h\mathbf{A}^{\text{ref}}\mathbf{M}_D)^k\bar{\mathbf{v}}_0$.

Definition F.1. Let $\mathbf{A}^{\text{ref}} := \mathbf{B}(\mathbf{G}_0^{\text{w}} + \mathbf{G}_0^{\text{ab}})\mathbf{B} + \mathbf{C}$, where $\mathbf{B}, \mathbf{G}^{\text{w}}, \mathbf{G}^{\text{ab}}, \mathbf{C}$ are defined in Definition C.1. Moreover, define the symmetric matrix

$$\mathbf{H} := \mathbf{M}_D^{1/2}\mathbf{A}^{\text{ref}}\mathbf{M}_D^{1/2} = \mathbf{M}_D^{1/2}(\mathbf{A}^{\text{ref}}\mathbf{M}_D)\mathbf{M}_D^{-1/2}.$$

Assumption F.2. Assume that $\gamma_\psi, \gamma_{\text{data}}, \gamma_P \geq 0$ with $\gamma_\psi + \gamma_{\text{data}} + \gamma_P < 1/2$. Moreover, we only consider initial vectors W_0 which satisfy the conditions (W1) – (W7) in Proposition E.3. Similar to Theorem 10, we further assume that

$$\begin{aligned} K_M^{-1} &\leq \lambda_{\min}(\mathbf{M}_D) \leq \lambda_{\max}(\mathbf{M}_D) \leq K_M \\ \psi_{D,p} &= O(1) \\ \psi_{D,q} &= O(m^{\gamma_\psi-1}) \\ h &\leq \lambda_{\max}(\mathbf{H})^{-1} \\ 0 < \varepsilon &< \frac{1/2 - (\gamma_\psi + \gamma_{\text{data}} + \gamma_P)}{3}, \end{aligned}$$

where $K_M > 0$ is a constant and \mathbf{H} depends on D and W_0 . Note that since $\text{eig}(\mathbf{M}_D) = \text{eig}(\mathbf{M}_{D,1}) \cup \text{eig}(\mathbf{M}_{D,-1})$ by construction of \mathbf{M}_D , the first condition is equivalent to

$$K_M^{-1} \leq \lambda_{\min}(\mathbf{M}_{D,\sigma}) \leq \lambda_{\max}(\mathbf{M}_{D,\sigma}) \leq K_M \text{ for } \sigma = \pm 1.$$

Lemma F.3. Let Assumption F.2 be satisfied. The matrix \mathbf{A}^{ref} is of the form

$$\mathbf{A}^{\text{ref}} = \begin{pmatrix} \mathbf{A}_1^{\text{ref}} & \\ & \mathbf{A}_2^{\text{ref}} \end{pmatrix}$$

with $0 \prec \mathbf{A}_1^{\text{ref}}, \mathbf{A}_2^{\text{ref}} \in \mathbb{R}^{2 \times 2}$ and

$$\lambda_{\min}(\mathbf{A}_1^{\text{ref}}) = \Theta(m), \quad \lambda_{\max}(\mathbf{A}_1^{\text{ref}}) = \Theta(m), \quad \lambda_{\min}(\mathbf{A}_2^{\text{ref}}) = \Theta(1), \quad \lambda_{\max}(\mathbf{A}_2^{\text{ref}}) = \Theta(1).$$

Proof Since $b_{i,0} = 0$ by initialization property (W1) in Proposition E.3, we have $\Sigma_{\sigma,ab,0} = \Sigma_{\sigma,b^2,0} = 0$. Since the distributions of Z_a, Z_w have densities by (Q1) in Assumption 2, we have $\text{Var}(Z_a), \text{Var}(Z_w) > 0$. This yields

$$\mathbf{G}_{\sigma,0}^{\text{w}} + \mathbf{G}_{\sigma,0}^{\text{ab}} = \begin{pmatrix} \Sigma_{\sigma,w^2,0} + \Sigma_{\sigma,a^2,0} & \\ & \Sigma_{\sigma,w^2,0} \end{pmatrix} \stackrel{(W5), (W6)}{=} \begin{pmatrix} \Theta(m) & \\ & \Theta(1) \end{pmatrix}.$$

Hence,

$$\mathbf{G}_0^{\text{w}} + \mathbf{G}_0^{\text{ab}} = \tilde{\mathbf{P}}(\tilde{\mathbf{G}}_0^{\text{w}} + \tilde{\mathbf{G}}_0^{\text{ab}})\tilde{\mathbf{P}} = \tilde{\mathbf{P}} \begin{pmatrix} \mathbf{G}_{1,0}^{\text{w}} + \mathbf{G}_{1,0}^{\text{ab}} & \\ & \mathbf{G}_{-1,0}^{\text{w}} + \mathbf{G}_{-1,0}^{\text{ab}} \end{pmatrix} \tilde{\mathbf{P}}$$

$$\begin{aligned}
 &= \tilde{\mathbf{P}} \begin{pmatrix} \Theta(m) & & & \\ & \Theta(1) & & \\ & & \Theta(m) & \\ & & & \Theta(1) \end{pmatrix} \tilde{\mathbf{P}} = \begin{pmatrix} \Theta(m) & & & \\ & \Theta(m) & & \\ & & \Theta(1) & \\ & & & \Theta(1) \end{pmatrix} \\
 &=: \begin{pmatrix} \mathbf{G}_1 & \\ & \mathbf{G}_2 \end{pmatrix}.
 \end{aligned}$$

We have seen in Definition C.1 that

$$\mathbf{B} = \begin{pmatrix} \hat{\mathbf{B}} & \\ & \hat{\mathbf{B}} \end{pmatrix}, \quad \hat{\mathbf{B}} = \begin{pmatrix} 1 & \alpha \\ \alpha & 1 \end{pmatrix}, \quad \mathbf{C} = \begin{pmatrix} 0 & 0 & & \\ 0 & 0 & & \\ & & 1 & 1 \\ & & 1 & 1 \end{pmatrix} =: \begin{pmatrix} \mathbf{0} & \\ & \hat{\mathbf{C}} \end{pmatrix}.$$

Using the previous results, we obtain

$$\mathbf{A}^{\text{ref}} = \begin{pmatrix} \hat{\mathbf{B}}\mathbf{G}_1\hat{\mathbf{B}} & \\ & \hat{\mathbf{B}}\mathbf{G}_2\hat{\mathbf{B}} + \hat{\mathbf{C}} \end{pmatrix} =: \begin{pmatrix} \mathbf{A}_1^{\text{ref}} & \\ & \mathbf{A}_2^{\text{ref}} \end{pmatrix}.$$

The matrix $\hat{\mathbf{B}}$ is fixed and invertible since $|\alpha| \neq 1$. Moreover, $\text{eig}(\hat{\mathbf{C}}) = \{0, 2\}$. This yields

$$\begin{aligned}
 \text{eig}(\mathbf{A}_1^{\text{ref}}) &= \text{eig}(\hat{\mathbf{B}}\mathbf{G}_1\hat{\mathbf{B}}) = \Theta(m) \\
 \text{eig}(\mathbf{A}_2^{\text{ref}}) &= \text{eig}(\hat{\mathbf{B}}\mathbf{G}_2\hat{\mathbf{B}} + \hat{\mathbf{C}}) = \Theta(1). \quad \blacksquare
 \end{aligned}$$

Proposition F.4. *Let Assumption F.2 be satisfied. We have*

$$\begin{aligned}
 h \sum_{k=0}^{\infty} \|(\mathbf{I}_4 - h\mathbf{A}^{\text{ref}}\mathbf{M})^k \bar{\mathbf{v}}_0\|_{\infty} &= O(m^{\varepsilon + \gamma_{\psi} - 1}) \\
 h \sum_{k=0}^{\infty} \|(\mathbf{I}_4 - h\mathbf{A}^{\text{ref}}\mathbf{M})^k\|_{\infty} &= O(1).
 \end{aligned}$$

Proof We divide the proof in multiple steps:

(1) *Investigate the initial vector:*

By definition, we have $\bar{\mathbf{v}}_0 = \mathbf{v}_0 - \mathbf{v}_D^{\text{opt}}$. Therefore,

$$\begin{aligned}
 |\bar{\mathbf{v}}_0| &\leq |\mathbf{v}_0| + |\mathbf{v}_D^{\text{opt}}| \\
 &\leq \begin{pmatrix} |p_{1,0}| \\ |p_{-1,0}| \\ |q_{1,0}| \\ |q_{-1,0}| \end{pmatrix} + \begin{pmatrix} \psi_{D,p} \\ \psi_{D,p} \\ \psi_{D,q} \\ \psi_{D,q} \end{pmatrix} = \begin{pmatrix} |\Sigma_{1,wa,0} + \alpha\Sigma_{-1,wa,0}| \\ |\Sigma_{-1,wa,0} + \alpha\Sigma_{1,wa,0}| \\ |\Sigma_{1,wb,0} + \alpha\Sigma_{-1,wb,0}| \\ |\Sigma_{-1,wb,0} + \alpha\Sigma_{1,wb,0}| \end{pmatrix} + \begin{pmatrix} \psi_{D,p} \\ \psi_{D,p} \\ \psi_{D,q} \\ \psi_{D,q} \end{pmatrix} \\
 &\stackrel{(\text{W1}), (\text{W7}), \text{F.2}}{\leq} \begin{pmatrix} O(m^{\varepsilon}) \\ O(m^{\varepsilon}) \\ 0 \\ 0 \end{pmatrix} + \begin{pmatrix} O(1) \\ O(1) \\ O(m^{\gamma_{\psi} - 1}) \\ O(m^{\gamma_{\psi} - 1}) \end{pmatrix} \leq \begin{pmatrix} O(m^{\varepsilon}) \\ O(m^{\varepsilon}) \\ O(m^{\gamma_{\psi} - 1}) \\ O(m^{\gamma_{\psi} - 1}) \end{pmatrix}.
 \end{aligned}$$

Thus, we can group

$$\bar{\mathbf{v}}_0 = \begin{pmatrix} \bar{\mathbf{v}}_{0,1} \\ \bar{\mathbf{v}}_{0,2} \end{pmatrix}$$

with $\bar{\mathbf{v}}_{0,1}, \bar{\mathbf{v}}_{0,2} \in \mathbb{R}^2$ and $\|\bar{\mathbf{v}}_{0,1}\|_{\infty} \leq O(m^{\varepsilon}), \|\bar{\mathbf{v}}_{0,2}\|_{\infty} \leq O(m^{\gamma_{\psi} - 1})$.

(2) *Diagonalization yields a simple bound:*

The matrix $\mathbf{A}^{\text{ref}}\mathbf{M}$ is similar to the symmetric matrix

$$\mathbf{H} := \mathbf{M}^{1/2}\mathbf{A}^{\text{ref}}\mathbf{M}^{1/2} = \mathbf{M}^{1/2}(\mathbf{A}^{\text{ref}}\mathbf{M})\mathbf{M}^{-1/2} \succ 0 .$$

The matrix \mathbf{H} can thus be orthogonally diagonalized as $\mathbf{H} = \mathbf{U}\mathbf{D}\mathbf{U}^\top$ with \mathbf{U} orthogonal and \mathbf{D} diagonal such that \mathbf{D} contains the eigenvalues of \mathbf{H} in descending order. Then, $\mathbf{I}_4 - h\mathbf{D}$ only contains non-negative entries due to the condition $h \leq \lambda_{\max}(\mathbf{H})^{-1}$ with its maximal entry being $1 - h\lambda_{\min}(\mathbf{H})$. Thus, $\|(\mathbf{I}_4 - h\mathbf{D})^k\|_2 = (1 - h\lambda_{\min}(\mathbf{H}))^k$. By applying $(\mathbf{I}_4 - h\mathbf{A}^{\text{ref}}\mathbf{M})\mathbf{M}^{-1/2} = \mathbf{M}^{-1/2} - h\mathbf{A}^{\text{ref}}\mathbf{M}^{1/2} = \mathbf{M}^{-1/2}(\mathbf{I}_4 - h\mathbf{H})$ inductively, we find $(\mathbf{I}_4 - h\mathbf{A}^{\text{ref}}\mathbf{M})^k\mathbf{M}^{-1/2} = \mathbf{M}^{-1/2}(\mathbf{I}_4 - h\mathbf{H})^k$. We can now compute

$$\begin{aligned} h \sum_{k=0}^{\infty} \|(\mathbf{I}_4 - h\mathbf{A}^{\text{ref}}\mathbf{M})^k\|_2 &= h \sum_{k=0}^{\infty} \|\mathbf{M}^{-1/2}(\mathbf{I}_4 - h\mathbf{H})^k\mathbf{M}^{1/2}\|_2 \\ &= h \sum_{k=0}^{\infty} \|\mathbf{M}^{-1/2}\mathbf{U}(\mathbf{I}_4 - h\mathbf{D})^k\mathbf{U}^\top\mathbf{M}^{1/2}\|_2 \\ &\leq \|\mathbf{M}^{-1/2}\|_2 \|\mathbf{M}^{1/2}\|_2 \cdot h \sum_{k=0}^{\infty} \|(\mathbf{I}_4 - h\mathbf{D})^k\|_2 \\ &= \text{cond}(\mathbf{M}^{1/2}) h \sum_{k=0}^{\infty} (1 - h\lambda_{\min}(\mathbf{H}))^k \\ &= \sqrt{\text{cond}(\mathbf{M})} \frac{h}{1 - (1 - h\lambda_{\min}(\mathbf{H}))} \\ &= \frac{\sqrt{\text{cond}(\mathbf{M})}}{\lambda_{\min}(\mathbf{H})} \leq O(1) , \end{aligned} \tag{11}$$

where $\lambda_{\min}(\mathbf{H}) \geq \Omega(1)$ since for $\mathbf{v} \in \mathbb{R}^4$, we have

$$\mathbf{v}^\top \mathbf{H} \mathbf{v} = (\mathbf{M}^{1/2}\mathbf{v})^\top \mathbf{A}^{\text{ref}}(\mathbf{M}^{1/2}\mathbf{v}) \geq \lambda_{\min}(\mathbf{A}^{\text{ref}})\mathbf{v}^\top \mathbf{M} \mathbf{v} \geq \lambda_{\min}(\mathbf{A}^{\text{ref}})\lambda_{\min}(\mathbf{M})\mathbf{v}^\top \mathbf{v} \tag{12}$$

where $\lambda_{\min}(\mathbf{A}^{\text{ref}})\lambda_{\min}(\mathbf{M}) = \Theta(1)$ by Assumption F.2 and Lemma F.3.

(3) *$\mathbf{A}^{\text{ref}}\mathbf{M}$ has 2 “large” eigenvalues:*

Let

$$\mathbf{M} = \begin{pmatrix} \mathbf{M}_{11} & \mathbf{M}_{12} \\ \mathbf{M}_{12}^\top & \mathbf{M}_{22} \end{pmatrix}, \quad \mathbf{M}^{1/2} = \begin{pmatrix} \tilde{\mathbf{M}}_{11} & \tilde{\mathbf{M}}_{12} \\ \tilde{\mathbf{M}}_{12}^\top & \tilde{\mathbf{M}}_{22} \end{pmatrix}$$

be the block decompositions of \mathbf{M} and $\mathbf{M}^{1/2}$ into 2×2 blocks. Then,

$$\mathbf{M}^{1/2}\mathbf{A}^{\text{ref}}\mathbf{M}^{1/2} = \begin{pmatrix} \tilde{\mathbf{M}}_{11}\mathbf{A}_1^{\text{ref}}\tilde{\mathbf{M}}_{11} + \tilde{\mathbf{M}}_{12}\mathbf{A}_2^{\text{ref}}\tilde{\mathbf{M}}_{12}^\top & * \\ * & * \end{pmatrix}$$

and by Cauchy’s interlacing theorem (cf. e.g. Corollary III.1.5 in (Bhatia, 2013)), the second largest eigenvalue $\lambda_2(\mathbf{H})$ of \mathbf{H} satisfies

$$\begin{aligned} \lambda_2(\mathbf{H}) &\geq \lambda_2(\tilde{\mathbf{M}}_{11}\mathbf{A}_1^{\text{ref}}\tilde{\mathbf{M}}_{11} + \tilde{\mathbf{M}}_{12}\mathbf{A}_2^{\text{ref}}\tilde{\mathbf{M}}_{12}^\top) = \lambda_{\min}(\tilde{\mathbf{M}}_{11}\mathbf{A}_1^{\text{ref}}\tilde{\mathbf{M}}_{11} + \tilde{\mathbf{M}}_{12}\mathbf{A}_2^{\text{ref}}\tilde{\mathbf{M}}_{12}^\top) \\ &\geq \lambda_{\min}(\tilde{\mathbf{M}}_{11}\mathbf{A}_1^{\text{ref}}\tilde{\mathbf{M}}_{11}) \geq \lambda_{\min}(\mathbf{A}_1^{\text{ref}})\lambda_{\min}(\tilde{\mathbf{M}}_{11})^2 \geq \lambda_{\min}(\mathbf{A}_1^{\text{ref}})\lambda_{\min}(\mathbf{M}^{1/2})^2 \\ &= \lambda_{\min}(\mathbf{A}_1^{\text{ref}})\lambda_{\min}(\mathbf{M}) \geq \Theta(m) . \end{aligned} \tag{13}$$

(4) *Lower components of eigenvectors to large eigenvalues are small:*

Let $\mathbf{w} = (\mathbf{w}_1, \mathbf{w}_2)^\top$ be an eigenvector of $\mathbf{A}^{\text{ref}}\mathbf{M}$ with eigenvalue $\lambda \geq \lambda_2(\mathbf{H}) \geq \Theta(m)$. The lower part of the identity $\lambda\mathbf{w} = \mathbf{A}^{\text{ref}}\mathbf{M}\mathbf{w}$ reads as

$$\lambda\mathbf{w}_2 = \mathbf{A}_2^{\text{ref}}\mathbf{M}_{12}^\top\mathbf{w}_1 + \mathbf{A}_2^{\text{ref}}\mathbf{M}_{22}\mathbf{w}_2 ,$$

which yields

$$\begin{aligned} \Theta(m)\|\mathbf{w}_2\|_2 &\leq \lambda\|\mathbf{w}_2\|_2 \leq \|\mathbf{A}_2^{\text{ref}}\|_2\|\mathbf{M}_{12}^\top\|_2\|\mathbf{w}_1\|_2 + \|\mathbf{A}_2^{\text{ref}}\|_2\|\mathbf{M}_{22}\|_2\|\mathbf{w}_2\|_2 \\ &\leq \Theta(1)\|\mathbf{w}_1\|_2 + \Theta(1)\|\mathbf{w}_2\|_2 \end{aligned}$$

and hence (for large m)

$$\|\mathbf{w}_2\|_2 \leq \frac{\Theta(1)}{\Theta(m) - \Theta(1)}\|\mathbf{w}_1\|_2 \leq O(m^{-1})\|\mathbf{w}_1\|_2 . \quad (14)$$

(5) *The first two eigenvectors of $\mathbf{A}^{\text{ref}}\mathbf{M}$ are “well-conditioned”:*

Let

$$\begin{aligned} \mathbf{U} &= (\mathbf{U}_1 \quad \mathbf{U}_2) = \begin{pmatrix} \mathbf{U}_{11} & \mathbf{U}_{12} \\ \mathbf{U}_{21} & \mathbf{U}_{22} \end{pmatrix}, \quad \mathbf{F} = (\mathbf{F}_1 \quad \mathbf{F}_2) := \mathbf{U}_1^\top \mathbf{M}^{1/2}, \\ \mathbf{W} &= \begin{pmatrix} \mathbf{W}_1 \\ \mathbf{W}_2 \end{pmatrix} := \mathbf{M}^{-1/2}\mathbf{U}_1 . \end{aligned}$$

The columns of \mathbf{W} are the eigenvectors of $\mathbf{A}^{\text{ref}}\mathbf{M}$ to the 2 largest eigenvalues:

$$\begin{aligned} \mathbf{A}^{\text{ref}}\mathbf{M}\mathbf{W} &= \mathbf{M}^{-1/2}\mathbf{M}^{1/2}\mathbf{A}^{\text{ref}}\mathbf{M}^{1/2}\mathbf{U}_1 = \mathbf{M}^{-1/2}\mathbf{U}\mathbf{D}\mathbf{U}^\top\mathbf{U}_1 = \mathbf{M}^{-1/2}\mathbf{U}\mathbf{D} \begin{pmatrix} \mathbf{I}_2 \\ \mathbf{0} \end{pmatrix} \\ &= \mathbf{M}^{-1/2}\mathbf{U}_1\mathbf{D}_1 = \mathbf{W}\mathbf{D}_1 , \end{aligned} \quad (15)$$

where \mathbf{D}_1 is the upper left 2×2 block of \mathbf{D} . Thus,

$$\begin{aligned} \|\mathbf{F}\|_2 &\leq \|\mathbf{U}_1^\top\|_2\|\mathbf{M}^{1/2}\|_2 = 1 \cdot \lambda_{\max}(\mathbf{M}^{1/2}) = \Theta(1) \\ \|\mathbf{W}\|_2 &\leq \|\mathbf{M}^{-1/2}\|_2\|\mathbf{U}_1\|_2 = \lambda_{\max}(\mathbf{M}^{-1/2}) \cdot 1 = \Theta(1) \\ \|\mathbf{W}_2\|_2 &\leq \|\mathbf{W}_2\|_F \stackrel{(14)}{\leq} O(m^{-1})\|\mathbf{W}_1\|_F \leq O(m^{-1})\|\mathbf{W}\|_2 \leq O(m^{-1}) . \end{aligned}$$

We want to show that \mathbf{W}_1^{-1} exists and $\|\mathbf{W}_1^{-1}\|_2$ is sufficiently small. Observe that $\mathbf{I}_2 = \mathbf{U}_1^\top\mathbf{U}_1 = \mathbf{F}\mathbf{W} = \mathbf{F}_1\mathbf{W}_1 + \mathbf{F}_2\mathbf{W}_2$ and

$$\|\mathbf{F}_2\mathbf{W}_2\|_2 \leq \|\mathbf{F}_2\|_2\|\mathbf{W}_2\|_2 \leq O(m^{-1}) \leq \frac{1}{2}$$

for large m . Hence, $\mathbf{F}_1\mathbf{W}_1 = \mathbf{I}_2 - \mathbf{F}_2\mathbf{W}_2$ is invertible with

$$(\mathbf{F}_1\mathbf{W}_1)^{-1} = \sum_{k=0}^{\infty} (\mathbf{F}_2\mathbf{W}_2)^k, \quad \|(\mathbf{F}_1\mathbf{W}_1)^{-1}\|_2 \leq \sum_{k=0}^{\infty} \|\mathbf{F}_2\mathbf{W}_2\|_2^k \leq 2 .$$

Since $\mathbf{F}_1\mathbf{W}_1$ has full rank, \mathbf{W}_1 and \mathbf{F}_1 must also have full rank. Hence, $(\mathbf{F}_1\mathbf{W}_1)^{-1} = \mathbf{W}_1^{-1}\mathbf{F}_1^{-1}$ and

$$\|\mathbf{W}_1^{-1}\|_2 \leq \|(\mathbf{F}_1\mathbf{W}_1)^{-1}\|_2\|\mathbf{F}_1\|_2 \leq O(1) .$$

(6) *Bound the sum for a “similar” initial vector:*

Note that for $\tilde{\mathbf{v}}_2 := \mathbf{W}_2 \mathbf{W}_1^{-1} \bar{\mathbf{v}}_{0,1}$, we have

$$\mathbf{W} \mathbf{W}_1^{-1} \bar{\mathbf{v}}_{0,1} = \begin{pmatrix} \mathbf{I}_2 \\ \mathbf{W}_2 \mathbf{W}_1^{-1} \end{pmatrix} \bar{\mathbf{v}}_{0,1} = \begin{pmatrix} \bar{\mathbf{v}}_{0,1} \\ \tilde{\mathbf{v}}_2 \end{pmatrix} \quad (16)$$

and $\tilde{\mathbf{v}}_2$ is “small”:

$$\|\tilde{\mathbf{v}}_2\|_2 \leq \|\mathbf{W}_2\|_2 \|\mathbf{W}_1^{-1}\|_2 \|\bar{\mathbf{v}}_{0,1}\|_2 \leq O(m^{-1})O(1)O(m^\varepsilon) = O(m^{\varepsilon-1}).$$

By Eq. (15), we have $\mathbf{A}^{\text{ref}} \mathbf{M} \mathbf{W} = \mathbf{W} \mathbf{D}_1$, where \mathbf{D}_1 is the upper left 2×2 block of \mathbf{D} . Therefore,

$$\begin{aligned} h \sum_{k=0}^{\infty} \|(\mathbf{I}_4 - h\mathbf{A}^{\text{ref}} \mathbf{M})^k \mathbf{W} \mathbf{W}_1^{-1} \bar{\mathbf{v}}_{0,1}\|_2 &= h \sum_{k=0}^{\infty} \|\mathbf{W} (\mathbf{I}_2 - h\mathbf{D}_1)^k \mathbf{W}_1^{-1} \bar{\mathbf{v}}_{0,1}\|_2 \\ &\leq \|\mathbf{W}\|_2 \|\mathbf{W}_1^{-1}\|_2 \|\bar{\mathbf{v}}_{0,1}\|_2 \cdot h \sum_{k=0}^{\infty} \|(\mathbf{I}_2 - h\mathbf{D}_1)^k\|_2, \end{aligned}$$

where

$$\|\mathbf{W}\|_2 \|\mathbf{W}_1^{-1}\|_2 \|\bar{\mathbf{v}}_{0,1}\|_2 \leq O(1)O(1)O(m^\varepsilon) = O(m^\varepsilon)$$

and we can compute the remaining sum similar to step (2):

$$h \sum_{k=0}^{\infty} \|(\mathbf{I}_2 - h\mathbf{D}_1)^k\|_2 = h \sum_{k=0}^{\infty} (1 - h\lambda_2(\mathbf{H}))^k \leq \frac{h}{1 - (1 - h\lambda_2(\mathbf{H}))} \stackrel{(13)}{\leq} O(m^{-1}).$$

(7) *Bound the original sum:*

Using $\bar{\mathbf{v}}_0 = \mathbf{W} \mathbf{W}_1^{-1} \bar{\mathbf{v}}_{0,1} + \begin{pmatrix} 0 \\ \bar{\mathbf{v}}_{0,2} - \tilde{\mathbf{v}}_2 \end{pmatrix}$, we obtain

$$\begin{aligned} &h \sum_{k=0}^{\infty} \|(\mathbf{I}_4 - h\mathbf{A}^{\text{ref}} \mathbf{M})^k \bar{\mathbf{v}}_0\|_2 \\ (16) \quad &\leq h \sum_{k=0}^{\infty} \|(\mathbf{I}_4 - h\mathbf{A}^{\text{ref}} \mathbf{M})^k \mathbf{W} \mathbf{W}_1^{-1} \bar{\mathbf{v}}_{0,1}\|_2 \\ &\quad + h \sum_{k=0}^{\infty} \left\| (\mathbf{I}_4 - h\mathbf{A}^{\text{ref}} \mathbf{M})^k \begin{pmatrix} 0 \\ \bar{\mathbf{v}}_{0,2} - \tilde{\mathbf{v}}_2 \end{pmatrix} \right\|_2 \\ &\leq h \sum_{k=0}^{\infty} \|(\mathbf{I}_4 - h\mathbf{A}^{\text{ref}} \mathbf{M})^k \mathbf{W} \mathbf{W}_1^{-1} \bar{\mathbf{v}}_{0,1}\|_2 \\ &\quad + h \sum_{k=0}^{\infty} \|(\mathbf{I}_4 - h\mathbf{A}^{\text{ref}} \mathbf{M})^k\|_2 \cdot (\|\bar{\mathbf{v}}_{0,2}\|_2 + \|\tilde{\mathbf{v}}_2\|_2) \\ (11), \text{ Step (5)} \quad &\leq O(m^{-1})O(m^\varepsilon) + O(1) (O(m^{\gamma_\psi-1}) + O(m^{\varepsilon-1})) \\ &\leq O(m^{\varepsilon+\gamma_\psi-1}). \quad \blacksquare \end{aligned}$$

Appendix G. Training Dynamics

In this section, we investigate how much the weights W_k change during training, which allows us to prove Theorem 10 at the end of this section. To this end, we first define important terms.

Definition G.1. For any sequence $(z_k)_{k \in \mathbb{N}_0}$, define

$$\Delta_k z := \max_{0 \leq l \leq k} |z_l - z_0| ,$$

where the supremum should be taken element-wise if z is a vector or a matrix. Moreover, let

$$\kappa_{u,k} := h \sum_{l=0}^k \|\mathbf{u}_l\|_\infty, \quad \tilde{\mathbf{Q}} := \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix}, \quad \mathbf{1}_3 := \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}, \quad \mathbf{1}_{3 \times 3} := \begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix} .$$

Now, we can state a general result, which resembles a first-order Taylor approximation:¹⁰

Proposition G.2. Let $k \in \mathbb{N}_0$, $\sigma \in \{\pm 1\}$ and $i \in I_\sigma$. Then, with $|\cdot|$ and \leq understood component-wise,

$$\begin{aligned} \Delta_k \boldsymbol{\theta}_i &\leq \kappa_{u,k} \tilde{\mathbf{Q}} |\boldsymbol{\theta}_{i,0}| + 2\kappa_{u,k}^2 e^{2\kappa_{u,k}} \|\boldsymbol{\theta}_{i,0}\|_\infty \mathbf{1}_3 \\ \Delta_k \boldsymbol{\Sigma}_\sigma &\leq \kappa_{u,k} (\tilde{\mathbf{Q}} |\boldsymbol{\Sigma}_{\sigma,0}| + |\boldsymbol{\Sigma}_{\sigma,0}| \tilde{\mathbf{Q}}) + 8\kappa_{u,k}^2 e^{4\kappa_{u,k}} \|\boldsymbol{\Sigma}_{\sigma,0}\|_\infty \mathbf{1}_{3 \times 3} . \end{aligned}$$

Proof The inequality

$$\|(\mathbf{A} + \mathbf{I})^2 - 2\mathbf{A} - \mathbf{I}\|_\infty \leq (\|\mathbf{A}\|_\infty + \|\mathbf{I}\|_\infty)^2 - 2\|\mathbf{A}\|_\infty - \|\mathbf{I}\|_\infty$$

for arbitrary matrices \mathbf{A} looks like an incorrect application of the triangle inequality due to the minus signs. However, it is correct since the subtracted terms exactly match terms in the expansion of the first term (since $\|\mathbf{I}\|_\infty = 1$):

$$\begin{aligned} \|(\mathbf{A} + \mathbf{I})^2 - 2\mathbf{A} - \mathbf{I}\|_\infty &= \|\mathbf{A}^2 + 2\mathbf{A} + \mathbf{I} - 2\mathbf{A} - \mathbf{I}\|_\infty = \|\mathbf{A}^2\|_\infty \\ &\leq \|\mathbf{A}\|_\infty^2 = \|\mathbf{A}\|_\infty^2 + 2\|\mathbf{A}\|_\infty + \|\mathbf{I}\|_\infty - 2\|\mathbf{A}\|_\infty - \|\mathbf{I}\|_\infty \\ &= (\|\mathbf{A}\|_\infty + \|\mathbf{I}\|_\infty)^2 - 2\|\mathbf{A}\|_\infty - \|\mathbf{I}\|_\infty . \end{aligned}$$

We can apply the same trick to obtain bounds on $|\boldsymbol{\theta}_{i,k} - \boldsymbol{\theta}_{i,0}|$ and $|\boldsymbol{\Sigma}_{\sigma,k} - \boldsymbol{\Sigma}_{\sigma,0}|$:¹¹ Define

$$\tilde{\mathbf{Q}}_k := h \sum_{l=0}^k \mathbf{Q}_{\sigma,l}, \quad \tilde{s}_k := h \sum_{l=0}^k \|\mathbf{Q}_{\sigma,l}\|_\infty .$$

Since

$$\boldsymbol{\theta}_{i,k} = (\mathbf{I}_3 + h\mathbf{Q}_{\sigma,k-1}) \cdots (\mathbf{I}_3 + h\mathbf{Q}_{\sigma,0}) \boldsymbol{\theta}_{i,0}$$

10. In the “first-order term”, the matrices are still sparse. “Higher-order” approximations are not useful for our purpose.

11. The bound on $\Delta_k \boldsymbol{\theta}_i$ and $\Delta_k \boldsymbol{\Sigma}_\sigma$ then follows since the bound is increasing in k .

$$\Sigma_{\sigma,k} = (\mathbf{I}_3 + h\mathbf{Q}_{\sigma,k-1}) \cdots (\mathbf{I}_3 + h\mathbf{Q}_{\sigma,0}) \Sigma_{\sigma,0} (\mathbf{I}_3 + h\mathbf{Q}_{\sigma,0}) \cdots (\mathbf{I}_3 + h\mathbf{Q}_{\sigma,k-1}),$$

we find with $1 + x \leq e^x$:

$$\begin{aligned} & \|\boldsymbol{\theta}_{i,k} - \tilde{\mathbf{Q}}_{k-1} \boldsymbol{\theta}_{i,0} - \boldsymbol{\theta}_{i,0}\|_{\infty} \\ & \leq (1 + h\|\mathbf{Q}_{\sigma,k-1}\|_{\infty}) \cdots (1 + h\|\mathbf{Q}_{\sigma,0}\|_{\infty}) \|\boldsymbol{\theta}_{i,0}\|_{\infty} - \tilde{s}_{k-1} \|\boldsymbol{\theta}_{i,0}\|_{\infty} - \|\boldsymbol{\theta}_{i,0}\|_{\infty} \\ & \leq (e^{\tilde{s}_{k-1}} - \tilde{s}_{k-1} - 1) \|\boldsymbol{\theta}_{i,0}\|_{\infty} \end{aligned}$$

and similarly

$$\begin{aligned} & \|\Sigma_{\sigma,k} - \tilde{\mathbf{Q}}_{k-1} \Sigma_{\sigma,0} - \Sigma_{\sigma,0} \tilde{\mathbf{Q}}_{k-1} - \Sigma_{\sigma,0}\|_{\infty} \\ & \leq (1 + h\|\mathbf{Q}_{\sigma,k-1}\|_{\infty}) \cdots (1 + h\|\mathbf{Q}_{\sigma,0}\|_{\infty}) \|\Sigma_{\sigma,0}\|_{\infty} (1 + h\|\mathbf{Q}_{\sigma,0}\|_{\infty}) \cdots (1 + h\|\mathbf{Q}_{\sigma,k-1}\|_{\infty}) \\ & \quad - (\tilde{s}_{k-1} \|\Sigma_{\sigma,0}\|_{\infty} + \|\Sigma_{\sigma,0}\|_{\infty} \tilde{s}_{k-1} + 1) \\ & \leq (e^{2\tilde{s}_{k-1}} - 2\tilde{s}_{k-1} - 1) \|\Sigma_{\sigma,0}\|_{\infty}. \end{aligned}$$

Observe that

$$e^x - x - 1 = \sum_{k=2}^{\infty} \frac{x^k}{k!} = x^2 \sum_{k=0}^{\infty} \frac{x^k}{(k+2)!} \stackrel{(k+2)! \geq 2k!}{\leq} \frac{1}{2} x^2 \sum_{k=0}^{\infty} \frac{x^k}{k!} = \frac{1}{2} x^2 e^x.$$

Obviously,

$$|\tilde{\mathbf{Q}}_k| \leq h \sum_{l=0}^k |\mathbf{Q}_{\sigma,l}| = h \begin{pmatrix} 0 & 0 & \sum_{l=0}^k |r_{\sigma,l}| \\ 0 & 0 & \sum_{l=0}^k |s_{\sigma,l}| \\ \sum_{l=0}^k |r_{\sigma,l}| & \sum_{l=0}^k |s_{\sigma,l}| & 0 \end{pmatrix} \leq \kappa_{u,k} \tilde{\mathbf{Q}}.$$

We also have $\tilde{s}_k \leq 2\kappa_{u,k}$ since

$$\|\mathbf{Q}_{\sigma,l}\|_{\infty} = \max_{i \in \{1, \dots, 3\}} \sum_{j=1}^3 |(\mathbf{Q}_{\sigma,l})_{ij}| = |r_{\sigma,l}| + |s_{\sigma,l}| \leq 2\|\mathbf{u}_l\|_{\infty}.$$

Aggregating the previous results and using $\kappa_{u,k-1} \leq \kappa_{u,k}$ yields

$$\begin{aligned} |\boldsymbol{\theta}_{i,k} - \boldsymbol{\theta}_{i,0}| & \leq |\tilde{\mathbf{Q}}_{k-1} \boldsymbol{\theta}_{i,0}| + \|\boldsymbol{\theta}_{i,k} - \tilde{\mathbf{Q}}_{k-1} \boldsymbol{\theta}_{i,0} - \boldsymbol{\theta}_{i,0}\|_{\infty} \mathbf{1}_3 \\ & \leq \kappa_{u,k} \tilde{\mathbf{Q}} |\boldsymbol{\theta}_{i,0}| + 2\kappa_{u,k}^2 e^{2\kappa_{u,k}} \|\boldsymbol{\theta}_{i,0}\|_{\infty} \mathbf{1}_3. \\ |\Sigma_{\sigma,k} - \Sigma_{\sigma,0}| & \leq |\tilde{\mathbf{Q}}_{k-1} \Sigma_{\sigma,0}| + |\Sigma_{\sigma,0} \tilde{\mathbf{Q}}_{k-1}| \\ & \quad + \|\Sigma_{\sigma,k} - \tilde{\mathbf{Q}}_{k-1} \Sigma_{\sigma,0} - \Sigma_{\sigma,0} \tilde{\mathbf{Q}}_{k-1} - \Sigma_{\sigma,0}\|_{\infty} \mathbf{1}_{3 \times 3} \\ & \leq \kappa_{u,k} (\tilde{\mathbf{Q}} |\Sigma_{\sigma,0}| + |\Sigma_{\sigma,0}| \tilde{\mathbf{Q}}) + 8\kappa_{u,k}^2 e^{4\kappa_{u,k}} \|\Sigma_{\sigma,0}\|_{\infty} \mathbf{1}_{3 \times 3}. \quad \blacksquare \end{aligned}$$

Corollary G.3. *Let Assumption F.2 be satisfied. If $\kappa_{u,k} \leq O(m^{\varepsilon + \gamma_{\psi} - 1})$ for some $k \in \mathbb{N}_0$ with bound independent of k , we have*

$$\Delta_k \boldsymbol{\theta}_i \leq O(m^{2\varepsilon + \gamma_{\psi}}) \begin{pmatrix} O(m^{-3/2}) \\ O(m^{-3/2}) \\ O(m^{-1}) \end{pmatrix},$$

$$\Delta_k \Sigma_\sigma \leq O(m^{2\varepsilon+\gamma_\psi}) \begin{pmatrix} O(m^{-1}) & O(m^{-1}) & O(1) \\ O(m^{-1}) & O(m^{-3/2}) & O(m^{-1}) \\ O(1) & O(m^{-1}) & O(m^{-1}) \end{pmatrix}$$

with a bound independent of k .

Proof Note that since $\varepsilon + \gamma_\psi < 1$, we have $e^{4\kappa_{u,k}} \leq O(1)$.

(a) By properties (W1), (W2), and (W3) in Proposition E.3, we have

$$|\boldsymbol{\theta}_{i,0}| = \begin{pmatrix} |a_{i,0}| \\ |b_{i,0}| \\ |w_{i,0}| \end{pmatrix} \leq \begin{pmatrix} m^\varepsilon \\ 0 \\ m^{\varepsilon-1/2} \end{pmatrix}.$$

We can now apply Proposition G.2 to obtain

$$\begin{aligned} |\boldsymbol{\theta}_{i,k} - \boldsymbol{\theta}_{i,0}| &\leq \kappa_{u,k} \tilde{\mathbf{Q}} |\boldsymbol{\theta}_{i,0}| + 2\kappa_{u,k}^2 e^{2\kappa_{u,k}} \|\boldsymbol{\theta}_{i,0}\|_\infty \mathbf{1}_3 \\ &\leq \kappa_{u,k} \begin{pmatrix} m^{\varepsilon-1/2} \\ m^{\varepsilon-1/2} \\ m^\varepsilon \end{pmatrix} + 2\kappa_{u,k}^2 e^{2\kappa_{u,k}} m^\varepsilon \mathbf{1}_3 \\ &\leq O(m^{2\varepsilon+\gamma_\psi}) \begin{pmatrix} O(m^{-3/2}) \\ O(m^{-3/2}) \\ O(m^{-1}) \end{pmatrix} + O(m^{\varepsilon+\gamma_\psi}) \begin{pmatrix} O(m^{-2}) \\ O(m^{-2}) \\ O(m^{-2}) \end{pmatrix} \\ &\stackrel{\varepsilon+\gamma_\psi \leq 1/2}{\leq} O(m^{2\varepsilon+\gamma_\psi}) \begin{pmatrix} O(m^{-3/2}) \\ O(m^{-3/2}) \\ O(m^{-1}) \end{pmatrix}. \end{aligned}$$

(b) By properties (W1), (W5), (W6) and (W7) in Proposition E.3, we have

$$|\Sigma_{\sigma,0}| = \begin{pmatrix} |\Sigma_{\sigma,a^2,0}| & |\Sigma_{\sigma,ab,0}| & |\Sigma_{\sigma,wa,0}| \\ |\Sigma_{\sigma,ab,0}| & |\Sigma_{\sigma,b^2,0}| & |\Sigma_{\sigma,wb,0}| \\ |\Sigma_{\sigma,wa,0}| & |\Sigma_{\sigma,wb,0}| & |\Sigma_{\sigma,w^2,0}| \end{pmatrix} = \begin{pmatrix} O(m) & 0 & O(m^\varepsilon) \\ 0 & 0 & 0 \\ O(m^\varepsilon) & 0 & O(1) \end{pmatrix}.$$

Since $\varepsilon \leq 1$, we can conclude $\|\Sigma_{\sigma,0}\|_\infty \leq O(m)$ and

$$\begin{aligned} \tilde{\mathbf{Q}} |\Sigma_{\sigma,0}| + |\Sigma_{\sigma,0}| \tilde{\mathbf{Q}} &= \begin{pmatrix} O(m^\varepsilon) & 0 & O(1) \\ O(m^\varepsilon) & 0 & O(1) \\ O(m) & 0 & O(m^\varepsilon) \end{pmatrix} + \begin{pmatrix} O(m^\varepsilon) & O(m^\varepsilon) & O(m) \\ 0 & 0 & 0 \\ O(1) & O(1) & O(m^\varepsilon) \end{pmatrix} \\ &= \begin{pmatrix} O(m^\varepsilon) & O(m^\varepsilon) & O(m) \\ O(m^\varepsilon) & 0 & O(1) \\ O(m) & O(1) & O(m^\varepsilon) \end{pmatrix}. \end{aligned}$$

We can now apply Proposition G.2 to obtain

$$\begin{aligned} &|\Sigma_{\sigma,k} - \Sigma_{\sigma,0}| \\ &\leq \kappa_{u,k} (\tilde{\mathbf{Q}} |\Sigma_{\sigma,0}| + |\Sigma_{\sigma,0}| \tilde{\mathbf{Q}}) + 8\kappa_{u,k}^2 e^{4\kappa_{u,k}} \|\Sigma_{\sigma,0}\|_\infty \mathbf{1}_{3 \times 3} \\ &\leq O(m^{2\varepsilon+\gamma_\psi}) \begin{pmatrix} O(m^{-1}) & O(m^{-1}) & O(1) \\ O(m^{-1}) & 0 & O(m^{-1}) \\ O(m^{-1}) & O(m^{-1}) & O(m^{-1}) \end{pmatrix} + O(m^{\varepsilon+\gamma_\psi-2}) \mathbf{1}_{3 \times 3} \end{aligned}$$

$$\varepsilon + \gamma_\psi \leq 1/2 \quad O(m^{2\varepsilon + \gamma_\psi}) \begin{pmatrix} O(m^{-1}) & O(m^{-1}) & O(1) \\ O(m^{-1}) & O(m^{-3/2}) & O(m^{-1}) \\ O(1) & O(m^{-1}) & O(m^{-1}) \end{pmatrix}. \quad \blacksquare$$

Remark G.4. We will prove in Proposition G.6 that the assumption of Corollary G.3 is satisfied. Although the first inequality of Corollary G.3 already provides bounds on the change of the individual weights, the second inequality is interesting as well because its bounds are stronger than what one would expect only from the individual weight bounds in the first inequality: For the sake of simplicity, pretend that $\varepsilon = \gamma_\psi = 0$. Then, for example, one could argue using the first inequality that

$$\begin{aligned} \Delta_k \Sigma_{\sigma, a^2} &= \max_{0 \leq l \leq k} \left| \sum_{i \in I_\sigma} (a_{i,l}^2 - a_{i,0}^2) \right| \leq \sum_{i \in I_\sigma} \max_{0 \leq l \leq k} |a_{i,l}^2 - a_{i,0}^2| \\ &= \sum_{i \in I_\sigma} \max_{0 \leq l \leq k} |a_{i,l} + a_{i,0}| \cdot |a_{i,l} - a_{i,0}| \leq \sum_{i \in I_\sigma} (|a_{i,0}| + \Delta_k a_i) \Delta_k a_i \\ &\leq \sum_{i \in I_\sigma} (O(1) + O(m^{-3/2})) O(m^{-3/2}) = |I_\sigma| O(m^{-3/2}) \leq O(m^{-1/2}), \end{aligned}$$

which is weaker than the bound $\Delta_k \Sigma_{\sigma, a^2} \leq O(m^{-1})$ obtained by the second inequality. These stronger bounds will be crucial in proving that the assumption $\kappa_{u,k} \leq O(m^{\varepsilon + \gamma_\psi - 1})$ of Corollary G.3 is satisfied. Also, note that for $\varepsilon = \gamma_\psi = 0$, the weakest bound

$$\Delta_k \Sigma_{\sigma, wa} \leq O(1)$$

cannot be improved: $\Sigma_{\sigma, wa} = p_\sigma$ is the slope of $f_{W, \tau, \sigma}$, which initially satisfies $|\Sigma_{\sigma, wa, 0}| \leq O(1)$ by (W7) and needs to converge to an p_σ^{opt} that is independent of $\Sigma_{\sigma, wa, 0}$ and also satisfies $|p_\sigma^{\text{opt}}| \leq \psi_{D,p} \leq O(1)$. Our proof works since $\Sigma_{\sigma, wa}$ only occurs in $h\mathbf{G}^{\text{wab}}$ with a small factor h , but neither in \mathbf{G}^{w} nor \mathbf{G}^{ab} .

We will soon use Corollary G.3 to prove its own assumption $\kappa_{u,k} \leq O(m^{\varepsilon + \gamma_\psi - 1})$. To this end, we first need a lemma that connects the reference system $\delta \bar{\mathbf{v}} = -h\mathbf{A}^{\text{ref}}\mathbf{M}\bar{\mathbf{v}}$ to the actual system $\delta \bar{\mathbf{v}} = -h\mathbf{A}\mathbf{M}\bar{\mathbf{v}}$.

Lemma G.5. For $m \geq 1$, let $\|\cdot\|$ denote an arbitrary vector norm on \mathbb{R}^m and its induced matrix norm. Let $k \in \mathbb{N}_0$, $\mathbf{K}_0, \dots, \mathbf{K}_{k-1} \in \mathbb{R}^{m \times m}$ and $\tilde{\mathbf{K}} \in \mathbb{R}^{m \times m}$. If

$$\delta_{k-1} := \sum_{l=0}^{k-1} \|\tilde{\mathbf{K}}^l\| \cdot \sup_{l \in \{0, \dots, k-1\}} \|\mathbf{K}_l - \tilde{\mathbf{K}}\| < 1,$$

where $\delta_{-1} := 0$, then each sequence $\mathbf{v}_0, \dots, \mathbf{v}_k \in \mathbb{R}^m$ with $\mathbf{v}_{l+1} = \mathbf{K}_l \mathbf{v}_l$ for all $l \in \{0, \dots, k-1\}$ satisfies

$$\sum_{l=0}^k \|\mathbf{v}_l\| \leq \frac{1}{1 - \delta_{k-1}} \sum_{l=0}^k \|\tilde{\mathbf{K}}^l \mathbf{v}_0\|.$$

Proof Clearly, for $l \in \{0, \dots, k-1\}$,

$$\mathbf{v}_{l+1} = \tilde{\mathbf{K}} \mathbf{v}_l + (\mathbf{K}_l - \tilde{\mathbf{K}}) \mathbf{v}_l$$

and hence, by induction on l ,

$$\mathbf{v}_l = \tilde{\mathbf{K}}^l \mathbf{v}_0 + \sum_{l'=0}^{l-1} \tilde{\mathbf{K}}^{l-1-l'} (\mathbf{K}_{l'} - \tilde{\mathbf{K}}) \mathbf{v}_{l'}$$

for all $l \in \{0, \dots, k\}$. Summing norms on both sides yields

$$\begin{aligned} \sum_{l=0}^k \|\mathbf{v}_l\| &\leq \sum_{l=0}^k \|\tilde{\mathbf{K}}^l \mathbf{v}_0\| + \sum_{l=0}^k \sum_{l'=0}^{l-1} \|\tilde{\mathbf{K}}^{l-1-l'} (\mathbf{K}_{l'} - \tilde{\mathbf{K}}) \mathbf{v}_{l'}\| \\ &= \sum_{l=0}^k \|\tilde{\mathbf{K}}^l \mathbf{v}_0\| + \sum_{l'=0}^{k-1} \sum_{l=l'+1}^k \|\tilde{\mathbf{K}}^{l-1-l'} (\mathbf{K}_{l'} - \tilde{\mathbf{K}}) \mathbf{v}_{l'}\| \\ &\leq \sum_{l=0}^k \|\tilde{\mathbf{K}}^l \mathbf{v}_0\| + \sum_{l'=0}^{k-1} \left(\sum_{l=0}^{k-1-l'} \|\tilde{\mathbf{K}}^l\| \right) \cdot \sup_{l \in \{0, \dots, k-1\}} \|\mathbf{K}_l - \tilde{\mathbf{K}}\| \cdot \|\mathbf{v}_{l'}\| \\ &\leq \sum_{l=0}^k \|\tilde{\mathbf{K}}^l \mathbf{v}_0\| + \delta_{k-1} \sum_{l'=0}^k \|\mathbf{v}_{l'}\|. \end{aligned}$$

Hence $(1 - \delta_{k-1}) \sum_{l=0}^k \|\mathbf{v}_l\| \leq \sum_{l=0}^k \|\tilde{\mathbf{K}}^l \mathbf{v}_0\|$ and since $\delta_{k-1} < 1$, the inequality is preserved when dividing by $1 - \delta_{k-1}$. \blacksquare

Proposition G.6. *Let Assumption F.2 be satisfied. We have*

$$\kappa_{u,k} \leq O(m^{\varepsilon+\gamma\psi^{-1}}),$$

where $\kappa_{u,k}$ was defined in Definition G.1 and the bound $O(m^{\varepsilon+\gamma\psi^{-1}})$ is independent of $k \in \mathbb{N}_0$.

Proof By Proposition C.5, we know that $\bar{\mathbf{v}}_{k+1} = (\mathbf{I}_4 - h\mathbf{A}_k \mathbf{M}_D) \bar{\mathbf{v}}_k$. We want to bound $\kappa_{u,k} = h \sum_{l=0}^k \|\mathbf{u}_l\|_\infty = h \sum_{l=0}^k \|\mathbf{B} \mathbf{M}_D \bar{\mathbf{v}}_l\|_\infty$ by comparing it to the reference system $\delta \bar{\mathbf{v}} = -h\mathbf{A}^{\text{ref}} \mathbf{M}_D \bar{\mathbf{v}}$ using Lemma G.5. Hence, we define

$$\begin{aligned} \delta_k &:= \sum_{l=0}^k \|(\mathbf{I}_4 - h\mathbf{A}^{\text{ref}} \mathbf{M}_D)^l\|_\infty \cdot \sup_{0 \leq l \leq k} \|(\mathbf{I}_4 - h\mathbf{A}^{\text{ref}} \mathbf{M}_D) - (\mathbf{I}_4 - h\mathbf{A}_l \mathbf{M}_D)\|_\infty \\ &= h \sum_{l=0}^k \|(I - h\mathbf{A}^{\text{ref}} \mathbf{M}_D)^l\|_\infty \cdot \sup_{0 \leq l \leq k} \|(\mathbf{A}_l - \mathbf{A}^{\text{ref}}) \mathbf{M}_D\|_\infty. \end{aligned} \quad (17)$$

For m large enough, we want to prove by induction that $\delta_k \leq 1/2$ for all $k \in \mathbb{N}_0$. Trivially, $\delta_{-1} = 0 \leq 1/2$. Now let $k \in \mathbb{N}_0$ with $\delta_{k-1} \leq 1/2$.

(1) By Lemma C.4, we have $\tilde{\mathbf{u}}_k = -\tilde{\mathbf{B}} \tilde{\mathbf{M}}_D \tilde{\mathbf{v}}_k$ and hence $\mathbf{u}_k = -\mathbf{B} \mathbf{M}_D \bar{\mathbf{v}}_k$. Thus,

$$\kappa_{u,k} = h \sum_{l=0}^k \|\mathbf{u}_l\|_\infty \leq \|\mathbf{B}\|_\infty \|\mathbf{M}_D\|_\infty \cdot h \sum_{l=0}^k \|\bar{\mathbf{v}}_l\|_\infty.$$

Because $\delta_{k-1} \leq 1/2$, we can apply Lemma G.5 and obtain

$$h \sum_{l=0}^k \|\bar{\mathbf{v}}_l\|_\infty \leq \frac{1}{1 - \delta_{k-1}} h \sum_{l=0}^k \|(\mathbf{I}_4 - h\mathbf{A}^{\text{ref}} \mathbf{M}_D)^l \bar{\mathbf{v}}_0\|_\infty$$

$$\begin{aligned} &\leq 2h \sum_{l=0}^{\infty} \|(\mathbf{I}_4 - h\mathbf{A}^{\text{ref}}\mathbf{M}_D)^l \bar{\mathbf{v}}_0\|_{\infty} \\ &\stackrel{\text{Proposition F.4}}{\leq} O(m^{\varepsilon+\gamma_{\psi}-1}). \end{aligned}$$

Norm equivalence (cf. Definition A.3) yields

$$\|\mathbf{M}_D\|_{\infty} \leq O(\|\mathbf{M}_D\|_2) = O(\lambda_{\max}(\mathbf{M}_D)) \stackrel{\text{Assumption F.2}}{=} O(1). \quad (18)$$

Hence, we can write

$$\kappa_{u,k} = O(m^{\varepsilon+\gamma_{\psi}-1}), \quad (19)$$

where, in accordance with Definition A.2, the constant in $O(m^{\varepsilon+\gamma_{\psi}-1})$ does not depend on the induction step k .

(2) Let us investigate the components of Eq. (17):

$$\begin{aligned} h \sum_{l=0}^k \|(\mathbf{I}_4 - h\mathbf{A}^{\text{ref}}\mathbf{M}_D)^l\|_{\infty} &\leq h \sum_{l=0}^{\infty} \|(\mathbf{I}_4 - h\mathbf{A}^{\text{ref}}\mathbf{M}_D)^l\|_{\infty} \stackrel{\text{Proposition F.4}}{=} O(1) \\ (\mathbf{A}_l - \mathbf{A}^{\text{ref}})\mathbf{M}_D &= \mathbf{B} \left((\mathbf{G}_l^{\text{w}} - \mathbf{G}_0^{\text{w}}) + (\mathbf{G}_l^{\text{ab}} - \mathbf{G}_0^{\text{ab}}) + h\mathbf{G}_l^{\text{wab}} \right) \mathbf{B}\mathbf{M}_D \\ \Rightarrow \|(\mathbf{A}_l - \mathbf{A}^{\text{ref}})\mathbf{M}_D\|_{\infty} &\stackrel{(18)}{=} O(1) \cdot (\|\mathbf{G}_l^{\text{w}} - \mathbf{G}_0^{\text{w}}\|_{\infty} + \|\mathbf{G}_l^{\text{ab}} - \mathbf{G}_0^{\text{ab}}\|_{\infty} + h\|\mathbf{G}_l^{\text{wab}}\|_{\infty}). \end{aligned}$$

First of all, for $0 \leq l \leq k$,

$$\|\mathbf{G}_l^{\text{w}} - \mathbf{G}_0^{\text{w}}\|_{\infty} \leq \max_{\sigma \in \{\pm 1\}} \Delta_k \Sigma_{\sigma, w^2} \stackrel{\text{Corollary G.3}}{\leq} O(m^{2\varepsilon+\gamma_{\psi}-1}).$$

Similarly, for $0 \leq l \leq k$,

$$\|\mathbf{G}_l^{\text{ab}} - \mathbf{G}_0^{\text{ab}}\|_{\infty} \leq \max_{\sigma \in \{\pm 1\}} (\Delta_k \Sigma_{\sigma, a^2} + \Delta_k \Sigma_{\sigma, ab} + \Delta_k \Sigma_{\sigma, b^2}) \leq O(m^{2\varepsilon+\gamma_{\psi}-1}).$$

Observe that

$$h|r_{\sigma,l}| \leq h\|\mathbf{u}_l\|_{\infty} \leq h \sum_{l'=0}^k \|\mathbf{u}_{l'}\|_{\infty} = \kappa_{u,k} \stackrel{(19)}{=} O(m^{\varepsilon+\gamma_{\psi}-1})$$

and similarly $h|s_{\sigma,l}| \leq O(m^{\varepsilon+\gamma_{\psi}-1})$. Thus, we find

$$\begin{aligned} h\|\mathbf{G}_l^{\text{wab}}\|_{\infty} &= \max_{\sigma \in \{\pm 1\}} h|r_{\sigma,l}\Sigma_{\sigma, wa,l} + s_{\sigma,l}\Sigma_{\sigma, wb,l}| \\ &= O(m^{\varepsilon+\gamma_{\psi}-1}) \cdot \left(\max_{\sigma \in \{\pm 1\}} |\Sigma_{\sigma, wa,l}| + |\Sigma_{\sigma, wb,l}| \right). \end{aligned}$$

Similar to the other calculations, we can compute for $0 \leq l \leq k$

$$\begin{aligned} |\Sigma_{\sigma, wa,l}| &\leq |\Sigma_{\sigma, wa,0}| + \Delta_k \Sigma_{\sigma, wa} \stackrel{(W7)}{\leq} O(m^{\varepsilon}) + O(m^{2\varepsilon+\gamma_{\psi}}) = O(m^{2\varepsilon+\gamma_{\psi}}) \\ |\Sigma_{\sigma, wb,l}| &\leq |\Sigma_{\sigma, wb,0}| + \Delta_k \Sigma_{\sigma, wb} \leq 0 + O(m^{2\varepsilon+\gamma_{\psi}-1}) = O(m^{2\varepsilon+\gamma_{\psi}-1}), \end{aligned}$$

which yields $h\|\mathbf{G}_l^{\text{wab}}\|_\infty \leq O(m^{3\varepsilon+2\gamma_\psi-1})$.

We can now revisit the beginning of step (2) to obtain

$$\|(\mathbf{A}_l - \mathbf{A}^{\text{ref}})\mathbf{M}_D\|_\infty \leq O(m^{3\varepsilon+2\gamma_\psi-1})$$

and

$$\begin{aligned} \delta_k &\stackrel{(17)}{=} h \sum_{l=0}^k \|(\mathbf{I}_4 - h\mathbf{A}^{\text{ref}}\mathbf{M}_D)^l\|_\infty \cdot \sup_{0 \leq l \leq k} \|(\mathbf{A}_l - \mathbf{A}^{\text{ref}})\mathbf{M}_D\|_\infty \\ &= O(1) \cdot O(m^{3\varepsilon+2\gamma_\psi-1}) = O(m^{3\varepsilon+2\gamma_\psi-1}) . \end{aligned}$$

We have shown that $\delta_{k-1} \leq 1/2$ implies $\delta_k \leq O(m^{3\varepsilon+2\gamma_\psi-1})$, where the constant in $O(m^{3\varepsilon+2\gamma_\psi-1})$ does not depend on k . Since $3\varepsilon + 2\gamma_\psi < 1$ by Assumption F.2, we have $\lim_{m \rightarrow \infty} m^{3\varepsilon+2\gamma_\psi-1} = 0$ and there exists $m_0 \in \mathbb{N}_0$ such that for all $m \geq m_0$ and $k \in \mathbb{N}_0$, $\delta_{k-1} \leq 1/2$ implies $\delta_k \leq 1/2$ and the induction works. Thus, for all $m \geq m_0$ and $k \in \mathbb{N}_0$, we know that $\delta_{k-1} \leq 1/2$ and we can apply step (1) to obtain

$$\kappa_{u,k} = O(m^{\varepsilon+\gamma_\psi-1}) . \quad \blacksquare$$

We can now prove our main theorem:

Proof of Theorem 10

Since $\gamma_\psi + \gamma_{\text{data}} + \gamma_P < 1/2$, we can assume without loss of generality that

$$0 < \varepsilon < \frac{1/2 - (\gamma_\psi + \gamma_{\text{data}} + \gamma_P)}{3} .$$

Moreover, if (W1) – (W7) from Proposition E.3 are satisfied, we have (similar to Eq. (12) in the proof of Proposition F.4)

$$\begin{aligned} \lambda_{\max}(\mathbf{H}) &= \lambda_{\max}(\mathbf{M}_D^{1/2} \mathbf{A}^{\text{ref}} \mathbf{M}_D^{1/2}) \leq \|\mathbf{M}_D^{1/2}\|_2^2 \lambda_{\max}(\mathbf{A}^{\text{ref}}) = \lambda_{\max}(\mathbf{M}_D) \lambda_{\max}(\mathbf{A}^{\text{ref}}) \\ &\stackrel{\text{Lemma F.3}}{\leq} K_M \Theta(m) . \end{aligned}$$

Hence, there exists a constant $C_{\text{lr}}^{-1} > 0$ with $\lambda_{\max}(\mathbf{H}) \leq C_{\text{lr}}^{-1} m$. For this choice, assumption (3) yields

$$h \leq C_{\text{lr}} m^{-1} \leq \frac{1}{\lambda_{\max}(\mathbf{H})} .$$

Therefore, Assumption F.2 is satisfied whenever the initialization satisfies (W1) – (W7).

By our choice of ε , we have

$$2\varepsilon + \gamma_\psi - 3/2 < -1 - \gamma_P - \gamma_{\text{data}} \leq -1 - \gamma_P . \quad (20)$$

For $\underline{x} := K_{\text{data}} m^{-\gamma_{\text{data}}}$, we then obtain using (W4), Corollary G.3 and Proposition G.6:

$$(|a_{i,0}| - |a_{i,0} - a_{i,k}|) \underline{x} \geq \left(\Omega(m^{-1-\gamma_P}) - O(m^{2\varepsilon+\gamma_\psi-3/2}) \right) \Theta(m^{-\gamma_{\text{data}}}) \stackrel{(20)}{=} \Omega(m^{-1-\gamma_P-\gamma_{\text{data}}})$$

$$|b_{i,k}| = |b_{i,k} - b_{i,0}| \leq O(m^{2\varepsilon + \gamma_\psi - 3/2}) .$$

Using (20) again, we find that there exists m_0 such that for all $m \geq m_0$ and all $i \in I, k \in \mathbb{N}_0$,

$$|b_{i,k}| < (|a_{i,0}| - |a_{i,0} - a_{i,k}|)\underline{x} ,$$

which, in terms of Lemma B.2, means $W_k \in \mathcal{S}_{W_0}(\underline{x})$. Since $\underline{x} \leq \underline{x}_D$ by assumption, we also have $W_k \in \mathcal{S}_{W_0}(\underline{x}_D)$. But then, Lemma B.2 tells us that $\nabla L_{D,\tau}(W_k) = \nabla L_D(W_k)$ and that $f_W|_{[\underline{x},\infty)} = f_{W,\tau,1}|_{[\underline{x},\infty)}$ as well as $f_W|_{(-\infty,-\underline{x}]} = f_{W,\tau,-1}|_{(-\infty,-\underline{x}]}$ are affine. Hence, $(W_k)_{k \in \mathbb{N}_0}$ satisfies the original gradient descent iteration

$$W_{k+1} = W_k - h \nabla L_D(W_k)$$

and (v) is satisfied. Moreover, by Corollary G.3, we obtain (i), (ii) and (iii) (up to a factor 2 in front of ε , which can be resolved by shrinking ε):

$$\begin{aligned} |a_{i,k} - a_{i,0}| &\leq O(m^{2\varepsilon + \gamma_\psi - 3/2}) \\ |b_{i,k} - b_{i,0}| &\leq O(m^{2\varepsilon + \gamma_\psi - 3/2}) \\ |w_{i,k} - w_{i,0}| &\leq O(m^{2\varepsilon + \gamma_\psi - 1}) . \end{aligned}$$

In order to find a similar bound for c , we recall from Lemma C.4 that $\delta c = h(\hat{s}_1 + \hat{s}_{-1})$, from Definition G.1 that $\kappa_{u,k} = h \sum_{l=0}^k \|\mathbf{u}_l\|_\infty$ and from Definition C.1 that

$$\mathbf{u} = \mathbf{B} \begin{pmatrix} \hat{r}_1 \\ \hat{r}_{-1} \\ \hat{s}_1 \\ \hat{s}_{-1} \end{pmatrix} \quad \text{and therefore} \quad \begin{pmatrix} \hat{r}_1 \\ \hat{r}_{-1} \\ \hat{s}_1 \\ \hat{s}_{-1} \end{pmatrix} = \mathbf{B}^{-1} \mathbf{u} ,$$

since \mathbf{B} is invertible due to $|\alpha| \neq 1$. Because \mathbf{B} is fixed, we therefore obtain

$$\begin{aligned} |c_k - c_0| &\leq \sum_{l=0}^{k-1} |c_{l+1} - c_l| = \sum_{l=0}^{k-1} |\delta c_l| = h \sum_{l=0}^{k-1} |\hat{s}_{1,l} + \hat{s}_{-1,l}| \leq O(\kappa_{u,k}) \\ &\stackrel{\text{Proposition G.6}}{\leq} O(m^{2\varepsilon + \gamma_\psi - 1}) , \end{aligned}$$

which shows (iv) after rescaling ε .

All of this holds under the assumption that $m \geq m_0$ and (W1) – (W7), where m_0 is independent of W_0, k, h . By Proposition E.3, the assumption holds with probability $\geq 1 - O(m^{-\gamma_P})$. \blacksquare

Appendix H. Multi-Dimensional Inputs

In the following, we investigate the case where

- the one-dimensional x values of D are projected onto a line in a d -dimensional input space ($d \in \mathbb{N}$), and
- a two-layer neural network with d -dimensional input is trained on this (degenerate) data set.

Hence, if we know that there exists a bound $B \in (0, \infty)$ with $p_X(x) \leq B$ for all $x \in \mathbb{R}$, then

$$p_{X+Y}(x) \leq B \int_{-\infty}^{\infty} p_Y(y) \, dy = B .$$

Moreover, if p_X and p_Y are symmetric, then p_{X+Y} is also symmetric:

$$\begin{aligned} p_{X+Y}(x) &= \int_{\mathbb{R}} p_X(x-y)p_Y(y) \, dy = \int_{\mathbb{R}} p_X(y-x)p_Y(-y) \, dy \\ &\stackrel{y':=-y}{=} \int_{\mathbb{R}} p_X((-x)-y')p_Y(y') \, dy' = p_{X+Y}(-x) . \end{aligned}$$

This directly yields (Q1) for Z_a .

(Q2) Since Z_a can be written as a linear combination of random variables that satisfy (Q2), Z_a must also satisfy (Q2) by the Minkowski inequality.

By Eq. (21), we obtain

$$f_{\tilde{W}_k}(zx) = f_{W_k}(x)$$

for all $k \in \mathbb{N}_0$ and $x \in \mathbb{R}$. Especially, we can apply Theorems 10, 12 and 22 and obtain that under the assumptions of these theorems, the kinks of $x \mapsto f_{\tilde{W}_k}(zx)$ do not cross the data points with high probability.

Since the assumptions of the theorems are (up to modifying constants) invariant under multiplying the x_j by a positive constant, we can also allow $\|z\|_2 \neq 1$ as long as $z \neq 0$.

Appendix I. Inconsistency Proofs

In this section, we give proofs of the inconsistency results in Section 6.

Proof of Corollary 24 Let $\gamma_\psi \in (0, 1/2)$ be sufficiently large such that $\frac{1}{2-2\gamma_\psi} \geq 1 - \varepsilon$. We can then choose $K_{\text{param}} > 0$ sufficiently large such that

$$m_n \leq K_{\text{param}} n^{\frac{1}{2-2\gamma_\psi}}$$

for all $n \geq 1$. Choose $K_{\text{data}} > 0$ such that $P_X^{\text{data}}([-K_{\text{data}}, K_{\text{data}}]) = 0$. Moreover, let $\gamma_{\text{data}} = 0$ and $\gamma_P > 0$ such that $\gamma_P + \gamma_{\text{data}} + \gamma_\psi < 1/2$.

Let C_{lr} be the corresponding constant from Theorem 22. Since $m_n \rightarrow \infty$ and $h_n < o(m_n^{-1})$, there exists an n_0 such that for all $n \geq n_0$, we have

$$h_n \leq C_{\text{lr}} m_n^{-1} .$$

By Theorem 22 we hence obtain for all $n \geq n_0$ that f_{W_k} is affine on $(-\infty, -K_{\text{data}}]$ and $[K_{\text{data}}, \infty)$ with probability $\geq 1 - C_P m_n^{-\gamma_P} \rightarrow 1$ ($n \rightarrow \infty$). But such a function satisfies $f_{W_k} \in \mathcal{F}_{K_{\text{data}}}$ and, because $P_X^{\text{data}}([-K_{\text{data}}, K_{\text{data}}]) = 0$, we obtain

$$R_{P^{\text{data}}}(f_{W_k}) \geq R_{P^{\text{data}}, K_{\text{data}}}^* = R_{P^{\text{data}}, 0}^* \stackrel{\text{(P4)}}{>} R_{P^{\text{data}}}^* ,$$

which yields inconsistency. \blacksquare

Proof of Corollary 26 Consider an NN as in Corollary 24 that is inconsistent on a distribution P^{data} on $\mathbb{R} \times \mathbb{R}$. Furthermore, fix an arbitrary vector $\mathbf{z} \in \mathbb{R}^d$ with $\|\mathbf{z}\|_2 = 1$.

For $(x, y) \sim P^{\text{data}}$, let \tilde{P}^{data} denote the distribution of $(x\mathbf{z}, y)$. It is easy to show that the optimal population risks satisfy

$$R_{\tilde{P}^{\text{data}}}^* = R_{P^{\text{data}}}^* . \quad (23)$$

Let $D \sim (P^{\text{data}})^n$, i.e., D consists of n i.i.d. data points $(x_j, y_j) \sim P^{\text{data}}$, then $\tilde{D} \sim (\tilde{P}^{\text{data}})^n$ with \tilde{D} defined in Remark H.1. Let \tilde{W}_0 be independent from \tilde{D} and initialized analogous to Assumption 2 as discussed in Remark H.1. Let

$$\tilde{W}_{k+1} = \tilde{W}_k - h_n \nabla L_{\tilde{D}}(\tilde{W}_k) .$$

Let $W_k := \mathbf{Z}\tilde{W}_k$. As shown in Remark H.1, W_0 satisfies Assumption 2 and $(W_k)_{k \in \mathbb{N}_0}$ arises from gradient descent on D :

$$W_{k+1} = W_k - h_n \nabla L_D(W_k) .$$

Moreover, $f_{\tilde{W}_k}(x\mathbf{z}) = f_{W_k}(x)$ for all $x \in \mathbb{R}, k \in \mathbb{N}_0$ and therefore

$$R_{\tilde{P}^{\text{data}}}(f_{\tilde{W}_k}) = R_{P^{\text{data}}}(f_{W_k}) \quad (24)$$

for all $k \in \mathbb{N}_0$. Since the NN with one-dimensional input is inconsistent on P^{data} , the NN with d -dimensional input must be inconsistent on \tilde{P}^{data} by (23) and (24). \blacksquare

In order to prove Corollary 25, we first show that if functions that are affine on $\mathbb{R} \setminus \{0\}$ cannot approach the Bayes risk $R_{P^{\text{data}}}^*$, there exists $\delta > 0$ such that the same holds for functions that are affine on $\mathbb{R} \setminus [-\delta, \delta]$.

Lemma I.1. *Let P^{data} be a bounded distribution on $\mathbb{R} \times \mathbb{R}$ satisfying (P1). For $\delta \geq 0$, define*

$$\mathcal{F}_\delta := \{f : \mathbb{R} \rightarrow \mathbb{R} \mid f \text{ affine on } (-\infty, -\delta) \text{ and } (\delta, \infty)\}, \quad R_{P^{\text{data}}, \delta}^* := \inf_{f \in \mathcal{F}_\delta} R_{P^{\text{data}}}(f) .$$

Then, $\lim_{\delta \searrow 0} R_{P^{\text{data}}, \delta}^* = R_{P^{\text{data}}, 0}^*$.

Proof In the case $\eta = \infty$, we obviously have $R_{P^{\text{data}}, \delta}^* = R_{P^{\text{data}}, 0}^*$ for all $\delta \geq 0$ with $P_X^{\text{data}}((-\delta, \delta)) = 0$ and we are done.

Obviously, $R_{P^{\text{data}}, \delta}^*$ is non-increasing in δ . In order to derive a contradiction, assume that $\lim_{\delta \searrow 0} R_{P^{\text{data}}, \delta}^* < R_{P^{\text{data}}, 0}^*$.

For $\delta' > 0, \sigma \in \{\pm 1\}$, consider $P_\sigma := P^{\text{data}}(\cdot \mid \sigma X > \delta')$. For sufficiently small δ' , P is well-defined and P_X is not only concentrated on a single x value due to (P1). Now, fix such a $\delta' > 0$.

For $f \in \mathcal{F}_\delta$, define $\mathbf{v}_\sigma(f) \in \mathbb{R}^{2 \times 2}$ as the slope and intercept of f on $\sigma(\delta, \infty)$. As in Definition 8, we can construct an invertible matrix $\mathbf{M}_{P_\sigma} = \mathbf{M}_{P_{\sigma, \delta}}$ and a vector $\mathbf{v}_{P_\sigma}^{\text{opt}} = \mathbf{v}_{P_{\sigma, \delta}}^{\text{opt}}$. Analogous to Remark D.6, we obtain for $\delta \leq \delta'$:

$$R_P(f) \geq (\mathbf{v}_\sigma(f) - \mathbf{v}_{P_\sigma}^{\text{opt}})^\top \mathbf{M}_{P_\sigma} (\mathbf{v}_\sigma(f) - \mathbf{v}_{P_\sigma}^{\text{opt}}) .$$

Hence, for $f \in \mathcal{F}_\delta$ with $\delta \leq \delta'$ and $R_{P^{\text{data}}}(f) \leq R_{P^{\text{data}},0}^*$, we obtain

$$\begin{aligned} R_{P^{\text{data}},0}^* &\geq R_{P^{\text{data}}}(f) \geq \frac{1}{2} \mathbb{E}_{(x,y) \sim P^{\text{data}}} \mathbb{1}_{\sigma(\delta', \infty)}(x) (y - f(x))^2 \\ &= \frac{1}{2} P_X^{\text{data}}(\sigma(\delta', \infty)) \cdot \mathbb{E}_{(x,y) \sim P_\sigma} (y - f(x))^2 \\ &\geq \frac{1}{2} P_X^{\text{data}}(\sigma(\delta', \infty)) \cdot (\mathbf{v}_\sigma(f) - \mathbf{v}_{P_\sigma}^{\text{opt}})^\top \mathbf{M}_{P_\sigma} (\mathbf{v}_\sigma(f) - \mathbf{v}_{P_\sigma}^{\text{opt}}). \end{aligned}$$

Since $\mathbf{M}_{P_\sigma} \succ 0$ and $P_X^{\text{data}}(\sigma(\delta', \infty)) > 0$, there has to exist a constant C such that $\|\mathbf{v}_\sigma(f)\|_\infty \geq C$ for all such f and $\sigma \in \{\pm 1\}$.

Now, pick $f \in \mathcal{F}_\delta$, $0 < \delta \leq \delta'$, with $R_{P^{\text{data}}}(f) \leq R_{P^{\text{data}},0}^*$ and let $f_0 \in \mathcal{F}_{\text{hsal}}$ be its affine continuation (i.e., $f_0(x) = f(x)$ for $|x| > \delta$). Then, $|f_0(x)| \leq C(1 + |x|)$ for $x > 0$ and therefore

$$\begin{aligned} R_{P^{\text{data}},0}^* &\leq R_{P^{\text{data}}}(f_0) \\ &= \frac{1}{2} \mathbb{E}_{(x,y) \sim P^{\text{data}}} \mathbb{1}_{[-\delta, \delta]}(x) (y - f_0(x))^2 + \frac{1}{2} \mathbb{E}_{(x,y) \sim P^{\text{data}}} \mathbb{1}_{[-\delta, \delta]^c}(x) (y - f_0(x))^2 \\ &\leq \frac{1}{2} \mathbb{E}_{(x,y) \sim P^{\text{data}}} \mathbb{1}_{[-\delta, \delta]}(x) 2(y^2 + f_0(x)^2) + R_{P^{\text{data}}}(f) \\ &\leq \underbrace{\mathbb{E}_{(x,y) \sim P^{\text{data}}} \mathbb{1}_{[-\delta, \delta]}(x) y^2}_{\rightarrow 0 \quad (\delta \rightarrow 0)} + \underbrace{P_X^{\text{data}}([-\delta, \delta]) \cdot (C(1 + |\delta|))^2}_{\rightarrow 0 \quad (\delta \rightarrow 0)} + R_{P^{\text{data}}}(f). \end{aligned}$$

Since we can choose $R_{P^{\text{data}}}(f)$ arbitrarily close to $R_{P^{\text{data}},\delta}^*$, it follows that $R_{P^{\text{data}},0}^* \leq \lim_{\delta \searrow 0} R_{P^{\text{data}},\delta}^*$, which contradicts our initial assumption. \blacksquare

Proof of Corollary 25 Since

$$\begin{aligned} \frac{1}{\gamma\eta} &< \frac{1}{2}, \\ 1 - \frac{1}{2\gamma} + \frac{1}{\eta\gamma} &= 1 - \frac{1 - \frac{2}{\eta}}{2\gamma} = \frac{2\gamma - \left(1 - \frac{2}{\eta}\right)}{2\gamma} < \frac{1}{2}, \end{aligned}$$

there exist some

$$\begin{aligned} \gamma_\psi &\geq \max \left\{ 0, 1 - \frac{1}{2\gamma} \right\} \\ \gamma_{\text{data}} &> \frac{1}{\gamma\eta} \\ \gamma_P &> 0 \end{aligned}$$

such that $\gamma_\psi + \gamma_{\text{data}} + \gamma_P < 1/2$. We then have

$$\begin{aligned} \gamma &\leq \frac{1}{2 - 2\gamma_\psi} \\ 1 - \gamma\eta\gamma_{\text{data}} &< 0 \end{aligned}$$

and therefore

$$\begin{aligned} m_n &\leq O\left(n^{\frac{1}{2-2\gamma_\psi}}\right) \\ O(nm_n^{-\eta\gamma_{\text{data}}}) &= o(1) \\ O(m_n^{-\gamma_P}) &= o(1). \end{aligned}$$

Let C_{lr} be the corresponding constant from Theorem 22. Since $m_n \rightarrow \infty$ and $h_n < o(m_n^{-1})$, there exists an n_0 such that for all $n \geq n_0$, we have

$$h_n \leq C_{\text{lr}} m_n^{-1}.$$

Hence, by Theorem 22, we obtain for all $n \geq n_0$ that $f_{W_k} \in \mathcal{F}_{K_{\text{data}} m_n^{-\gamma_{\text{data}}}}$ with probability $\geq 1 - C_P(m_n^{-\gamma_P} + nm_n^{-\eta\gamma_{\text{data}}}) \rightarrow 1$ for $n \rightarrow \infty$. By assumption (P4), we have $R_{P_{\text{data}},0}^* > R_{P_{\text{data}}}^*$ and by Lemma I.1, there exists $\delta > 0$ such that $R_{P_{\text{data}},\delta}^* > R_{P_{\text{data}}}^*$. For n sufficiently large, we have $K_{\text{data}} m_n^{-\gamma_{\text{data}}} \leq \delta$ and therefore

$$R_{P_{\text{data}}}(f_{W_k}) \geq R_{P_{\text{data}},\delta}^* > R_{P_{\text{data}}}^*$$

with probability $\geq 1 - C_P(m_n^{-\gamma_P} + nm_n^{-\eta\gamma_{\text{data}}}) \rightarrow 1$ for $n \rightarrow \infty$. This shows inconsistency. \blacksquare

Appendix J. Miscellaneous

In this section, we prove a fact that has been mentioned in the main paper.

Lemma J.1. *Let $D = ((x_1, y_1), \dots, (x_n, y_n)) \in (\mathbb{R} \times \mathbb{R})^n$ with $n \geq 1$ and $x_j \neq 0$ for all j . Then, by adding three points to D , we can achieve that $\mathbf{M}_{D,\sigma}$ is invertible for both $\sigma \in \{\pm 1\}$ and that $\psi_{D,q} = 0$.*

Proof By adding at most one point to D , we can ensure that both D_1 and D_{-1} are nonempty. Now consider the case of D_1 (D_{-1} can be handled analogously). For $x' := 1 + \max_{(x,y) \in D_1} x$ and $y' \in \mathbb{R}$ yet to be specified, consider the data set $\tilde{D} := D \cup \{(x', y')\}$. Since the kernels of two different matrices $\mathbf{M}_{x_j} \succeq 0$ and $\mathbf{M}_{x'} \succeq 0$ only intersect in zero, we have

$$\mathbf{M}_{\tilde{D},1} = \frac{1}{n+1} \left(\mathbf{M}_{x'} + \sum_{(x,y) \in D_1} \mathbf{M}_x \right) \succ 0,$$

i.e., $\mathbf{M}_{\tilde{D},1}$ is invertible. Moreover, we have

$$\begin{aligned} \begin{pmatrix} p_{\tilde{D},1}^{\text{opt}} \\ q_{\tilde{D},1}^{\text{opt}} \end{pmatrix} &= \mathbf{v}_{\tilde{D},1}^{\text{opt}} = \mathbf{M}_{\tilde{D},1}^{-1} \mathbf{u}_{\tilde{D},1}^0 = \frac{1}{n+1} \left(\mathbf{u}_{(x',y')} + \sum_{(x,y) \in D_1} \mathbf{u}_{(x,y)}^0 \right) \\ &= \frac{y'}{n+1} \mathbf{M}_{\tilde{D},1}^{-1} \begin{pmatrix} x' \\ 1 \end{pmatrix} + \frac{n}{n+1} \mathbf{M}_{\tilde{D},1}^{-1} \mathbf{u}_{D,1}^0 \end{aligned}$$

We need to show that we can choose y' such that $q_{\tilde{D},1}^{\text{opt}} = 0$. Assume the contrary, which means that there exists $z \in \mathbb{R}$ with

$$\mathbf{M}_{\tilde{D},1}^{-1} \begin{pmatrix} x' \\ 1 \end{pmatrix} = \begin{pmatrix} z' \\ 0 \end{pmatrix}$$

or, equivalently,

$$\begin{pmatrix} x' \\ 1 \end{pmatrix} = \mathbf{M}_{\tilde{D},1} \begin{pmatrix} z' \\ 0 \end{pmatrix} = \frac{z'}{n+1} \left(\begin{pmatrix} x' \\ 1 \end{pmatrix} + \sum_{(x,y) \in D_1} \begin{pmatrix} x \\ 1 \end{pmatrix} \right).$$

Since D_1 is nonempty by assumption and all $(x, y) \in D_1$ satisfy $x < x'$, we obtain the desired contradiction.

Overall, we can therefore satisfy $M_{D,\sigma} \succ 0$ for both $\sigma \in \{\pm 1\}$ by adding at most point to D and we can then satisfy $\psi_{D,q} = 0$ by adding at most two more points to D . \blacksquare

Appendix K. NTK Relation Proofs

Proposition K.1. *Let $W \in \mathbb{R}^{3m+1}$ be a parameter vector that induces the same activation pattern as W_0 on D , i.e.,*

$$\text{sgn}(a_i x_j + b_i) = \text{sgn}(a_{i,0} x_j + b_{i,0})$$

for all $i \in I, j \in J$, such that $f_{W,\tau,\text{sgn}(x_j)}(x_j) = f_W(x_j)$ for all $j \in J$. By the proof of Lemma B.2, this is satisfied e.g. for $W \in \mathcal{S}_{W_0}(\underline{x}_D)$. Then, following the notation from Section 8 and Definition C.1, we have

$$\begin{aligned} \mathbf{K} &= \mathbf{X} \mathbf{A} \mathbf{X}^\top \\ \mathbf{M}_D &= \frac{1}{n} \mathbf{X}^\top \mathbf{X} \\ \bar{\mathbf{v}} &= (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \bar{\mathbf{f}}, \end{aligned}$$

where $h = 0$ is used in the definition of \mathbf{A} .

Proof We first show $\mathbf{K}_{jk} = [\mathbf{X} \mathbf{A} \mathbf{X}^\top]_{jk}$ for $j, k \in J$. Recall from Section 8 and Definition C.1 that

$$\mathbf{X} = \begin{pmatrix} x_1 & 0 & 1 & 0 \\ \vdots & \vdots & \vdots & \vdots \\ x_{n'} & 0 & 1 & 0 \\ 0 & x_{n'+1} & 0 & 1 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & x_n & 0 & 1 \end{pmatrix},$$

$$\mathbf{A} = \begin{pmatrix} \Sigma_{1,a^2} + \Sigma_{1,w^2} & 0 & \Sigma_{1,ab} & 0 \\ 0 & \Sigma_{-1,a^2} + \Sigma_{-1,a^2} & 0 & \Sigma_{-1,ab} \\ \Sigma_{1,ab} & 0 & 1 + \Sigma_{1,b^2} + \Sigma_{1,w^2} & 1 \\ 0 & \Sigma_{-1,ab} & 1 & 1 + \Sigma_{-1,b^2} + \Sigma_{-1,w^2} \end{pmatrix}.$$

We consider several cases:

- Case 1: $x_j > 0$ and $x_k < 0$. Then, since no neuron is activated for both x_j and x_k ,

$$\begin{aligned} K_{jk} &= \sum_{i \in I} \left(\frac{\partial f_W(x_j)}{\partial a_i} \frac{\partial f_W(x_k)}{\partial a_i} + \frac{\partial f_W(x_j)}{\partial b_i} \frac{\partial f_W(x_k)}{\partial b_i} + \frac{\partial f_W(x_j)}{\partial w_i} \frac{\partial f_W(x_k)}{\partial w_i} \right) \\ &\quad + \frac{\partial f_W(x_j)}{\partial c} \frac{\partial f_W(x_k)}{\partial c} \\ &= 0 + 1 \cdot 1 = 1 = \begin{pmatrix} x_j \\ 0 \\ 1 \\ 0 \end{pmatrix}^\top \mathbf{A} \begin{pmatrix} 0 \\ x_k \\ 0 \\ 1 \end{pmatrix} = [(\mathbf{XAX})^\top]_{jk} . \end{aligned}$$

- Case 2: $x_j, x_k > 0$. Then,

$$\begin{aligned} K_{jk} &= \sum_{i \in I_1} \left(\frac{\partial f_W(x_j)}{\partial a_i} \frac{\partial f_W(x_k)}{\partial a_i} + \frac{\partial f_W(x_j)}{\partial b_i} \frac{\partial f_W(x_k)}{\partial b_i} + \frac{\partial f_W(x_j)}{\partial w_i} \frac{\partial f_W(x_k)}{\partial w_i} \right) \\ &\quad + \frac{\partial f_W(x_j)}{\partial c} \frac{\partial f_W(x_k)}{\partial c} \\ &= \sum_{i \in I_1} (w_i^2 x_j x_k + w_i^2 + (a_i x_j + b_i)(a_i x_k + b_i)) + 1 \\ &= \Sigma_{1,w^2}(x_j x_k + 1) + \Sigma_{1,a^2} x_j x_k + \Sigma_{1,ab}(x_j + x_k) + \Sigma_{1,b^2} + 1 \\ &= \begin{pmatrix} x_j \\ 0 \\ 1 \\ 0 \end{pmatrix}^\top \mathbf{A} \begin{pmatrix} x_k \\ 0 \\ 1 \\ 0 \end{pmatrix} = [(\mathbf{XAX})^\top]_{jk} . \end{aligned}$$

- Case 3: $x_j, x_k < 0$. This can be handled just like Case 2.

The identity $\mathbf{M}_D = \frac{1}{n} \mathbf{X}^\top \mathbf{X}$ is easy to verify using the definitions of \mathbf{M}_D and \mathbf{X} .

It remains to show the last identity. For $j \in J$ and $\sigma = \text{sgn}(x_j)$, we have

$$(\mathbf{X}\mathbf{v})_j = p_\sigma x_j + q_\sigma \stackrel{\text{Lemma C.4}}{=} f_{W,\tau,\sigma}(x_j) = f_W(x_j) .$$

Therefore,

$$(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \begin{pmatrix} f_W(x_1) \\ \vdots \\ f_W(x_n) \end{pmatrix} = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{X}\mathbf{v} = \mathbf{v} . \quad (25)$$

Using Definition C.1 and Definition 8, we also obtain

$$\begin{aligned} (\mathbf{X}^\top \mathbf{X})^{-1} &= (n\mathbf{M}_D)^{-1} = (n\tilde{\mathbf{P}}\tilde{\mathbf{M}}_D\tilde{\mathbf{P}})^{-1} = \frac{1}{n} \tilde{\mathbf{P}}\tilde{\mathbf{M}}_D^{-1}\tilde{\mathbf{P}} \\ \tilde{\mathbf{M}}_D^{-1} &= \begin{pmatrix} \mathbf{M}_{D,1} & \\ & \mathbf{M}_{D,-1} \end{pmatrix}^{-1} = \begin{pmatrix} \mathbf{M}_{D,1}^{-1} & \\ & \mathbf{M}_{D,-1}^{-1} \end{pmatrix} \\ \frac{1}{n} \tilde{\mathbf{P}}\mathbf{X}^\top \begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix} &= \frac{1}{n} \tilde{\mathbf{P}} \begin{pmatrix} \sum_{j \in J_1} x_j y_j \\ \sum_{j \in J_{-1}} x_j y_j \\ \sum_{j \in J_1} y_j \\ \sum_{j \in J_{-1}} y_j \end{pmatrix} = \frac{1}{n} \begin{pmatrix} \sum_{j \in J_1} x_j y_j \\ \sum_{j \in J_1} y_j \\ \sum_{j \in J_{-1}} x_j y_j \\ \sum_{j \in J_{-1}} y_j \end{pmatrix} = \begin{pmatrix} \hat{\mathbf{u}}_{D,1}^0 \\ \hat{\mathbf{u}}_{D,-1}^0 \end{pmatrix} \end{aligned}$$

and therefore

$$\begin{aligned}
 (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix} &= \tilde{\mathbf{P}} \begin{pmatrix} \mathbf{M}_{D,1}^{-1} & \\ & \mathbf{M}_{D,-1}^{-1} \end{pmatrix} \begin{pmatrix} \hat{\mathbf{u}}_{D,1}^0 \\ \hat{\mathbf{u}}_{D,-1}^0 \end{pmatrix} = \tilde{\mathbf{P}} \begin{pmatrix} \mathbf{v}_{D,1}^{\text{opt}} \\ \mathbf{v}_{D,-1}^{\text{opt}} \end{pmatrix} = \tilde{\mathbf{P}} \tilde{\mathbf{v}}^{\text{opt}} \\
 &= \mathbf{v}^{\text{opt}} .
 \end{aligned} \tag{26}$$

Subtracting (26) from (25) yields the desired identity

$$(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \bar{\mathbf{f}} = \bar{\mathbf{v}} . \quad \blacksquare$$

Appendix L. Experimental Details for Section 10

In Table L.1, we provide the detailed experimental results for the plots shown in Section 10 and some more results. In particular, we provide results for two other bias initializations:

- Random point combination (**He+5**): This is one of the variants of the initialization methods proposed by Steinwart (2019), given by $b_i = -\sum_{k=1}^5 \lambda_k \langle \mathbf{a}_i, \mathbf{x}_{j_k} \rangle$, where j_k are random indices and λ is randomly drawn from a simplex.
- Negative random point kink initialization (X_{k-}): Here, we choose $b_i = -\langle \mathbf{a}_i, \mathbf{x} \rangle$, where \mathbf{x} is a random data point. If b_i is positive, we update $b_i \leftarrow -b_i$ and $\mathbf{a}_i \leftarrow -\mathbf{a}_i$, such that b_i becomes negative and the kink hyperplane does not change.

We also provide results for the data-generating distribution $P_{\text{ex}}^{\text{data}}$ used in Figure 9: We sample

$$\begin{aligned}
 x &= \sqrt{\frac{196}{105}} \sigma \beta , \\
 y &= \frac{1}{0.727} \left(\cos \left(7\pi \left(x - \frac{x}{\sqrt{1+x^2}} \right) \right) + \frac{1}{5} x^2 + \frac{1}{10} \varepsilon + 0.074 \right) ,
 \end{aligned}$$

where $\sigma \sim \mathcal{U}\{-1, 1\}$, $\beta \sim \text{Beta}(5, 2)$ and $\varepsilon \sim \mathcal{N}(0, 1)$ are independent. The factor $\sqrt{196/105}$ is chosen such that $\text{Var}(x) = 1$. The parameter 5 in the Beta distribution ensures $\eta = 5$ in the sense of Assumption 16. The distribution of y is designed to yield visually interesting results, make y approximately normalized and ensure $q_{P^{\text{data}}, \pm 1}^{\text{opt}} \approx 0$.

References

Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning*, 2019.

Raman Arora, Amitabh Basu, Poorya Mianjy, and Anirbit Mukherjee. Understanding deep neural networks with rectified linear units. In *International Conference on Learning Representations*, 2018.

Param/Opt/Bias-init	$P_{\text{ex}}^{\text{data}}$	P_1^{data}	P_2^{data}	P_4^{data}	P_8^{data}	P_{16}^{data}	P_{32}^{data}	P_{64}^{data}
He/SGD/Zero	9288 ± 53	1020 ± 95	1266 ± 67	1968 ± 25	2424 ± 21	2877 ± 26	3293 ± 23	3846 ± 48
He/SGD/PyTorch	7226 ± 42	500 ± 3	1017 ± 7	1638 ± 20	2319 ± 19	2929 ± 19	3317 ± 21	3859 ± 44
He/SGD/ $U(1)$	6986 ± 38	482 ± 4	928 ± 5	1540 ± 6	2251 ± 10	2753 ± 10	3268 ± 28	3843 ± 35
He/SGD/ $U_+(1)$	8057 ± 21	507 ± 5	1102 ± 8	2325 ± 11	3556 ± 14	4484 ± 15	5315 ± 16	5750 ± 17
He/SGD/ $U_-(1)$	3768 ± 65	233 ± 4	637 ± 5	1300 ± 5	1948 ± 8	2445 ± 9	2686 ± 9	2896 ± 14
He/SGD/ $U_{k_-}(1)$	2385 ± 58	109 ± 3	612 ± 6	1320 ± 6	1910 ± 7	2402 ± 9	2661 ± 9	2882 ± 15
He/SGD/ X_{k_-}	1641 ± 11	115 ± 3	576 ± 5	1310 ± 6	1996 ± 9	2554 ± 9	2862 ± 9	3111 ± 15
He/SGD/He+5	6404 ± 32	490 ± 4	983 ± 11	1674 ± 9	2540 ± 13	3117 ± 11	3450 ± 12	3940 ± 17
He/SGDM/Zero	7678 ± 417	554 ± 6	897 ± 5	1930 ± 14	3490 ± 154	3496 ± 71	3694 ± 46	3691 ± 45
He/SGDM/PyTorch	3044 ± 68	291 ± 6	821 ± 4	1550 ± 8	2659 ± 33	3408 ± 83	3539 ± 41	3703 ± 43
He/SGDM/ $U(1)$	2549 ± 58	127 ± 3	576 ± 5	1270 ± 6	1982 ± 7	2555 ± 9	2858 ± 11	3138 ± 15
He/SGDM/ $U_+(1)$	3595 ± 54	175 ± 3	581 ± 4	1192 ± 6	2052 ± 8	2714 ± 11	3053 ± 11	3425 ± 14
He/SGDM/ $U_-(1)$	1649 ± 7	45 ± 2	364 ± 5	1243 ± 6	1948 ± 8	2423 ± 9	2657 ± 11	2782 ± 14
He/SGDM/ $U_{k_-}(1)$	1644 ± 7	29 ± 1	412 ± 9	1313 ± 7	1911 ± 7	2385 ± 9	2631 ± 11	2769 ± 15
He/SGDM/ X_{k_-}	1520 ± 5	32 ± 1	331 ± 6	1264 ± 7	1998 ± 9	2511 ± 17	2785 ± 9	2982 ± 17
He/SGDM/He+5	3108 ± 97	230 ± 7	797 ± 6	1463 ± 7	2360 ± 10	2946 ± 11	3215 ± 11	3360 ± 18
He/Adam/Zero	1554 ± 6	72 ± 2	449 ± 7	1232 ± 5	1779 ± 7	2496 ± 9	2973 ± 10	3265 ± 15
He/Adam/PyTorch	1544 ± 6	53 ± 2	430 ± 7	1234 ± 6	1786 ± 8	2458 ± 9	2963 ± 10	3266 ± 16
He/Adam/ $U(1)$	1550 ± 6	41 ± 3	405 ± 6	1233 ± 6	1791 ± 7	2219 ± 9	2632 ± 9	2918 ± 15
He/Adam/ $U_+(1)$	1632 ± 15	87 ± 8	536 ± 10	1160 ± 7	1874 ± 10	2533 ± 13	3055 ± 11	3395 ± 19
He/Adam/ $U_-(1)$	1521 ± 6	19 ± 1	279 ± 4	1141 ± 7	1733 ± 7	2061 ± 8	2377 ± 9	2584 ± 13
He/Adam/ $U_{k_-}(1)$	1512 ± 5	13 ± 1	268 ± 4	1160 ± 7	1713 ± 7	2057 ± 8	2376 ± 9	2582 ± 13
He/Adam/ X_{k_-}	1498 ± 4	13 ± 1	266 ± 4	1162 ± 7	1771 ± 8	2160 ± 8	2533 ± 9	2788 ± 15
He/Adam/He+5	1534 ± 5	50 ± 2	425 ± 6	1249 ± 7	1785 ± 7	2339 ± 9	2838 ± 11	3134 ± 16
NTK/SGD/Zero	7813 ± 41	529 ± 3	821 ± 4	1186 ± 4	1651 ± 6	2057 ± 8	2357 ± 8	2633 ± 11
NTK/SGD/PyTorch	7371 ± 30	501 ± 3	818 ± 4	1182 ± 5	1644 ± 6	2055 ± 7	2348 ± 8	2644 ± 12
NTK/SGD/ $U(1)$	6909 ± 26	367 ± 5	705 ± 4	1182 ± 5	1689 ± 7	2123 ± 8	2421 ± 8	2810 ± 15
NTK/SGD/ $U_+(1)$	7416 ± 25	386 ± 9	823 ± 4	1287 ± 9	1793 ± 11	2194 ± 15	2454 ± 13	3002 ± 23
NTK/SGD/ $U_-(1)$	5903 ± 31	378 ± 4	691 ± 4	1125 ± 5	1672 ± 6	2102 ± 9	2408 ± 8	2692 ± 13
NTK/SGD/ $U_{k_-}(1)$	5618 ± 33	166 ± 4	629 ± 4	1124 ± 5	1662 ± 6	2096 ± 9	2403 ± 8	2688 ± 13
NTK/SGD/ X_{k_-}	4505 ± 35	159 ± 5	661 ± 5	1136 ± 4	1663 ± 6	2088 ± 7	2399 ± 8	2693 ± 13
NTK/SGD/He+5	6659 ± 26	446 ± 4	786 ± 4	1189 ± 5	1658 ± 6	2084 ± 8	2353 ± 8	2677 ± 13
NTK/SGDM/Zero	5761 ± 230	500 ± 4	826 ± 4	1239 ± 7	1727 ± 8	2113 ± 11	2458 ± 13	2694 ± 17
NTK/SGDM/PyTorch	2061 ± 68	297 ± 13	815 ± 5	1215 ± 5	1728 ± 8	2132 ± 10	2484 ± 13	2722 ± 21
NTK/SGDM/ $U(1)$	1768 ± 42	268 ± 10	707 ± 5	1136 ± 5	1670 ± 6	2103 ± 8	2417 ± 9	2624 ± 13
NTK/SGDM/ $U_+(1)$	2954 ± 45	368 ± 9	732 ± 4	1150 ± 6	1715 ± 6	2166 ± 8	2454 ± 8	2629 ± 13
NTK/SGDM/ $U_-(1)$	1602 ± 39	121 ± 9	572 ± 7	1133 ± 5	1643 ± 7	2055 ± 9	2378 ± 9	2616 ± 15
NTK/SGDM/ $U_{k_-}(1)$	1564 ± 6	68 ± 2	508 ± 7	1128 ± 5	1638 ± 7	2049 ± 8	2377 ± 9	2611 ± 14
NTK/SGDM/ X_{k_-}	1552 ± 6	70 ± 2	523 ± 7	1157 ± 4	1637 ± 6	2054 ± 7	2388 ± 8	2630 ± 14
NTK/SGDM/He+5	2126 ± 77	256 ± 13	784 ± 5	1196 ± 5	1689 ± 7	2106 ± 9	2447 ± 8	2644 ± 12
NTK/Adam/Zero	1605 ± 7	59 ± 3	461 ± 8	1219 ± 5	1655 ± 6	2124 ± 7	2491 ± 10	2747 ± 14
NTK/Adam/PyTorch	1600 ± 9	47 ± 2	440 ± 7	1216 ± 5	1677 ± 6	2145 ± 8	2568 ± 9	2818 ± 15
NTK/Adam/ $U(1)$	1628 ± 7	67 ± 4	466 ± 8	1132 ± 6	1705 ± 7	2175 ± 9	2591 ± 9	2852 ± 14
NTK/Adam/ $U_+(1)$	1664 ± 8	78 ± 5	530 ± 9	1237 ± 8	1846 ± 9	2471 ± 12	2857 ± 11	3108 ± 17
NTK/Adam/ $U_-(1)$	1580 ± 7	44 ± 2	364 ± 6	1093 ± 6	1594 ± 6	1984 ± 7	2334 ± 8	2539 ± 13
NTK/Adam/ $U_{k_-}(1)$	1559 ± 7	24 ± 1	322 ± 5	1085 ± 6	1585 ± 6	1974 ± 7	2333 ± 9	2539 ± 12
NTK/Adam/ X_{k_-}	1528 ± 6	23 ± 1	348 ± 5	1116 ± 5	1625 ± 6	2023 ± 8	2360 ± 8	2583 ± 14
NTK/Adam/He+5	1594 ± 6	43 ± 2	447 ± 7	1203 ± 4	1691 ± 7	2159 ± 9	2589 ± 9	2852 ± 15

Table L.1: For different combinations of data distribution, parameterization, optimizer and bias initialization, this table shows the mean RMSE over 100 runs and the estimated standard deviation of the mean estimator, multiplied by 10^4 and rounded to the nearest integer, respectively. Cell colors are interpolated on a logarithmic scale from white (best result in column) to red (worst result in column). The bias initialization methods He+5 and X_{k_-} and the distribution $P_{\text{ex}}^{\text{data}}$ are explained in Section L, while the other terms are explained in Section 10.

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 *International Conference on Machine Learning*, 2019a.

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 *Neural Information Processing Systems*, 2019b.

Sanjeev Arora, Simon S. Du, Zhiyuan Li, Ruslan Salakhutdinov, Ruosong Wang, and Dingli Yu. Harnessing the power of infinitely wide deep nets on small-data tasks. In *International*

- Conference on Learning Representations*, 2019c.
- Mikhail Belkin, Daniel J. Hsu, and Partha Mitra. Overfitting or perfect fitting? Risk bounds for classification and regression rules that interpolate. In *Neural Information Processing Systems*, 2018.
- Rajendra Bhatia. *Matrix Analysis*, volume 169. Springer Science & Business Media, 2013.
- Avrim Blum and Ronald L. Rivest. Training a 3-node neural network is NP-complete. In *Neural Information Processing Systems*, 1989.
- Digvijay Boob, Santanu S. Dey, and Guanghui Lan. Complexity of training relu neural network. *Discrete Optimization*, 2020.
- Yuan Cao and Quanquan Gu. Generalization error bounds of gradient descent for learning over-parameterized deep ReLU networks. In *AAAI Conference on Artificial Intelligence*, 2020.
- Zixiang Chen, Yuan Cao, Difan Zou, and Quanquan Gu. How much over-parameterization is sufficient to learn deep ReLU networks? In *International Conference on Learning Representations*, 2021.
- Lenaïc Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. In *Neural Information Processing Systems*, 2019.
- François Chollet and others. Keras, 2015.
- Luc Devroye, László Györfi, and Gábor Lugosi. *A Probabilistic Theory of Pattern Recognition*, volume 31. Springer Science & Business Media, 1996.
- Simon Du, Jason Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. In *International Conference on Machine Learning*, 2019a.
- 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.
- András Faragó and Gábor Lugosi. Strong universal consistency of neural network classifiers. *IEEE Transactions on Information Theory*, 39(4):1146–1151, 1993.
- Kenji Fukumizu and Shun-ichi Amari. Local minima and plateaus in hierarchical structures of multilayer perceptrons. *Neural Networks*, 13(3):317–327, 2000.
- Rong Ge, Runzhe Wang, and Haoyu Zhao. Mildly overparametrized neural nets can memorize training data efficiently. *arXiv:1909.11837*, 2019.
- Gene H. Golub and Charles F. Van Loan. Matrix Computations. *The Johns Hopkins University Press, Baltimore, USA*, 1989.

- Marco Gori and Alberto Tesi. On the problem of local minima in backpropagation. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 14(1):76–86, 1992.
- Yury Gorishniy, Ivan Rubachev, Valentin Khruikov, and Artem Babenko. Revisiting deep learning models for tabular data. In *Neural Information Processing Systems*, 2021.
- László Györfi, Michael Kohler, Adam Krzyzak, and Harro Walk. *A Distribution-Free Theory of Nonparametric Regression*. Springer Science & Business Media, 2002.
- Thomas Hamm and Ingo Steinwart. Adaptive learning rates for support vector machines working on data with low intrinsic dimension. *The Annals of Statistics*, 49(6):3153–3180, 2021.
- Fengxiang He, Bohan Wang, and Dacheng Tao. Piecewise linear activations substantially shape the loss surfaces of neural networks. In *International Conference on Learning Representations*, 2019.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification. In *IEEE Conference on Computer Vision*, 2015.
- David Holzmüller. Convergence analysis of neural networks. Master’s thesis, University of Stuttgart, 2019.
- Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural Tangent Kernel: Convergence and generalization in neural networks. In *Neural Information Processing Systems*, 2018.
- 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*, 2019.
- Arlind Kadra, Marius Lindauer, Frank Hutter, and Josif Grabocka. Regularization is all you need: Simple neural nets can excel on tabular data. *arXiv:2106.11189*, 2021.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- Herbert KH Lee. Consistency of posterior distributions for neural networks. *Neural Networks*, 13(6):629–642, 2000.
- Yuanzhi Li and Yingyu Liang. Learning overparameterized neural networks via stochastic gradient descent on structured data. In *Neural Information Processing Systems*, 2018.
- Tengyuan Liang, Alexander Rakhlin, and Xiyu Zhai. On the multiple descent of minimum-norm interpolants and restricted lower isometry of kernels. In *Conference on Learning Theory*, 2020.
- Kyle Luther and H. Sebastian Seung. Sample variance decay in randomly initialized ReLU networks. *arXiv:1902.04942*, 2019.

- Nicole Mücke and Ingo Steinwart. Empirical Risk Minimization in the Interpolating Regime with Application to Neural Network Learning. *arXiv:1905.10686*, 2019.
- Alexander Rakhlin and Xiyu Zhai. Consistency of interpolation with laplace kernels is a high-dimensional phenomenon. In *Conference on Learning Theory*, 2019.
- Itay Safran and Ohad Shamir. Spurious local minima are common in two-layer ReLU neural networks. In *International Conference on Machine Learning*, 2018.
- Andrew M. Saxe, James L. McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. In *International Conference on Learning Representations*, 2014.
- Ravid Shwartz-Ziv and Amitai Armon. Tabular data: Deep learning is not all you need. *Information Fusion*, 81:84–90, 2022.
- Zhao Song and Xin Yang. Quadratic suffices for over-parametrization via matrix chernoff bound. *arXiv:1906.03593*, 2019.
- Eduardo D. Sontag and Héctor J. Sussmann. Backpropagation can give rise to spurious local minima even for networks without hidden layers. *Complex Systems*, 3(1):91–106, 1989.
- Daniel Soudry and Yair Carmon. No bad local minima: Data independent training error guarantees for multilayer neural networks. *arXiv:1605.08361*, 2016.
- Ingo Steinwart. A sober look at neural network initializations. *arXiv:1903.11482*, 2019.
- Ingo Steinwart and Andreas Christmann. *Support Vector Machines*. Springer Science & Business Media, 2008.
- Halbert White. Connectionist nonparametric regression: Multilayer feedforward networks can learn arbitrary mappings. *Neural networks*, 3(5):535–549, 1990.
- Francis Williams, Matthew Trager, Daniele Panozzo, Claudio Silva, Denis Zorin, and Joan Bruna. Gradient dynamics of shallow univariate relu networks. In *Neural Information Processing Systems*, 2019.
- Yuan Yao, Lorenzo Rosasco, and Andrea Caponnetto. On early stopping in gradient descent learning. *Constructive Approximation*, 26(2):289–315, 2007.
- Chulhee Yun, Suvrit Sra, and Ali Jadbabaie. Small nonlinearities in activation functions create bad local minima in neural networks. In *International Conference on Learning Representations*, 2019.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*, 2017.
- Difan Zou and Quanquan Gu. An improved analysis of training over-parameterized deep neural networks. In *Neural Information Processing Systems*, 2019.

Difan Zou, Yuan Cao, Dongruo Zhou, and Quanquan Gu. Stochastic gradient descent optimizes over-parameterized deep relu networks. *arXiv:1811.08888*, 2018.