

Deep Linear Networks can Benignly Overfit when Shallow Ones Do

Niladri S. Chatterji

Computer Science Department, Stanford University, 353 Jane Stanford Way, Stanford, CA 94305.

NILADRI@CS.STANFORD.EDU

Philip M. Long

Google, 1600 Amphitheatre Parkway, Mountain View, CA, 94043.

PLONG@GOOGLE.COM

Editor: Ohad Shamir

Abstract

We bound the excess risk of interpolating deep linear networks trained using gradient flow. In a setting previously used to establish risk bounds for the minimum ℓ_2 -norm interpolant, we show that randomly initialized deep linear networks can closely approximate or even match known bounds for the minimum ℓ_2 -norm interpolant. Our analysis also reveals that interpolating deep linear models have exactly the same conditional variance as the minimum ℓ_2 -norm solution. Since the noise affects the excess risk only through the conditional variance, this implies that depth does not improve the algorithm’s ability to “hide the noise”. Our simulations verify that aspects of our bounds reflect typical behavior for simple data distributions. We also find that similar phenomena are seen in simulations with ReLU networks, although the situation there is more nuanced.

Keywords: generalization, benign overfitting, implicit bias, interpolation, neural networks, regression

1. Introduction

Recent empirical studies (Zhang et al., 2017; Belkin et al., 2019a) have brought to light the surprising phenomenon that overparameterized neural network models trained with variants of gradient descent generalize well despite perfectly fitting noisy data. This seemingly violates the once widely accepted principle that learning algorithms should trade off between some measure of the regularity of a model, and its fit to the data. To understand this, a rich line of research has emerged to establish conditions under which extreme overfitting—fitting the data perfectly—is benign in simple models (see Belkin et al., 2018; Hastie et al., 2022; Bartlett et al., 2020). Another closely connected thread of research to understand generalization leverages the recognition that training by gradient descent engenders an implicit bias (see Neyshabur et al., 2015; Soudry et al., 2018; Ji and Telgarsky, 2019). These results can be paraphrased as follows: training until the loss is driven to zero will produce a model that, among models that interpolate the data, minimizes some data-independent regularity criterion.

Our paper continues this study of benign overfitting but with a more complex model class, deep linear networks. Deep linear networks are often studied theoretically (see, e.g., Saxe et al., 2014; Arora et al., 2018), because some of the relevant characteristics of deep learning in the presence of nonlinearities are also present in linear networks but in a setting

that is more amenable to analysis. The analyses of linear networks have included a number of results on implicit bias (see, e.g., Azulay et al., 2021; Min et al., 2021). Recently, one of these analyses (Azulay et al., 2021), of two-layer networks trained by gradient flow with a “balanced” initialization, was leveraged in an analysis of benign overfitting (Chatterji et al., 2022). (For a mapping $x \rightarrow xWv$ parameterized by a hidden layer $W \in \mathbb{R}^{d \times m}$ and an output layer $v \in \mathbb{R}^{m \times 1}$, initial values of v and W are balanced if $vv^\top = W^\top W$.) Min et al. (2021) analyzed implicit bias in two-layer linear networks under more general conditions including the unbalanced case.

In this paper, we analyze benign overfitting in deep linear networks of arbitrary depth trained by gradient flow. Our first main result is a bound on the excess risk. The bound is in terms of some characteristics of the joint distribution of the training data previously used to analyze linear regression with the standard parameterization, including notions of the effective rank of the covariance matrix, and it holds under similar conditions on the data distribution. Another key quantity used in the bound concerns the linear map Θ computed by the network after training—it is the norm of the projection of this map onto the subspace orthogonal to the span of the training examples. This norm can further be bounded in terms of its value at initialization, and a quantity that reflects how rapidly training converged. In contrast with previous analyses on two-layer networks (Chatterji et al., 2022), this analysis holds whether this initialization is balanced or not.

Our second main result is a high-probably risk bound that holds for networks in which the first and last layers are initialized randomly, and the middle layers are all initialized to the identity. Our bound holds whenever the scale of the initialization of the first layer is small enough, and the scale of the initialization of the last layer is large enough. This includes the extreme case where the first layer is initialized to zero. As the scale of the initialization of the first layer goes to zero, our bound approaches the known bound for the minimum ℓ_2 -norm interpolator with the standard parameterization. Our final main theoretical result illustrates our bounds using a simple covariance matrix used in previous work (Bartlett et al., 2020; Chatterji and Long, 2022) which might be viewed as a canonical case where overfitting is benign for linear regression with the standard parameterization.

These bounds were obtained in the absence of a precise characterization of the implicit bias of gradient flow for deep linear networks, or a closed-form formula for the model produced.

A key point of our analysis is that the projection of the linear map Θ computed by the interpolating network onto the span of the rows of the design matrix X is exactly equal to minimum ℓ_2 -norm interpolant Θ_{ℓ_2} . The risk of Θ naturally decomposes into contributions from this projection and $\Theta_{X^\perp} = \Theta - \Theta_{\ell_2}$. We can use previous analyses of Θ_{ℓ_2} to bound the former.

Figure 1 contains plots from simulation experiments where the excess risk of a deep linear model increases with the scale of the initialization of the first layer, as in the upper bounds of our analysis. A similar effect is also seen when the first layer is initialized at a unit scale, and the scale of the initialization of the last layer varies. In both cases, we also see that as the function computed by the network at initialization approaches the zero function, the trained model approaches the minimum ℓ_2 -norm interpolant. Figure 2 includes plots of analogous experiments with networks with ReLU nonlinearities. As in the linear case the excess risk increases with the scale of the initialization of the first layer, but we do not see a

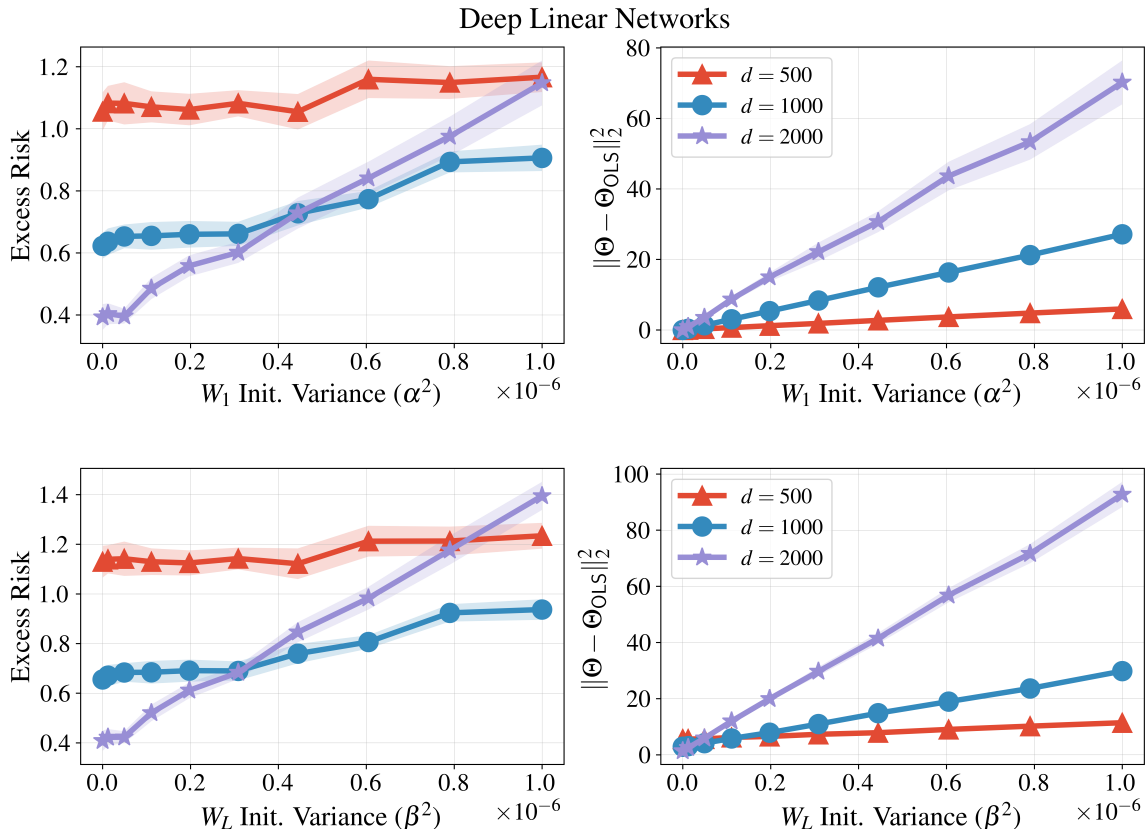


Figure 1: Three-layer linear networks trained by gradient descent on data generated by an underlying linear model. The model is trained on $n = 100$ points drawn from the generative model $y = x\Theta^* + \omega$, where $x \sim \mathcal{N}(0, \Sigma)$ and $\omega \sim \mathcal{N}(0, 1)$. The excess risk is defined as $\mathbb{E}_x [\|x\Theta - x\Theta^*\|^2]$. We empirically find that when the initialization variance of either the first layer (α^2) or the last layer (β^2) is close to zero, the final solution is close to the minimum ℓ_2 -norm interpolator and suffers small excess risk. While when the initialization variance is large, that is, when the network is initialized away from the origin, the excess risk is larger due to the component of the final solution outside the span of the data. Additional details in Section 7.

significant increase in excess risk with the scale of the initialization of the last layer. More details of the experiments are described in Section 7.

Intuitively, the harm from overfitting arises from fitting the noise, and the effect of fitting the noise is analyzed in the conditional variance of the estimator. In the setting studied here, as in linear regression with the standard parameterization, the conditional variance is entirely determined by the projection of Θ onto the span of the rows of the data matrix X which is equal to Θ_{ℓ_2} . Thus, when learning deep linear networks with quadratic loss, aspects of training that affect the inductive bias, such as the initialization, architecture,

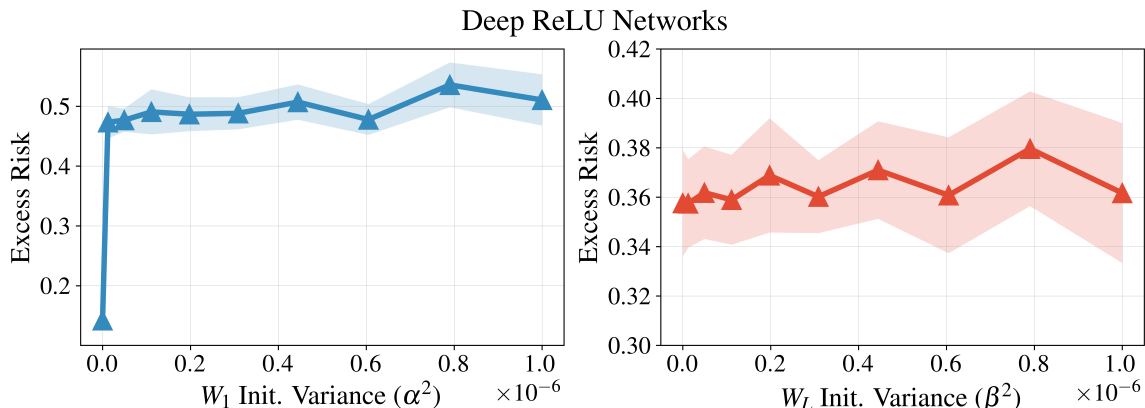


Figure 2: Three-layer ReLU networks trained by gradient descent on data generated by an underlying two-layer ReLU teacher network. The model is trained on $n = 500$ points drawn from the generative model $y = f^*(x) + \omega$, where f^* is a two-layer ReLU network with width 50, $x \sim \mathcal{N}(0, I_{10 \times 10})$ and $\omega \sim \mathcal{N}(0, 1)$. The excess risk is defined as $\mathbb{E}_x [\|f(x) - f^*(x)\|^2]$. In ReLU models we find that the risk scales differently as we scale the initialization variance of the first layer (α^2) and that of the last layer (β^2). When we scale α^2 , similar to deep linear models we find that risk is smaller for smaller values of α^2 . However, this is not the case when we scale β^2 . This highlights a surprising asymmetry in the role played by the initialization scales of the different layers in ReLU networks. For additional details about the experiment see Section 7.

etc., do not affect this variance term—no matter how they are chosen, the distribution of the variance term is determined by Θ_{ℓ_2} . To see an effect of implicit bias in deep linear networks on the consequence of fitting the noise, we must analyze a loss function other than the quadratic loss.

Our upper bounds reveal no benefit in representing linear transformations by deep networks, and, in our simulations, we see no benefit with random initialization. This is because non-zero random initialization usually contributes additional error to the bias as the random initialization is typically a poor guess for the regression function. (In rare cases it could reduce the bias, though, if by chance it approximates the regression function.)

Our analysis also leverages the effect of imbalanced initialization on implicit bias—our treatment partially extends the results by Min et al. (2021) from the two-layer case to the deep case, and then combines them with our general risk bound.

Organization. In Section 2 we describe our problem setting and our assumptions. Then in Section 3 we present our main results and in Sections 4, 5 and 6 we prove these results. We provide additional simulations and simulation details in Section 7. We conclude with a discussion in Section 8. In Appendix A we highlight other related work on benign overfitting, implicit bias, and on linear networks. Finally, in Appendix B we present omitted technical details.

2. Preliminaries

This section includes notational conventions and a description of the setting.

2.1 Notation

Given a vector v , let $\|v\|$ denote its Euclidean norm. Given a matrix M , let $\|M\|$ denote its Frobenius norm and let $\|M\|_{op}$ denote its operator norm. For any $j \in \mathbb{N}$, we denote the set $\{1, \dots, j\}$ by $[j]$. We will use c, c', c_1, c_x, \dots to denote positive absolute constants, which may take different values in different contexts.

2.2 Setting

We analyze linear regression with d inputs and q outputs from n examples. Throughout the paper we assume that $d > n$. Although we assume throughout that the input dimension d is finite, it is straightforward to extend our results to infinite d .

Let $X \in \mathbb{R}^{n \times d}$ be the data matrix, and $Y \in \mathbb{R}^{n \times q}$ be the response matrix, and let $x_1, \dots, x_n \in \mathbb{R}^{1 \times d}$ be the rows of X and $y_1, \dots, y_n \in \mathbb{R}^{1 \times q}$ be the rows of Y .

For random $(x, y) \in \mathbb{R}^{1 \times d} \times \mathbb{R}^{1 \times q}$, let

$$\Theta^* \in \arg \min_{\Theta \in \mathbb{R}^{d \times q}} \mathbb{E}_{(x,y)} [\|y - x\Theta\|^2]$$

be an arbitrary optimal linear regressor. We let $\Omega = Y - X\Theta^* \in \mathbb{R}^{n \times q}$ be the noise matrix.

Define the *excess risk* of an estimate $\Theta \in \mathbb{R}^{d \times q}$ to be

$$\text{Risk}(\Theta) := \mathbb{E}_{x,y} [\|y - x\Theta\|^2 - \|y - x\Theta^*\|^2],$$

where x, y are test samples that are independent of Θ .

Denote the second moment matrix of the covariates by $\Sigma := \mathbb{E}[x^\top x] \in \mathbb{R}^{d \times d}$ with eigenvalues $\lambda_1 \geq \dots \geq \lambda_d \geq 0$. We will use the following definitions of the “effective rank” that Bartlett et al. (2020) previously used in the analysis of the excess risk of the minimum ℓ_2 -norm interpolant.

Definition 1. Given any $j \in [d]$, define $s_j := \sum_{i>j} \lambda_i$ and

$$r_j := \frac{s_j}{\lambda_{j+1}} \quad \text{and} \quad R_j := \frac{s_j^2}{\sum_{i>j} \lambda_i^2}.$$

We define the index k below. The value of k shall help determine what we consider the “tail” of the covariance matrix.

Definition 2. For a large enough constant b (that will be fixed henceforth), define

$$k := \min\{j \geq 0 : r_j \geq bn\},$$

where the minimum of the empty set is defined as ∞ .

We are now ready to introduce the assumptions of our paper.

Assumptions. Let c_x and c_y denote absolute constants.

(A.1) The samples $(x_1, y_1), \dots, (x_n, y_n)$ are drawn i.i.d.

(A.2) The covariates x and responses y are mean-zero.

(A.3) The covariates x satisfy $x = \Sigma^{1/2}u$, where u is isotropic and has components that are independent c_x -sub-Gaussian random variables, that is, for all $\phi \in \mathbb{R}^d$

$$\mathbb{E} \left[\exp \left(\phi^\top u \right) \right] \leq \exp \left(c_x \|\phi\|^2 / 2 \right).$$

(A.4) The difference $y - x\Theta^*$ is c_y -sub-Gaussian, conditionally on x ; that is, for all $\phi \in \mathbb{R}^q$

$$\mathbb{E}_y \left[\exp \left(\phi^\top (y - x\Theta^*) \right) \mid x \right] \leq \exp \left(c_y \|\phi\|^2 / 2 \right)$$

(note that this implies that $\mathbb{E}[y \mid x] = x\Theta^*$ and $\mathbb{E}[\|y - x\Theta^*\|^2] \leq cq$).

(A.5) Almost surely, the projection of the data X on the space orthogonal to any eigenvector of Σ spans a space of dimension n .

All the constants going forward may depend on the values of c_x and c_y . The assumptions made here are standard in the benign overfitting literature (see Bartlett et al., 2020; Chatterji et al., 2022). They are satisfied for example in the case where x is a mean-zero Gaussian whose covariance Σ has full rank, $d > n$, and the noise $y - x\Theta^*$ is independent and Gaussian.

2.3 Deep Linear Models

We analyze linear models represented by deep linear networks with m hidden units at each layer. We denote the weight matrices by W_1, \dots, W_L , where $W_1 \in \mathbb{R}^{m \times d}$, $W_2, \dots, W_{L-1} \in \mathbb{R}^{m \times m}$, and $W_L \in \mathbb{R}^{q \times m}$. The standard representation of the network's linear transformation, denoted by $\Theta \in \mathbb{R}^{d \times q}$, is

$$\Theta = (W_L \cdots W_1)^\top \in \mathbb{R}^{d \times q}.$$

Define P_X to be the projection onto the row span of X , that is, $P_X := X^\top (XX^\top)^{-1}X$. Let

$$\Theta_X := P_X \Theta \quad \text{and} \quad \Theta_{X^\perp} := (I - P_X) \Theta.$$

For n datapoints $(x_1, y_1), \dots, (x_n, y_n)$, where $x_i \in \mathbb{R}^{1 \times d}$ and $y_i \in \mathbb{R}^{1 \times q}$, the training loss is given by

$$\mathcal{L}(\Theta) := \sum_{i=1}^n \|y_i - x_i \Theta\|^2 = \|Y - X\Theta\|^2.$$

We will analyze the generalization properties of deep linear models trained with gradient flow, that is, for all $j \in [L]$,

$$\frac{dW_j^{(t)}}{dt} = -\nabla_{W_j^{(t)}} \mathcal{L}(\Theta^{(t)}).$$

We study the following random initialization scheme in our paper.

Definition 3. (Random initialization) Given $\alpha, \beta > 0$, the entries of the first layer $W_1^{(0)}$ and the last layer $W_L^{(0)}$ are initialized using i.i.d. draws from $\mathbf{N}(0, \alpha^2)$ and $\mathbf{N}(0, \beta^2)$ respectively. The remaining layers $W_2^{(0)}, \dots, W_{L-1}^{(0)}$ are initialized to the identity I_m .

A similar initialization scheme has been studied previously (Zou et al., 2020). Our analysis will show that starting from random initialization the scale of the network grows in a controlled manner which is captured by the following definition.

Definition 4. We say that training is perpetually Λ bounded if, for all $t \geq 0$ and all $S \subseteq [L]$,

$$\prod_{j \in S} \|W_j^{(t)}\|_{op} \leq \Lambda.$$

In our subsequent analysis, this notion of perpetually Λ bounded shall allow us to control the behavior of the network in the null space of the data matrix X .

2.4 The Minimum ℓ_2 -norm Interpolant

It will be helpful to compare the generalization of the deep linear model with the result of applying the minimum ℓ_2 -norm interpolant resulting from the standard parameterization.

Definition 5. For any $X \in \mathbb{R}^{n \times d}$ and $Y \in \mathbb{R}^{n \times q}$, define $\Theta_{\ell_2} = X^\top (XX^\top)^{-1}Y$.

Under Assumption (A.5), the matrix XX^\top is full rank and therefore Θ_{ℓ_2} is well defined. As previously noted, the excess risk of this canonical interpolator has been studied in prior work (see Bartlett et al., 2020; Tsigler and Bartlett, 2020).

3. Main Results

In this section, we present our excess risk bounds. Our first result applies to any deep linear model trained until interpolation. Second, we shall specialize this result to the case where the model is randomly initialized. Lastly, we present an excess risk bound for a randomly initialized network in a setting with a spiked covariance matrix.

3.1 Excess Risk bound for Deep Linear Models

The following theorem is an excess risk bound for *any* deep linear model trained until it interpolates in terms of the rate of convergence of its training, along with the effective ranks of the covariance matrix.

Theorem 6. Under Assumptions (A.1)-(A.5), there is an absolute constant $c > 0$ such that, for all $\delta < 1/2$ and all depths $L > 1$, the following holds. With probability at least $1 - c\delta$, if $\Theta = \lim_{t \rightarrow \infty} \Theta^{(t)}$ for a perpetually Λ -bounded training process for which $\lim_{t \rightarrow \infty} \mathcal{L}(\Theta^{(t)}) = 0$, and $n \geq c \max\{r_0, k, \log(1/\delta)\}$, then

$$\text{Risk}(\Theta) = \text{Bias}(\Theta_{\ell_2}) + \text{Variance}(\Theta_{\ell_2}) + \Xi,$$

where

$$\begin{aligned} \text{Bias}(\Theta_{\ell_2}) &\leq \frac{cs_k}{n} \|\Theta^*\|^2, \\ \text{Variance}(\Theta_{\ell_2}) &\leq cq \log(q/\delta) \left(\frac{k}{n} + \frac{n}{R_k} \right), \end{aligned}$$

and

$$\Xi = \frac{cs_k}{n} \|\Theta_{X^\perp}\|^2 \leq \frac{cs_k}{n} \left[\|\Theta_{X^\perp}^{(0)}\| + L\Lambda^2 \sqrt{\lambda_1 n} \int_{t=0}^{\infty} \sqrt{\mathcal{L}(\Theta^{(t)})} dt \right]^2.$$

This bound shows that the *conditional bias* of the estimator Θ is upper bounded by the conditional bias of the minimum ℓ_2 -norm interpolant $\text{Bias}(\Theta_{\ell_2})$ plus Ξ , which is the additional bias incurred by the component of Θ outside the row span of X . This additional term Ξ depends not only on the eigenvalues of the covariance matrix but also on the specifics of the optimization procedure such as the initial linear model $(\Theta^{(0)})$, the size of the weights throughout training (Λ) and the rate of decay of the loss.

Interestingly, the *conditional variance* of the interpolator Θ , is in fact identical to the conditional variance of the minimum ℓ_2 -norm interpolant. This follows because, as we will show in the proof, the component of the interpolator Θ in the row span of X is in fact equal to Θ_{ℓ_2} , and the conditional variance depends only on this component within the row span of X . The variance captures the effect of perfectly fitting the noise in the data, and our analysis shows that the harm incurred by fitting the noise is unaffected by parameterizing a linear model as a deep linear model.

Essentially matching lower bounds (up to constants) on the variance term are known (Bartlett et al., 2020).

3.2 Excess Risk Bound under Random Initialization

Our next main result establishes a high-probability bound on the excess risk, and in particular on Ξ , when the network is trained after a random initialization (see Definition 3).

Theorem 7. *Under Assumptions (A.1)-(A.5), there is an absolute constant $c > 0$ such that, for all $\delta < 1/2$, if*

- *the initialization scales β and α satisfy $\beta \geq c \max \left\{ 1, \frac{\lambda_1^{1/4} \sqrt{Ln} (\sqrt{\|\Theta^*\| \lambda^{1/4} + q^{1/4}})}{\sqrt{s_k}} \right\}$ and $\alpha \leq 1$;*
- *the width $m \geq c \max \left\{ d + q + \log(1/\delta), \frac{L^2 \alpha^2 \lambda_1 s_0 n^2 q \log(n/\delta)}{\beta^2 s_k^2} \right\}$;*
- *the network is trained using random initialization as described in Definition 2.3;*
- *the number of samples satisfies $n \geq c \max\{r_0, k, \log(1/\delta)\}$,*

then, with probability at least $1 - c\delta$,

$$\text{Risk}(\Theta) \leq \text{Bias}(\Theta_{\ell_2}) + \text{Variance}(\Theta_{\ell_2}) + \Xi,$$

where

$$\begin{aligned} \text{Bias}(\Theta_{\ell_2}) &\leq \frac{cs_k}{n} \|\Theta^*\|^2, \\ \text{Variance}(\Theta_{\ell_2}) &\leq cq \log(q/\delta) \left(\frac{k}{n} + \frac{n}{R_k} \right), \\ &\equiv \leq \frac{c\alpha^2 s_k}{n} \left[q\beta^2 + \frac{L^2(\alpha + 1/L)^4 \lambda_1 n^2}{s_k^2} \left(\lambda_1 \|\Theta^*\|^2 + q + \frac{\alpha^2 \beta^2 s_0 q \log(n/\delta)}{m} \right) \right]. \end{aligned}$$

Note that the bound on Ξ of Theorem 7 can be made arbitrarily small by decreasing α while keeping the other parameters fixed. When $\alpha = 0$, our bound shows that the model has the same risk as the minimum ℓ_2 -norm interpolant.

Recall from the simulation in Figure 1 that as the initialization of the last layer β approaches 0, the model produced by gradient descent gets closer to the minimum ℓ_2 -norm interpolant. Our bound on Ξ does not approach 0 as $\beta \rightarrow 0$, and we do not know how to prove that this happens in general with high probability.

Regarding the role of overparameterization, we find that one component of our bound on Ξ gets smaller as the width m is increased. However, our bound gets larger as we increase depth L .

As mentioned earlier, the bound on the conditional variance, which captures the effect of fitting the noise, is sharp up to constants, however we do not know whether the upper bound on the conditional bias, and specifically Ξ in Theorem 7, can be improved. It is also unclear whether conditions on β and m can be relaxed.

Next, to facilitate the interpretation of our bounds, we apply Theorem 7 in a canonical setting where benign overfitting occurs for the minimum ℓ_2 -norm interpolant.

Definition 8 ((k, ε) -spike model). *For $0 < \varepsilon < 1$ and $k \in \mathbb{N}$, a (k, ε) -spike model is a setting where the eigenvalues of Σ are $\lambda_1 = \dots = \lambda_k = 1$ and $\lambda_{k+1} = \dots = \lambda_d = \varepsilon$.*

The (k, ε) -spike model is a setting where there are k high variance directions, and many $(d - k)$ low variance directions that can be used to “hide” the energy of the noise. Note that, in this model, if $d \geq cn$ and $n \geq ck$ for a large enough constant c , then k satisfies the requirement of Definition 2, since $r_k = \varepsilon(d - k)/\varepsilon = d - k \geq bn$. Since this covariance matrix has full rank, it may be used in one of the concrete settings where all of our assumptions are satisfied described at the end of Section 2.2.

Corollary 9. *Under Assumptions (A.1)-(A.5), there is an absolute constant $c > 0$, such that, for any $0 < \varepsilon < 1$ and $k \in \mathbb{N}$, if Σ is an instance of the (k, ε) -spike model, for any input dimension d , output dimension q , depth $L > 1$, and number of samples n , there are initialization scales $\alpha > 0$ and $\beta > 0$ such that the following holds. For all $\delta < 1/2$, if*

- the width $m \geq c(d + q + \log(1/\delta))$;
- the network is trained as described in Section 2.3;
- the input dimension $d \geq cn$;
- the number of samples $n \geq c \max\{k + \varepsilon d, \log(1/\delta)\}$,

then, with probability at least $1 - c\delta$,

$$\text{Risk}(\Theta) \leq c \left(\frac{\varepsilon d \|\Theta^*\|^2 + qk \log(q/\delta)}{n} + \frac{nq \log(q/\delta)}{d} \right). \quad (1)$$

Corollary 9 gives the simple bound obtained by a choice of parameters that includes a sufficiently small value of α . For larger values of α the bound of Theorem 7 may behave differently in the case of the (k, ε) -spike model. We find that if we regard $\|\Theta^*\|^2$ as a constant then, the excess risk approaches zero if

$$\frac{\varepsilon d}{n} \rightarrow 0, \quad \frac{qk \log(q/\delta)}{n} \rightarrow 0 \quad \text{and} \quad \frac{nq}{d} \rightarrow 0,$$

which recovers the known sufficient conditions for the minimum ℓ_2 -norm interpolant to benignly overfit in this setting. One example is where

$$q = 5, k = 5, \delta = 1/100, d = n^2, \varepsilon = 1/n^2,$$

and $n \rightarrow \infty$.

4. Proof of Theorem 6

The proof of Theorem 6 needs some lemmas, which we prove first. Throughout this section the assumptions of Theorem 6 are in force.

A key point is that the projection of any interpolator onto the row span of X , including the model output by training a deep linear network, is the minimum ℓ_2 -norm interpolant.

Lemma 10. *For any interpolator Θ , $\Theta_X = P_X \Theta = \Theta_{\ell_2}$.*

Proof. Since Θ interpolates the data

$$Y = X\Theta = X(\Theta_X + \Theta_{X^\perp}) = X\Theta_X. \quad (2)$$

Recall that $\Theta_X = X^\top (XX^\top)^{-1} X\Theta = P_X \Theta$, where P_X projects onto the row span of X . Continuing, we get that

$$\begin{aligned} \Theta_{\ell_2} &= X^\top (XX^\top)^{-1} Y \\ &= X^\top (XX^\top)^{-1} X\Theta_X \\ &= P_X \Theta_X = P_X P_X \Theta = P_X \Theta = \Theta_X. \end{aligned}$$

■

Using the formula for the minimum ℓ_2 -norm interpolant, we can now write down an expression for the excess risk.

Lemma 11. *The excess risk of any interpolator Θ of the data satisfies*

$$\text{Risk}(\Theta) \leq c \text{Tr} \left((\Theta^* - \Theta_{X^\perp})^\top B (\Theta^* - \Theta_{X^\perp}) \right) + cq \log(q/\delta) \text{Tr}(C)$$

with probability at least $1 - \delta$ over the noise matrix $\Omega = Y - X\Theta^*$, where

$$\begin{aligned} B &:= \left(I - X^\top (XX^\top)^{-1} X \right) \Sigma \left(I - X^\top (XX^\top)^{-1} X \right) \quad \text{and} \\ C &:= (XX^\top)^{-1} X \Sigma X^\top (XX^\top)^{-1}. \end{aligned}$$

Proof. We have $y - x\Theta^*$ is conditionally mean-zero given x , thus

$$\begin{aligned} \text{Risk}(\Theta) &= \mathbb{E}_{x,y} [\|y - x\Theta\|^2] - \mathbb{E}_{x,y} [\|y - x\Theta^*\|^2] \\ &= \mathbb{E}_{x,y} [\|y - x\Theta^* + x(\Theta^* - \Theta)\|^2] - \mathbb{E}_{x,y} [\|y - x\Theta^*\|^2] \\ &= \mathbb{E}_x [\|x(\Theta^* - \Theta)\|^2]. \end{aligned} \tag{3}$$

Since Θ interpolates the data, by Lemma 10 we know that

$$\Theta = \Theta_{\ell_2} + \Theta_{X^\perp} = X^\top (XX^\top)^{-1} Y + \Theta_{X^\perp}.$$

Now because $Y = X\Theta^* + \Omega$ we find that

$$\begin{aligned} \text{Risk}(\Theta) &= \mathbb{E}_x \left[\left\| x \left(I - X^\top (XX^\top)^{-1} X \right) (\Theta^* - \Theta_{X^\perp}) - x X^\top (XX^\top)^{-1} \Omega \right\|^2 \right] \\ &\leq 2\mathbb{E}_x \left[\left\| x \left(I - X^\top (XX^\top)^{-1} X \right) (\Theta^* - \Theta_{X^\perp}) \right\|^2 \right] + 2\mathbb{E}_x \left[\left\| x X^\top (XX^\top)^{-1} \Omega \right\|^2 \right] \\ &\leq 2\mathbb{E}_x \left[x \left(I - X^\top (XX^\top)^{-1} X \right) (\Theta^* - \Theta_{X^\perp}) (\Theta^* - \Theta_{X^\perp})^\top \left(I - X^\top (XX^\top)^{-1} X \right) x^\top \right] \\ &\quad + 2\mathbb{E}_x \left[x X^\top (XX^\top)^{-1} \Omega \Omega^\top (XX^\top)^{-1} X x^\top \right] \\ &\stackrel{(i)}{=} 2\text{Tr} \left((\Theta^* - \Theta_{X^\perp})^\top \left(I - X^\top (XX^\top)^{-1} X \right) \mathbb{E}_x \left[x^\top x \right] \left(I - X^\top (XX^\top)^{-1} X \right) (\Theta^* - \Theta_{X^\perp}) \right) \\ &\quad + 2\text{Tr} \left(\Omega^\top (XX^\top)^{-1} X \mathbb{E}_x \left[x^\top x \right] X^\top (XX^\top)^{-1} \Omega \right) \\ &\stackrel{(ii)}{=} 2\text{Tr} \left((\Theta^* - \Theta_{X^\perp})^\top B (\Theta^* - \Theta_{X^\perp}) \right) + 2\text{Tr} \left(\Omega^\top C \Omega \right). \end{aligned}$$

where (i) follows by using the cyclic property of the trace, and (ii) follows by the definition of the matrices B and C .

Let $\omega_1, \dots, \omega_q$ denote the columns of the error matrix Ω . Then

$$\text{Tr}(\Omega^\top C \Omega) = \sum_{i=1}^q \omega_i^\top C \omega_i.$$

Invoking (Bartlett et al., 2020, Lemma S.2) bounds each term in the sum by $cq \log(q/\delta) \text{Tr}(C)$ with probability at least $1 - \delta/q$. A union bound completes the proof. \blacksquare

To work on the first term in the upper bound of the excess risk, we would like an upper bound on $\|\Theta_{X^\perp}^{(t)}\|$. Toward this end, we first establish a high-probability bound on $\|X\|_{op}$.

Lemma 12. *There is a constant $c > 0$ such that for any $\delta \in (0, 1)$, if $n \geq c \max\{r_0, \log(\frac{1}{\delta})\}$, with probability at least $1 - \delta$, $\|X\|_{op} \leq c\sqrt{\lambda_1 n}$.*

Proof. By (Koltchinskii and Lounici, 2017, Lemma 9), with probability at least $1 - \delta$

$$\begin{aligned} \|X\|_{op} &= \sqrt{\|X^\top X\|_{op}} \\ &\leq \sqrt{n \left(\|\Sigma\|_{op} + \left\| \frac{1}{n} X^\top X - \Sigma \right\|_{op} \right)} \\ &\leq \sqrt{n \left(\|\Sigma\|_{op} + \|\Sigma\|_{op} \max \left\{ \sqrt{\frac{r_0}{n}}, \sqrt{\frac{\log(1/\delta)}{n}}, \frac{r_0}{n}, \frac{\log(1/\delta)}{n} \right\} \right)}. \end{aligned}$$

Recalling that $n \geq c \max\{r_0, \log(1/\delta)\}$, this implies that, with probability at least $1 - \delta$, $\|X\|_{op} \leq c\sqrt{n\|\Sigma\|_{op}}$. \blacksquare

Next, we will calculate a formula for the time derivative of $\Theta^{(t)}$. Its definition will make use of products of matrices before and after a given layer.

Definition 13. For $j \in [L]$ define $A_j = \prod_{k=L}^{j+1} W_k^{(t)}$ and $B_j = \prod_{k=j-1}^1 W_k^{(t)}$.

Now we are ready for our lemma giving the time derivative of $\Theta^{(t)}$.

Lemma 14. *At any time $t \geq 0$,*

$$\frac{d\Theta^{(t)}}{dt} = - \sum_{j=1}^L B_j^\top B_j X^\top (X\Theta^{(t)} - Y) A_j A_j^\top.$$

Proof. Let us suppress the superscript (t) to ease notation. The gradient flow dynamics is defined as

$$\frac{dW_j}{dt} = -\nabla_{W_j} \mathcal{L}(\Theta),$$

where

$$\nabla_{W_j} \mathcal{L}(\Theta) = (W_L \cdots W_{j+1})^\top (X\Theta - Y)^\top (W_{j-1} \cdots W_1 X^\top)^\top. \quad (4)$$

So by the chain rule of differentiation,

$$\begin{aligned}
 \frac{d\Theta}{dt} &= \frac{d(W_1^\top \cdots W_L^\top)}{dt} \\
 &= \sum_{j=1}^L \left(W_1^\top \cdots W_{j-1}^\top \right) \frac{dW_j^\top}{dt} \left(W_{j+1}^\top \cdots W_L^\top \right) \\
 &= \sum_{j=1}^L \left(W_1^\top \cdots W_{j-1}^\top \right) \left(-\nabla_{W_j} \mathcal{L}(\Theta) \right)^\top \left(W_{j+1}^\top \cdots W_L^\top \right) \\
 &= - \sum_{j=1}^L \left(W_1^\top \cdots W_{j-1}^\top \right) \left(W_{j-1} \cdots W_1 X^\top \right) (X\Theta - Y) (W_L \cdots W_{j+1}) \left(W_{j+1}^\top \cdots W_L^\top \right) \\
 &= - \sum_{j=1}^L B_j^\top B_j X^\top (X\Theta - Y) A_j A_j^\top.
 \end{aligned}$$

■

Toward the goal of proving a high-probability bound on $\|\Theta_{X^\perp}^{(t)}\|$, we next bound its rate of growth.

Lemma 15. *There is a constant $c > 0$ such that, if $n \geq c \max\{r_0, \log(1/\delta)\}$, with probability at least $1 - \delta$, if training is perpetually Λ bounded, then, for all $t \geq 0$,*

$$\frac{1}{2} \frac{d\|\Theta_{X^\perp}^{(t)}\|^2}{dt} \leq \sum_{j=2}^L \text{Tr} \left(\Theta^{(t)\top} P_{X^\perp} B_j^\top B_j X^\top (X\Theta^{(t)} - Y) A_j A_j^\top \right).$$

Proof. Given matrices A and B , we let $A \cdot B = \text{Tr}(A^\top B)$ denote the matrix inner product.

By the chain rule,

$$\begin{aligned}
 \frac{1}{2} \frac{d\|\Theta_{X^\perp}^{(t)}\|^2}{dt} &= \Theta_{X^\perp}^{(t)} \cdot \frac{d\Theta_{X^\perp}^{(t)}}{dt} \\
 &= \Theta_{X^\perp}^{(t)} \cdot P_{X^\perp} \frac{d\Theta^{(t)}}{dt} \\
 &\stackrel{(i)}{=} \Theta_{X^\perp}^{(t)} \cdot P_{X^\perp} \left(- \sum_{j=1}^L B_j^\top B_j X^\top (X\Theta^{(t)} - Y) A_j A_j^\top \right) \\
 &= - \sum_{j=1}^L \Theta_{X^\perp}^{(t)} \cdot P_{X^\perp} \left(B_j^\top B_j X^\top (X\Theta^{(t)} - Y) A_j A_j^\top \right), \tag{5}
 \end{aligned}$$

where (i) follows by the formula derived in Lemma 14.

Let us consider a particular term in the sum above,

$$\begin{aligned}
 & \Theta_{X^\perp}^{(t)} \cdot P_{X^\perp} \left(B_j^\top B_j X^\top (X\Theta^{(t)} - Y) A_j A_j^\top \right) \\
 &= \text{Tr} \left(\Theta_{X^\perp}^{(t)\top} P_{X^\perp} B_j^\top B_j X^\top (X\Theta^{(t)} - Y) A_j A_j^\top \right) \\
 &= \text{Tr} \left(\Theta^{(t)\top} P_{X^\perp}^\top P_{X^\perp} B_j^\top B_j X^\top (X\Theta^{(t)} - Y) A_j A_j^\top \right) \\
 &= \text{Tr} \left(\Theta^{(t)\top} P_{X^\perp} B_j^\top B_j X^\top (X\Theta^{(t)} - Y) A_j A_j^\top \right). \tag{6}
 \end{aligned}$$

In the case where $j = 1$, the RHS is equal to

$$\text{Tr} \left(\Theta^{(t)\top} P_{X^\perp} X^\top (X\Theta^{(t)} - Y) \left(\prod_{k=L}^2 W_k^{(t)} \right) \left(\prod_{k=L}^2 W_k^{(t)} \right)^\top \right) = 0,$$

since $P_{X^\perp} X^\top = 0$.

In the case $j > 1$, we have

$$\begin{aligned}
 & \Theta_{X^\perp}^{(t)} \cdot P_{X^\perp} \left(B_j^\top B_j X^\top (X\Theta^{(t)} - Y) A_j A_j^\top \right) \\
 &= \text{Tr} \left(\Theta^{(t)\top} P_{X^\perp} B_j^\top B_j X^\top (X\Theta^{(t)} - Y) A_j A_j^\top \right)
 \end{aligned}$$

completing the proof. ■

Lemma 16. *There is a constant $c > 0$ such that, if $n \geq c \max\{r_0, \log(1/\delta)\}$, with probability at least $1 - \delta$, if training is perpetually Λ bounded, then, for all $t \geq 0$,*

$$\|\Theta_{X^\perp}^{(t)}\| \leq \|\Theta_{X^\perp}^{(0)}\| + c(L-1)\Lambda^2 \sqrt{\lambda_1 n} \int_{s=0}^t \sqrt{\mathcal{L}(\Theta^{(s)})} \, ds.$$

Proof. Let us consider one of the terms in the RHS of Lemma 15. We have

$$\begin{aligned}
 & \text{Tr} \left(\Theta^{(t)\top} P_{X^\perp} B_j^\top B_j X^\top (X\Theta^{(t)} - Y) A_j A_j^\top \right) \\
 &= \text{Tr} \left(\Theta_{X^\perp}^{(t)\top} B_j^\top B_j X^\top (X\Theta^{(t)} - Y) A_j A_j^\top \right) \\
 &\leq \|\Theta_{X^\perp}^{(t)}\| \left\| B_j^\top B_j X^\top (X\Theta^{(t)} - Y) A_j A_j^\top \right\| \\
 &\stackrel{(i)}{\leq} \|\Theta_{X^\perp}^{(t)}\| \left\| B_j^\top B_j \right\|_{op} \left\| A_j A_j^\top \right\|_{op} \|X\|_{op} \|X\Theta^{(t)} - Y\| \\
 &\leq \|\Theta_{X^\perp}^{(t)}\| \left(\prod_{k=1}^{j-1} \|W_k^{(t)}\|_{op}^2 \right) \left(\prod_{k=j+1}^L \|W_k^{(t)}\|_{op}^2 \right) \|X\|_{op} \|X\Theta^{(t)} - Y\| \\
 &= \|\Theta_{X^\perp}^{(t)}\| \left(\prod_{k \neq j} \|W_k^{(t)}\|_{op} \right)^2 \|X\|_{op} \|X\Theta^{(t)} - Y\| \\
 &\stackrel{(ii)}{\leq} \|\Theta_{X^\perp}^{(t)}\| \|X\|_{op} \Lambda^2 \sqrt{\mathcal{L}(\Theta^{(t)})},
 \end{aligned}$$

where (i) follows since for any matrices $\|AB\| \leq \|A\|_{op}\|B\|$, and (ii) follows since training is perpetually Λ bounded.

Summing over layers $j = 2, \dots, L$, we get that,

$$\frac{1}{2} \frac{d\|\Theta_{X^\perp}^{(t)}\|^2}{dt} \leq (L-1)\|\Theta_{X^\perp}^{(t)}\| \|X\|_{op} \Lambda^2 \sqrt{\mathcal{L}(\Theta^{(t)})}.$$

Now note that,

$$\frac{1}{2} \frac{d\|\Theta_{X^\perp}^{(t)}\|^2}{dt} = \frac{\|\Theta_{X^\perp}^{(t)}\| d\|\Theta_{X^\perp}^{(t)}\|}{dt} \leq (L-1)\|\Theta_{X^\perp}^{(t)}\| \|X\|_{op} \Lambda^2 \sqrt{\mathcal{L}(\Theta^{(t)})},$$

which in turn implies that, when $\|\Theta_{X^\perp}^{(t)}\| > 0$, we have

$$\frac{d\|\Theta_{X^\perp}^{(t)}\|}{dt} \leq (L-1) \|X\|_{op} \Lambda^2 \sqrt{\mathcal{L}(\Theta^{(t)})}.$$

If, for all $s \in [0, t]$, we have $\|\Theta_{X^\perp}^{(s)}\| \neq 0$, then by integrating this differential inequality we conclude that

$$\|\Theta_{X^\perp}^{(t)}\| - \|\Theta_{X^\perp}^{(0)}\| \leq (L-1) \|X\|_{op} \Lambda^2 \int_{s=0}^t \sqrt{\mathcal{L}(\Theta^{(s)})} ds. \quad (7)$$

Otherwise, if $T = \sup\{s : \|\Theta_{X^\perp}^{(s)}\| = 0\}$,

$$\|\Theta_{X^\perp}^{(t)}\| \leq (L-1) \|X\|_{op} \Lambda^2 \int_{s=T}^t \sqrt{\mathcal{L}(\Theta^{(s)})} ds,$$

which implies (7).

Applying Lemma 12 which is a high probability upper bound on $\|X\|_{op}$ completes the proof. \blacksquare

Armed with these lemmas, we are now ready to prove the first of our main results.

Proof of Theorem 6. Combining Lemma 11 with Lemmas 6 and 11 by Bartlett et al. (2020) to bound $\text{Tr}(C)$ we get, with probability at least $1 - c\delta$,

$$\text{Risk}(\Theta) \leq \text{Tr} \left((\Theta^* - \Theta_{X^\perp})^\top B (\Theta^* - \Theta_{X^\perp}) \right) + \underbrace{cq \log(q/\delta) \left(\frac{k}{n} + \frac{n}{R_k} \right)}_{=:\text{Variance}(\Theta_{\ell_2})}.$$

We begin by bounding the first term in the RHS above. Let $\theta_1^*, \dots, \theta_q^*$ be the columns of Θ^* and $\theta_{X^\perp, 1}, \dots, \theta_{X^\perp, q}$ be the columns of Θ_{X^\perp} . By invoking (Chatterji et al., 2022, Eq. 54) for each of the q outputs, and applying a union bound, we find that with probability at

least $1 - cq(\delta/q) = 1 - c\delta$,

$$\begin{aligned} \text{Tr} \left((\Theta^* - \Theta_{X^\perp})^\top B (\Theta^* - \Theta_{X^\perp}) \right) &= \sum_{i=1}^q (\theta_i^* - \theta_{X^\perp, i})^\top B (\theta_i^* - \theta_{X^\perp, i}) \\ &\leq \frac{cS_k}{n} \sum_{i=1}^q \|\theta_i^* - \theta_{X^\perp, i}\|^2 \\ &= \frac{cS_k}{n} \|\Theta^* - \Theta_{X^\perp}\|^2 \\ &\leq \frac{2cS_k}{n} (\|\Theta^*\|^2 + \|\Theta_{X^\perp}\|^2). \end{aligned}$$

Define $\text{Bias}(\Theta_{\ell_2}) := \frac{2cS_k}{n} \|\Theta^*\|^2$ and let $\Xi := \frac{2cS_k}{n} \|\Theta_{X^\perp}\|^2$. The bound on Ξ follows by invoking Lemma 16. \blacksquare

5. Proof of Theorem 7

The assumptions of Theorem 7 are in force throughout this section. Before starting its proof, we establish some lemmas.

Definition 17. *For a large enough absolute constant c , we say that the network enjoys a δ -good initialization if*

$$\begin{aligned} \alpha/c &< \sigma_{\min}(W_1^{(0)}) \leq \sigma_{\max}(W_1^{(0)}) < c\alpha, \\ \beta/c &< \sigma_{\min}(W_L^{(0)}) \leq \sigma_{\max}(W_L^{(0)}) < c\beta, \end{aligned}$$

and

$$\mathcal{L}(\Theta^{(0)}) < c \left(\|Y\|^2 + \frac{\alpha^2 \beta^2 q \|X\|^2 \log(n/\delta)}{m} \right).$$

The following proposition is proved in Appendix B. It guarantees that for wide networks, optimization is successful starting from random initialization.

Proposition 18. *There is a constant c such that, given any $\delta \in (0, 1)$, if the initialization scales α and β , along with the network width m , satisfy*

$$\begin{aligned} m &\geq c \max \left\{ d + q + \log(1/\delta), \frac{L^2 \alpha^2 \|X\|_{op}^2 \|X\|^2 q \log(n/\delta)}{\beta^2 \sigma_{\min}^4(X)} \right\}, \\ \beta &\geq c \max \left\{ 1, \sqrt{\frac{L \|X\|_{op} \|Y\|}{\sigma_{\min}^2(X)}} \right\}, \\ \alpha &\leq 1, \end{aligned}$$

then with probability at least $1 - \delta$:

1. the initialization is δ -good;
2. training is perpetually $c(\alpha + 1/L)\beta$ bounded;

3. for all $t > 0$, we have that

$$\begin{aligned} \mathcal{L}(\Theta^{(t)}) &< \mathcal{L}(\Theta^{(0)}) \exp\left(-\frac{\beta^2 \sigma_{\min}^2(X)}{4e} \cdot t\right) \\ &\leq c \left(\|Y\|^2 + \frac{\alpha^2 \beta^2 q \|X\|^2 \log(n/\delta)}{m} \right) \exp\left(-\frac{\beta^2 \sigma_{\min}^2(X)}{4e} \cdot t\right). \end{aligned}$$

The reader may notice that the roles of α and β in Proposition 18 are asymmetric. We focused on that case that α is small because the updates of W_1 are in the span of the rows of X , which is not necessarily the case for the other layers, including W_L . This means that the scale of W_1 in the null space of X remains the same as it was at initialization, so that a small scale at initialization pays dividends throughout training.

The next lemma shows that the projection of the model computed by the network onto the null space of X is the same as the model obtained by projecting the first layer weights, and combining them with the other layers.

Lemma 19. For all $t \geq 0$,

$$\Theta_{X^\perp}^{(t)} = \left(W_L^{(t)} \cdots W_2^{(t)} W_{1,X^\perp}^{(t)} \right)^\top,$$

where

$$W_{1,X^\perp}^{(t)} := W_1^{(t)}(I - P_X).$$

Proof. By definition

$$\Theta^{(t)} = (W_L^{(t)} \cdots W_1^{(t)})^\top = \left(W_1^{(t)} \right)^\top \cdots \left(W_L^{(t)} \right)^\top \in \mathbb{R}^{d \times q}.$$

Therefore,

$$\Theta_{X^\perp}^{(t)} = (I - P_X)\Theta^{(t)} = (I - P_X)(W_1^{(t)})^\top \cdots (W_L^{(t)})^\top = (W_{1,X^\perp}^{(t)})^\top \cdots (W_L^{(t)})^\top. \quad \blacksquare$$

The subsequent lemma shows that the projection of the first layer onto the null space of X does not change during training.

Lemma 20. For all $t \geq 0$, $W_{1,X^\perp}^{(t)} = W_{1,X^\perp}^{(0)}$.

Proof. We have

$$\begin{aligned} \frac{dW_{1,X^\perp}^{(t)}}{dt} &= \frac{dW_1^{(t)}(I - P_X)}{dt} \\ &= \left(\frac{dW_1^{(t)}}{dt} \right) (I - P_X) \\ &= - \left((W_L \cdots W_2)^\top (X\Theta - Y)^\top X \right) (I - P_X) \quad (\text{by using Eq. (4)}) \\ &= - \left((W_L \cdots W_2)^\top (X\Theta - Y)^\top \right) (X(I - P_X)) \\ &= - \left((W_L \cdots W_2)^\top (X\Theta - Y)^\top \right) (0) \\ &= 0. \quad \blacksquare \end{aligned}$$

By using the previous two lemmas regarding the first layer weights W_1 we can now prove an alternate bound on $\|\Theta_{X^\perp}^{(t)}\|$. In contrast to the previous bound that we derived in Lemma 16, here the initial scale of W_1 plays a role in controlling the growth in $\|\Theta_{X^\perp}^{(t)}\|$.

Lemma 21. *There is constant $c > 0$ such that, if training is perpetually Λ bounded, then, for all $t \geq 0$,*

$$\|\Theta_{X^\perp}^{(t)}\| \leq \|\Theta_{X^\perp}^{(0)}\| + cL\|W_1^{(0)}\|_{op}\|X\|_{op}\Lambda^2 \int_{s=0}^t \sqrt{\mathcal{L}(\Theta^{(s)})} ds.$$

Proof. Let us once again consider one of the terms in the RHS of Lemma 15. We have

$$\begin{aligned} & \text{Tr} \left(\Theta_{X^\perp}^{(t)\top} P_{X^\perp} B_j^\top B_j X^\top (X\Theta^{(t)} - Y) A_j A_j^\top \right) \\ &= \text{Tr} \left(\Theta_{X^\perp}^{(t)\top} P_{X^\perp} (W_1^{(t)})^\top \left(\prod_{k=j-1}^2 W_k^{(t)} \right)^\top B_j X^\top (X\Theta^{(t)} - Y) A_j A_j^\top \right) \\ &= \text{Tr} \left(\Theta_{X^\perp}^{(t)\top} W_{1,X^\perp}^\top \left(\prod_{k=j-1}^2 W_k^{(t)} \right)^\top B_j X^\top (X\Theta^{(t)} - Y) A_j A_j^\top \right). \end{aligned}$$

Continuing by using the fact that for any matrices $\|AB\| \leq \|A\|_{op}\|B\|$, we get that

$$\begin{aligned} & \Theta_{X^\perp}^{(t)} \cdot P_{X^\perp} \left(B_j^\top B_j X^\top (X\Theta^{(t)} - Y) A_j A_j^\top \right) \\ & \leq \|\Theta_{X^\perp}^{(t)}\| \|W_{1,X^\perp}^{(t)}\|_{op} \left\| \left(\prod_{k=j-1}^2 W_k^{(t)} \right)^\top B_j \right\|_{op} \|A_j A_j^\top\|_{op} \|X\|_{op} \|X\Theta^{(t)} - Y\| \\ & \stackrel{(i)}{\leq} \|\Theta_{X^\perp}^{(t)}\| \|W_{1,X^\perp}^{(t)}\|_{op} \|X\|_{op} \Lambda^2 \sqrt{\mathcal{L}(\Theta^{(t)})} \\ & \stackrel{(ii)}{\leq} \|\Theta_{X^\perp}^{(t)}\| \|W_{1,X^\perp}^{(0)}\|_{op} \|X\|_{op} \Lambda^2 \sqrt{\mathcal{L}(\Theta^{(t)})} \\ & \leq \|\Theta_{X^\perp}^{(t)}\| \|W_1^{(0)}\|_{op} \|X\|_{op} \Lambda^2 \sqrt{\mathcal{L}(\Theta^{(t)})}, \end{aligned}$$

since $\|W_{1,X^\perp}^{(0)}\|_{op} \leq \|W_1^{(0)}\|_{op}$, where (i) follows since training is Λ perpetually bounded and so

$$\left\| \left(\prod_{k=j-1}^2 W_k^{(t)} \right)^\top B_j \right\|_{op} \|A_j A_j^\top\|_{op} \leq \left(\prod_{k \neq \{1,j\}} \|W_k^{(t)}\|_{op} \right) \left(\prod_{k \neq \{j\}} \|W_k^{(t)}\|_{op} \right) \leq \Lambda^2$$

and (ii) follows since by Lemma 20, $W_{1,X^\perp}^{(t)} = W_{1,X^\perp}^{(0)}$.

Summing over layers $j = 2, \dots, L$, we get that,

$$\frac{1}{2} \frac{d\|\Theta_{X^\perp}^{(t)}\|^2}{dt} \leq (L-1) \|\Theta_{X^\perp}^{(t)}\| \|W_1^{(0)}\|_{op} \|X\|_{op} \Lambda^2 \sqrt{\mathcal{L}(\Theta^{(t)})}.$$

Thus, we have that

$$\frac{1}{2} \frac{d\|\Theta_{X^\perp}^{(t)}\|^2}{dt} = \frac{\|\Theta_{X^\perp}^{(t)}\| d\|\Theta_{X^\perp}^{(t)}\|}{dt} \leq (L-1) \|\Theta_{X^\perp}^{(t)}\| \|W_1^{(0)}\|_{op} \|X\|_{op} \Lambda^2 \sqrt{\mathcal{L}(\Theta^{(t)})}$$

which in turn implies that, when $\|\Theta_{X^\perp}^{(t)}\| \neq 0$, we have

$$\frac{d\|\Theta_{X^\perp}^{(t)}\|}{dt} \leq (L-1) \|W_1^{(0)}\|_{op} \|X\|_{op} \Lambda^2 \sqrt{\mathcal{L}(\Theta^{(t)})}.$$

Therefore, by integrating this differential inequality as in the proof of Lemma 16, we conclude that

$$\|\Theta_{X^\perp}^{(t)}\| - \|\Theta_{X^\perp}^{(0)}\| \leq (L-1) \|W_1^{(0)}\|_{op} \|X\|_{op} \Lambda^2 \int_{s=0}^t \sqrt{\mathcal{L}(\Theta^{(s)})} ds.$$

■

We also need a lemma that bounds the Frobenius norm of the data matrix X .

Lemma 22. *There is a constant $c > 0$ such that for any $\delta \in (0, 1)$, if $n \geq c \log(1/\delta)$, then with probability at least $1 - \delta$, $\|X\| \leq c\sqrt{ns_0}$.*

Proof. The rows of X are n i.i.d. draws from a distribution, where each sample can be written as $x_i = \Sigma^{1/2} u_i$, where u_i has components that are independent c_x -sub-Gaussian random variables. Define $u_{\text{stacked}} := (u_1, u_2, \dots, u_n) \in \mathbb{R}^{dn}$ to be concatenation of the vectors u_1, \dots, u_n and define $\Sigma_{\text{stacked}}^{1/2} \in \mathbb{R}^{dn \times dn}$ to be a block diagonal matrix with $\Sigma^{1/2} \in \mathbb{R}^{d \times d}$ repeated n times along its diagonal. Then,

$$\|X\|^2 = \sum_{i=1}^n \|x_i\|^2 = \sum_{i=1}^n \|\Sigma^{1/2} u_i\|^2 = \left\| \Sigma_{\text{stacked}}^{1/2} u_{\text{stacked}} \right\|^2.$$

Now, u_{stacked} is an isotropic, c_x -sub-Gaussian random vector. Therefore, by applying (Vershynin, 2018, Theorem 6.3.2) we know that the sub-Gaussian norm (Vershynin, 2018, Definition 2.5.3) of $\|X\| = \|\Sigma_{\text{stacked}}^{1/2} u_{\text{stacked}}\|$ is

$$\left\| \|\Sigma_{\text{stacked}}^{1/2} u_{\text{stacked}}\| - c_1 \sqrt{n \text{Tr}(\Sigma)} \right\|_{\psi_2} = \left\| \|X\| - c_1 \sqrt{ns_0} \right\|_{\psi_2} \leq c_2 \sqrt{\lambda_1}.$$

Therefore, by Hoeffding's bound (Vershynin, 2018, Proposition 2.5.2) we get that

$$\mathbb{P} \left[\left| \|X\| - c_1 \sqrt{ns_0} \right| \geq \eta \right] \leq 2 \exp(-c_3 \eta^2 / \lambda_1).$$

Setting $\eta^2 = ns_0 / \lambda_1 = nr_0$ and noting that $n \geq \log(1/\delta) \geq \log(1/\delta) / r_0$ completes the proof. ■

Finally, we have a simple lemma that bounds the Frobenius norm of the responses Y .

Lemma 23. *There is a constant $c > 0$ such that for any $\delta \in (0, 1)$, if $n \geq c \log(1/\delta)$, then with probability at least $1 - \delta$, $\|Y\| \leq c(\|X\|_{op} \|\Theta^*\| + \sqrt{qn})$.*

Proof. Note that $Y = X\Theta^* + \Omega$, and therefore

$$\|Y\| \leq \|X\Theta^*\| + \|\Omega\| \leq \|X\|_{op}\|\Theta^*\| + \|\Omega\|, \quad (8)$$

where the last inequality follows since for any matrices $\|AB\| \leq \|A\|_{op}\|B\|$. Now each entry in $\Omega \in \mathbb{R}^{n \times q}$ is a zero-mean and c_y -sub-Gaussian. Therefore, by Bernstein's bound (Ver-shynin, 2018, Theorem 2.8.1),

$$\mathbb{P} [\|\Omega\|^2 - \mathbb{E} [\|\Omega\|^2] \geq qn] \leq 2 \exp(-c_1qn).$$

Now $\mathbb{E} [\|\Omega\|^2] = n\mathbb{E} [\|y - x\Theta^*\|^2] \leq c_2qn$, by Assumption (A.4), and $2 \exp(-qn) \leq \delta$ since $n \geq c \log(1/\delta) \geq c \log(1/\delta)/q$. Thus, with probability at least $1 - \delta$

$$\|\Omega\|^2 \leq c_3qn.$$

Combining this with Eq. (8) completes the proof. ■

With all of the pieces in place we are now ready to prove the theorem.

Proof of Theorem 7. Define a “good event” \mathcal{E} as the intersection of the following events:

- \mathcal{E}_1 , the excess risk bound stated in Theorem 6 holds.
- \mathcal{E}_2 , the bounds stated in Proposition 18 hold.
- \mathcal{E}_3 , $\|X\|_{op} \leq c\sqrt{\lambda_1 n}$.
- \mathcal{E}_4 , $\|X\| \leq c\sqrt{s_0 n}$.
- \mathcal{E}_5 , $\sigma_{\min}(X) \geq \frac{\sqrt{s_k}}{c}$.
- \mathcal{E}_6 , $\|Y\| \leq c(\|X\|_{op}\|\Theta^*\| + \sqrt{qn})$.

Now, Theorem 6 and Proposition 18 each hold with probability at least $1 - c\delta$. Lemma 12 implies that the event \mathcal{E}_3 holds with probability at least $1 - \delta$. By Lemma 22, the event \mathcal{E}_4 holds with probability at least $1 - \delta$. For \mathcal{E}_5 , notice that

$$\sigma_{\min}(X) = \sqrt{\sigma_{\min}(XX^\top)} = \sqrt{\sigma_{\min}(X_{:k}X_{:k}^\top + X_{k:}X_{k:}^\top)} \geq \sqrt{\sigma_{\min}(X_{k:}X_{k:}^\top)} = \sigma_{\min}(X_{k:}),$$

where $X_{:k}$ are the first k columns of X and $X_{k:}$ are the last $d - k$ columns of X . Since $n \geq c \log(1/\delta)$, by (Bartlett et al., 2020, Lemma 9) we know that with probability at least $1 - \delta$

$$\sigma_{\min}(X) \geq \sigma_{\min}(X_{k:}) \geq \sqrt{\frac{s_k}{c_1} \left(1 - \frac{c_1 n}{r_k}\right)} \geq \frac{\sqrt{s_k}}{c} \quad (\text{since } r_k \geq bn \text{ by Definition 2}).$$

Finally, by Lemma 23 event \mathcal{E}_6 holds with probability at least $1 - \delta$. Therefore, by a union bound the good event \mathcal{E} holds with probability at least $1 - c'\delta$. Let us assume that this event occurs going forward in the proof.

Proposition 18 guarantees that the training process is $c_2(\alpha + 1/L)\beta$ -perpetually bounded and the loss converges to zero. Therefore, by applying Theorem 6, the risk is bounded by

$$\text{Risk}(\Theta) \leq \text{Bias}(\Theta_{\ell_2}) + \text{Variance}(\Theta_{\ell_2}) + \Xi,$$

where

$$\begin{aligned} \text{Bias}(\Theta_{\ell_2}) &\leq \frac{cs_k}{n} \|\Theta^*\|^2, \\ \text{Variance}(\Theta_{\ell_2}) &\leq cq \log(q/\delta) \left(\frac{k}{n} + \frac{n}{R_k} \right), \\ &\equiv \leq \frac{cs_k}{n} \left[\|\Theta_{X^\perp}^{(0)}\| + L\alpha(\alpha + 1/L)^2 \beta \sqrt{\lambda_1 n} \int_{t=0}^{\infty} \sqrt{\mathcal{L}(\Theta^{(t)})} dt \right]^2. \end{aligned} \quad (9)$$

In the rest of the proof we shall bound the term Ξ .

For this, we would like to apply Proposition 18, which we can, since $\alpha \leq 1$,

$$\begin{aligned} \beta &\geq c_3 \max \left\{ 1, \frac{\lambda_1^{1/4} \sqrt{Ln} \left(\sqrt{\|\Theta^*\|} \lambda^{1/4} + q^{1/4} \right)}{\sqrt{s_k}} \right\} \\ &\geq c_4 \max \left\{ 1, \sqrt{\frac{L\|X\|_{op}\|Y\|}{\sigma_{\min}^2(X)}} \right\} \end{aligned} \quad (\text{by events } \mathcal{E}_3, \mathcal{E}_5 \text{ and } \mathcal{E}_6),$$

and

$$\begin{aligned} m &\geq c_5 \max \left\{ d + q + \log(1/\delta), \frac{L^2 \alpha^2 \lambda_1 s_0 n^2 q \log(n/\delta)}{\beta^2 s_k^2} \right\} \\ &\geq c_6 \max \left\{ d + q + \log(1/\delta), \frac{L^2 \alpha^2 \|X\|_{op}^2 \|X\|^2 q \log(n/\delta)}{\beta^2 \sigma_{\min}^4(X)} \right\} \end{aligned} \quad (\text{by events } \mathcal{E}_3 \text{ and } \mathcal{E}_4).$$

Thus, by Proposition 18 we know that for all $t > 0$,

$$\begin{aligned} \mathcal{L}(\Theta^{(t)}) &< c_7 \left(\|Y\|^2 + \frac{\alpha^2 \beta^2 q \|X\|^2 \log(n/\delta)}{m} \right) \exp \left(-\frac{\beta^2 \sigma_{\min}^2(X)}{4e} \cdot t \right) \\ &< c_8 \left(\lambda_1 n \|\Theta^*\|^2 + qn + \frac{\alpha^2 \beta^2 q s_0 n \log(n/\delta)}{m} \right) \exp(-c_9 \beta^2 s_k \cdot t) \quad (\text{by events } \mathcal{E}_3\text{-}\mathcal{E}_6). \end{aligned}$$

Integrating the RHS above we get that

$$\int_{t=0}^{\infty} \sqrt{\mathcal{L}(\Theta^{(t)})} dt \leq c_{10} \frac{\sqrt{(\lambda_1 \|\Theta^*\|^2 + q)n} + \alpha \beta \sqrt{\frac{s_0 q n \log(n/\delta)}{m}}}{\beta^2 s_k}. \quad (10)$$

Proposition 18 also guarantees that the initialization is δ -good. That is, $\|W_1^{(0)}\|_{op} \leq c_{11}\alpha$ and $\|W_L^{(0)}\| \leq c_{11}\beta$. So,

$$\begin{aligned}
 \|\Theta_{X^\perp}^{(0)}\| &= \|(I - P_X)\Theta^{(0)}\| \leq \|(I - P_X)\|_{op}\|\Theta^{(0)}\| \\
 &\leq \|\Theta^{(0)}\| \\
 &= \|W_L^{(0)}W_{L-1}^{(0)} \cdots W_1^{(0)}\| \\
 &= \|W_L^{(0)}W_1^{(0)}\| \quad (\text{since } W_2^{(0)} = \cdots = W_{L-1}^{(0)} = I) \\
 &\leq \min \left\{ \|W_1^{(0)}\|_{op}\|W_L^{(0)}\|, \|W_1^{(0)}\| \|W_L^{(0)}\|_{op} \right\} \\
 &\stackrel{(i)}{\leq} c_{12}\sqrt{\min\{q, d\}}\alpha\beta \\
 &\leq c_{12}\sqrt{q}\alpha\beta,
 \end{aligned} \tag{11}$$

where (i) follows since the initialization was good, and the ranks of $W_1^{(0)}$ and $W_L^{(0)}$ are bounded by d and q respectively.

Plugging the bounds obtained in Eqs. (10) and (11) into Eq. (9) we have that

$$\begin{aligned}
 &\equiv \leq \frac{c_{13}s_k}{n} \left[\sqrt{q}\alpha\beta + L\alpha(\alpha + 1/L)^2\beta^2\sqrt{\lambda_1 n} \left(\frac{\sqrt{(\lambda_1\|\Theta^\star\|^2 + q)n} + \alpha\beta\sqrt{\frac{s_0qn \log(n/\delta)}{m}}}{\beta^2 s_k} \right) \right]^2 \\
 &\leq \frac{c_{13}s_k}{n} \left[\sqrt{q}\alpha\beta + \frac{L\alpha(\alpha + 1/L)^2\sqrt{\lambda_1 n} \left(\sqrt{(\lambda_1\|\Theta^\star\|^2 + q)n} + \alpha\beta\sqrt{\frac{s_0qn \log(n/\delta)}{m}} \right)}{s_k} \right]^2 \\
 &\leq \frac{c_{14}\alpha^2 s_k}{n} \left[q\beta^2 + \frac{L^2(\alpha + 1/L)^4\lambda_1 n \left((\lambda_1\|\Theta^\star\|^2 + q)n + \frac{\alpha^2\beta^2 s_0qn \log(n/\delta)}{m} \right)}{s_k^2} \right] \\
 &\leq \frac{c_{14}\alpha^2 s_k}{n} \left[q\beta^2 + \frac{L^2(\alpha + 1/L)^4\lambda_1 n^2}{s_k^2} \left(\lambda_1\|\Theta^\star\|^2 + q + \frac{\alpha^2\beta^2 s_0q \log(n/\delta)}{m} \right) \right].
 \end{aligned}$$

This completes our proof. ■

6. Proof of Corollary 9

When Σ is an instance of the (k, ε) -spike model we find that

$$r_0 = s_0/\lambda_1 = k + \varepsilon(d - k), \quad s_k = \varepsilon(d - k) \quad \text{and} \quad R_k = (d - k). \tag{12}$$

First, for a large enough c_1 , we set

$$\beta = c_1 \max \left\{ 1, \frac{\lambda_1^{1/4} \sqrt{Ln} \left(\sqrt{\|\Theta^\star\| \lambda_1^{1/4} + q^{1/4}} \right)}{\sqrt{s_k}} \right\} = c_1 \max \left\{ 1, \frac{\sqrt{Ln} \left(\sqrt{\|\Theta^\star\| + q^{1/4}} \right)}{\sqrt{\varepsilon(d - k)}} \right\}.$$

Given this choice of β , for any q, n, k, d, L , if $\alpha > 0$ is chosen to be small enough then,

$$m \geq c_1(d + q + \log(1/\delta)) = c_1 \max \left\{ d + q + \log(1/\delta), \frac{L^2 \alpha^2 \lambda_1 s_0 n^2 q \log(n/\delta)}{\beta^2 s_k^2} \right\}.$$

Also by the assumption on the number of samples,

$$n \geq c_2 \max \{k + \varepsilon d, \log(1/\delta)\} \geq c_2 \max \{r_0, k, \log(1/\delta)\}.$$

We are now in position to invoke Theorem 7. By this theorem we get that,

$$\begin{aligned} \text{Risk}(\Theta) &\leq \frac{c_3 \varepsilon (d - k)}{n} \|\Theta^*\|^2 + c_3 q \log(q/\delta) \left(\frac{k}{n} + \frac{n}{d - k} \right) + \Xi \\ &\leq \frac{c \varepsilon d}{n} \|\Theta^*\|^2 + c q \log(q/\delta) \left(\frac{k}{n} + \frac{n}{d} \right) + \Xi \quad (\text{since } d \geq c_4 k). \end{aligned}$$

Recall from above that the upper bound on Ξ scales with α^2 . Thus, for small enough α it is a lower order term.

7. Additional Simulations and Details

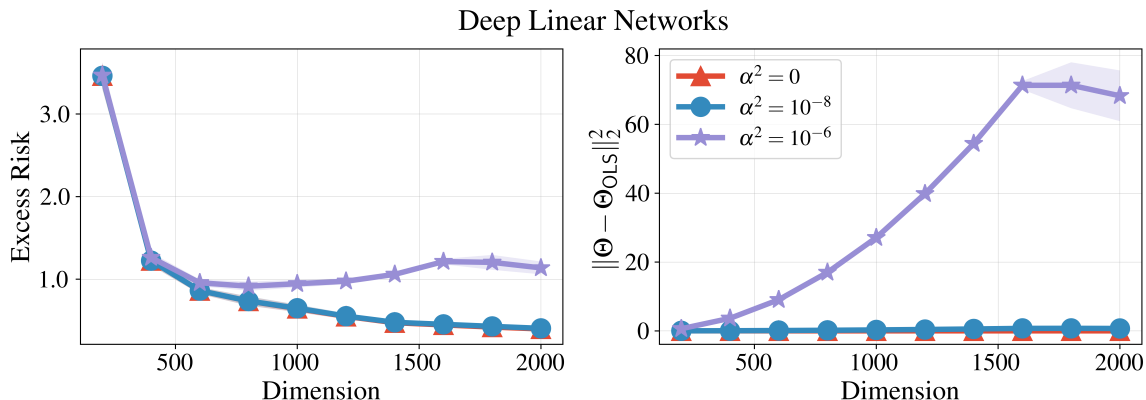


Figure 3: Excess risk and distance from the minimum ℓ_2 -norm interpolator of three-layer linear networks trained by gradient descent on data generated by an underlying linear model as the input dimension varies. The model is trained on $n = 100$ points drawn from the generative model $y = x\Theta^* + \omega$, where $x \sim \mathcal{N}(0, \Sigma)$ and $\omega \sim \mathcal{N}(0, 1)$. The excess risk is defined as $\mathbb{E}_x [\|x\Theta - x\Theta^*\|^2]$. In line with our theory, we find that when the initialization scale is small, final solution is close to the minimum ℓ_2 -norm interpolator and the resulting excess risk is small.

Inspired by our theory, we ran simulations to study the excess risk of several linear networks and ReLU networks as a function of both the initialization scale and dimension.¹ In line

1. Code at <https://github.com/niladri-chatterji/Benign-Deep-Linear>

with our theoretical upper bounds, we find that for deep linear networks as the initialization scale of either the first layer (α) or the last layer (β) is large, the excess risk of the model is larger (see Figure 1). In deep ReLU networks (see Figure 2), we find an asymmetry in the roles of α and β . The excess risk increases when we increase α , but is largely unaffected by the scale of the initialization of the final layer β .

In all of our figures we report the average over 20 runs. We also report the 95% confidence interval assuming that the statistic of interest follows a Gaussian distribution.

Setup for deep linear models. For Figures 1 and 3 the generative model for the underlying data was $y = x\Theta^* + \omega$, where

1. $\Theta^* \in \mathbb{R}^{d \times 3}$ is drawn uniformly over the set of matrices with unit Frobenius norm. The output dimension $q = 3$;
2. the covariates $x \sim \mathbf{N}(0, \Sigma)$, where the eigenvalues of Σ are as follows: $\lambda_1 = \dots = \lambda_{10} = 1$ and $\lambda_{11} = \dots = \lambda_d = 0.01$;
3. the noise ω is drawn independently from $\mathbf{N}(0, 1)$.

For these figures the number of samples $n = 100$ across all experiments. All of the models are trained on the squared loss with full-batch gradient descent with step-size 10^{-4} , until the training loss is smaller than 10^{-7} .

We train models that have 2 hidden layers ($L = 3$). The width of the middle layers m is set to be $10(d + q)$, where d is the input dimension and q is the output dimension.

For the top half of Figure 1 and Figure 3 when we vary the initialization scale of the first layer α , we initialize all of the middle layers to the identity, and initialize entries of the last layer with i.i.d. draws from $\mathbf{N}(0, 1)$.

For the bottom half of Figure 1 when we vary the initialization scale of the last layer α , we initialize all of the middle layers to the identity, and initialize entries of the first layer with i.i.d. draws from $\mathbf{N}(0, 1)$.

Setup for deep ReLU models. For Figure 2 the generative model for the underlying data was $y = f^*(x) + \omega$, where

1. $f^*(x)$ is a two-layer feedforward ReLU network with width 10 and output dimension 3 which was randomly initialized according to LeCun initialization;
2. the covariates $x \sim \mathbf{N}(0, I_{10 \times 10})$;
3. the noise ω is drawn independently from $\mathbf{N}(0, 1)$.

The networks are trained on $n = 500$ samples. Again, all of the models are trained on the squared loss with full-batch gradient descent with step-size 10^{-4} , until the training loss is smaller than 10^{-7} .

We train models that have $L = 3$ layers. The width of the middle layers (m) is set to be 50.

For left half of Figure 2 when we vary the initialization scale of the first layer α , we initialize all of the middle layers to the identity, and initialize entries of the last layer with i.i.d. draws from $\mathbf{N}(0, 1)$.

For the right half of Figure 2 when we vary the initialization scale of the last layer α , we initialize all of the middle layers to the identity, and initialize entries of the first layer with i.i.d. draws from $N(0, 1)$.

8. Discussion

We have provided upper bounds on the excess risk for deep linear networks that interpolate the data with respect to the quadratic loss, and presented simulation studies that verify that the some aspects of our bounds reflect typical behavior.

As mentioned in the introduction, our analysis describes a variety of conditions under which the generalization behavior of interpolating deep linear networks is similar, or the same, as the behavior of the minimum ℓ_2 -norm interpolant with the standard parameterization. Among other things, this motivates study of loss functions other than the quadratic loss used in this work. The softmax loss would be a natural choice.

Looking at our proofs, it appears that the only way that a deep linear parameterization can promote benign overfitting is for the function computed by the network at initialization to approximate the regression function. (Formalizing this with a lower bound, possibly in the case of random initialization, or with an arbitrary initialization and a randomly chosen regression function Θ^* , is a potential topic for further research.) The benefits of a good approximation to the regression function at initialization has been explored in the case of two-layer linear networks (Chatterji et al., 2022). Extending this analysis to deep networks is a potential subject for further study.

We focused on a particular random initialization scheme in this paper, it is possible to study other initialization schemes as well. For example, we believe that, if the width m of the network is somewhat larger, a similar analysis should go through without our simplifying assumption that W_2, \dots, W_{L-1} are initialized exactly to the identity, and instead are initialized randomly.

Recently, Mallinar et al. (2022) established conditions under which interpolation with the minimum ℓ_2 -norm interpolator is “tempered”, achieving risk within a constant factor of the Bayes risk. Here we show that the risk of interpolating deep linear networks is (nearly) equal to the risk of the minimum ℓ_2 -norm interpolator, this implies that when the minimum ℓ_2 -norm interpolator is tempered, so is the output of the deep linear model. We hope that our techniques lay the groundwork for other results about tempered overfitting.

While here we analyzed the network obtained by the continuous-time gradient flow it is straightforward to use our techniques to obtain similar results for gradient descent with small enough step-size at the expense of a more involved analysis.

As mentioned after the statement of Theorem 7, its bounds could potentially be improved. (We have not attempted to prove any lower bounds in this work.)

In Figure 1 we found that when the scale of the random initialization of either the first (α) or the last layer (β) goes to zero the trained model approach the minimum ℓ_2 -norm interpolator. Understanding why and when this happens is an avenue for future research.

Finally, examining the extent to which the effects described here carry over when non-linearities are present is a natural next step.

Acknowledgments

We thank anonymous reviewers for their careful reading of an earlier version of this paper and their valuable feedback.

Appendix A. Additional Related Work

In this appendix, we describe a wider variety of related work.

A.1 Benign Overfitting and Double Descent

This subsection includes descriptions of some of the most closely related work that we know. For a wider sample, we point the interested reader to a couple of surveys (Bartlett et al., 2021; Belkin, 2021).

Papers have studied the excess risk of the minimum ℓ_2 -norm interpolant (Bartlett et al., 2020; Hastie et al., 2022; Muthukumar et al., 2020; Bartlett and Long, 2021), which is obtained as a result of minimizing the squared loss using gradient descent with no explicit regularization. While these previous papers directly analyzed the closed form expression of the minimum ℓ_2 -norm interpolant, followup work (Negrea et al., 2020; Koehler et al., 2021; Chinot and Lerasle, 2020; Chinot et al., 2022) employed tools from uniform convergence to analyze its excess risk. Prior work (Kobak et al., 2020; Wu and Xu, 2020; Tsigler and Bartlett, 2020) analyzed ridge regression with small or even negative regularization, and identified settings where using zero or even negative regularization can be optimal.

Techniques have also been developed to upper bound the excess risk of the sparsity-inducing minimum ℓ_1 -norm interpolant (Koehler et al., 2021; Li and Wei, 2021; Wang et al., 2022; Donhauser et al., 2022). Furthermore, lower bounds on the excess risk that show that sparsity can be incompatible with benign overfitting, and that the excess risk of sparse interpolators maybe exponentially larger than that of dense interpolators have also been derived (Chatterji and Long, 2022).

Kernel ridgeless regression has been actively studied (Liang and Rakhlin, 2020; Mei and Montanari, 2019). Careful theoretical analysis and simulations have revealed that kernel “ridgeless” regression can lead to multiple descent curves (Liang et al., 2020). A handful of papers have analyzed the risk of random features models (Mei and Montanari, 2019; Li et al., 2021b).

Several papers have also studied benign overfitting in linear classification of the canonical maximum ℓ_2 -margin classifier (Montanari et al., 2019; Deng et al., 2022; Chatterji and Long, 2021; Hsu et al., 2021; Muthukumar et al., 2021; Wang and Thrampoulidis, 2021; Cao et al., 2021), the maximum ℓ_1 -margin classifier (Liang and Sur, 2022), and classifiers obtained by minimizing polynomially-tailed classification losses (Wang et al., 2021b). Results have also been obtained on data that is linearly separable with two-layer leaky ReLU networks (Frei et al., 2022), and with two-layer convolutional networks with smooth nonlinearities (Cao et al., 2022).

Furthermore, this phenomenon has been studied in nearest neighbor models (Belkin et al., 2018), latent factor models (Bunea et al., 2022) and the Nadaraya-Watson estimator (Belkin et al., 2019b) with a singular kernel.

Shamir (2022) highlighted the importance of the choice of the loss function, since an interpolator that benignly overfits with respect one loss function may not with respect to another one.

A.2 Implicit Bias

In the closely related problem of matrix factorization, several papers (Gunasekar et al., 2017; Arora et al., 2019) established conditions under which the solution of gradient flow converges to the minimum nuclear norm solution. This rank minimization behavior of gradient flow was also shown to approximately hold for ReLU networks (Timor et al., 2023).

Papers have also studied the implicit bias of gradient descent for *diagonal* linear networks and found that it depends on the scale of the initialization and step-size (Woodworth et al., 2020; Yun et al., 2021; Nacson et al., 2022). They found that depending on the scale of the initialization and step-size the network converges to the minimum ℓ_2 -norm solution (kernel regime), or to the minimum ℓ_1 -norm solution (rich regime) or to a solution that interpolates between these two norms. Gunasekar et al. (2018b); Jagadeesan et al. (2022) studied linear convolutional networks and found that under this parameterization of a linear model, gradient descent implicitly minimizes norms of the Fourier transform of the predictor. Techniques have also been developed to study the implicit bias of mirror descent (Gunasekar et al., 2018a; Li et al., 2022).

For linear classifiers, minimizing exponentially-tailed losses including the logistic loss leads to the maximum ℓ_2 -margin classifier (Soudry et al., 2018; Ji and Telgarsky, 2019; Nacson et al., 2019). Ji and Telgarsky (2018) found that this is also the case when the linear classifier is parameterized as a deep linear classifier.

As a counterpoint to this line of research, Arora et al. (2019) and Razin and Cohen (2020) raised the possibility that the implicit bias of deep networks may be unexplainable by a simple function such as a norm. Li et al. (2021a) showed that under certain conditions gradient flow with infinitesimal initialization is equivalent to a simple heuristic rank minimization algorithm. Vardi and Shamir (2021) showed that it might be impossible altogether to capture the implicit bias of even two-layer ReLU networks using any functional form.

A.3 Optimization of Deep Linear Networks

Several papers have studied deep linear networks as a means to understand the benefit of overparameterization while optimizing nonlinear networks. Arora et al. (2018) argued that depth promotes a form of implicit acceleration when performing gradient descent on deep linear networks. While other papers (Du and Hu, 2019) showed that for wide enough networks that are randomly initialized by Gaussians the loss converges at a linear rate. Other papers have analyzed the convergence rate under other initialization schemes such as orthogonal initialization (Hu et al., 2020) and near-identity initialization (Bartlett et al., 2019; Zou et al., 2020). Rates of convergence for accelerated methods such as Polyak’s heavy ball method have also been established (Wang et al., 2021a).

Saxe et al. (2014) analyzed the effect of initialization and step-size in training deep linear networks. Kawaguchi (2016) identified a number of properties of the loss landscape of deep linear networks, including the absence of suboptimal local minima; Achour et al. (2021) characterized global minima, strict saddles, and non-strict saddles for these landscapes.

Appendix B. Proof of Proposition 18

In this appendix, we prove Proposition 18. Its proof uses some technical lemmas, which we derive first.

It is useful to recall the definition of a δ -good initialization from above.

Definition 17. *For a large enough absolute constant c , we say that the network enjoys a δ -good initialization if*

$$\begin{aligned}\alpha/c &< \sigma_{\min}(W_1^{(0)}) \leq \sigma_{\max}(W_1^{(0)}) < c\alpha, \\ \beta/c &< \sigma_{\min}(W_L^{(0)}) \leq \sigma_{\max}(W_L^{(0)}) < c\beta,\end{aligned}$$

and

$$\mathcal{L}(\Theta^{(0)}) < c \left(\|Y\|^2 + \frac{\alpha^2 \beta^2 q \|X\|^2 \log(n/\delta)}{m} \right).$$

The next two lemmas shall be useful in showing that our initialization scheme leads to a δ -good initialization. First, to bound the singular values of the weight matrices at initialization we apply the following result from Vershynin (2010).

Lemma 24. *There exists a constant c such that given any $\delta \in (0, 1)$ if $m \geq c(d + q + \log(1/\delta))$, then with probability at least $1 - \delta/2$*

$$\begin{aligned}\alpha/2 &< \sigma_{\min}(W_1^{(0)}) \leq \sigma_{\max}(W_1^{(0)}) < 2\alpha \quad \text{and} \\ \beta/2 &< \sigma_{\min}(W_L^{(0)}) \leq \sigma_{\max}(W_L^{(0)}) < 2\beta.\end{aligned}$$

Proof. By (Vershynin, 2010, Corollary 5.35) we have that with probability at least $1 - 4e^{-\eta^2/2}$

$$\begin{aligned}\frac{\alpha}{\sqrt{m}} (\sqrt{m} - \sqrt{d} - \eta) &\leq \sigma_{\min}(W_1^{(0)}) \leq \sigma_{\max}(W_1^{(0)}) \leq \frac{\alpha}{\sqrt{m}} (\sqrt{m} + \sqrt{d} + \eta) \quad \text{and} \\ \frac{\beta}{\sqrt{m}} (\sqrt{m} - \sqrt{q} - \eta) &\leq \sigma_{\min}(W_L^{(0)}) \leq \sigma_{\max}(W_L^{(0)}) \leq \frac{\beta}{\sqrt{m}} (\sqrt{m} + \sqrt{q} + \eta).\end{aligned}$$

So since $m \geq c(d + q + \log(1/\delta))$, where c is a large enough constant, by picking $\eta = \sqrt{m}/8$ we get that with probability at least $1 - \delta/2$

$$\begin{aligned}\frac{\alpha}{2} &< \sigma_{\min}(W_1^{(0)}) \leq \sigma_{\max}(W_1^{(0)}) < 2\alpha \\ \frac{\beta}{2} &< \sigma_{\min}(W_L^{(0)}) \leq \sigma_{\max}(W_L^{(0)}) < 2\beta,\end{aligned}$$

completing the proof. ■

The next lemma shows that the loss is controlled at initialization.

Lemma 25. *There is a positive constant c such that, for any $\delta \in (0, 1)$, provided that $m \geq c(d + q + \log(1/\delta))$, with probability at least $1 - \delta/2$*

$$\mathcal{L}(\Theta^{(0)}) < c \left(\|Y\|^2 + \frac{\alpha^2 \beta^2 q \|X\|^2 \log(n/\delta)}{m} \right).$$

Proof. The lemma directly follows by invoking (Zou et al., 2020, Proposition 3.3). \blacksquare

The next lemma shows that if the weights remain close to their initial values then the loss decreases at a certain rate.

Lemma 26. *If $\beta \geq 2$ and there was a good initialization (see Definition 17), at any time $t > 0$ if, for all $j \in [L]$*

$$\|W_j^{(t)} - W_j^{(0)}\|_{op} < \frac{1}{2L}$$

then

$$\frac{d\mathcal{L}(\Theta^{(t)})}{dt} < -\frac{\beta^2 \sigma_{\min}^2(X)}{4e} \mathcal{L}(\Theta^{(t)}).$$

Proof. By the chain rule, we have that

$$\frac{d\mathcal{L}(\Theta^{(t)})}{dt} = \nabla_{\Theta} \mathcal{L}(\Theta^{(t)}) \cdot \frac{d\Theta}{dt} = -\nabla_{\Theta} \mathcal{L}(\Theta^{(t)}) \cdot \nabla_{\Theta} \mathcal{L}(\Theta^{(t)}) = -\|\nabla_{\Theta} \mathcal{L}(\Theta^{(t)})\|^2.$$

Further, observe that

$$\begin{aligned} \|\nabla_{\Theta} \mathcal{L}(\Theta^{(t)})\|^2 &= \sum_{j=1}^L \|\nabla_{W_j} \mathcal{L}(\Theta^{(t)})\|^2 \\ &\geq \|\nabla_{W_1} \mathcal{L}(\Theta^{(t)})\|^2 \\ &= \|(W_L^{(t)} \dots W_2^{(t)})^\top (X\Theta - Y)^\top X\|^2. \end{aligned}$$

Continuing by applying (Zou et al., 2020, Lemma B.3) to the RHS of the inequality above we get that,

$$\begin{aligned} \|\nabla_{\Theta} \mathcal{L}(\Theta^{(t)})\|^2 &\geq \sigma_{\min}^2(X) \|X\Theta - Y\|^2 \left(\prod_{k \neq 1} \sigma_{\min}^2(W_k^{(t)}) \right) \\ &= \sigma_{\min}^2(X) \mathcal{L}(\Theta^{(t)}) \left(\prod_{k \neq 1} \sigma_{\min}^2(W_k^{(t)}) \right) \\ &\stackrel{(i)}{\geq} \sigma_{\min}^2(X) \mathcal{L}(\Theta^{(t)}) \left(\prod_{k \neq 1} \left(\sigma_{\min}(W_k^{(0)}) - \|W_k^{(t)} - W_k^{(0)}\|_{op} \right)^2 \right) \\ &\stackrel{(ii)}{>} \sigma_{\min}^2(X) \mathcal{L}(\Theta^{(t)}) \left(\frac{\beta}{2} - \frac{1}{2L} \right)^2 \left(1 - \frac{1}{2L} \right)^{2(L-2)} \\ &= \frac{\beta^2}{4} \sigma_{\min}^2(X) \mathcal{L}(\Theta^{(t)}) \left(1 - \frac{1}{\beta L} \right)^2 \left(1 - \frac{1}{2L} \right)^{2(L-2)} \\ &\geq \frac{\beta^2}{4} \sigma_{\min}^2(X) \mathcal{L}(\Theta^{(t)}) \left(1 - \frac{1}{2L} \right)^{2L-2} \\ &\geq \frac{\beta^2 \sigma_{\min}^2(X)}{4e} \mathcal{L}(\Theta^{(t)}), \end{aligned}$$

where (i) follows since on a good initialization for all $j \in [L]$, $\sigma_{\min}(W_j^{(0)}) \geq \min\{\beta/2, 1\} \geq 1$ and by assumption $\|W_j^{(t)} - W_j^{(0)}\|_{op} < 1/(2L)$. Inequality (ii) follows since $\beta \geq 2$, and there was a good initialization (which implies that the event in Lemma 24 occurs). ■

The next lemma shows that if the weight matrices remain close throughout the path of gradient flow then the loss decreases.

Lemma 27. *If $\beta \geq 2$ and there was a good initialization (see Definition 17), given any $t > 0$ if for all $0 \leq s < t$ and for all $j \in [L]$*

$$\|W_j^{(s)} - W_j^{(0)}\|_{op} < \frac{1}{2L}$$

then,

$$\mathcal{L}(\Theta^{(t)}) < \mathcal{L}(\Theta^{(0)}) \exp\left(-\frac{\beta^2 \sigma_{\min}^2(X)}{4e} \cdot t\right).$$

Proof. By assumption $\|W_j^{(s)} - W_j^{(0)}\|_{op} < 1/(2L)$ for all $0 \leq s < t$. Thus, by invoking Lemma 26 we know that for all $0 \leq s < t$,

$$\frac{d\mathcal{L}(\theta^{(s)})}{ds} < -\frac{\beta^2 \sigma_{\min}^2(X)}{4e} \mathcal{L}(\Theta^{(s)}).$$

which implies

$$\int_0^t \frac{1}{\mathcal{L}(\Theta^{(s)})} \frac{d\mathcal{L}(\theta^{(s)})}{ds} ds < -\int_0^t \frac{\beta^2 \sigma_{\min}^2(X)}{4e} ds.$$

Integrating both sides we get that

$$\mathcal{L}(\Theta^{(t)}) < \mathcal{L}(\Theta^{(0)}) \exp\left(-\frac{\beta^2 \sigma_{\min}^2(X)}{4e} \cdot t\right).$$

■

We also need the lemma that controls the growth of the operator norm of $W_j^{(t)}$.

Lemma 28. *There is a positive absolute constant c such that if*

- $\beta \geq c \max\left\{1, \sqrt{\frac{L\|X\|_{op}\|Y\|}{\sigma_{\min}^2(X)}}\right\}$ and
- $m \geq c \frac{L^2 \alpha^2 \|X\|_{op}^2 \|X\|^2 q \log(n/\delta)}{\beta^2 \sigma_{\min}^4(X)},$

then on a good initialization, given any $t > 0$ if for all $0 \leq s < t$

$$\|W_j^{(s)} - W_j^{(0)}\|_{op} < \frac{1}{2L}$$

then

$$\|W_j^{(t)} - W_j^{(0)}\|_{op} < \frac{1}{2L}.$$

Proof. Applying (Zou et al., 2020, Lemma A.1) with $A = B = I$ we have, for all $s \geq 0$,

$$\left\| \nabla_{W_j} \mathcal{L}(\Theta^{(s)}) \right\|^2 \leq 2e \|X\|_{op}^2 \mathcal{L}(\Theta^{(s)}).$$

By the definition of gradient flow,

$$W_j^{(t)} - W_j^{(0)} = - \int_0^t \nabla \mathcal{L}_{W_j}(\Theta^{(s)}) \, ds,$$

which implies that,

$$\begin{aligned} \|W_j^{(t)} - W_j^{(0)}\|_{op} &\leq \int_0^t \|\nabla \mathcal{L}_{W_j}(\Theta^{(s)})\|_{op} \, ds \\ &\leq \int_0^t \sqrt{2e} \|X\|_{op} \sqrt{\mathcal{L}(\Theta^{(s)})} \, ds \\ &\stackrel{(i)}{<} \sqrt{2e} \|X\|_{op} \int_0^t \sqrt{\mathcal{L}(\Theta^{(0)})} \exp\left(-\frac{\beta^2 \sigma_{\min}^2(X)}{4e} \cdot s\right) \, ds \\ &= \sqrt{2e} \|X\|_{op} \sqrt{\mathcal{L}(\Theta^{(0)})} \int_0^t \exp\left(-\frac{\beta^2 \sigma_{\min}^2(X)}{4e} \cdot s\right) \, ds \\ &= \frac{4e \sqrt{2e} \|X\|_{op} \sqrt{\mathcal{L}(\Theta^{(0)})}}{\beta^2 \sigma_{\min}^2(X)} \left[1 - \exp\left(-\frac{\beta^2 \sigma_{\min}^2(X)t}{4e}\right)\right] \\ &< \frac{30 \|X\|_{op} \sqrt{\mathcal{L}(\Theta^{(0)})}}{\beta^2 \sigma_{\min}^2(X)}, \end{aligned}$$

where (i) follows by applying Lemma 27.

On a good initialization we have that

$$\mathcal{L}(\Theta^{(0)}) < c_1 \left(\|Y\|^2 + \frac{\alpha^2 \beta^2 q \|X\|^2 \log(n/\delta)}{m} \right).$$

Plugging this into the previous inequality we find that

$$\begin{aligned} \|W_j^{(t)} - W_j^{(0)}\|_{op} &< \frac{30 \|X\|_{op}}{\beta^2 \sigma_{\min}^2(X)} \sqrt{c_1 \left(\|Y\|^2 + \frac{\alpha^2 \beta^2 q \|X\|^2 \log(n/\delta)}{m} \right)} \\ &\leq c_2 \left(\frac{\|X\|_{op} \|Y\|}{\beta^2 \sigma_{\min}^2(X)} + \frac{\alpha \|X\|_{op} \|X\|}{\beta \sigma_{\min}^2(X)} \sqrt{\frac{q \log(n/\delta)}{m}} \right). \end{aligned}$$

So for our lemma to be satisfied it suffices if

$$c_2 \frac{\alpha \|X\|_{op} \|X\|}{\beta \sigma_{\min}^2(X)} \sqrt{\frac{q \log(n/\delta)}{m}} < \frac{1}{4L} \Leftrightarrow m > c_3 \frac{L^2 \alpha^2 \|X\|_{op}^2 \|X\|^2 q \log(n/\delta)}{\beta^2 \sigma_{\min}^4(X)}$$

and

$$c_2 \frac{\|X\|_{op} \|Y\|}{\beta^2 \sigma_{\min}^2(X)} < \frac{1}{4L} \Leftrightarrow \beta > c_4 \sqrt{\frac{L \|X\|_{op} \|Y\|}{\sigma_{\min}^2(X)}}$$

■

Armed with these lemmas, we are now ready to prove the main result of this appendix. Recall its statement from above.

Proposition 18. *There is a constant c such that, given any $\delta \in (0, 1)$, if the initialization scales α and β , along with the network width m , satisfy*

$$\begin{aligned} m &\geq c \max \left\{ d + q + \log(1/\delta), \frac{L^2 \alpha^2 \|X\|_{op}^2 \|X\|^2 q \log(n/\delta)}{\beta^2 \sigma_{\min}^4(X)} \right\}, \\ \beta &\geq c \max \left\{ 1, \sqrt{\frac{L \|X\|_{op} \|Y\|}{\sigma_{\min}^2(X)}} \right\}, \\ \alpha &\leq 1, \end{aligned}$$

then with probability at least $1 - \delta$:

1. the initialization is δ -good;
2. training is perpetually $c(\alpha + 1/L)\beta$ bounded;
3. for all $t > 0$, we have that

$$\begin{aligned} \mathcal{L}(\Theta^{(t)}) &< \mathcal{L}(\Theta^{(0)}) \exp \left(-\frac{\beta^2 \sigma_{\min}^2(X)}{4e} \cdot t \right) \\ &\leq c \left(\|Y\|^2 + \frac{\alpha^2 \beta^2 q \|X\|^2 \log(n/\delta)}{m} \right) \exp \left(-\frac{\beta^2 \sigma_{\min}^2(X)}{4e} \cdot t \right). \end{aligned}$$

Proof. We will first prove that Part 1 of the proposition holds with probability $1 - \delta$ and then prove the other two parts assuming that the initialization was good.

Proof of Part 1. By invoking Lemmas 24 and 25 we have that a good initialization holds with probability at least $1 - \delta$.

Proof of Parts 2 and 3. Assume that the initialization was good. We claim that, for all $t > 0$,

- for all $j \in [L]$, $\|W_j^{(t)} - W_j^{(0)}\|_{op} < \frac{1}{2L}$ and
- $\mathcal{L}(\Theta^{(t)}) < \mathcal{L}(\Theta^{(0)}) \exp \left(-\frac{\beta^2 \sigma_{\min}^2(X)}{4e} \cdot t \right)$.

Assume for contradiction that this does not hold. Since $\|W_j^{(t)} - W_j^{(0)}\|_{op} - \frac{1}{2L}$ and $\mathcal{L}(\Theta^{(t)}) - \mathcal{L}(\Theta^{(0)}) \exp \left(-\frac{\beta^2 \sigma_{\min}^2(X)}{4e} \cdot t \right)$ are continuous functions of t , by the Intermediate Value Theorem there is a least value $t_0 > 0$ such that one of these quantities equals zero. But Lemmas 27 and 28 contradict this. This proves the first bound on the loss in Part 3. For the second bound on the loss, we note that on a good initialization

$$\mathcal{L}(\Theta^{(0)}) < c' \left(\|Y\|^2 + \frac{\alpha^2 \beta^2 q \|X\|^2 \log(n/\delta)}{m} \right).$$

Finally, we show that the training is perpetually $c(\alpha + 1/L)\beta$ bounded. Recall that this means that, for all $t \geq 0$, for all $S \subseteq [L]$,

$$\prod_{k \in S} \left\| W_k^{(t)} \right\|_{op} \leq c(\alpha + 1/L)\beta.$$

This follows since on a good initialization, $\|W_1^{(0)}\|_{op} \leq 2\alpha$, $\|W_L^{(0)}\|_{op} \leq 2\beta$ and for all $j \in \{2, \dots, L-1\}$, $\|W_j^{(0)}\|_{op} = 1$. Further, by Lemma 28 we know that for all $t \geq 0$, $\|W_j^{(t)} - W_j^{(0)}\|_{op} \leq 1/(2L)$. Putting these two facts together, proves that the process is perpetually $c(\alpha + 1/L)\beta$ bounded. ■

References

- El Mehdi Achour, François Malgouyres, and Sébastien Gerchinovitz. The loss landscape of deep linear neural networks: a second-order analysis. *arXiv preprint arXiv:2107.13289*, 2021.
- Sanjeev Arora, Nadav Cohen, and Elad Hazan. On the optimization of deep networks: implicit acceleration by overparameterization. In *International Conference on Machine Learning (ICML)*, 2018.
- Sanjeev Arora, Nadav Cohen, Wei Hu, and Yuping Luo. Implicit regularization in deep matrix factorization. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- Shahar Azulay, Edward Moroshko, Mor Shpigel Nacson, Blake Woodworth, Nathan Srebro, Amir Globerson, and Daniel Soudry. On the implicit bias of initialization shape: beyond infinitesimal mirror descent. In *International Conference on Machine Learning (ICML)*, 2021.
- Peter Bartlett and Philip Long. Failures of model-dependent generalization bounds for least-norm interpolation. *Journal of Machine Learning Research (JMLR)*, 2021.
- Peter Bartlett, David Helmbold, and Philip Long. Gradient descent with identity initialization efficiently learns positive-definite linear transformations by deep residual networks. *Neural Computation*, 2019.
- Peter Bartlett, Philip Long, Gábor Lugosi, and Alexander Tsigler. Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences (PNAS)*, 2020.
- Peter Bartlett, Andrea Montanari, and Alexander Rakhlin. Deep learning: a statistical viewpoint. *Acta Numerica*, 2021.
- Mikhail Belkin. Fit without fear: remarkable mathematical phenomena of deep learning through the prism of interpolation. *Acta Numerica*, 2021.

- Mikhail Belkin, Daniel Hsu, and Partha Mitra. Overfitting or perfect fitting? Risk bounds for classification and regression rules that interpolate. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences (PNAS)*, 2019a.
- Mikhail Belkin, Alexander Rakhlin, and Alexandre Tsybakov. Does data interpolation contradict statistical optimality? In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2019b.
- Florentina Bunea, Seth Strimas-Mackey, and Marten Wegkamp. Interpolation under latent factor regression models. *Journal of Machine Learning Research (JMLR)*, 2022.
- Yuan Cao, Quanquan Gu, and Mikhail Belkin. Risk bounds for over-parameterized maximum margin classification on sub-Gaussian mixtures. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- Yuan Cao, Zixiang Chen, Mikhail Belkin, and Quanquan Gu. Benign overfitting in two-layer convolutional neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- Niladri Chatterji and Philip Long. Finite-sample analysis of interpolating linear classifiers in the overparameterized regime. *Journal of Machine Learning Research (JMLR)*, 2021.
- Niladri Chatterji and Philip Long. Foolish crowds support benign overfitting. *Journal of Machine Learning Research (JMLR)*, 2022.
- Niladri Chatterji, Philip Long, and Peter Bartlett. The interplay between implicit bias and benign overfitting in two-layer linear networks. *Journal of Machine Learning Research (JMLR)*, 2022.
- Geoffrey Chinot and Matthieu Lerasle. On the robustness of the minimum ℓ_2 interpolator. *arXiv preprint arXiv:2003.05838*, 2020.
- Geoffrey Chinot, Matthias Löffler, and Sara van de Geer. On the robustness of minimum norm interpolators and regularized empirical risk minimizers. *Annals of Statistics*, 2022.
- Zeyu Deng, Abla Kammoun, and Christos Thrampoulidis. A model of double descent for high-dimensional binary linear classification. *Information and Inference*, 2022.
- Konstantin Donhauser, Nicolò Ruggeri, Stefan Stojanovic, and Fanny Yang. Fast rates for noisy interpolation require rethinking the effect of inductive bias. In *International Conference on Machine Learning (ICML)*, 2022.
- Simon Du and Wei Hu. Width provably matters in optimization for deep linear neural networks. In *International Conference on Machine Learning (ICML)*, 2019.

- Spencer Frei, Niladri Chatterji, and Peter Bartlett. Benign overfitting without linearity: neural network classifiers trained by gradient descent for noisy linear data. In *Conference on Learning Theory (COLT)*, 2022.
- Suriya Gunasekar, Blake Woodworth, Srinadh Bhojanapalli, Behnam Neyshabur, and Nathana Srebro. Implicit regularization in matrix factorization. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- Suriya Gunasekar, Jason Lee, Daniel Soudry, and Nathan Srebro. Characterizing implicit bias in terms of optimization geometry. In *International Conference on Machine Learning (ICML)*, 2018a.
- Suriya Gunasekar, Jason Lee, Daniel Soudry, and Nathan Srebro. Implicit bias of gradient descent on linear convolutional networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018b.
- Trevor Hastie, Andrea Montanari, Saharon Rosset, and Ryan Tibshirani. Surprises in high-dimensional ridgeless least squares interpolation. *Annals of Statistics*, 2022.
- Daniel Hsu, Vidya Muthukumar, and Ji Xu. On the proliferation of support vectors in high dimensions. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2021.
- Wei Hu, Lechao Xiao, and Jeffrey Pennington. Provable benefit of orthogonal initialization in optimizing deep linear networks. In *International Conference on Learning Representations (ICLR)*, 2020.
- Meena Jagadeesan, Ilya Razenshteyn, and Suriya Gunasekar. Inductive bias of multi-channel linear convolutional networks with bounded weight norm. In *Conference on Learning Theory (COLT)*, 2022.
- Ziwei Ji and Matus Telgarsky. Gradient descent aligns the layers of deep linear networks. In *International Conference on Learning Representations (ICLR)*, 2018.
- Ziwei Ji and Matus Telgarsky. The implicit bias of gradient descent on nonseparable data. In *Conference on Learning Theory (COLT)*, 2019.
- Kenji Kawaguchi. Deep learning without poor local minima. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.
- Dmitry Kobak, Jonathan Lomond, and Benoit Sanchez. The optimal ridge penalty for real-world high-dimensional data can be zero or negative due to the implicit ridge regularization. *Journal of Machine Learning Research (JMLR)*, 2020.
- Frederic Koehler, Lijia Zhou, Danica Sutherland, and Nathan Srebro. Uniform convergence of interpolators: Gaussian width, norm bounds and benign overfitting. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- Vladimir Koltchinskii and Karim Lounici. Concentration inequalities and moment bounds for sample covariance operators. *Bernoulli*, 2017.

- Yue Li and Yuting Wei. Minimum ℓ_1 -norm interpolators: precise asymptotics and multiple descent. *arXiv preprint arXiv:2110.09502*, 2021.
- Zhiyuan Li, Yuping Luo, and Kaifeng Lyu. Towards resolving the implicit bias of gradient descent for matrix factorization: greedy low-rank learning. In *International Conference on Learning Representations (ICLR)*, 2021a.
- Zhiyuan Li, Tianhao Wang, Jason Lee, and Sanjeev Arora. Implicit bias of gradient descent on reparametrized models: on equivalence to mirror descent. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- Zhu Li, Zhi-Hua Zhou, and Arthur Gretton. Towards an understanding of benign overfitting in neural networks. *arXiv preprint arXiv:2106.03212*, 2021b.
- Tengyuan Liang and Alexander Rakhlin. Just interpolate: Kernel “ridgeless” regression can generalize. *Annals of Statistics*, 2020.
- Tengyuan Liang and Pragya Sur. A precise high-dimensional asymptotic theory for boosting and min- ℓ_1 -norm interpolated classifiers. *Annals of Statistics*, 2022.
- 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 (COLT)*, 2020.
- Neil Mallinar, James Simon, Amirhesam Abedsoltan, Parthe Pandit, Mikhail Belkin, and Preetum Nakkiran. Benign, tempered, or catastrophic: a refined taxonomy of overfitting. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- Song Mei and Andrea Montanari. The generalization error of random features regression: precise asymptotics and the double descent curve. *Communications on Pure and Applied Mathematics*, 2019.
- Hancheng Min, Salma Tarmoun, René Vidal, and Enrique Mallada. On the explicit role of initialization on the convergence and implicit bias of overparametrized linear networks. In *International Conference on Machine Learning (ICML)*, 2021.
- Andrea Montanari, Feng Ruan, Youngtak Sohn, and Jun Yan. The generalization error of max-margin linear classifiers: high-dimensional asymptotics in the overparametrized regime. *arXiv preprint arXiv:1911.01544*, 2019.
- Vidya Muthukumar, Kailas Vodrahalli, Vignesh Subramanian, and Anant Sahai. Harmless interpolation of noisy data in regression. *IEEE Journal on Selected Areas in Information Theory*, 2020.
- Vidya Muthukumar, Adhyayan Narang, Vignesh Subramanian, Mikhail Belkin, Daniel Hsu, and Anant Sahai. Classification vs regression in overparameterized regimes: does the loss function matter? *Journal of Machine Learning Research (JMLR)*, 2021.
- Mor Shpigel Nacson, Nathan Srebro, and Daniel Soudry. Stochastic gradient descent on separable data: exact convergence with a fixed learning rate. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2019.

- Mor Shpigel Nacson, Kavya Ravichandran, Nathan Srebro, and Daniel Soudry. Implicit bias of the step size in linear diagonal neural networks. In *International Conference on Machine Learning (ICML)*, 2022.
- Jeffrey Negrea, Gintare Karolina Dziugaite, and Daniel Roy. In defense of uniform convergence: generalization via derandomization with an application to interpolating predictors. In *International Conference on Machine Learning (ICML)*, 2020.
- Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. In search of the real inductive bias: on the role of implicit regularization in deep learning. In *International Conference on Learning Representations (ICLR) Workshop*, 2015.
- Noam Razin and Nadav Cohen. Implicit regularization in deep learning may not be explainable by norms. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- Andrew Saxe, James McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. In *International Conference on Learning Representations (ICLR)*, 2014.
- Ohad Shamir. The implicit bias of benign overfitting. In *Conference on Learning Theory (COLT)*, 2022.
- Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The implicit bias of gradient descent on separable data. *Journal of Machine Learning Research (JMLR)*, 2018.
- Nadav Timor, Gal Vardi, and Ohad Shamir. Implicit regularization towards rank minimization in ReLU networks. In *International Conference on Algorithmic Learning Theory (ALT)*, 2023.
- Alexander Tsigler and Peter Bartlett. Benign overfitting in ridge regression. *arXiv preprint arXiv:2009.14286*, 2020.
- Gal Vardi and Ohad Shamir. Implicit regularization in ReLU networks with the square loss. In *Conference on Learning Theory (COLT)*, 2021.
- Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- Roman Vershynin. *High-dimensional probability: an introduction with applications in data science*. Cambridge University Press, 2018.
- Guillaume Wang, Konstantin Donhauser, and Fanny Yang. Tight bounds for minimum ℓ_1 -norm interpolation of noisy data. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2022.
- Jun-Kun Wang, Chi-Heng Lin, and Jacob Abernethy. A modular analysis of provable acceleration via Polyak’s momentum: training a wide ReLU network and a deep linear network. In *International Conference on Learning Representations (ICLR)*, 2021a.

- Ke Wang and Christos Thrampoulidis. Benign overfitting in binary classification of Gaussian mixtures. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021.
- Ke Alexander Wang, Niladri Chatterji, Saminul Haque, and Tatsunori Hashimoto. Is importance weighting incompatible with interpolating classifiers? In *International Conference on Learning Representations (ICLR)*, 2021b.
- Blake Woodworth, Suriya Gunasekar, Jason Lee, Edward Moroshko, Pedro Savarese, Itay Golan, Daniel Soudry, and Nathan Srebro. Kernel and rich regimes in overparametrized models. In *Conference on Learning Theory (COLT)*, 2020.
- Denny Wu and Ji Xu. On the optimal weighted ℓ_2 regularization in overparameterized linear regression. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- Chulhee Yun, Shankar Krishnan, and Hossein Mobahi. A unifying view on implicit bias in training linear neural networks. In *International Conference on Learning Representations (ICLR)*, 2021.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations (ICLR)*, 2017.
- Difan Zou, Philip Long, and Quanquan Gu. On the global convergence of training deep linear resnets. In *International Conference on Learning Representations (ICLR)*, 2020.