



# Singular Value Perturbation and Deep Network Optimization

Rudolf H. Riedi<sup>1</sup> · Randall Balestrieri<sup>2</sup> · Richard G. Baraniuk<sup>3</sup>

Received: 21 November 2021 / Revised: 19 July 2022 / Accepted: 8 October 2022 /

Published online: 27 November 2022

© The Author(s) 2022

## Abstract

We develop new theoretical results on matrix perturbation to shed light on the impact of architecture on the performance of a deep network. In particular, we explain analytically what deep learning practitioners have long observed empirically: the parameters of some deep architectures (e.g., residual networks, ResNets, and Dense networks, DenseNets) are easier to optimize than others (e.g., convolutional networks, ConvNets). Building on our earlier work connecting deep networks with continuous piecewise-affine splines, we develop an exact local linear representation of a deep network layer for a family of modern deep networks that includes ConvNets at one end of a spectrum and ResNets, DenseNets, and other networks with skip connections at the other. For regression and classification tasks that optimize the squared-error loss, we show that the optimization loss surface of a modern deep network is piecewise quadratic in the parameters, with local shape governed by the singular values of a matrix that is a function of the local linear representation. We develop new perturbation results for how the singular values of matrices of this sort behave as we add a fraction of the identity and multiply by certain diagonal matrices. A direct application of our perturbation results explains analytically why a network with skip connections (such as a ResNet or DenseNet) is easier to optimize than a ConvNet: thanks to its more stable singular values and smaller condition number, the local loss surface of such a network is less erratic, less eccentric, and features local minima that are more

---

We dedicate this paper to Prof. Ron DeVore to celebrate his 80th birthday and his lifetime of achievements in approximation theory.

---

Communicated by Albert Cohen.

---

✉ Richard G. Baraniuk  
richb@rice.edu

<sup>1</sup> Mathematics Department, HES-SO The University of Applied Sciences of Western Switzerland, Fribourg, Switzerland

<sup>2</sup> Meta AI, FAIR, New York City, USA

<sup>3</sup> Department of Electrical and Computer Engineering, Rice University, Houston, TX 77005, USA

accommodating to gradient-based optimization. Our results also shed new light on the impact of different nonlinear activation functions on a deep network's singular values, regardless of its architecture.

**Keywords** Deep learning · Optimization landscape · Perturbation theory · Singular values

**Mathematics Subject Classification** 68T07 · 47A55


## 1 Introduction

Deep learning has significantly advanced our ability to address a wide range of difficult inference and approximation problems. Today's machine learning landscape is dominated by deep (neural) networks (DNs) [19], which are compositions of a large number of simple parameterized linear and nonlinear operators, known as *layers*. An all-too-familiar story of late is that of plugging a DN into an application as a black box, learning its parameter values using gradient descent optimization with copious training data, and then significantly improving performance over classical task-specific approaches.

Over the past decade, a menagerie of DN architectures has emerged, including Convolutional Neural Networks (ConvNets) that feature affine convolution operations [18], Residual Networks (ResNets) that extend ConvNets with skip connections that jump over some layers [12], Dense Networks (DenseNets) with several parallel skip connections [15], and beyond. A natural question for the practitioner is: Which architecture should be preferred for a given task? Approximation capability does not offer a point of differentiation, because, as their size (number of parameters) grows, these and most other DN models attain universal approximation capability [6].

Practitioners know that DNs with skip connections, such as ResNets and DenseNets, are much preferred over ConvNets, because empirically their gradient descent learning converges faster and more stably to a better solution. In other words, it is not *what* a DN can approximate that matters, but rather *how it learns* to approximate. Empirical studies [21] have indicated that this is because the so-called *loss landscape* of the objective function navigated by gradient descent as it optimizes the DN parameters is much smoother for ResNets and DenseNets as compared to ConvNets (see Fig. 1). However, to date there has been no analytical work in this direction.

In this paper, we provide the first analytical characterization of the local properties of the DN loss landscape that enables us to quantitatively compare different DN architectures. The key is that, since the layers of all modern DNs (ConvNets, ResNets, and DenseNets included) are continuous piecewise-affine (CPA) spline mappings [1, 2], *the loss landscape is a continuous piecewise function of the DN parameters*. In particular, for regression and classification problems where the DN approximates a function by minimizing the  $\ell_2$ -norm squared-error, we show that the loss landscape is a continuous piecewise quadratic function of the DN parameters. The local eccentricity of this landscape and width of each local minimum basin is governed by the *singular values* of a matrix that is a function of not only the DN parameters but also the DN



**Fig. 1** Optimization loss landscape of two deep networks (DNs) along a 2D slice of their parameter space (from [21]). (left) Convolutional neural network (ConvNet) with no skip connections. (right) Residual network (ResNet) with skip connections. *This paper develops new matrix perturbation theory tools to understand these landscapes and in particular explain why skip connections produce landscapes that are less erratic, less eccentric, and feature local minima that are more accommodating to gradient-based optimization.* Used with permission

architecture. This enables us to quantitatively compare different DN architectures in terms of their singular values.

Let us introduce some notation to elaborate on the above programme. We study state-of-the-art DNs whose layers  $f_k$  comprise a nonlinear, continuous piecewise-linear *activation function*  $\phi$  that is applied element-wise to the output of an affine transformation. We focus on the ubiquitous class of activations  $\phi(t) = \max(\eta t, t)$ , which yields the so-called rectified linear unit (ReLU) for  $\eta = 0$ , leaky-ReLU for small  $\eta \geq 0$ , and absolute value for  $\eta = -1$  [11]. In an abuse of notation that is standard in the machine learning community,  $\phi$  can also operate on vector inputs by applying the above activation function to each coordinate of the vector input separately. Focusing for this introduction on ConvNets and ResNets (we deal with DenseNets in Sect. 3), denote the input to layer  $k$  by  $z_k$  and the output by  $z_{k+1}$ . Then we can write

$$z_{k+1} = f_k(z_k) := \phi_k(W_k z_k + b_k) + \rho z_k, \quad (1)$$

where  $W_k$  is an  $n \times n$  *weight matrix* and  $b_k$  is a vector of offsets.<sup>1</sup> Typically  $W_k$  is a (strided) convolution matrix. The choice  $\rho = 0$  corresponds to a ConvNet, while  $\rho = 1$  corresponds to a ResNet.

Previous work [1, 2] has demonstrated that the operator  $f_k$  is a collection of continuous piecewise-affine (CPA) splines that partition the layer's input space into polytopal regions and fit a different affine transformation on each region with the constraint that the overall mapping is continuous. This means that, locally around the input  $z_k$  and the DN parameters  $W_k, b_k$ , (1) can be written as

$$z_{k+1} = D_k W_k z_k + D_k b_k + \rho z_k, \quad (2)$$

<sup>1</sup> Without loss of generality, we assume that all  $W_k$  are square. Rectangular  $W_k$  are easily handled by extending with zeros.

where the diagonal matrix  $D_k$  contains 1s at the positions where the corresponding entries of  $W_k z_k + b_k$  are positive and  $\eta$  where they are negative.

In supervised learning, we are given a set of  $G$  labeled training data pairs  $\{x^{(g)}, y^{(g)}\}_{g=1}^G$ , and we tune the DN parameters  $W_k, b_k$  such that, when datum  $x^{(g)}$  is input to the DN, the output  $\hat{y}^{(g)}$  is close to the true label  $y^{(g)}$  as measured by some loss function  $L$ . In this paper, we focus on the  $\ell_2$ -norm squared-error loss averaged over a subset of the training data (called a *mini-batch*)

$$L := \frac{1}{G} \sum_{g=1}^G \|y^{(g)} - \hat{y}^{(g)}\|_2^2. \quad (3)$$

Standard practice is to use some flavor of gradient descent to iteratively reduce  $L$  by differentiating with respect to  $W_k, b_k$ . For ease of exposition, we focus our analysis first on the case of *fully stochastic gradient descent* that involves only a single data point per gradient step (i.e.,  $G = 1$ ), meaning that we iteratively minimize

$$L^{(g)} := \|y^{(g)} - \hat{y}^{(g)}\|_2^2 \quad (4)$$

for different choices of  $g$ . We then extend our theoretical results to arbitrary  $G > 1$  in Sect. 3.6. Let  $z_k^{(g)}$  denote the input to the  $k$ -th layer when the DN input is  $x^{(g)}$ . Using the CPA spline formulation (2), it is easy to show (see (46)) that, for fixed  $y^{(g)}$ , the loss function  $L^{(g)}$  is continuous and *piecewise-quadratic* in  $W_k, b_k$ .

The squared-error loss  $L^{(g)}$  is ubiquitous in *regression* tasks, where  $y^{(g)}$  is real-valued, but also relevant for classification tasks, where  $y^{(g)}$  is discrete-valued. Indeed, recent empirical work [16] has demonstrated that state-of-the-art DNs trained for *classification* tasks using the squared-error loss perform as well or better than the same networks trained using the more popular cross-entropy loss.<sup>2</sup>

We can characterize the local geometric properties of this piecewise-quadratic loss surface as follows. First, without loss of generality, we simplify our notation. For the rest of the paper, we label the layer of interest by  $k = 0$  and suppress the subscript 0 (i.e.,  $W_0 \rightarrow W$ ); we assume that there are  $p$  subsequent layers  $(f_1, \dots, f_p)$  between layer 0 and the output  $\hat{y}$ . Optimizing the weights  $W$  of layer  $f$  with training datum  $x^{(g)}$  that produces layer input  $z^{(g)}$  requires the analysis of the DN output  $\hat{y}$ , due to the chain rule calculation of the gradient of  $L$  with respect to  $W$ . (As we discuss below, there is no need to analyze the optimization of  $b$ .) For further simplicity, we suppress the superscript  $(g)$  that enumerates the training data whenever possible. We thus write the DN output as

$$\hat{y} = M(\rho)DWz + B \quad (5)$$

<sup>2</sup> We note in passing that one can use the techniques developed in this paper to show that the cross-entropy loss is continuous and piecewise-cross-entropic in  $W_k, b_k$ . Analysis of the local Hessian of this piecewise loss would lead to the same conclusion as with the squared-error loss, namely that the loss landscapes of ResNets/DenseNets are strictly better behaved than those of ConvNets. We leave the details of this analysis for future work.

where

$$M(\rho) = D_p W_p \cdots D_1 W_1 + \rho \text{Id} \quad (6)$$

collects the combined effect of applying  $W_k$  in subsequent layers  $f_k, k = 1, \dots, p$ , and  $B$  collects the combined offsets. Note that  $z$  reflects the influence of the training datum combined with the action of the layers preceding  $f$  and can be interpreted as the input to the shallower network consisting of layers  $f, f_1, \dots, f_p$ . This justifies using the index 0 for the layer under consideration.

Using this notation and fixing  $b$  as well as the input  $z$ , the piecewise-quadratic loss function  $L$  can be written locally as a linear term in  $W$  plus a quadratic term in  $W$ , which can be written as a quadratic form featuring matrix  $Q$  (see Lemma 8)

$$z^T W^T D^T M(\rho)^T M(\rho) D W z = w^T Q^T Q w. \quad (7)$$

Here,  $w$  denotes the columnized vector version of the matrix  $W$ .

The semi-axes of the ellipsoidal level sets of the local quadratic loss (7) are determined by the *singular values* of the matrix  $Q$ , which we can write as  $s_i \cdot |z[j]|$  according to Corollary 9, with  $s_i = s_i(MD)$  the  $i$ -th singular value of the linear mapping  $MD$  and  $z[j]$  the  $j$ -th entry in the vector  $z$ . The eccentricity of the ellipsoidal level sets is governed by the *condition number* of  $Q$  which, therefore, factors as

$$\kappa(Q) = \frac{\max_i s_i(MD)}{\min_i^* s_i(MD)} \cdot \frac{\max_j |z[j]|}{\min_j^* |z[j]|} = \kappa(MD) \cdot \kappa(\text{diag}(z)). \quad (8)$$

Here,  $\min^*$  denotes the minimum after discarding all vanishing elements (cf. (53)). In this paper we focus on the effect of the DN architecture on the condition number of the matrix  $MD$  representing the action of the subsequent layers, which is

$$\kappa(MD) = \frac{\max_i s_i(MD)}{\min_i^* s_i(MD)}. \quad (9)$$

However, the factorization in (8) provides insights into the role played by the training datum  $z$  (see Corollary 9). Extensions of these results to batches ( $G > 1$ ) are found in Sect. 3.6. (We note at this point that the optimization of the offset vector  $b$  has no effect on the shape of the loss function.)

For a fixed set of training data (i.e., fixed  $x$ ), the difference between the loss landscapes of a ConvNet and a ResNet is determined solely by  $\rho$  in (6). Therefore, we can make a fair, quantitative comparison between the loss landscapes of the ConvNet and ResNet architectures by studying how the singular values of the linear mapping  $M(\rho)D$  evolve as we move infinitesimally from a ConvNet ( $\rho = 0$ ) towards a ResNet ( $\rho = 1$ ).

Addressing how the singular values and condition number of the linear mapping  $M(\rho)D$  (and hence the DN optimization landscape) change when passing from  $\rho = 0$  towards  $\rho = 1$  requires a nontrivial extension of *matrix perturbation theory*. Our focus

and main contributions in this paper lie in exploring this under-explored territory.<sup>3</sup> Our analysis technique is new, and so we dedicate most of the paper to its development, which is potentially of independent interest.

We briefly summarize our key theoretical findings from Sects. 2–6 and our penultimate result in Theorem 22. We provide concrete bounds on the growth or decay of the singular values of any square matrix  $M_o = M(0)$  when it is perturbed to  $M(\rho) = M_o + \rho \text{Id}$  by adding a multiple of the identity  $\rho \text{Id}$  for  $0 \leq \rho \leq 1$ . The bounds are in terms of the largest and smallest eigenvalues of the symmetrized matrix  $M_o^T + M_o$ . The behavior of the largest singular value under this perturbation can be controlled quite accurately in the sense that random matrices tend to follow the upper bound we provide quite closely with high probability. For all other singular values, we require that they be non-degenerate, meaning that they have multiplicity 1 for all but finitely many  $\rho$ , a property that holds for most matrices in a generic sense (see Lemma 12). Further, we establish bounds for arbitrary singular values that do not hinge on the assumption of multiplicity 1 but that are somewhat less tight.

Our new perturbation results enable us to draw three rigorous conclusions regarding the optimization landscape of a DN. First, with regards to network architecture, we establish that the addition of skip connections improves the condition number of a DN's optimization landscape. To this end, we point out that *no* guarantees on the piecewise-quadratic optimization landscape's condition number are available for ConvNets; in particular,  $\kappa(M(0))$  can take any value, even if the entries of  $M(0)$  are bounded. However, as we move from ConvNets ( $\rho = 0$ ) towards networks with skip connections such as ResNets and DenseNets ( $\rho > 0$ ), bounds on  $\kappa(M(\rho))$  become available. For an appropriate random initialization of the weights and a wide network architecture (large  $n$ ), we show that the condition number of a ResNet ( $\rho = 1$ ) is bounded above asymptotically by just 3 (see (121), (122), and (128)). Second, for the particular case of DNs employing absolute value activation functions, we show that the relevant condition number is easier to control than for ReLU activations, that is, absolute value requires less restrictive assumptions on the entries of  $M$  in order to bound the condition number. For absolute value, we also prove that the condition number of  $M(1)$  (ResNet) is smaller than that of  $M(0)$  (ConvNet) with probability asymptotically equal to 1. Third, with regards to the influence of the training data, we prove that extreme values of the data negatively affect the performance of linear least squares and that this impact is tempered when applying mini-batch learning due to its inherent averaging.

In the context of DN optimization, these results demonstrate that the local loss landscape of a DN with skip connections (e.g., ResNet, DenseNet) is better conditioned than that of a ConvNet and thus less erratic, less eccentric, and with local minima that are more accommodating to gradient-based optimization, particularly when the weights in  $W$  are small. This is typically the case at the initialization of the optimization [10] and often remains true throughout training [9]. In particular, our results also provide new insights into the best magnitude to randomly initialize the weights at the beginning of learning in order to reap the maximum benefit from a skip connection architecture.

<sup>3</sup> A closely related situation occurs in Tikhonov regularization, where  $\text{Id}$  is added to  $M^T M$  in order to regularize the ill-posed problem  $Mx = b$ .

This paper is organized as follows. Section 2 introduces our notation and provides a review of existing useful results on the perturbation of singular values. Section 3 overviews our prior work establishing that CPA DNs can be viewed as locally linear and leverages this essential property to identify the singular values relevant for understanding the shape of the quadratic loss surface used in gradient-based optimization. Section 4 develops the perturbation results needed to capture the influence on the singular values when adding a skip connection, i.e., when passing from  $M_o$  to  $M_o + \text{Id}$ . In particular, we provide bounds on the condition number of  $M_o + \rho \text{Id}$  in terms of the largest singular value of  $M_o$  and the largest and the smallest eigenvalue of  $M_o^T + M_o$ . Section 5 combines the findings of Sects. 3 and 4 to show that a layer with a skip connection and weights bounded by an appropriate constant will have a bounded condition number. Clearly, this applies to weights drawn from a uniform distribution. Section 6 is dedicated to a ResNet layer with random weights of bounded standard deviation. Here, we establish bounds on the condition number of the corresponding ResNet and on the probability with which they hold, explicitly in terms of the width  $n$  of the network. Further, we provide convincing numerical evidence that the asymptotic bounds apply in practice for widths as modest as  $n \geq 5$ . We conclude in Sect. 7 with a synthesis of our results and perspectives on future research directions. All proofs are provided within the main paper. Various empirical experiments sprinkled throughout the paper support our theoretical results.

## 2 Background

### 2.1 Notation

**Singular Values** Let  $A$  be a  $n \times n$ -matrix, and let  $\|u\|$  denote the  $\ell_2$ -norm of the vector  $u$ . All products are to be understood as matrix multiplications, even if the factors are vectors. Vectors are always column vectors.

Let  $u_i$  denote a unit-eigenvector of  $A^T A$  to its *eigenvalue*  $\lambda_i = \lambda_i(A^T A)$  such that the vectors  $u_i$  ( $i = 1, \dots, n$ ) form an orthonormal basis.

The *singular values* of  $A$  are defined as the square roots of the eigenvalues of  $A^T A$

$$s_i = s_i(A) = \sqrt{\lambda_i(A^T A)} = \|Au_i\|. \quad (10)$$

Note that the singular values are always real and non-negative valued, since

$$0 \leq \|Au_i\|^2 = u_i^T A^T A u_i = \lambda_i u_i^T u_i = \lambda_i. \quad (11)$$

For reasons of definiteness, we take the eigenvalues of  $A^T A$ , and thus the singular values of  $A$  to be ordered by size:  $s_1 \geq \dots \geq s_n$ .

**Singular Vectors** The vectors  $u_i$  are called *right singular vectors* of  $A$ . The *left singular vectors* of  $A$  are denoted  $v_i$  and are defined as an orthonormal basis with the property that

$$Au_i = s_i v_i \quad \text{and} \quad A^T v_i = s_i u_i. \quad (12)$$

Such a basis of left-singular vectors always exists.

**SVD** Collecting the right singular vectors  $u_i$  as columns into an orthogonal matrix  $U$ , the left singular vectors  $v_i$  into an orthogonal matrix  $V$ , and the singular values  $s_i$  into a diagonal matrix  $\Sigma$  in corresponding order, we obtain the Singular Value Decomposition (SVD) of  $A = V\Sigma U^T$ .

**Eigenvectors** Note that  $A$  may have eigenvectors that differ from the singular vectors. In particular, eigenvectors are not always orthogonal onto each other. Clearly, if  $A$  is symmetric then the eigenvectors and right-singular vectors coincide and every *eigenvalue*  $r_i$  of  $A$  corresponds to a singular value  $s_i = |r_i|$ . However, to avoid confusion, we will refrain from referring to the eigenvalues and the singular values of the same matrix whenever possible.

**Smooth Matrix Deformations** Consider a family of  $n \times n$  matrices  $A(\rho)$ , where the variable  $\rho$  is assumed to lie in some interval  $I$ . We say that the family  $A(\rho)$  is  $C^m(I)$  if all its entries are  $C^m(I)$ , with  $m = 0$  corresponding to being continuous and  $m = 1$  to continuously differentiable. If  $A$  is  $C^m(I)$ , then so is  $A^T A$ .

We denote the matrix of derivatives of the entries of  $A$  by  $A'(\rho) = \frac{d}{d\rho} A(\rho)$ .

**Multiplicity of Singular Values** Denote by  $S_i \subset I$  the set of  $\rho$  values for which  $s_i$  is a *simple* singular value:

$$S_i = \{\rho \in I : s_i(\rho) \neq s_j(\rho) \text{ for all } j \neq i\}. \quad (13)$$

For later use, let

$$S_i^* = \{\rho \in S_i : s_i(\rho) \neq 0\}. \quad (14)$$

For completeness, we mention an obvious fact that follows from the SVD:

$$s_i(A) = s_i(A^T). \quad (15)$$

## 2.2 Singular Values Under Additive Perturbation

A classical result that will prove to be key in our study is due to Wielandt and Hoffman [14]. It uses the notion of the Frobenius norm  $\|A\|_F$  of the  $n \times n$ -matrix  $A$ . Denoting the entries of  $A$  by  $A[i, j]$  and its singular values by  $s_i(A)$ , it is well known that

$$\|A\|_F^2 = \sum_{i,j=1}^n |A[i, j]|^2 = \text{trace}(A^T A) = \sum_{i=1}^n s_i^2(A). \quad (16)$$

**Proposition 1** (Wielandt and Hoffman) *Let  $A$  and  $\tilde{A}$  be symmetric  $n \times n$ -matrices with eigenvalues ordered by size. Then*

$$\sum_{i=1}^n \left| \lambda_i(A) - \lambda_i(\tilde{A}) \right|^2 \leq \|A - \tilde{A}\|_F^2. \quad (17)$$

This result can easily be strengthened to read as follows.



**Corollary 2** (Wielandt and Hoffman) *Let  $A$  and  $\tilde{A}$  be any  $n \times n$ -matrices with singular values ordered by size. Then*

$$\sum_{i=1}^n \left| s_i(A) - s_i(\tilde{A}) \right|^2 \leq \|A - \tilde{A}\|_F^2. \quad (18)$$

Note that this result, as most others, holds actually also for rectangular matrices with the usual adjustments.

**Proof** We establish Corollary 2. The  $2n$  eigenvalues of the symmetric matrix

$$H_A = \begin{pmatrix} 0 & A \\ A^T & 0 \end{pmatrix} \quad (19)$$

are  $s_i(A)$  with eigenvector  $\frac{1}{\sqrt{2}} \begin{pmatrix} v_i \\ u_i \end{pmatrix}$  and  $-s_i(A)$  with eigenvector  $\frac{1}{\sqrt{2}} \begin{pmatrix} -v_i \\ u_i \end{pmatrix}$ . Also,  $\|H_A\|_F^2 = 2\|A\|_F^2$ . Therefore, when applying Proposition 1 with  $A$  replaced by  $H_A$  and  $\tilde{A}$  replaced by  $H_{\tilde{A}}$ , we run through the left-hand side of (18) twice and in turn obtain twice the right-hand side.

Corollary 2 immediately implies a well-known result that will be key: the singular values of a matrix depend continuously on its entries, as we state next.

**Lemma 3** *If  $A(\rho)$  is  $C^0(I)$ , then all its singular values  $s_i(\rho)$  depend continuously on  $\rho$ . Consequently, all  $S_i$  in (13) and  $S_i^*$  in (14) are then open.*

**Proof** To establish continuity, choose  $\rho_1$  and  $\rho_2$  arbitrary and apply Corollary 2 (in slight abuse of notation) with  $A = A(\rho_1)$  and  $\tilde{A} = A(\rho_2)$ .

As will become clear in the sequel, the main difficulty in establishing further results on the regularities of the singular values with elementary arguments and estimates lies in dealing with multiple zeros, i.e., with multiple singular values.

Also useful are some facts about the special role of the largest singular value.

**Proposition 4** (Operator norm) *Let  $A$  be a  $n \times n$ -matrix. Then, its largest singular value is equal to its operator norm, and we have the following inequalities:*

$$s_1(A) = \|A\|_{\text{Op}} := \max_{|x|=1} |Ax| \leq \|A\|_F \leq \sqrt{\text{rank}(A)} \cdot \|A\|_{\text{Op}}. \quad (20)$$

*In particular, if  $A$  is rank 1, then both norms coincide. Also, being a norm, the largest singular value satisfies*

$$s_1(A + \tilde{A}) \leq s_1(A) + s_1(\tilde{A}) \quad (21)$$

*for any matrices  $A$  and  $\tilde{A}$ .*

**Proof** The first equality follows from the fact that  $s_1$  is the largest semi-axis of the ellipsoid that is the image of the unit-sphere. The first inequality follows from (16). The last inequality follows again from considering the aforementioned ellipsoid.

Another proof of the continuity of singular values is found in the well-known interlacing properties of singular values, which we adapt to our notation and summarize for the convenience of the reader (see, e.g. [5, Theorem 6.1.3-5.]).

**Proposition 5** (Interlacing) *Let  $A$  and  $\tilde{A}$  be  $n \times n$ -matrices with singular values ordered by size.*

- (Weyl additive perturbation) *We have:*

$$s_{i+j-1}(A) \leq s_i(\tilde{A}) + s_j(A - \tilde{A}). \quad (22)$$

- (Cauchy interlacing by deletion) *Let  $\tilde{A}$  be obtained from  $A$  by deleting  $m$  rows or  $m$  columns. Then*

$$s_i(A) \geq s_i(\tilde{A}) \geq s_{i+m}(A). \quad (23)$$

From (22) applied with  $j = 1$ , we note that the singular values are in fact uniformly continuous. Indeed, we have for any  $i$  that

$$|s_i(A) - s_i(\tilde{A})| \leq s_1(A - \tilde{A}) \leq \|A - \tilde{A}\|_F. \quad (24)$$

### 2.3 Singular Values Under Smooth Perturbation

We now turn to the study of how singular values of  $A(\rho)$  change as a function of  $\rho$ , thereby leveraging the concept of differentiability.

This can be done in several ways; the following computation is particularly simple and well known. Let  $N(\rho)$  denote a family of symmetric  $\mathcal{C}^1(I)$  matrices. Dropping all indices and variables  $\rho$  for ease of reading, the eigenvalues and eigenvectors satisfy

$$Nu = \lambda u, \quad u^T u = 1. \quad (25)$$

Denoting the derivatives with respect to  $\rho$  by  $N'$ ,  $\lambda'$  and  $u'$  and assuming they exist, we obtain

$$N'u + Nu' = \lambda'u + \lambda u', \quad u'^T u = u^T u' = 0. \quad (26)$$

We then left-multiply the first equality by  $u^T$  and use the second equality to find

$$u^T N'u + u^T Nu' = \lambda' u^T u + \lambda u^T u' = \lambda'. \quad (27)$$

Finally, as  $N$  is symmetric we see that  $u^T Nu' = (Nu)^T u' = \lambda u^T u' = 0$ , leading to

$$\lambda' = u^T N'u. \quad (28)$$

In quantum mechanics, the last equation is known as the Hellmann–Feynman Theorem and was discovered independently by several authors in the 1930s. It is known, though seldom mentioned, that the formula may fail when eigenvalues coincide. For

further studies in this direction, we refer the interested reader to works in the field of degenerate perturbation theory.

For a proof of the Hellmann–Feynman Theorem, it is natural to consider the function  $\Psi : \mathbb{R}^{n+2} \rightarrow \mathbb{R}^{n+1}$

$$\Psi(\rho, u, \lambda) = \begin{pmatrix} Nu - \lambda u \\ u^T u - 1 \end{pmatrix}, \quad (29)$$

where  $u$  is an  $n$ -dimensional column vector.

The Implicit Function Theorem applied to  $\Psi$  then yields the following result.

**Theorem 6** (Hellmann–Feynman) *Assume  $N(\rho)$  forms a family of symmetric  $\mathcal{C}^1(I)$  matrices. Then, as a function of  $\rho$ , each of its eigenvalues  $\lambda_i = \lambda_i(\rho)$  is continuously differentiable where it is simple. Its derivative reads as*

$$\lambda'_i(\rho) = u_i^T N' u_i. \quad (30)$$

*Its corresponding right singular vector  $u_i(\rho)$  can be chosen such that it becomes  $\mathcal{C}^1(S_i)$  as well.*

If in addition to the assumptions of Theorem 6 the family  $N$  is actually  $\mathcal{C}^m(I)$ , then  $\lambda_i(\rho)$  and  $u_i(\rho)$  are  $\mathcal{C}^m(S_i)$ . From Hellmann–Feynmann we may conclude the following.

**Corollary 7** *Let  $A(\rho)$  be  $\mathcal{C}^1(I)$ . Assume that  $S_i^*$  consists of isolated points only. Then,  $s_i(A)$  and its left singular vector  $v_i$  are  $\mathcal{C}^1(S_i^*)$  with*

$$s'_i(\rho) = u_i(\rho)^T A'(\rho)^T v_i(\rho). \quad (31)$$


Simple examples show that the above results cannot be improved without stronger assumptions (see Example 1 below and Fig. 2).

**Proof** Apply Theorem 6 to the family  $N = A^T A$ . Note that  $N' = A'^T A + A^T A'$  and apply the chain rule to  $\lambda_i = (s_i(\rho))^2$  to find the first equality of the following formula:

$$s'_i = \frac{1}{2s_i} u_i^T (A'^T A + A^T A') u_i = \frac{1}{2s_i} 2u_i^T A'^T A u_i = u_i^T A'^T v_i. \quad (32)$$

Also, we have  $u_i^T A'^T A u_i = (u_i^T A'^T A u_i)^T = u_i^T A^T A' u_i$ , since this is a real-valued number. This implies the second equality in the formula above. Finally note that  $A u_i = s_i v_i$  to complete the proof.

*Example 1.* The diagonal family  $M = \text{diag}(-4 \quad -2) + \rho \text{Id}$  is  $\mathcal{C}^\infty(\mathbb{R})$ . When ordered by size, the eigenvalues of  $N = M^T M$  are  $\max((\rho - 4)^2, (\rho - 2)^2)$  and  $\min((\rho - 4)^2, (\rho - 2)^2)$ . They are not differentiable where they coincide. In particular, the matrix norm of  $N$  is not differentiable at  $\rho = 3$ . The singular values of  $M$  are  $s_1 = \max(|\rho - 4|, |\rho - 2|)$  and  $s_2 = \min(|\rho - 4|, |\rho - 2|)$ ; they are not differentiable at their respective zeros. Also, the left singular vectors  $v_i$  change sign at the zero of  $s_i$  and so are not even continuous there (see Fig. 2 for a similar example).



**Fig. 2** Singular values of two families of matrices of the form  $M(\rho) = M_0 + \rho \text{Id}$  for  $0 \leq \rho \leq 1$ . **a**  $M_0 = \text{diag}(-0.4 \ -0.2 \ -0.1)$  demonstrates that even the largest singular value may not be increasing or everywhere differentiable. **b** The entries of the  $15 \times 15$  matrix  $M_0$  are drawn from a zero-mean Gaussian distribution with  $\sigma = 0.3$ . This example demonstrates that the singular values may be non-convex and may have several extrema even without coinciding with others

### 3 Singular Values in Deep Learning


In this section, we establish that the squared-error loss landscape of a deep network (DN) is governed locally by the singular values of the matrix of weights of a layer and that the difference between certain network architectures can be interpreted as a perturbation of this matrix. Combining this key insight with the perturbation theory of singular values for DNs, which will be developed in the sections to follow, will enable us to characterize the dependence of the loss landscape on the underlying architecture. Doing so, we will ultimately provide analytical insight into numerous phenomena that have been observed only empirically so far, such as the fact that so called ResNets [13] are easier to optimize than ConvNets [18].

#### 3.1 The Continuous Piecewise-Affine Structure of Deep Networks

The most popular DNs contain activation nonlinearities that are continuous piecewise-affine (CPA), such as the ReLU, leaky-ReLU, absolute value, and max-pooling nonlinearities [11]. To be more precise, let  $f_k$  denote a layer of a DN. We will assume that it can be written as an affine map  $W_k z_k + b_k$  followed by the activation  $\phi_k$

$$z_{k+1} = f_k(z_k) = \phi_k(W_k z_k + b_k). \quad (33)$$

The entries of the matrix  $W_k$  are called *weights*; the constant additive term  $b_k$  is called the *bias*. More general settings can easily be employed; however, the above will be sufficient for our purpose. The above-mentioned, commonly used nonlinearities can be written in this form using



**Fig. 3** Iterative partitioning of input space by three consecutive layers of a toy deep network (DN). The first layer maps  $\mathbb{R}^2$  to  $\mathbb{R}^3$ , producing three cuts and seven regions, with one lying outside the visualized region

$$\phi_k(t) = \max(t, \eta t), \quad (34)$$

where the max is taken coordinate by coordinate. Clearly, any activation of the form (34) is continuous and piecewise linear, rendering the corresponding layer CPA (See Fig. 3). The following choices for  $\eta$  correspond to ReLU ( $\eta = 0$ ) [10], leaky-ReLU ( $0 < \eta < 1$ ) [25] and absolute value [4] ( $\eta = -1$ ) (see also [1, 2]). In summary:

$$\begin{cases} \eta > 0 & \text{leaky-ReLU} \\ \eta = 0 & \text{ReLU} \\ \eta = -1 & \text{absolute value} \end{cases} \quad (35)$$

For example, in two dimensions we obtain  $\phi_k^{(\text{Abs})}(t_1, t_2) = (|t_1|, |t_2|)$ . We will consider only nonlinearities that can be written in the form (34). Such nonlinearities are in fact linear in each hyper-octant of space, and can be written as

$$\phi_k(t) = D_k t,$$

where  $D_k$  is a diagonal matrix with entries that depend on the signs of the coordinates of the input  $t = W_k z_k + b_k$ . The diagonal entries of  $D_k$  equal either 1 or  $\eta$  depending on the hyper-octant.

Consequently, the input space of the layer  $f_k$  is divided into *regions* that depend on  $W_k$  and  $b_k$  in which  $f_k$  is affine and can be rewritten as

$$z_{k+1} = f_k(z_k) = \phi_k(W_k z_k + b_k) = (D_k W_k) z_k + D_k b_k. \quad (36)$$

By the linear nature of  $z_k \mapsto W_k z_k$ , these regions of the input space ( $z_k$ -space) are piecewise linear. Concatenating several layers into a multi-layer  $F(z_1) = f_p \circ \dots \circ f_1(z_1)$  then leads to a subdivision of its input space into piecewise-linear regions, in which  $F$  is affine, a result due to [1, 2]. In each linear region, this multi-layer can be written as

$$z_{p+1} = D_p W_p \dots D_1 W_1 z_1 + B_p, \quad (37)$$

where all constant terms are collected in

$$B_p = D_p W_p \dots D_2 W_2 b_1 + D_p W_p \dots D_3 W_3 b_2 + \dots + b_p. \quad (38)$$

A multi-layer can be considered to be an entire DN or a part of one.

### 3.2 From ConvNets to ResNets and DenseNets, and Learning

A DN with an architecture  $F(z_1) = f_p \circ \dots \circ f_1(z_1)$  as described above is called a *Convolutional Network* (ConvNet) [18] whenever the  $W_k$  are circulant-block-circulant; otherwise it is called a multi-layer-perceptron. Without loss of generality, we will focus on ConvNets and their extensions, particularly ResNets and DenseNets, in this paper. While such architectures are responsible for the recent cresting wave of deep learning successes, they are known to be intricate to effectively employ when the number of layers  $p$  large. This motivated research on alternative architectures.

When adding the input  $z_k$  to the output  $\phi_k(W_k z_k + b_k)$  of a convolutional layer (a skip connection), one obtains a *residual layer* which can be written in each region of linearity as

$$z_{k+1} = f_k(z_k) = \phi_k(W_k z_k + b_k) + z_k = (D_k W_k + \text{Id})z_k + D_k b_k. \quad (39)$$

A DN in which all (or most) layers are residual is called a *Residual Network* (ResNet) [13]. *Densely Connected Networks* (DenseNets) [15] allow the residual connection to bypass not just a single but multiple layers

$$z_{p+1} = (D_p W_p \dots D_1 W_1 + \text{Id})z_1 + B_p. \quad (40)$$

All of the above DN architectures are universal approximators. Learning the per-layer weights of any DN architecture is performed by solving an optimization problem that involves a training data set, a differentiable loss function such as the squared-error or cross-entropy, and a weight update policy such as gradient descent. Hence, the key benefit of one architecture over another only emerges when considering the learning problem, and in particular gradient-based optimization.

### 3.3 Learning

We now consider the learning process for one particular layer of a DN. As we make clear below, for the comparison of different network architectures, it is sufficient to study each layer independently of the others. For ease of exposition, our analysis in Sects. 3.3–3.5 also focuses on a single training data point, which corresponds to learning via fully stochastic gradient descent ( $G = 1$  in the Introduction). We extend our analysis to the general case of mini-batch gradient descent ( $G > 1$ ) in Sect. 3.6.

The layer being learned is  $f_0$ ; however, for clarity we drop the index 0 and use the notation

$$z_1 = f(z) = \phi(Wz + b) = DWz + b. \quad (\text{trained layer}) \quad (41)$$

The layers following  $f$  are the only ones relevant in the study of learning layer  $f$ . In fact, the previous layers transform the DN input into a feature map that, for the layer in question, can be seen as a fixed input. This part of the network will be denoted by

$F = f_p \circ \cdots \circ f_1$ . In each partition region, it takes the form

$$\hat{y} = MDWz + B. \quad (\text{trained plus subsequent layers}) \quad (42)$$

Here,  $B$  collects all of the constants,  $D$  represents the nonlinearity of the trained layer in the current region, and  $M$  collect the action of the subsequent layers  $F$ . We use the parameter  $\rho$  to distinguish between ConvNets and ResNets, with  $\rho = 0$  corresponding to convolution layers and  $\rho = 1$  to residual ones. The matrix  $M = M(\rho)$  representing the subsequent layers  $F$  becomes

$$M = M(\rho) = D_1 W_1 + \rho \text{Id} \quad (\text{one layer}) \quad (43)$$

$$M = M(\rho) = D_p W_p \cdots D_1 W_1 + \rho \text{Id}. \quad (\text{multi-layer}) \quad (44)$$

As a shorthand, we will often write  $M(\rho) = M_o + \rho \text{Id}$ , where

$$M_o = M(0) = D_p W_p \cdots D_1 W_1. \quad (45)$$

When saying “let  $M_o$  be arbitrary” we mean that no restrictions on the factors in (45) are imposed except for the  $D_i$  to be diagonal and the  $W_i$  to be square. When choosing  $D_p = \cdots = D_1 = W_p = \cdots = W_2 = \text{Id}$  and  $W_1$  arbitrary, we have  $M_o = W_1$  arbitrary. So, there is no ambiguity in our usage of the term “arbitrary”.

### 3.4 Quadratic Loss

We wish to quantify the advantage in training of a ResNet architecture over a ConvNet one. We consider a quadratic loss function such as the squared-error as mentioned in the Introduction (cf. (4); we drop the superscript  $(g)$ ):

$$L^{(g)} := \|y - \hat{y}\|_2^2 = (y - \hat{y})^T (y - \hat{y}) = \|y\|_2^2 - 2y^T \hat{y} + \|\hat{y}\|_2^2. \quad (46)$$

In the above, we used that  $y^T \hat{y} = (y^T \hat{y})^T = \hat{y}^T y$ , since this is a real number.

Recalling that the DN is a continuous piecewise-affine operator and fixing all parameters except for  $W$ , we can locally write  $\hat{y} = MDWz + B$  for all  $W$  in some region  $\mathcal{R}$  of  $F$  (cf. (4);  $z$  is the input to  $f$  produced by the datum  $x^{(g)}$ ). Writing  $\|\hat{y}\|_2^2 = \hat{y}^T \hat{y}$ , we find that the loss can be viewed as an affine term  $P(W)$  with  $P$  a polynomial of degree 1 plus a quadratic term  $\|MDWz\|_2^2$

$$L^{(g)} = P(W) + \|MDWz\|_2^2 \quad (W \in \mathcal{R}). \quad (47)$$

We stress once more that all parameters are fixed except for  $W$ . Denote by  $W_o = W_o(\rho, \mathcal{R})$  the minimum of the function  $h(W) = P(W) + \|MDWz\|_2^2$  ( $W$  arbitrary), which may very well lie outside  $\mathcal{R}$ . Writing  $\Delta W = W - W_o$  and developing  $h$  into its Taylor series at  $W_o$ , we find  $h(W) = h(W_o) + \|MD\Delta Wz\|_2^2$ , since all first order derivatives of  $h$  vanish at  $W_o$ .

It is worthwhile noting that the region  $\mathcal{R}$  does not change when changing the value of  $\rho$  in  $M = M(\rho)$ , since  $D\Delta Wz$  does not depend on  $\rho$ . From this, writing  $c = h(W_o)$  for short, we obtain for any  $\rho$

$$L^{(g)} = c + \|M(\rho)D\Delta Wz\|_2^2 \quad (W \in \mathcal{R}). \quad (48)$$

This means that the loss surface in each region  $\mathcal{R}$  of  $F$  is a part of an exact ellipsoid whose shape is determined by the second-order expression  $\|M(\rho)D\Delta Wz\|_2^2$  as a function of  $\Delta W$ . The value  $c$  should not be interpreted as the minimum of the loss function, since  $W_o$  may lie outside the region  $\mathcal{R}$ .

Before commenting further on the loss, however, let us untangle the influence of the data  $z$  (which remains unchanged during training) from the weights  $W$  in (47) (analogous for  $\Delta W$  in (48)). When listing the variables  $W[i, j]$  row by row in a “flattened” vector  $w$ , the term (47) becomes a quadratic form defined by the following matrix  $Q$ .

**Lemma 8** *In the quadratic term (47) of the MSE, we have the following matrices:  $W$  is  $n \times n$ ,  $MD$  is  $n \times n$ , and  $z$  is  $n \times 1$ . Denote the flattened version of  $W$  by  $w$  ( $n^2 \times 1$ ) and the singular values of  $MD$  by  $s_i = s_i(MD)$ . Let  $Q = \text{diag}(z) \otimes (MD)$  meaning that*

$$Q = \begin{pmatrix} z[1]MD & 0 & \dots & 0 \\ 0 & z[2]MD & \dots & 0 \\ & & \ddots & \\ 0 & \dots & 0 & z[n]MD \end{pmatrix}. \quad (49)$$

Then, we have (cf. (47))

$$\|MDWz\|_2^2 = w^T Q^T Q w, \quad (50)$$

and that the singular values of  $Q$  factor as

$$s_{i,j}(Q) = s_i(MD) \cdot |z[j]|, \quad (51)$$

with  $1 \leq i, j \leq n$ .

**Proof** Direct computation easily verifies (50). For (51), note that the characteristic function of  $A = Q^T Q$  is as follows, where in slight abuse of notation  $\text{Id}$  denotes the identity in the appropriate dimensions

$$\det(A - \lambda \text{Id}) = \prod_{j=1}^n \det\left(\left\{z[j]^2(MD)^T(MD)\right\} - \lambda \text{Id}\right). \quad (52)$$

Its zeros are easily identified as  $\lambda_{i,j} = z[j]^2 \lambda_i((MD)^T(MD)) = z[j]^2 s_i^2(MD)$ , establishing (51).



### 3.5 Condition Number

We return now to the shape of the loss surface of the trained layer  $f$ . We point out that the loss surface is exactly equal to a portion of an ellipsoid whose geometry is completely characterized by the spectrum of the matrix  $Q$  due to (48) and Lemma 8. Its curvature and higher-dimensional eccentricity has a considerable influence on the numerical behavior of gradient-based optimization. In essence, the larger the eccentricity, the less accurately the gradient points towards the maximal change of the loss when making a non-infinitesimal change in  $W$ .

The condition number  $\kappa$  of the matrix  $Q$  provides a measure of its higher-dimensional eccentricity and, therefore, of the difficulty of optimization by linear least squares. It is defined as the ratio of the largest singular value  $s_1$  to the smallest non-zero singular value. Denoting the latter by  $s_*$  for convenience, the condition number of an arbitrary matrix  $A$  is given by

$$\kappa(A) := \frac{s_1(A)}{\min\{s_i(A) : s_i(A) > 0\}} = \frac{s_1(A)}{\min_i^* s_i(A)} = \frac{s_1(A)}{s_*(A)}. \quad (53)$$

Having identified  $\kappa(Q)$  as a most relevant quantity towards understanding the performance of linear least squares, we can apply Lemma 8 to identify and separate the influence of data and subsequent layers on the loss landscape as follows. For an efficient notation we recall that the singular values of  $\text{diag}(z)$  are  $|z[j]|$  and that vanishing  $z[j]$  do not contribute to  $\kappa(\text{diag}(z))$  (cf. (53)).

**Corollary 9** (The Data Factor) *The condition number of the matrix  $Q$  of Lemma 8 factors as*

$$\kappa(Q) = \kappa(MD) \cdot \kappa(\text{diag}(z)). \quad (54)$$

Our main interest lies in the impact of the architecture of the  $p$  layers following  $f$  (adding a residual link or not), leaving the rest of the DN as is. Since the input data  $z$  does not depend on this architectural choice, the performance of the linear least-squares optimization we focus on here is governed by  $\kappa(MD) = \kappa(M(\rho)D)$ .

Consequently, we are faced with two types of perturbation settings: First, the evolution of the loss landscape of DNs as their architecture is continuously interpolated from a ConvNet ( $\rho = 0$ ) to a ResNet ( $\rho = 1$ ), that is, when moving from  $M(0)$  to  $M(1)$ . In this case, since the region  $\mathcal{R}$  of  $F$  does not depend on  $\rho$ , such an analysis is meaningful for understanding the optimization of the loss via (48). Second, the effect of multiplying  $M$  by a diagonal matrix  $D$ . Before studying the perturbations in the subsequent sections, we make a remark regarding the use of mini-batches.

### 3.6 Learning Using Mini-Batches

In deep learning, one works typically with mini-batches of data  $x^{(g)}$  ( $g = 1, \dots, G$ ) that produce the inputs  $z^{(g)}$  to the trained layer  $f$ . The loss of a mini-batch is obtained by averaging the losses of the individual data points (cf. (3), (47) and (50)). Using the

superscript  $(g)$  in a self-explanatory way and letting  $w$  be as in Lemma 8, we have

$$L_{\text{batch}} = \frac{1}{G} \sum_{g=1}^G P^{(g)}(W) + \frac{1}{G} \sum_{g=1}^G \|Q^{(g)}w\|_2^2, \quad W \in \bigcap_{g=1}^G \mathcal{R}^{(g)}. \quad (55)$$

Letting  $z^{(g)}[k]$  denote the  $k$ -th coordinate of the input to the trained layer produced by the  $g$ -th datum  $x^{(g)}$  we can rewrite the mini-batch loss in a form analogous to (47), which enables us to extend our analysis of the shape of the loss surface to mini-batches.

**Lemma 10** *There exists a linear polynomial  $P_{\text{batch}}$  and a symmetric matrix  $Q_{\text{batch}}$  such that*

$$L_{\text{batch}} = P_{\text{batch}}(W) + w^T Q_{\text{batch}}^T Q_{\text{batch}} w. \quad (56)$$

*The singular values of  $Q_{\text{batch}}$  can be bounded as follows:*

$$s_{i,j}^2(Q_{\text{batch}}) \leq \frac{1}{G} \sum_{g=1}^G |z^{(g)}[j]|^2 s_1^2(M^{(g)} D^{(g)}), \quad (57)$$

$$s_{i,j}^2(Q_{\text{batch}}) \geq \frac{1}{G} \sum_{g=1}^G |z^{(g)}[j]|^2 s_n^2(M^{(g)} D^{(g)}). \quad (58)$$

Notably, for any  $j$  with  $z^{(g)}[j] = 0$  ( $1 \leq g \leq G$ ), we have that  $s_{i,j}(Q_{\text{batch}}) = 0$  for  $1 \leq i \leq n$ .

**Proof** The linear polynomial is obviously  $P_{\text{batch}}(W) = (1/G) \sum_{g=1}^G P^{(g)}(W)$ . To clarify the choice of  $Q_{\text{batch}}$ , set  $A = (1/G) \sum_{g=1}^G (Q^{(g)})^T Q^{(g)}$ . Then,  $L_{\text{batch}} = P_{\text{batch}}(W) + w^T A w$ , and  $A$  is symmetric and positive semi-definite.

Notably, the matrix  $A$  inherits the diagonal block structure of the  $Q^{(g)}$  as given in (49). In a slight abuse of notation, we can write  $A = \text{diag}(A_1 \dots, A_n)$ , where the  $j$ -th block is a symmetric  $n \times n$  matrix

$$A_j = \frac{1}{G} \sum_{g=1}^G z^{(g)}[j]^2 (M^{(g)} D^{(g)})^T (M^{(g)} D^{(g)}). \quad (59)$$

Following the argumentation of the proof of Lemma 8, meaning that we exploit the fact that the blocks do not interact, it becomes apparent that the eigenvalues of  $A$  are those of the blocks  $A_j$ . In short, we can list the eigenvalues of  $A$  as  $\lambda_{i,j}(A) = \lambda_i(A_j) \geq 0$ , where we arrange eigenvalues by size in descending order within each block. Arranging these eigenvalues by block into a diagonal matrix  $\Lambda$ , there is an orthogonal matrix  $U$  such that  $A = U^T \Lambda U$ . Now arranging their square roots  $\sqrt{\lambda_{i,j}(A)}$  in the same order into the diagonal matrix  $S$ , we can set  $Q_{\text{batch}} = U^T S U$ . Then,  $A = Q_{\text{batch}}^T Q_{\text{batch}}$  and  $w^T A w = \|Q_{\text{batch}} w\|_2^2$  as desired. While we could choose  $Q_{\text{batch}}$  to be any matrix of the form  $V^T S U$  with  $V$  orthogonal and find identical singular values, we chose a symmetric matrix for reasons of simplicity and convenience.

Clearly,  $s_{i,j}^2(Q_{\text{batch}}) = \lambda_{i,j}(A) = \lambda_i(A_j)$ . For the upper bound of these singular values, we proceed as follows

$$s_{i,j}^2(Q_{\text{batch}}) = \lambda_i(A_j) \leq \|A_j\|_{\text{Op}} \leq \frac{1}{G} \sum_{g=1}^G z^{(g)}[j]^2 \|M^{(g)} D^{(g)}\|_{\text{Op}}^2. \quad (60)$$

From this, (57) follows easily. For the lower bound (58) let  $u$  denote the normalized eigenvector to the smallest eigenvalue  $\lambda_{n,j}$  of  $A_j$ . Using (59) we obtain

$$s_{i,j}^2(Q_{\text{batch}}) = \lambda_i(A_j) \geq \lambda_n(A_j) = u^T A_j u = \frac{1}{G} \sum_{g=1}^G z^{(g)}[j]^2 \|(M^{(g)} D^{(g)})u\|_2^2. \quad (61)$$

From this, using  $\|(M^{(g)} D^{(g)})u\|_2 \geq s_n(M^{(g)} D^{(g)})\|u\|_2 = s_n(M^{(g)} D^{(g)})$ , the bound (58) follows.

For a convenient bound of the condition number of  $Q_{\text{batch}}$ , we set  $\bar{z}[j] = \sqrt{(1/G) \sum_{g=1}^G |z^{(g)}[j]|^2}$ , which can be interpreted as the  $j$ -th mean square feature of the batch. From Lemma 10 we can conclude that  $\bar{z}[j] = 0$  implies  $s_{i,j}(Q_{\text{batch}}) = 0$  for all  $i$ . Since vanishing singular values are disregarded in the computation of the condition number as defined in this paper, we obtain the following result.

**Corollary 11** *In the notation of Lemma 10, assume that there are constants  $C$  and  $c > 0$  such that for  $1 \leq g \leq G$*

$$s_1(M^{(g)} D^{(g)}) \leq C \quad \text{and} \quad s_n(M^{(g)} D^{(g)}) \geq c. \quad (62)$$

*Assume that there is at least one non-vanishing  $\bar{z}[j]$ , and denote by  $\min_j^* \bar{z}[j]$  the minimum over all non-vanishing  $\bar{z}[j]$ . Then*

$$\kappa(Q_{\text{batch}}) \leq \frac{C}{c} \cdot \frac{\max_j \bar{z}[j]}{\min_j^* \bar{z}[j]}. \quad (63)$$

## 4 Perturbation of a DN's Loss Landscape When Moving from a ConvNet to a ResNet

In this section we exploit differentiability of the family  $M(\rho)$  in order to study the perturbation of its singular values when moving from ConvNets with  $\rho = 0$  to ResNets with  $\rho = 1$  (see Sect. 3.2).

We start with a general result that will prove most useful.

**Lemma 12** *Let  $N(\rho)$  be a family of symmetric  $n \times n$  matrices with entries that are polynomials in  $\rho$ . Let  $\mathcal{E}$  denote the set of  $\rho \in \mathbb{R}$  for which  $N(\rho)$  has multiple eigenvalues. Then the following dichotomy holds:*

- *Either  $\mathcal{E}$  is finite or  $\mathcal{E} = \mathbb{R}$ .*

Also, if the entries of  $N$  are real analytical functions of  $\rho$  in some open set containing the compact interval  $[\rho_1, \rho_2]$ , then the set of  $\rho \in [\rho_1, \rho_2]$  with multiple eigenvalues of  $N(\rho)$  is either finite or all of  $[\rho_1, \rho_2]$ .

Clearly, this result extends easily to the multiplicity the singular values of any family of square matrices  $A(\rho)$  by considering  $N(\rho) = A^T A$ .

**Proof** Fixing  $\rho$  for a moment, the eigenvalues of  $N(\rho)$  are the zeros of the characteristic polynomial  $h(\lambda) = \det(N(\rho) - \lambda \text{Id})$ . This polynomial  $h$  has multiple zeros  $\lambda$  if and only if its discriminant is zero. It is well known that the discriminant  $\chi_h$  of a polynomial  $h$  is itself a polynomial in the coefficients of  $h$ .

If the entries of  $N$  are polynomials in  $\rho$ , then the discriminant of  $h$  is, as a function of  $\rho$ , in fact a polynomial  $\chi_h(\rho)$ . In other words, the set  $\mathcal{E}$  is the set of zeros of the polynomial  $\chi_h(\rho)$ . This set is finite unless the polynomial  $\chi_h(\rho)$  vanishes identically, in which case  $\mathcal{E}$  becomes all of  $\mathbb{R}$ .

Finally, if the entries of  $N$  are real analytical, then so is  $\chi_h(\rho)$ , and the identity theorem implies the claim.

Combining with Lemma 3 and Theorem 6 we obtain the following.

**Corollary 13** *Let  $N$  be as in Lemma 12. In the real analytical case, assume that  $\rho_1 = 0$  and  $\rho_2 = \rho$  for some fixed  $\rho$ . Assume that  $N(0)$  has no multiple eigenvalues. Then  $\mathcal{E}$  is finite. Thus, for any  $i$  the set  $[0, \rho] \setminus S_i$  is finite and*

$$\lambda_i(\rho) - \lambda_i(0) = \int_0^\rho \lambda'_i(t) dt = \int_0^\rho u_i(t)^T N'(t) u_i(t) dt. \quad (64)$$

After these general remarks, let us now consider  $M(\rho) = M_o + \rho \text{Id}$  (45) and use (64) to bound the growth of the singular values of  $M$ . To this end, we note the following.

**Lemma 14** (Derivative) *Let  $M_o$  be arbitrary (see (45)). Writing  $u_i = u_i(\rho)$  for short, the derivative of the eigenvalue  $\lambda_i(\rho)$  of  $M^T M$  becomes, on  $S_i$  (see (13))*

$$\lambda'_i(\rho) = u_i^T (M_o^T + M_o) u_i + 2\rho = 2u_i(\rho)^T M_o u_i(\rho) + 2\rho. \quad (65)$$

On  $S_i^*$  (see (14)) we have

$$s'_i(\rho) = \lambda'_i(\rho)/(2s_i) = u_i^T v_i = \cos(\alpha_i(\rho)), \quad (66)$$

where  $\alpha_i(\rho)$  denotes the angle between the  $i$ -th left and right singular vectors.

Using (66) it is possible to compute the derivative  $s'_i(\rho)$  numerically with ease.

**Proof** Apply Theorem 6 to  $N(\rho) = M^T M$ . Use  $N' = (M')^T M + M^T (M') = M + M^T = M_o^T + M_o + 2\text{Id}$  and  $u_i^T u_i = 1$  to obtain

$$\lambda'_i(\rho) = u_i^T (M^T + M) u_i = u_i^T (M_o^T + M_o) u_i + 2\rho. \quad (67)$$

Also, note that  $u_i^T M_o^T u_i = (u_i^T M_o^T u_i)^T = u_i^T M_o u_i$ , since this is a real-valued number. This establishes (65). In a similar way

$$\lambda'_i(\rho) = u_i^T (M^T + M) u_i = 2(u_i^T M u_i) = 2s_i u_i^T v_i. \quad (68)$$

The chain rule implies (66). Alternatively, (66) follows from Corollary 7.

**Definition 15** For a family of matrices of the form  $M(\rho) = M_o + \rho \text{Id}$ , denote the smallest eigenvalue of the symmetric matrix  $M_o^T + M_o$  by  $\underline{\tau}$  and its largest by  $\bar{\tau}$ .

Note that, due to  $\underline{\tau} \leq \bar{\tau}$ , we have

$$\|M_o^T + M_o\|_{\text{Op}} = \max(|\underline{\tau}|, |\bar{\tau}|) = \max(-\underline{\tau}, \bar{\tau}) \leq 2s_1(M_o). \quad (69)$$

Using Corollary 13 with  $N = M^T M$  and (65) from Lemma 14, we have the following immediate result.

**Theorem 16** Assume that  $M_o$  has no multiple singular values but is otherwise arbitrary (see (45)). Then, for all  $i$ ,

$$\rho \underline{\tau} + \rho^2 \leq s_i(\rho)^2 - s_i(0)^2 \leq \rho \bar{\tau} + \rho^2. \quad (70)$$

We mention that a random matrix with iid uniform or Gaussian entries has almost surely no multiple singular values (see Theorem 34).

We continue with some observations on the growth of singular values that leverage the simple form of  $M(\rho)$  (see (45)) explicitly and that apply even in the case of  $M_o$  having multiple zeros.

Obviously, (70) holds also if for any  $\rho$  the matrix  $M(\rho)$  is known to have no multiple singular values. Even without any such knowledge, we have the following useful consequence of Weyl's additive perturbation bound.

**Lemma 17** Let  $M_o$  be arbitrary (see (45)). Then we have

$$\rho - s_1(M_o) \leq s_i(M(\rho)) \leq \rho + s_1(M_o). \quad (71)$$

**Proof** Apply (24) with  $A = M(\rho) = M_o + \rho \text{Id}$  and  $\tilde{A} = \rho \text{Id}$  to obtain

$$|s_i(M(\rho)) - \rho| \leq s_1(A - \tilde{A}) = s_1(M_o). \quad (72)$$

For a comparison, note that (70) implies that for any  $i$

$$\sqrt{s_i^2(M_o) + \rho \underline{\tau} + \rho^2} \leq s_i(\rho) \leq \sqrt{s_i^2(M_o) + \rho \bar{\tau} + \rho^2} \quad (73)$$

whenever it holds. Since  $\bar{\tau} \leq 2s_1(M_o)$  and  $s_i(0) \leq s_1(0)$ , the upper bound is clearly sharper than that of (71).

Continuing our observations, the following fact is obvious due to the simple form of  $M(\rho)$ .

**Lemma 18** (Eigenvector) *Let  $M_o$  be arbitrary (see (45)). Assume that, for some  $\rho_o$ , the vector  $w = u_i(\rho_o)$  is a right singular vector as well as an eigenvector of  $M(\rho_o)$  with eigenvalue  $t$ . Then, for all  $\rho$ , the vector  $w$  is an eigenvector and right singular vector of  $M(\rho)$  with eigenvalue  $t + \rho - \rho_o$  and singular value  $|t + \rho - \rho_o|$ .*

Note that the singular value of the fixed vector  $w$  may change order, as other singular values may increase or decrease at a different speed.

Further, we have that  $|s'_i(\rho)| \leq 1$  on  $S_i^*$  which is obvious from geometry and from (66). The eigenvectors of  $M_o$ , if they exist, achieve extreme growth or decay.

**Lemma 19** (Maximum Growth) *Let  $M_o$  be arbitrary (see (45)). Assume that  $s'_i(\rho_o) = 1$  for some  $\rho_o$ . Then,  $w = u_i(\rho_o)$  is eigenvector and right singular vector of  $M(\rho)$  for all  $\rho$  with eigenvalue  $t + \rho - \rho_o$  and singular value  $|t + \rho - \rho_o|$  for  $t = s_i(\rho_o) > 0$ .*

*The analogous results holds for  $s'_i(\rho_o) = -1$  with  $t = -s_i(\rho_o) < 0$ .*

Note that the singular value of the fixed vector  $w$  may change order, since other singular values may increase or decrease at a different speed.

Of interest are also the zeros of the singular values since differentiability may fail there (c.f. Lemma 14).

**Lemma 20** (Non-vanishing Singular Values) *Let  $M = M_o + \rho \text{Id}$ . Assume that  $s_i(\rho_o) = 0$  for some  $\rho_o$ . Then,  $w = u_i(\rho_o)$  is an eigenvector and right singular vector of  $M(\rho)$  for all  $\rho$  with eigenvalue  $\rho - \rho_o$  and singular value  $|\rho - \rho_o|$ . Consequently, if  $M_o$  has no eigenvectors, then  $s_i(\rho) \neq 0$  for all  $\rho$  and  $i$ .*

Theorem 16 shows that under mild conditions all singular values should be expected to increase when adding a multiple of the identity. However, when concerned only with the condition number, the following result suffices, requiring no technical assumptions at all.

**Lemma 21** *Let  $M_o$  be arbitrary (see (45)). Then*

$$s_n(0)^2 + \rho \underline{\tau} + \rho^2 \leq s_n(\rho)^2 \leq s_1(\rho)^2 \leq s_1(0)^2 + \rho \bar{\tau} + \rho^2. \quad (74)$$

**Proof** The argument consists of a direct computation, leveraging the special form of  $M(\rho) = M_o + \rho \text{Id}$ . For simplicity, let us denote the normalized eigenvector of  $N(\rho) = M^T M$  with eigenvalue  $\lambda_i(\rho) = s_i(\rho)^2$  by  $u = u_i(\rho)$ . Then,

$$\begin{aligned} \lambda_i(\rho) &= u^T N(\rho) u = u^T (M_o^T M_o + \rho (M_o^T + M_o) + \rho^2 \text{Id}) u \\ &= u^T (M_o^T M_o) u + \rho u^T (M_o^T + M_o) u + \rho^2 \\ &\leq \lambda_1(0) + \rho \bar{\tau} + \rho^2. \end{aligned} \quad (75)$$

Similarly, we can estimate the second to last expression from below by  $\lambda_n(0) + \rho \underline{\tau} + \rho^2$ . Since  $i$  is arbitrary, the claim follows.

The advantage of (71) is its simple form, while (74) is sharper (cf. (73)). As we will see, certain random matrices of interest tend to achieve the upper bound (74) quite closely with high probability. Combining this observation with Lemmata 17 and 21, we obtain the following.

**Theorem 22** *Let  $M_o$  be arbitrary (see (45)). Assume that  $s_1(M_o) < 1$ . Then*

$$\kappa(M_o + \text{Id}) \leq \frac{\sqrt{1 + \bar{\tau} + s_1(M_o)^2}}{1 - s_1(M_o)} \leq \frac{1 + s_1(M_o)}{1 - s_1(M_o)}. \quad (76)$$

*Alternatively, assume that  $\underline{\tau} > -1$ . Then*

$$\kappa(M_o + \text{Id}) \leq \sqrt{\frac{1 + \bar{\tau} + s_1(M_o)^2}{1 + \underline{\tau} + s_n(M_o)^2}} \leq \frac{1 + s_1(M_o)}{\sqrt{1 + \underline{\tau}}}. \quad (77)$$

We list all four bounds due to their different advantages. Due to (69), the first bounds in both (76) as well as (77) are in general sharper than the second. Indeed, for random matrices as used in Sect. 6.2 we have typically  $\bar{\tau} \simeq \sqrt{2}s_1(M_o) < 2s_1(M_o)$ . The second bounds are simpler and sufficient in certain deterministic settings.

The first bound of (77) is the tightest of all four bounds for the random matrices of interest in this paper. On the other hand, since we typically have  $-\underline{\tau} \simeq \sqrt{2}s_1(M_o)$  for such matrices, the stated sufficient condition for (76) is somewhat easier to meet than that of (77).

**Proof** From Lemmata 21 and 17 we obtain the first inequality of (76); the second part follows from  $\bar{\tau} \leq 2s_1(0)$ . Using Lemma 21 we obtain (77).

## 5 Influence of the Activation Nonlinearity on a Deep Network’s Loss Landscape

As we put forward in Sect. 3, each layer of a DN can be represented locally as a product of matrices  $D_i W_i$ , where  $W_i$  contains the weights of the layer while the diagonal  $D_i$  with entries equal to  $\eta$  or 1 (compare (35) and (36)) captures the effect of the activation nonlinearity.

In the last section we explored how the singular values of a matrix  $M_o$  change when adding a fraction of the identity, i.e., when moving from a ConvNet to a ResNet. In this section we explore the effect of the activation nonlinearities on the singular values. To this end, let  $\mathcal{D}(m, n, \eta)$  denote the set of all  $n \times n$  diagonal matrices with exactly  $m$  diagonal entries equal to  $\eta$  and  $n - m$  entries equal to 1:

$$\mathcal{D}(m, n, \eta) = \{D = \text{diag}(\lambda_1, \dots, \lambda_n) : \lambda_i \in \{\eta, 1\}, \text{trace}(D) = m\eta + n - m\}. \quad (78)$$

We will always assume that  $0 \leq m \leq n$ . Note that  $\mathcal{D}(0, n, \eta) = \{\text{Id}\}$  and  $\mathcal{D}(n, n, \eta) = \{\eta \text{Id}\}$ .

### 5.1 Absolute Value Activation

Consider the case  $\eta = -1$ , corresponding to an activation of the type “absolute value” (see (35)).

**Proposition 23** (Absolute Value) *Let  $M_o$  be arbitrary (see (45)). For any  $D \in \mathcal{D}(m, n, -1)$ , we have*

$$s_i(M) = s_i(MD) = s_i(DM). \quad (79)$$

**Proof** This is a simple consequence of the SVD. Since  $D$  is orthogonal due to  $\eta = -1$ ,  $DV$  and  $DU$  are orthogonal as well. Multiplication by  $D$  only changes the singular vectors, not the singular values:  $DM = D(V\Sigma U^T) = (DV)\Sigma U^T$  and  $MD = (V\Sigma U^T)D = V\Sigma(DU)^T$ .

## 5.2 ReLU Activation

The next simple case is the ReLU activation, corresponding to  $\eta = 0$  in (35). As we state next, the singular values of a linear map  $M$  decrease when the map is combined with a ReLU activation, from left or right.

**Lemma 24** (ReLU Decreases Singular Values) *Let  $M_o$  be arbitrary (see (45)). For any  $D \in \mathcal{D}(m, n, 0)$  we have*

$$s_i(M) \geq s_i(MD) \quad \text{and} \quad s_i(M) \geq s_i(DM). \quad (80)$$

Note that a similar result holds for leaky ReLU ( $\eta \neq 0$ ), however only for  $i = 1$ . This follows from the fact that  $s_1(MD) \leq s_1(M)s_1(D)$  for any two matrices  $M$  and  $D$ .

**Proof** This is a simple consequence of the Cauchy interlacing property by deletion (see (23)).

As a consequence, we have for ReLU that  $|s_i(M) - s_i(DM)| = s_i(M) - s_i(DM)$  and the analogous for  $MD$ . This motivates the computation of the difference of traces which equals the sum of squares of the singular values (recall (16)).

**Proposition 25** (ReLU) *Let  $M_o$  be arbitrary (see (45)). Let  $\rho$  be arbitrary and write  $M$  instead of  $M(\rho)$  for short. For any  $D \in \mathcal{D}(m, n, 0)$ , the following holds: Let  $k_i$  ( $i = 1, \dots, m$ ) be such that  $D[k_i k_i] = 0$ . Denote by  $s_i$  the singular values of  $M$  and by  $\tilde{s}_i$  those of  $MD$ . Then*

$$\sum_{i=1}^n s_i^2 - \sum_{i=1}^n \tilde{s}_i^2 = \sum_{i=1}^n |s_i^2 - \tilde{s}_i^2| = \|M\|_F^2 - \|MD\|_F^2 = \sum_{i=1}^m \sum_{j=1}^n M[jk_i]^2 \quad (81)$$

$$\sum_{i=1}^n |s_i - \tilde{s}_i|^2 \leq \sum_{i=1}^m \sum_{j=1}^n M[jk_i]^2. \quad (82)$$

**Proof** Write  $\tilde{M} = MD$  for short. The following equality simplifies our computation and holds in this special setting, because  $D$  amounts to an orthogonal projection enabling us to appeal to Pythagoras:  $\|M\|_F^2 = \|\tilde{M}\|_F^2 + \|M - \tilde{M}\|_F^2$ . This can also be



verified by direct computation. Lemma 24 implies

$$\sum_{i=1}^n |s_i^2 - \tilde{s}_i^2| = \sum_{i=1}^n s_i^2 - \sum_{i=1}^n \tilde{s}_i^2 = \|M\|_F^2 - \|\tilde{M}\|_F^2 = \|M - \tilde{M}\|_F^2 = \sum_{i=1}^m \sum_{j=1}^n M[j, k_i]^2, \quad (83)$$

yielding (81). Using Corollary 2 with  $A = M$  and  $\tilde{A} = \tilde{M}$ , we obtain (82).

Bounds on the effect of nonlinearities at the input can be obtained easily from bounds at the output using the following fact.

**Proposition 26** (Input–Output via Transposition) *For  $D$  diagonal,*

$$s_i(MD) = s_i(DM^T). \quad (84)$$

**Proof** This follows from the fact that matrix transposition does not alter singular values.

## 6 Bounding the Weights of a ResNet Bounds its Condition Number

In this section, we combine the results of the preceding two sections in order to quantify how bounds on the entries of the matrix  $W_i$  translate into bounds on the condition number. As we have seen, this is of relevance in the context of optimizing DNs, especially when comparing ConvNets to ResNets (recall Sect. 3.2).

To this end, we will proceed in three steps:

1. Bounds on the weights  $W_i[j, k]$  translate into bounds of  $M(0)$  (ConvNet)
2. Perturbation when adding a residual link, moving from  $M(0)$  (ConvNet) to  $M(1)$  (ResNet)
3. Perturbation by the activation nonlinearity of the trained layer (passing from  $M(1)$  to  $M(1)D$ )

We treat the two cases of deterministic weights (hard bounds) and random weights (high probability bounds) separately.

### 6.1 Hard Bounds

We start with a few facts on how hard bounds on the entries of a matrix are inherited by the entries and singular values of perturbed versions of the matrix. We use the term “hard bounds” for the following results on weights  $W_i$  bounded by a constant in order to distinguish them from “high probability bounds” for random weights  $W_i$  with bounded standard deviation. Clearly, for uniform random weights we can apply both types of results.

**Lemma 27** *Assume that all entries of the  $n \times n$  matrix  $W_i$  are bounded by the same constant:  $|W_i[j, k]| \leq c$ . Then*

$$s_1(W_i) \leq cn, \quad (85)$$

*with equality, for example, for the matrix  $W_i$  with all entries equal to  $c$ .*

**Proof** The bound is a simple consequence of (20):  $s_1(W_i) \leq \|W_i\|_F \leq cn$ . If all entries of  $W_i$  are equal to  $c$ , then the unit vector  $u_1 = (1/\sqrt{n})(1, \dots, 1)^T$  achieves  $W_i u_1 = cn u_1$ .

**Lemma 28** Let  $W_i$  ( $i = 1, \dots, p$ ) be matrices as in Lemma 27. Assume that  $c \leq 1/n$ . Let  $D_i$  ( $i = 1, \dots, p$ ) be diagonal matrices with entries bounded by 1. Then all the entries of  $M_o = M(0) = D_p W_p \cdots D_1 W_1$  are bounded by  $c$  as well.

**Proof** Each entry of  $D_i W_i$  is bounded by  $c$  and each entry of  $(D_2 W_2)(D_1 W_1)$  is bounded by  $nc^2 = (nc)c \leq c$ . Iterating the argument establishes the second claim.

Recall that  $M(0) = M_o$  corresponds to a multi-layer  $F = f_p, \dots, f_1$  with arbitrary nonlinearities. Let us add a skip-residual link by choosing  $\rho = 1$  in (44).

**Lemma 29** Let  $W_i$  be as in Lemma 28 and recall (44). Assume that  $c < 1/n$ . Then

$$0 < 1 - cn \leq s_n(M(1)) \leq s_1(M(1)) \leq 1 + cn < 2 \quad (86)$$

and

$$\kappa(M(1)) \leq \frac{1 + cn}{1 - cn}. \quad (87)$$

**Proof** To establish (86) apply Lemma 17 with  $i = 1$  and  $i = n$  and  $\rho = 1$ , and note that  $s_1(M_o) \leq cn$  by Lemmata 27 and 28. The bound (87) follows immediately.

Proposition 23 implies the following bound for the absolute value activation.

**Theorem 30** (Hard Bound, Absolute Value) Under the assumptions of Lemma 29 and for  $D \in \mathcal{D}(m, n, -1)$ , we have

$$\kappa(M(1)D) \leq \frac{1 + cn}{1 - cn}. \quad (88)$$

For a residual multi-layer (44) with ReLU nonlinearities, we find the following.

**Theorem 31** (Hard Bound, ReLU) Under the assumptions of Lemma 29 and for  $D \in \mathcal{D}(m, n, 0)$ , the following statements hold:

Provided  $2cn < 1/(1 + n)$ , we have

$$\kappa(M(1)D) \leq \frac{1 + cn}{\sqrt{1 - 2cn(1 + n)}}. \quad (89)$$

Provided  $m \leq n(1 - cn)^2$ , we have

$$\kappa(M(1)D) \leq \frac{1 + cn}{1 - cn - \sqrt{m/n}}. \quad (90)$$

Provided  $2cn < 1/(1 + 2m)$ , we have

$$\kappa(M(1)D) \leq \frac{1 + cn}{\sqrt{1 - 2cn(1 + 2m)}}. \quad (91)$$

**Proof** Applying Proposition 25 with  $\rho = 1$ , i.e.,  $M = M(1)$ , we have by (81) that

$$\sum_{i=1}^n |s_i^2 - \tilde{s}_i^2| = \|M(1) - M(1)D\|_F^2 \leq m \cdot (c^2n + 2c + 1). \quad (92)$$

Note that this bound cannot be made arbitrarily small by choosing  $c$  small and that it amounts to at least  $m$ . However, the bulk of the sum of deviations stems from some  $m$  singular values of  $M$  being changed to zero by  $D$ , compensating the constant term  $m$  in the above bound. More precisely,  $s_i \geq 1 - cn$  by (86) for all  $i$  while  $\tilde{s}_i = 0$  for  $i = n - m + 1, \dots, n$ . Thus,

$$\sum_{i=n-m+1}^n |s_i^2 - \tilde{s}_i^2| = \sum_{i=n-m+1}^n s_i^2 \geq m \cdot (1 - cn)^2, \quad (93)$$

implying that for any  $1 \leq j \leq n - m$

$$|s_j^2 - \tilde{s}_j^2| \leq \sum_{i=1}^{n-m} |s_i^2 - \tilde{s}_i^2| \leq m \cdot (c^2n + 2c + 1) - m \cdot (1 - cn)^2 \leq 2mc(1 + n) \quad (94)$$

(using  $cn - cn^2 < 0$ ). Therefore,

$$\begin{aligned} \tilde{s}_j^2 &\geq s_j^2 - 2mc(1 + n) \geq (1 - cn)^2 - 2mc(1 + n) = 1 - c(2n + 2m(1 + n) - cn^2) \\ &\geq 1 - 2c(n^2 + n) \end{aligned} \quad (95)$$

(using  $m(1 + n) \leq n^2 - 1 \leq n^2$ ). Since this bound holds for all non-zero singular values of  $\tilde{M}$ , we obtain

$$\tilde{s}_* \geq \sqrt{1 - 2c(n^2 + n)}. \quad (96)$$

Finally, we have  $\tilde{s}_1 \leq s_1 \leq 1 + cn$  by Lemma 24 and (86). This proves (89). Alternatively, we can estimate as follows:

$$\tilde{s}_j^2 \geq 1 - c(2n + 2m(1 + n) - cn^2) \geq 1 - cn(2 + 2m(1 + 1/n) - cn) \geq 1 - 2cn(1 + 2m) \quad (97)$$

using  $1/n \leq 1$ , which proves (91). At last, we can use (82) instead of (81). Analogous to before, we find

$$\sum_{i=1}^n |s_i - \tilde{s}_i|^2 \leq \|M(1) - M(1)D\|_F^2 \leq m \cdot (c^2n + 2c + 1), \quad (98)$$

implying for any  $1 \leq j \leq n - m$  that  $|s_j - \tilde{s}_j|^2 \leq 2mc(1 + n)$ . Therefore, provided  $2cn < 1/(1 + n)$  or  $2c(1 + n) < 1/n$ , we obtain (90) by observing that

$$\tilde{s}_j \geq s_j - \sqrt{2mc(1 + n)} \geq (1 - cn) - \sqrt{2mc(1 + n)} \geq 1 - cn - \sqrt{m/n}. \quad (99)$$

For leaky-ReLU, the situation is more intricate.

**Theorem 32** (Hard Bound, Leaky-ReLU) *Let  $W_i$  be as in Lemma 28 and recall (44). Assume that  $c < 1/n$ . Let  $D \in \mathcal{D}(m, n, \eta)$  with  $0 \leq \eta \leq 1$ .*

*Provided  $n \geq 5$  and  $\eta < 1$  large enough to ensure  $(1 - cn)^2 > 4(1 - \eta)\sqrt{m}$ , we have*

$$\kappa(M(1)D) \leq \frac{1 + cn}{\sqrt{(1 - cn)^2 - 4(1 - \eta)\sqrt{m}}}. \quad (100)$$

*Alternatively, if  $\eta < 1$  is such that  $1 - cn > (1 - \eta)m(3c + 1)$ , then we have*

$$\kappa(M(1)D) \leq \frac{1 + cn}{1 - cn - (1 - \eta)m(3c + 1)}. \quad (101)$$

This situation corresponds to a ResNet (44) with leaky-ReLU nonlinearities, because  $\eta$  can take any value between 0 and 1 (see (35)).

**Proof** We proceed along the lines of the arguments of Theorem 31. To this end, we need to generalize the results of Proposition 25.

In order to replace (81) we write for short  $M = M(1) = M_o + \text{Id}$  with singular values  $s_i$  and  $\tilde{M} = M(1)D$  with singular values  $\tilde{s}_i$ . Applying Proposition 1 with  $A = M^T M$  and  $\tilde{A} = \tilde{M}^T \tilde{M} = (MD)^T MD = DAD$ , we obtain

$$\sum_{i=1}^n |s_i^2 - \tilde{s}_i^2|^2 \leq \|A - \tilde{A}\|_F^2. \quad (102)$$

Since all entries of  $M_o$  are bounded by  $c$  due to Lemma 28 and using  $cn < 1$ , it follows that the entries of  $A$  are bounded as follows

$$|A[i, j]| \leq \begin{cases} (n - 2)c^2 + 2(1 + c)c \leq 3c & \text{if } i \neq j \\ (n - 1)c^2 + (1 + c)^2 \leq 3c + 1 & \text{if } i = j. \end{cases} \quad (103)$$

Using these bounds, direct computation shows that  $A - \tilde{A} = A - DAD$  contains  $2m(n - m)$  terms off the diagonal bounded by  $3c(1 - \eta)$ ,  $m$  terms in the diagonal bounded by  $(3c + 1)(1 - \eta^2)$ , further  $m^2 - m$  terms off the diagonal bounded by  $3c(1 - \eta^2)$ , and the remaining  $(n - m)^2$  terms equal to zero. Using  $cn < 1$  whenever apparent as well as  $(1 - \eta^2) \leq 2(1 - \eta)$ , we can estimate as follows:

$$\begin{aligned} \|A - \tilde{A}\|_F^2 &\leq (1 - \eta)^2 18c^2 m(n - m) + (1 - \eta^2)^2 \left( 9c^2(m^2 - m) + (3c + 1)^2 m \right) \\ &\leq (1 - \eta)^2 18c^2 m(n - m) + 4(1 - \eta)^2 \left( 9c^2 m^2 + 6cm + m \right) \\ &\leq (1 - \eta)^2 (42cm + 18c^2 m^2 + 4m) \leq 4(1 - \eta)^2 m(15/n + 1). \end{aligned} \quad (104)$$

We conclude that, for  $n \geq 5$ , we have  $|s_j^2 - \tilde{s}_j^2| \leq (1 - \eta)4\sqrt{m}$ . Together with  $s_j \geq (1 - cn)$ , we obtain (100).

Alternatively, we may strive at adapting (82) (which is based on Corollary 2) for  $\eta \neq 0$ . Recomputing (82) using (18) with  $A = M$  and  $\tilde{A} = \tilde{M}$  for non-zero  $\eta$ , we obtain

$$\sum_{i=1}^n |s_i - \tilde{s}_i|^2 \leq \|M - \tilde{M}\|_F^2 \leq (1 - \eta)^2 m(c^2 n + 2c + 1). \quad (105)$$

We obtain  $|s_j - \tilde{s}_j| \leq (1 - \eta)\sqrt{m(c^2 n + 2c + 1)} \leq (1 - \eta)\sqrt{m(3c + 1)}$  and (101).

We note that the earlier Theorem 22 will not lead to strong results in this context. However, it will prove effective in the random setting of the next section.

## 6.2 Bounds with High Probability for Random Matrices

While the results of the last section provide absolute guarantees, they are quite restrictive on the matrix  $M_o$  since they cover the worst cases. Working with the average behavior of a randomly initialized matrix  $M_o$  enables us to relax the restrictions from bounding the actual values of the weights to only bounding their standard deviation.

The strongest results in this context concern the operator norm and exploit cancellations in the entries of  $M_o$  (see [3, Proposition 3.1], [23, 24]). Most useful for this study is the well-known result due to Tracy and Widom [23, 24]. Translated into our setting, it reads as follows.

**Theorem 33** (Tracy-Widom) *Let  $X_{i,j}$  ( $1 \leq i, j$ ) be an infinite family of iid copies of a Uniform or Gaussian random variable  $X$  with zero mean and variance 1.*

*Let  $K_n$  be a sequence of  $n \times n$  matrices with entries  $K_n[i, j] = X_{i,j}$  ( $1 \leq i, j \leq n$ ).*

*Let  $\bar{t}_n$  denote the largest eigenvalue of  $K_n^T + K_n$  and let  $\underline{t}_n$  denote its smallest.*

*Then, as  $n \rightarrow \infty$ , both  $\bar{t}_n$  and  $-\underline{t}_n$  lie with high probability in the interval*

$$\left[ \sqrt{8n} - O(n^{-1/6}), \sqrt{8n} + O(n^{-1/6}) \right]. \quad (106)$$

*More precisely, the distributions of*

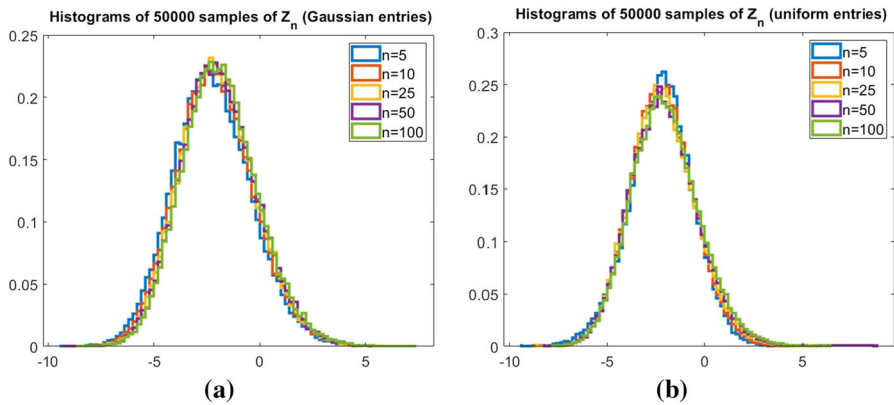
$$Z_n := (\bar{t}_n - \sqrt{8n}) n^{1/6} \quad (107)$$

$$\tilde{Z}_n := (-\underline{t}_n - \sqrt{8n}) n^{1/6} \quad (108)$$

*converge both (individually) to a Tracy-Widom law  $P_{TW}$ . The Tracy-Widom law is well concentrated at values close to 0.*

The random variables  $Z_n$  and  $\tilde{Z}_n$  are equal in distribution due to the symmetry of the random variable  $X$ . However, they are not independent.

It is easy to verify that the theory of Wigner matrices and the law of Tracy-Widom apply here and to establish Theorem 33. In particular,  $K_n^T + K_n$  is symmetric with iid entries above the diagonal and iid entries in the diagonal, all with finite non-zero second moment. The normalization of  $\bar{t}_n$  leading to  $Z_n$  (see (107)) is quite sharp, meaning that the distributions of  $Z_n$  are close to their limit starting at values as low as  $n = 5$  (see Fig. 4).



**Fig. 4** Empirical demonstration of the Tracy-Widom law. Note the collapsing histograms of  $Z_n = (\bar{t}_n - \sqrt{8n})n^{1/6}$  for  $n = 5, 10, 25, 50, 100$  where  $\bar{t}_n$  denotes the largest eigenvalue of the  $n \times n$  matrix of the form  $K_n + K_n^T$  and where the entries of  $K_n$  are drawn from a **a** Gaussian or **b** uniform distribution with zero mean and variance 1. Note that the right-side tails  $P[Z_n > t]$  decay very quickly, leading to negligible probabilities at moderate values of  $t$

The Tracy-Widom distributions are well studied and documented, with a right tail that decays exponentially fast [7]. Moreover, they are well concentrated at values close to 0. Combined with the fast convergence, this implies that  $\bar{t}_n$  is close to  $\sqrt{8n}$  with high probability even for modest  $n$ .

With the appropriate adjustments, Theorem 33 provides tight control on  $\bar{\tau}$  and  $\underline{\tau}$  (see Definition 15). In order to exploit this control toward a bound on the growth of singular values via Theorem 16, the following proves useful.

**Theorem 34** Assume that the  $n^2$  entries of the random matrix  $M_o$  are jointly continuous with some joint probability density function (e.g., Uniform or jointly Gaussian entries). Then, almost surely,  $M_o$  possesses no multiple singular value, and the growth of the singular values of  $M_o + \rho \text{Id}$  is bounded as in (70).

**Proof** As in the proof of Lemma 12, we argue that the discriminant  $\chi$  of the characteristic polynomial of the  $n \times n$  matrix  $M_o^T M_o$  is itself a polynomial of the entries of  $M_o$ . Viewing  $M_o$  as a point in  $\mathbb{R}^{n^2}$  defined by its entries, the matrices  $M_o$  with multiple singular values form a set that is identical to the zero-set of the polynomial  $\chi$  and form, therefore, a set of Lebesgue measure 0. The statement is now obvious.

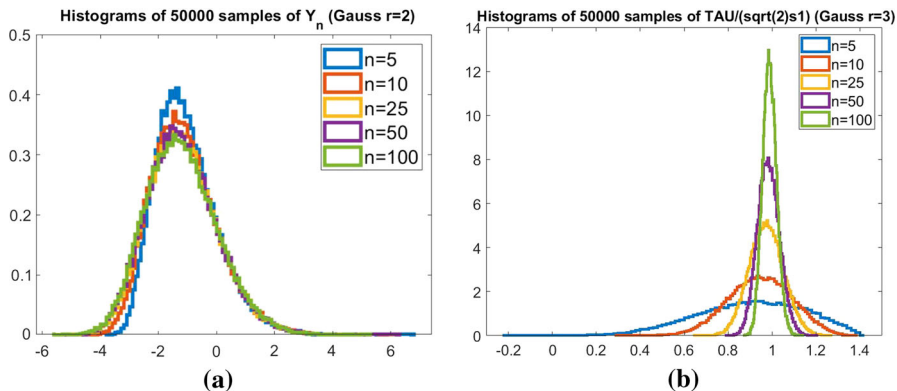
A related useful result reads as follows (see [5, Theorem 6.2.6 and Section 6.2.3] and the references therein).

**Theorem 35** Let  $K_n$  be as in Theorem 33, but Gaussian. Then, as  $n \rightarrow \infty$

$$\frac{s_1(K_n)}{\sqrt{n}} \xrightarrow{\text{a.s.}} 2. \quad (109)$$

Moreover, letting  $\mu_n = (\sqrt{n-1} + \sqrt{n})^2$  and  $q_n = \sqrt{\mu_n}(1/\sqrt{n-1} + 1/\sqrt{n})^{1/3}$ , the distribution of

$$Y_n := (s_1(K_n)^2 - \mu_n)/q_n \quad (110)$$



**Fig. 5** **a** Empirical demonstration of the convergence of  $Y_n$  in (110) to a Tracy-Widom law. Note that the right-side tails  $P[Y_n > t]$  decay very quickly, leading to negligible probabilities at moderate values of  $t$  and for  $n$  as small as 5 similar to  $Z_n$  from Fig. 4. **b** Empirical demonstration of the convergence of  $\bar{\tau}/(\sqrt{2}s_1)$  to 1: The ratio lies indeed with high probability between 0.6 and 1.4 for  $n$  as small as 10

converges narrowly to a Tracy-Widom law.

Comparing with Theorem 33, we should expect that  $s_1(K_n) \simeq 2\sqrt{n} \simeq \bar{\tau}_n/\sqrt{2}$  (compare to Fig. 5).

Combining Theorem 22 (with  $M_0 = W_1$ ) and Theorem 33, and using that (for  $n \geq 2$ )  $\mu_n \leq 4n$  and  $q_n \leq \sqrt{(4n)2^{1/3}(n-1)^{-1/6}} \leq \sqrt{(4n)2^{1/3}(n/2)^{-1/6}} \leq 2^{3/2}n^{1/3}$ , we obtain the following.

**Lemma 36** (High Probability Bound) *In the notation of Theorem 33, let  $r > 1$  and  $n \geq 2$ , set*

$$\sigma_n = \frac{1}{r\sqrt{8n}}, \quad (111)$$

and let  $W_1 = \sigma_n K_n$ . Denote the largest and smallest eigenvalues of  $W_1^T + W_1$  by  $\bar{\tau}_n = \bar{\tau}_n(r)$  and  $\underline{\tau}_n = \underline{\tau}_n(r)$  (compare with (69)).

(i) We have

$$\bar{\tau}_n(r) = \frac{1}{r} \left( 1 + Z_n \frac{1}{\sqrt{8n^{2/3}}} \right) \quad - \underline{\tau}_n(r) = \frac{1}{r} \left( 1 + \tilde{Z}_n \frac{1}{\sqrt{8n^{2/3}}} \right) \quad (112)$$

and

$$P[\underline{\tau}_n(r) > -1] = P[\tilde{Z}_n \leq (r-1)\sqrt{8n^{2/3}}].$$

(ii) In the Gaussian case, we have that  $s_1(W_1) \rightarrow 1/(\sqrt{2}r)$  almost surely. Also

$$s_1(W_1)^2 = \frac{1}{2r^2} \left( \frac{\mu_n}{4n} + n^{-2/3} \frac{q_n}{4n^{1/3}} Y_n \right) \leq \frac{1}{2r^2} \left( 1 + n^{-2/3} \frac{1}{\sqrt{2}} Y_n \right) \quad (113)$$

(iii) For any realization of  $W_1$  with  $s_1(W_1) < 1$

$$\kappa(W_1 + \text{Id}) \leq \frac{\sqrt{s_1^2(W_1) + \bar{\tau}_n + 1}}{1 - s_1(W_1)} \leq \frac{1 + s_1(W_1)}{1 - s_1(W_1)}. \quad (114)$$

(iv) For any realization of  $W_1$  with  $\underline{\tau}_n > -1$  we have the slightly tighter bounds

$$0 < s_n^2(W_1) + \underline{\tau}_n + 1 \leq s_n^2(W_1 + \text{Id}) \leq s_1^2(W_1 + \text{Id}) \leq s_1^2(W_1) + \bar{\tau}_n + 1. \quad (115)$$

The condition number of  $W_1 + \text{Id}$  is then bounded as

$$\kappa_n := \kappa(W_1 + \text{Id}) \leq \sqrt{\frac{s_1^2(W_1) + \bar{\tau}_n + 1}{\underline{\tau}_n + 1}} \leq \frac{s_1(W_1) + 1}{\sqrt{\underline{\tau}_n + 1}}. \quad (116)$$

When comparing (i) and (ii), it becomes apparent that the stated sufficient condition in (iii) is more easily satisfied than that of (iv).

For a single residual layer (43) with weights  $W_1$  and a ResNet with absolute value nonlinearities, we obtain the following.

**Corollary 37** (High Probability Bound, Absolute Value) *Let  $W_1$  be as in Lemma 36. Let  $D, D_1 \in \mathcal{D}(m, n, -1)$  for some fixed  $m$ . Set  $M_o = D_1 W_1$  and recall Definition 15.*

*(i) In the Gaussian case, we have  $s_1(M_o) \rightarrow 1/(\sqrt{2}r)$  almost surely.*

*(ii) For any realization of  $M_o$  for which  $s_1(M_o) < 1$ , we have*

$$\kappa((M_o + \text{Id})D) \leq \frac{\sqrt{s_1^2(M_o) + \bar{\tau} + 1}}{1 - s_1(M_o)} \leq \frac{1 + s_1(M_o)}{1 - s_1(M_o)}. \quad (117)$$

*(iii) For any realization of  $W_1$  with  $\underline{\tau} > -1$  the following bound holds*

$$\kappa((M_o + \text{Id})D) \leq \sqrt{\frac{s_1^2(M_o) + \bar{\tau} + 1}{\underline{\tau} + 1}} \leq \frac{s_1(M_o) + 1}{\sqrt{\underline{\tau} + 1}}. \quad (118)$$

**Proof** Note that  $M_o = D_1 W_1$  is a random matrix with the same properties as  $W_1$  itself, since  $D_1$  is deterministic and changes only the sign of some of the entries of  $W_1$ . We may, therefore, apply Lemma 36 to  $M_o$  (replacing  $W_1$ ).


Also,  $D$  does not alter the singular values:  $s_i(M_o) = s_i(W_1)$  and  $s_i((M_o + \text{Id})D) = s_i(M_o + \text{Id})$ .

### Probabilities of Exceptional Events

- Clarification regarding  $\underline{\tau}$  vs.  $\underline{\tau}_n$  in the setting of Corollary 37:

While we have  $s_1(M_o) = s_1(W_1)$ , one should note that  $\bar{\tau}$  (largest eigenvalue of  $M_o^T + M_o$ ) may differ from  $\bar{\tau}_n$  (largest eigenvalue of  $W_1^T + W_1$ ) for any particular realization. However, since  $M_o$  and  $W_1$  are equal in distribution due to the special form of  $D_1$ , we have that  $\bar{\tau}$  and  $\bar{\tau}_n$ , as well as  $-\underline{\tau}$  and  $-\underline{\tau}_n$  are all equal in distribution with distribution well concentrated around  $1/r$ .





**Fig. 6** Empirical demonstration of the efficiency of the bounds (115) and (116) using  $r = 2$ . The figures show histograms of the fractions **a**  $[s_n^2(W_1 + \text{Id})]/[s_n^2(0) + \tau_n + 1]$ , **b**  $[s_1^2(W_1) + \tau_n + 1]/[s_1^2(W_1 + \text{Id})]$ , and **c**  $\sqrt{(s_1^2(W_1) + \tau_n + 1)/(\tau_n + 1)}/\kappa(W_1 + \text{Id})$ . Note that all values are above 1 and actually quite close to 1

- Exception  $E_{\tau_n} = E_{\tau_n(r)} = \{\tau_n(r) \leq -1\}$ :

We have

$$P[E_{\tau_n}] \rightarrow 0 \text{ as } n \rightarrow \infty \text{ for } r > 1. \quad (119)$$

Moreover,  $P[E_{\tau_n}] \approx 0$  for values as small as  $r \geq 2$  and  $n \geq 5$ . We are able to support this claim through the following computations and simulations: Clearly, on the one hand,  $P[\tau_n > -1] = P[\tilde{Z}_n < (r-1)\sqrt{8n^{2/3}}] \rightarrow 1$  for  $r > 1$ . On the other hand, for  $r \geq 2$  and  $n \geq 5$ , we have  $(r-1)\sqrt{8n^{2/3}} > 8$  and may estimate  $P[\tau_n > -1] = P[\tilde{Z}_n < (r-1)\sqrt{8n^{2/3}}] \geq P[\tilde{Z}_n < 8] = P[Z_n < 8]$ . From Fig. 4 it appears that  $P[Z_n < 8] = 1$  for  $n \geq 5$  for all practical purposes. This is confirmed by our simulations in the following sense: Every single one of the 500,000 random matrices simulated with  $r = 2$  satisfied  $\tau_n > -1$ .

- Exception  $E_{s_1(W_1)} = \{s_1(W_1) \geq 1\}$ :

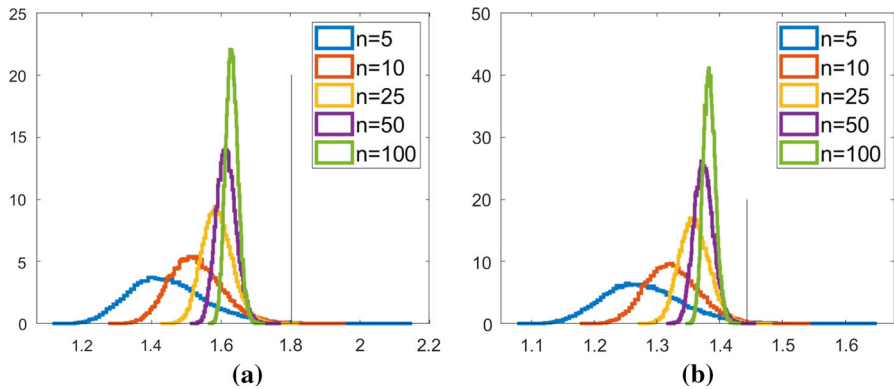
We have that

$$P[E_{s_1(W_1)}] \rightarrow 0 \text{ as } n \rightarrow \infty \text{ for } r > 1. \quad (120)$$

Moreover,  $P[E_{s_1(W_1)}] \approx 0$  for values as small as  $r \geq 2$  and  $n \geq 5$ . According to (113) and Theorem 35, we have  $P[s_1(W_1)^2 \geq 1] \leq P[Y_n \geq \sqrt{2}(2r^2 - 1)n^{2/3}] \rightarrow 0$ . As demonstrated in Fig. 5,  $Y_n$  converges rapidly to a Tracy-Widom law whence  $P[E_{s_1(W_1)}]$  is small for  $r \geq 2$  and modest  $n$ . In the setting of Corollary 37, the probability of the event  $E_{s_1(M_o)} = \{s_1(M_o) \geq 1\}$  tends to zero rapidly in a similar fashion.

The bounds (115) and (116) are surprisingly efficient. The first upper bound of (116), for example, is only about 10% larger than  $\kappa(W_1 + \text{Id})$  for large  $n$  (see Fig. 6).

We see an explanation for this efficiency of the bounds to be rooted in the fact that the singular values of random matrices tend to be as widely spread as possible, meaning that the largest singular values grow as fast as possible while the smallest singular values grow as slowly as possible. While this maximal spreading is well known for symmetric random matrices called “Wigner matrices,” we are not aware of any reports of this behavior for non-symmetric matrices such as  $W_1$ .



**Fig. 7** Empirical demonstration of the asymptotic bound (121). Histograms of the  $\kappa(W_1 + \text{Id})$  for **a**  $r = 2$  and **b**  $r = 3$  together with the upper bound in form of a deterministic expression from (121) in black. Note that the values of  $\kappa(W_1 + \text{Id})$  lie asymptotically below the deterministic expression as indicated in (121). Also, recall that  $\kappa(W_1 + \text{Id})$  and  $\kappa(D_1 W_1 + \text{Id})$  are equal in distribution in the setting of Corollary 37

For design purposes, an asymptotic bound for the condition number  $\kappa(M(1)D)$  in terms of the network parameters might be more useful than those in terms of the random realization as given in Corollary 37.

**Theorem 38** (High Probability Bound in the Limit) *With the notation and assumptions of Theorem 33 and Lemma 36, the upper bounds given there converge in distribution to constants as follows*

$$\sqrt{\frac{s_1^2(W_1) + \bar{\tau}_n + 1}{\bar{\tau}_n + 1}} \xrightarrow{\text{distr}} \sqrt{\frac{r^2 + r + 1/2}{r^2 - r}} \quad (121)$$

and

$$\frac{\sqrt{s_1^2(W_1) + \bar{\tau}_n + 1}}{1 - s_1(W_1)} \xrightarrow{\text{distr}} \frac{\sqrt{r^2 + r + 1/2}}{r - 1/\sqrt{2}}. \quad (122)$$

For a demonstration of the deterministic expression bounding the condition number as in (121), see Fig. 7. For  $r = 2$ , for example, the right-hand side of (121) becomes 1.8028, while the right-hand side of (122) becomes the slightly larger 1.9719.

**Proof** Since  $Z_n$  converges in distribution to a real-valued random variable  $Z$ , we can conclude that  $Z_n/n^{2/3}$  converges in distribution to 0. Indeed, for any  $\varepsilon > 0$  and any  $n_o \geq 1$ , we have for  $n \geq n_o$  that

$$P \left[ |Z_n/n^{2/3}| \geq \varepsilon \right] \leq P \left[ |Z_n| \geq \varepsilon n_o^{2/3} \right] \xrightarrow{n \rightarrow \infty} P \left[ |Z| \geq \varepsilon n_o^{2/3} \right]. \quad (123)$$

From this we obtain for any  $n_o \geq 1$  that

$$\limsup_{n \rightarrow \infty} P \left[ |Z_n/n^{2/3}| \geq \varepsilon \right] \leq P_{TW} \left[ |Z| \geq \varepsilon n_o^{2/3} \right], \quad (124)$$

and letting  $n_o \rightarrow \infty$  we conclude that  $P[|Z_n/n^{2/3}| \geq \varepsilon] \rightarrow 0$ . This establishes that  $Z_n/n^{2/3} \xrightarrow{\text{distr}} 0$ . Similarly, we obtain  $\tilde{Z}_n/n^{2/3} \xrightarrow{\text{distr}} 0$ .

By the continuous mapping theorem, and since  $r > 1$ , we obtain

$$\sqrt{\frac{s_1^2(W_1) + \bar{\tau}_n + 1}{\tau_n + 1}} \xrightarrow{\text{distr}} \sqrt{\frac{1/(2r^2) + 1/r + 1}{-1/r + 1}}. \quad (125)$$

Similarly,

$$\frac{\sqrt{s_1^2(W_1) + \bar{\tau}_n + 1}}{1 - s_1(W_1)} \xrightarrow{\text{distr}} \frac{\sqrt{1/(2r^2) + 1/r + 1}}{1 - 1/(\sqrt{2}r)}. \quad (126)$$

Simple algebra completes the proof.

Our simulations suggest that actually  $\bar{\tau}_n/s_1(W_1)$  converges in distribution to  $\sqrt{2}$  as  $n \rightarrow \infty$ , implying that we should see  $\bar{\tau}_n \simeq \sqrt{2}s_1(W_1)$  for most matrices in real world applications.

We turn now to the condition number of the relevant matrix  $(D_1 W_1 + \text{Id})D$  for a ResNet with ReLU activations that are represented by  $D_1 \in \mathcal{D}(\tilde{m}, n, 0)$  and  $D \in \mathcal{D}(m, n, 0)$ . We address the case of a single residual layer following the trained layer (cf. (43) with  $p = 1$ ).

To this end, we let

$$\nu = m + \tilde{m} - \frac{m \tilde{m}}{n}, \quad (127)$$

which can be considered as a measure of the extent to which these two nonlinearities affect the output in combination.

**Theorem 39** (High Probability Bound, ReLU) *Let  $D_1 \in \mathcal{D}(\tilde{m}, n, 0)$  and  $D \in \mathcal{D}(m, n, 0)$  with  $1 \leq m + \tilde{m}$ . Then,  $1 \leq \nu \leq n$ . Choose  $\theta$  such that  $4 < \theta \leq 2\sqrt{n}$ , and let  $\sigma_n$  and  $W_1$  be as in Lemma 36, with  $r$  defined as below and with Gaussian entries for  $W_1$ .*

(1) *Choose  $r > (2 + 2\nu + \theta)$ . Then, with probability at least  $1 - 2\exp(-\theta^2/2) - P[E_{s_1(W_1)}]$ , we have*

$$\kappa((D_1 W_1 + \text{Id})D) \leq 2 \frac{1 + 1/r}{(1 - (2 + 2\nu + \theta)/r)}. \quad (128)$$

(2) *Choose  $r' > 1$  and  $r > (r' + r'\nu + \theta)$ . Then, with probability at least  $1 - 2\exp(-\theta^2/2) - P[E_{\tau_n(r')}]$ , we have*

$$\kappa((D_1 W_1 + \text{Id})D) \leq 2 \frac{1 + s_1(W_1)}{(1 - (r' + r'\nu + \theta)/r)}. \quad (129)$$

We recall that, according to our simulations,  $P[E_{s_1(W_1)}] = P[s_1(W_1)^2 \geq 1] \leq P[Y_n \geq \sqrt{2}(2r^2 - 1)n^{2/3}]$  and  $P[E_{\tau_n(r')}] = P[Z_n > (r' - 1)\sqrt{8}n^{2/3}]$  are negligible for  $r, r' \geq 2$  and modest  $n$ .

We mention further that the restriction  $4 \leq \theta \leq 2\sqrt{n}$  is solely for simplicity of the representation and can be dropped at the cost of more complicated formulas (see proof). In the given form, useful bounds on the probability are obtained even with modest  $n$  as low as  $n = 7$ .

Clearly, the bound (128) on the condition number  $\kappa$  can be forced arbitrarily close to 2 by choosing  $r$  sufficiently large. In fact, bounds can be given that come arbitrarily close to 1, as becomes apparent from (149) and (143) in the proof.

**Proof** Let  $A = W_1 + \text{Id}$  with singular values  $s_i$  and  $\tilde{A} = (D_1 W_1 + \text{Id})D$  with singular values  $\tilde{s}_i$ . Clearly,  $r > r' > 1$ .

We proceed in several steps, addressing first (129).

(i) *Sure bound on the sum of perturbations.* We cannot apply Proposition 25 as we did in Theorem 31, since we need to take into account also the effect of the left-factor  $D_1$ . Due to Corollary 2 we have

$$\sum_{i=1}^n |s_i - \tilde{s}_i|^2 \leq \|\tilde{A} - A\|_2^2 = \|(D_1 W_1 D - W_1) + (D - \text{Id})\|_2^2. \quad (130)$$

We find  $k = n(m + \tilde{m}) - m\tilde{m}$  non-zero terms in the first matrix  $(D_1 W_1 D - W_1)$  that are of the form  $W_1[i, j]$ . The second matrix  $(D - \text{Id})$  is diagonal with exactly  $m$  non-zero terms in the diagonal that are all equal to  $-1$ . Thus, the Frobenius norm  $\|\tilde{A} - A\|_2^2$  consists of exactly  $k$  terms of the form  $W_1[i, j]^2$ , which we collect into the variable  $V_1$ ; several terms of the form  $2W_1[i, j]$ , which we collect into the variable  $V_2$ ; and finally  $m$  terms equal to 1. Note that  $V_2$  consists of at most  $m$  terms, depending on how many diagonal positions are non-zero in both  $(D_1 W_1 D - W_1)$  and  $(D - \text{Id})$ . We obtain

$$\sum_{i=1}^n |s_i - \tilde{s}_i|^2 \leq \|\tilde{A} - A\|_2^2 = V_1 + V_2 + m. \quad (131)$$

By elementary calculus we have  $n \leq k \leq n^2$ . Together with  $m, \tilde{m} \leq n$  and  $m + \tilde{m} \geq 1$ , this implies that  $1 \leq v \leq n$ .

(ii) *Concentration of the values of  $V_1$ ,  $V_2$  and  $\underline{\tau}_n$ .* The variable  $V_2$  is zero-mean Gaussian with a variance of at most  $4m\sigma_n^2$ . To formulate where the values of  $V_2$  concentrate, we will use a parameter  $\theta$ . A standard result says that

$$P[V_2 > \theta \cdot 2\sigma_n \sqrt{m}] \leq P[X > \theta \cdot \sigma_n] \leq \exp(-\theta^2/2). \quad (132)$$

Since  $W_1[i, j]/\sigma_n$  are independent standard normal variables, the variable  $Y = V_1/\sigma_n^2$  follows a Chi-square distribution with  $k$  degrees of freedom. One of the sharpest bounds known for the tail of a Chi-squared distribution is that of Massart and Laurent [17, Lemma 1, pg. 1325]. A corollary of their bound yields

$$P[Y \geq k + 2\sqrt{kq} + 2q] \leq \exp(-q). \quad (133)$$

Choosing  $q = \theta^2/2$  for symmetry with  $V_2$ , we obtain

$$P[V_1 > \sigma_n^2(k + 2\theta\sqrt{k/2} + \theta^2)] \leq \exp(-\theta^2/2). \quad (134)$$

Finally, since  $-\underline{\tau}_n$  and  $\bar{\tau}_n$  are equal in distribution, we obtain

$$P[-\underline{\tau}_n > r'\sigma_n\sqrt{8n}] = P[-\underline{\tau}_n > r'/r] = P[Z_n > (r' - 1)\sqrt{8n}^{2/3}] = P[E_{\underline{\tau}_n(r')}], \quad (135)$$

which is negligibly small for  $r' \geq 2$ .

In summary, using (132), (134) and (135) we find that, with probability at least  $1 - 2\exp(-\theta^2/2) - P[E_{\underline{\tau}_n(r)}]$ , we have simultaneously (115) as well as

$$\begin{aligned} V_2 &\leq \theta \cdot 2\sigma_n\sqrt{m} \\ V_1 &\leq \sigma_n^2(k + 2\theta\sqrt{k/2} + \theta^2) \\ -\underline{\tau}_n &\leq r'\sigma_n\sqrt{8n} = r'/r. \end{aligned} \quad (136)$$

(iii) *Sure bound on individual perturbations.* As in earlier proofs, we note that the bulk of the sum of deviations  $|s_i - \tilde{s}_i|^2$  stem from some  $m$  singular values of  $A$  being changed to zero by  $D$ , compensating the constant term  $m$  on the right. More precisely,  $s_i^2 \geq s_n^2 \geq 1 + \underline{\tau}_n$  by (115) while  $\tilde{s}_i = 0$  for  $i = n - m + 1, \dots, n$ . Thus,

$$\sum_{i=n-m+1}^n |s_i - \tilde{s}_i|^2 \geq m(1 + \underline{\tau}_n), \quad (137)$$

implying for any  $1 \leq j \leq n - m$  that

$$\begin{aligned} |s_j - \tilde{s}_j|^2 &\leq \sum_{i=1}^{n-m} |s_i - \tilde{s}_i|^2 \leq V_1 + V_2 + m - m \cdot (1 + \underline{\tau}_n) \\ &= V_1 + V_2 - m\underline{\tau}_n. \end{aligned} \quad (138)$$

(iv) *Probabilistic bound on the perturbation.* Combining the three bounds (136) that hold simultaneously with the indicated probability, we have

$$\begin{aligned} |s_j - \tilde{s}_j|^2 &\leq \sigma_n^2(k + 2\theta\sqrt{k/2} + \theta^2) + \theta \cdot 2\sigma_n\sqrt{m} + mr'\sigma_n\sqrt{8n} \\ &= \sigma_n\sqrt{8n} \left[ \sigma_n\sqrt{8n} \frac{k + \sqrt{2k}\theta + \theta^2}{8n} + 2\theta\sqrt{\frac{m}{8n}} + mr' \right]. \end{aligned} \quad (139)$$

Here, we factored  $\sigma_n\sqrt{8n}$ , since it simplifies to  $1/r$ .

(v) *Simplifications for convenience.* There are many ways to simplify this bound. For convenience of presentation, here is one. Assuming that  $\theta \leq 2\sqrt{n}$ , we have  $\theta \leq 2\sqrt{k}$  and

$$\frac{k + \sqrt{2k}\theta + \theta^2}{8n} \leq \frac{k + \sqrt{8k} + 4k}{8n} \leq \frac{8k}{8n} = v. \quad (140)$$

Using  $\sigma_n \sqrt{8n} v \leq 1, m/n \leq 1$  as well as  $v = m + \tilde{m}(1 - m/n) \geq m$  and  $1 + \theta/\sqrt{2} \leq \theta$ , we obtain

$$|s_j - \tilde{s}_j|^2 \leq \sigma_n \sqrt{8n} \left[ 1 + \theta/\sqrt{2} + r'v \right] \leq \sigma_n \sqrt{8n} [r'v + \theta]. \quad (141)$$

Noting that  $\sqrt{a} - \sqrt{b} = (a - b)/(\sqrt{a} + \sqrt{b}) \geq (a - b)/(1 + \sqrt{b}) \geq (a - b)/2$  whenever  $0 < a, b < 1$  and using (115),  $s_j \geq s_n$  and (135), we have

$$\begin{aligned} \tilde{s}_j &\geq s_j - |s_j - \tilde{s}_j| \geq \sqrt{1 - r'\sigma_n \sqrt{8n}} - \sqrt{\sigma_n \sqrt{8n}(r'v + \theta)} \\ &\geq \frac{1 - \sigma_n \sqrt{8n}(r' + r'v + \theta)}{1 + \sqrt{\sigma_n \sqrt{8n}(r'v + \theta)}} \geq \frac{1}{2} \left( 1 - \sigma_n \sqrt{8n}(r' + r'v + \theta) \right). \end{aligned} \quad (142)$$

Replacing  $\sigma_n \sqrt{8n}$  by  $1/r$ , this becomes

$$\tilde{s}_j \geq \frac{1 - (r' + r'v + \theta)/r}{1 + \sqrt{(r'v + \theta)/r}} \geq \frac{1}{2} (1 - (r' + r'v + \theta)/r). \quad (143)$$

This simple form is for convenience. Tighter bounds can be chosen along the lines above. This takes care of the denominator of  $\kappa$ . For the numerator, we estimate using well-known rules of operator norms (or alternatively Lemma 24 and (71)):

$$\tilde{s}_1 \leq (\|D_1\| \|W_1\| + \|\text{Id}\|) \|D\| \leq \|W_1\| + 1 = s_1(W_1) + 1. \quad (144)$$

(vi) *Let us now turn to (128).* We may proceed as in steps (iii) through (v) with minor modifications. First, replace (135) by

$$P[s_1(W_1) \geq 1/r] = P[s_1(W_1)^2 \geq 1/r^2] \leq P[Y_n \geq \sqrt{2}n^{2/3}] \rightarrow 0. \quad (145)$$

Thus, we have

$$\begin{aligned} V_2 &\leq \theta \cdot 2\sigma_n \sqrt{m} \\ V_1 &\leq \sigma_n^2(k + 2\theta\sqrt{k/2} + \theta^2) \\ s_1(W_1) &\leq 1/r = \sigma_n \sqrt{8n} \end{aligned} \quad (146)$$

with probability at least  $1 - 2 \exp(-\theta^2/2) - P[s_1(W_1) \geq 1/r] \geq 1 - 2 \exp(-\theta^2/2) - P[Y_n \geq \sqrt{2}n^{2/3}]$ , where we use the last expression for convenience of presentation, as it relates to the exceptional event  $E_{s_1(W_1)}$ .

Replacing  $s_i^2 \geq s_n^2 \geq 1 + \underline{\epsilon}_n$  by  $s_i^2 \geq s_n^2 \geq (1 - s_1(W_1))^2 \geq 1 - 2s_1(W_1)$ , we obtain the sure bound similar to (iii)

$$|s_j - \tilde{s}_j|^2 \leq V_1 + V_2 + 2ms_1(W_1) \quad (147)$$

and from this, computing as in (iv) and (v), the bound

$$|s_j - \tilde{s}_j|^2 \leq \sigma_n \sqrt{8n} \left[ 1 + \theta/\sqrt{2} + 2v \right] \leq \sigma_n \sqrt{8n} [2v + \theta], \quad (148)$$

which holds with probability at least  $1 - 2\exp(-\theta^2/2) - P[s_1(W_1) \geq 1]$ .

Noting that  $a - \sqrt{b} = (a^2 - b)/(a + \sqrt{b}) \geq (a^2 - b)/(a + 1) \geq (a^2 - b)/2$  for  $0 < a, b < 1$  and replacing  $\sigma_n \sqrt{8n}$  by  $1/r$ , we obtain, using  $s_j \geq s_n \geq 1 - s_1(W_1) \geq 1 - 1/r$

$$\begin{aligned} \tilde{s}_j &\geq s_j - |s_j - \tilde{s}_j| \geq (1 - 1/r) - \sqrt{(2v + \theta)/r} \\ &\geq \frac{(1 - 1/r)^2 - (2v + \theta)/r}{1 - 1/r + 1} \geq \frac{1}{2} \left( 1 - (2 + 2v + \theta)/r + 1/r^2 \right). \end{aligned} \quad (149)$$

Note that the very first bound  $(1 - 1/r) - \sqrt{(2v + \theta)/r}$  in (149) can be made arbitrarily close to 1 by increasing  $r$ . Dropping the last term  $1/r^2$  for convenience and noting that

$$\tilde{s}_1 \leq \|W_1\| + 1 = s_1(W_1) + 1 \leq 1 + 1/r \quad (150)$$

in the setting of step (vi), the proof is complete.

A similar argument can be given for Uniform random entries instead of Gaussian ones using the Central Limit Theorem and the Berry-Esseen Theorem instead of the tail behavior of Gaussian and Chi-Square variables. The statements and arguments will be more involved.

Also, similarly to the hard bounds it is possible to find bounds on the condition number for leaky-ReLU, implying again that choosing  $\eta$  close to 1 enables us to bound the condition number.

### 6.3 Advantages of Residual Layers

**Volatility of the Piecewise-Quadratic Loss Landscape** The preceding sections provide results on the condition number  $\kappa$  of the relevant matrix  $M(1)D$  for ResNets (recall Corollary 9). They establish conditions under which  $\kappa(M(1)D)$  can be bounded. To be more specific, let us provide some concrete examples assuming just one subsequent layer after the trained one, i.e., set  $p = 1$  in (44).

Consider first a *deterministic setting*. Assume that the entries of the  $n \times n$  weight matrix  $W_1$  are bounded by  $1/(3n)$ . Then the condition number of the relevant matrix  $\kappa(M(1)D)$  amounts to at most 2 for an absolute value activation nonlinearity according to (87). If in addition not more than  $1/4$  of the outputs of the trained layer are affected by the activation, then the condition number  $\kappa(M(1)D)$  is bounded by 8 for a ReLU activation according to (90). Under more technical assumptions, bounds on the condition number are also available for the leaky-ReLU activation. Further bounds are available without restrictions on the outputs but with more severe restrictions on the entries. Note that no a priori bounds on  $\kappa(M(0)D)$  are available, even if a bound on the entries of  $M(0)$  is imposed, as the matrices  $M(0) = \text{diag}(1/n, 1/(2n), \dots, 1/n^2)$  and similar examples demonstrate.

Consider now a *random weight setting* and assume that the standard deviation of the entries of the weight matrix  $W_1$  are bounded on the order of  $1/\sqrt{n}$ . Then the condition number  $\kappa(M(1)D)$  can be bounded for the absolute value activation with known probability using (116). For the ReLU activation, similar bounds are available as for the absolute value. Again, as in the deterministic case, the restrictions on the standard deviation of the weights become more severe the more outputs are affected by the ReLU (see (129)).

To the best of our knowledge, similar bounds are not available for the random matrix  $M(0)D$  corresponding to a ConvNet. Rather, the relevant condition number increases with  $n$ , as we put forward in the proof of the next result (see Theorem 42).

To better appreciate the impact of this difference between ConvNets and ResNets, recall that the loss surface is *piecewise* quadratic, meaning that, close to the optimal values  $W$ , regions with quite different activities of the nonlinearities and thus with different  $D$  and  $D_k$  may exist. Thus, lacking any bounds, the condition number  $\kappa(M(0)D)$  of a ConvNet may very well exhibit large fluctuations, resulting in a volatile shape for the loss landscape. For ResNets, the relevant condition number  $\kappa(M(1)D)$  is not only bounded via Theorem 39, but the bounds can also easily be rendered independent of the activities of the activation nonlinearities, i.e., independent of  $D$  and  $D_1$ . Indeed, obviously we can replace  $v$  everywhere by  $n$  and the statements still hold (cf. (149)).

**Corollary 40** (High Probability Bound, ReLU, for all regions) *Choose  $\theta$  such that  $4 < \theta \leq 2\sqrt{n}$ , and  $r > (2 + 2n + \theta)$ . Let  $\sigma_n$  and  $W_1$  be as in Lemma 36, with Gaussian entries for  $W_1$ . Set  $C = 1 + 1/r$  and  $c = (1 - (2 + 2n + \theta)/r)/2$ .*

*Then, for any  $D_1 \in \mathcal{D}(\tilde{m}, n, 0)$  and  $D \in \mathcal{D}(m, n, 0)$  with  $1 \leq m + \tilde{m}$ , with probability at least  $1 - 2\exp(-\theta^2/2) - P[E_{s_1(W_1)}]$ , we have*

$$c \leq s_i((D_1 W_1 + \text{Id})D) \leq C \quad \text{and} \quad \kappa((D_1 W_1 + \text{Id})D) \leq \frac{C}{c}. \quad (151)$$

The guarantees provided for the relevant matrix for a ResNet in terms of (151) assure that the resulting piecewise-quadratic loss surface cannot change shape too drastically.

**Influence of the Training Data** As mentioned earlier, we focus in our results and discussions mainly on how the weights and nonlinearities of the layers, i.e.,  $M(\rho)D$  influence the loss shape, with the exception of Sect. 3.6. This makes sense, since the influence of the data is literally factored out via Corollary 9 with a contribution to the condition number that is independent of the design choice of ConvNet versus ResNet.

Using mini-batches for learning affects the loss shape in two ways. First, it leads to smaller regions as put forward in Sect. 3.6 (cf. (55)) and thus to potentially increased volatility of the loss landscape. Second, Corollary 11 implies that using mini-batches alleviates the influence of data and features  $z^{(g)}$  (input to the trained layer  $f$ ) in the sense that the “data factor” of the condition number  $\kappa(Q)$  has a less severe impact, since feature values are averaged over  $G$  data points.

As mentioned above, these two kinds of influences of the data on the condition number are the same for both ConvNets and ResNets. Nevertheless, some difference prevails. For ConvNets, the “layer factor”  $\kappa(M(0)D)$  of the condition number  $\kappa(Q)$



leads mini-batches inducing even more loss landscape volatility due to the smaller regions. For ResNets, the uniform bound (151) carries over to mini-batches, and hence the smaller regions do not cause increased condition numbers and loss landscape volatility. *Again, no such guarantees are available for ConvNets.*

**Eccentricity of the Loss Landscape** Intuition from our singular values results says that *the loss landscape of a ResNet should be less eccentric than that of a ConvNet.* We now provide a theoretical result as well as a numerical validation to support this point of view.

**Lemma 41** *Let  $W_1$  be as in Lemma 36, and let  $D_1 \in \mathcal{D}(m, n, -1)$  for some  $m$ . Let  $M_o = D_1 W_1$ . Then, asymptotically as  $n \rightarrow \infty$*

$$P[\kappa(M_o) \geq (1 + s_1(M_o))/(1 - s_1(M_o)) > 0] \rightarrow 1. \quad (152)$$

Recalling that absolute value nonlinearities do not change singular values (recall Corollary 37), we obtain the following.

**Theorem 42** (Residual Links Improve Condition Number) *Let  $W_1$  be as in Lemma 36, and let  $D, D_1 \in \mathcal{D}(m, n, -1)$  for some  $m$ , corresponding to an absolute value activation nonlinearity. Consider a realization of  $W_1$  such that (152) holds. Then, adding a residual link improves the condition number of the matrix relevant for learning*

$$\kappa(D_1 W_1 D) \geq \kappa((D_1 W_1 + \text{Id})D). \quad (153)$$

By Lemma 41, the condition on  $W_1$  is satisfied with probability asymptotically equal to 1.

**Proof** To establish Lemma 41, it suffices to observe that  $s_1(M_o)$  converges almost surely to  $1/(\sqrt{2}r) < 1$ , as well as (117) and that  $\kappa = \kappa(D_1 W_1)$  is asymptotically of the order  $n$ .

To provide a rigorous argument, note that  $D_1 W_1$  is a Gaussian random matrix just like  $W_1$  itself, since  $D_1$  only changes some of the signs of  $W_1$ 's entries. Write  $\kappa = \kappa(D_1 W_1)$  for short.

Applying a result due to Edelman [8, Theorem 6.1] to the matrix  $D_1 W_1$  (multiplying a matrix by a constant does not change its condition number), we find that  $\kappa/n$  converges in distribution to a random variable  $K$  with density (for  $t > 0$ )

$$f_K(t) = \frac{2t + 4}{t^3} \exp(-2/t - 2/t^2). \quad (154)$$

Note that  $f_K$  is  $\mathcal{C}^1(\mathbb{R}^+)$  and  $f_K(t) \rightarrow 0$  as  $t \rightarrow 0$ . It follows that  $P[K < \varepsilon] \rightarrow 0$  as  $\varepsilon \rightarrow 0$ . We conclude that

$$P[\kappa < n\varepsilon] = P[\kappa/n < \varepsilon] \rightarrow P[K < \varepsilon] \quad (155)$$

as  $n \rightarrow \infty$ .

Now let  $\delta > 0$  be arbitrary. Choose  $\varepsilon > 0$  small enough such that  $P[K < \varepsilon] < \delta/2$ , and choose  $n$  large enough such that  $P[\kappa < n\varepsilon] < \delta$ . Increase  $n$  if necessary and choose  $r' < r$  close enough to 1 such that  $n\varepsilon = (1 + 1/r')/(1 - 1/r')$ .

Since  $D_1 W_1$  has the same properties as  $W_1$ , it follows that  $s_1(M_o)$  is equal in distribution to  $s_1(W_1)$  from (113). Then, since  $Y_n$  converges in distribution, we have that

$$\begin{aligned} P[s_1(M_o) > 1/r'] &= P[s_1(M_o)^2 > (1/r')^2] \leq P[1 + Y_n \frac{1}{\sqrt{2} \cdot n^{2/3}} > 2(r/r')^2] \\ &= P[Y_n > (2(r/r')^2 - 1)\sqrt{2} \cdot n^{2/3}] \rightarrow 0 \end{aligned} \quad (156)$$

as  $n \rightarrow \infty$ . Note that if  $s_1(M_o) \leq 1/r'$  then  $(1 + s_1(M_o))/(1 - s_1(M_o)) \leq (1 + 1/r')/(1 - 1/r')$ . As we just showed, the probability of this happening tends to 1 as  $n \rightarrow \infty$ .

Similarly, as in our other proofs, we note that  $P[\kappa < n\varepsilon \text{ or } (1 + s_1(M_o))/(1 - s_1(M_o)) > (1 + 1/r')/(1 - 1/r')]$  is at most equal to the sum of both probabilities and thus at most  $2\delta$  as  $n \rightarrow \infty$ . We conclude that

$$P[\kappa \geq n\varepsilon \text{ and } (1 + s_1(M_o))/(1 - s_1(M_o)) < n\varepsilon] \quad (157)$$

is asymptotically at least  $1 - 2\delta$ . With  $\delta > 0$  arbitrary, this establishes Lemma 41.

The theorem follows now from (117).

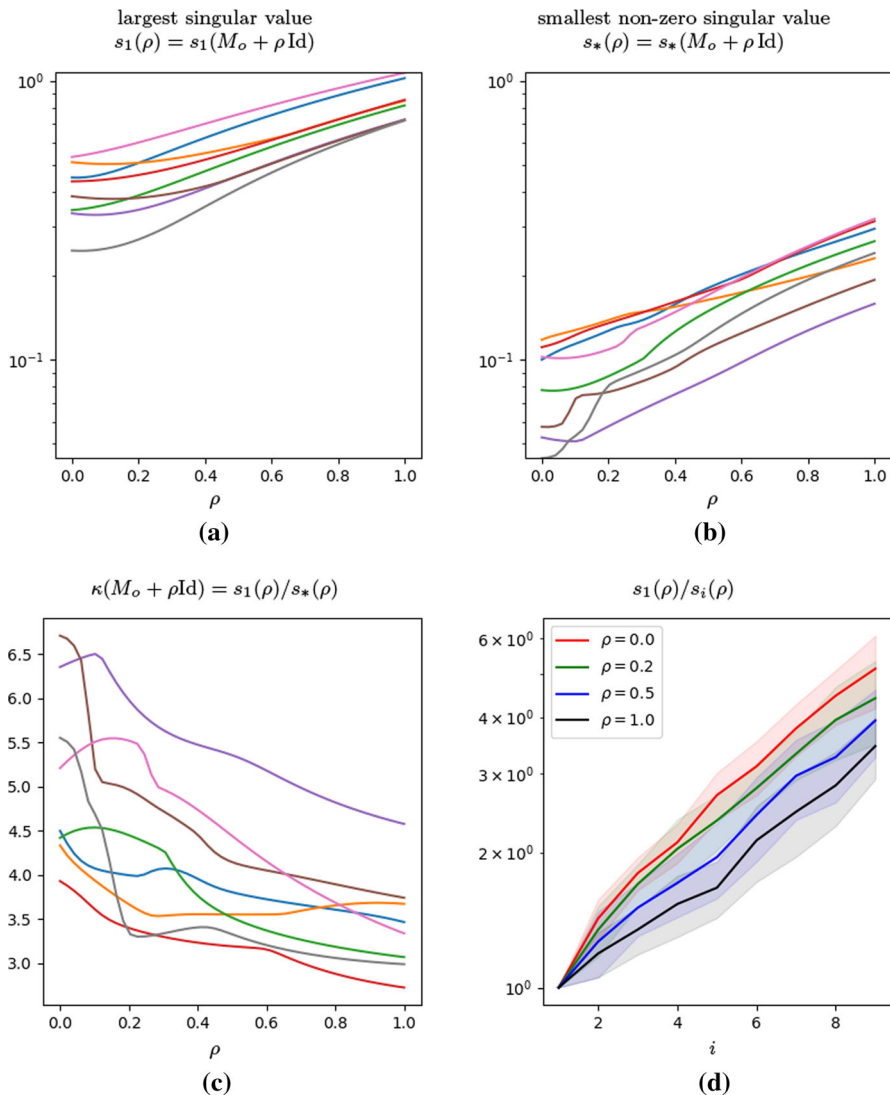
**Numerical Validation** To conclude the paper, we report on the results of a numerical experiment that validates our theoretical results on the singular values of deep networks and the conclusion that the loss landscape of a ResNet is more stable and less eccentric than that of a ConvNet. We constructed a toy four-layer deep network

$$f = f_2 \circ f_1 \circ f_0 \circ f_{-1} \quad (158)$$

with  $f_k$  of the form (1) with the nonlinear activation function  $\phi_k$  the leaky-ReLU with  $\eta = 0.1$  in (35). For  $f_{-1}, f_0, f_2$ , we set  $\rho = 0$  in (1), while for  $f_1$  we let  $\rho$  roam as a free parameter in the range  $\rho \in [0, 1]$  that includes as endpoints a ConvNet layer ( $\rho = 0$ ) and a ResNet layer ( $\rho = 1$ ). The input  $x$  was two-dimensional, and all other inputs, outputs, and bias vectors were 6-dimensional, making the weight matrix  $W_2$  of size  $6 \times 2$  and all other weight matrices of size  $6 \times 6$ . We generated the entries of the input  $x$ , random weights  $W_k$ , and random biases  $b_k$  as iid zero-mean unit Gaussian random variables and tasked the network with predicting the following six statistics of  $x$  by minimizing the  $\ell_2$  norm of the prediction error (squared error):

$$y = [x[1], x[2], \text{mean}(x), \text{max}(x), \text{min}(x), \|x\|_2]^T.$$

For each random  $x, W_k, b_k$  and each value of  $\rho \in [0, 1]$ , we computed the singular values of the matrix  $Q$  from (50) in Lemma 8 that determine the local shape of the piecewise-quadratic loss landscape for the 0-th layer. We conducted 8 realizations of the above experiment and summarize our results in Fig. 8.



**Fig. 8** Singular values of the matrix  $Q$  from (50) in Lemma 8 corresponding to the weight matrix  $W_0$  from the 0-th layer of the four-layer toy deep network (158) using leaky-ReLU activation functions. The singular values are a function of the coefficient  $\rho$  of the subsequent adjustable ResNet layer 1;  $\rho = 0$  corresponds to a ConvNet layer, while  $\rho = 1$  corresponds to a pure ResNet layer. We plot **a** the largest singular value  $s_1$ , **b** the smallest nonzero singular value, and **c** the condition number  $\kappa$  as a function of  $\rho$  for each of the 8 runs of the experiment in a different color. We also plot **d** the average and standard deviation of  $s_1/s_i$  as a function of  $i$  for four values of  $\rho$ . The vertical axes of **a**, **b**, and **d** are on a log scale. We observe that the singular values of the matrix  $Q$  increase steadily with  $\rho$  and the condition number decreases steadily with  $\rho$ , meaning that the eccentricity of the loss landscape also decreases steadily with  $\rho$

Figure 8 not only validates the theory we have developed above, but it also enables us to draw the following concrete conclusions for gradient-based optimization of the family of DNs we have considered. Recall that the loss surface of these networks consists of pieces of (hyper) ellipsoids according to (35) in Sect. 3.4. Independent of the location ( $W$ ) on the ellipsoid, the more eccentric the ellipsoid, the more the gradient points away from the local optimum. Figure 8a, b indicate for increasing  $\rho$  that, since the singular values increase, the image of the unit ball under  $Q$  becomes wider, and thus the level sets of the local loss surface become more narrow [8]. Moreover, Fig. 8c, d show that the local loss surface becomes less eccentric as  $\rho$  increases (a sphere would have a ratio equal to 1). Like our theory, these numerical results strongly recommend ResNets over ConvNets from an optimization perspective.

## 7 Conclusions

In this paper, we have developed new theoretical results on the perturbation of the singular values of a class of matrices relevant to the optimization of an important family of DNs. Applied to deep learning, our results show that the transition from a ConvNet to a network with skip connections (e.g., ResNet, DenseNet) leads to an increase in the size of the singular values of a local linear representation and a decrease in the condition number. For regression and classification tasks employing the classical squared-error loss function, the implications are that the piecewise-quadratic loss surface of a ResNet or DenseNet are less erratic, less eccentric and features local minima that are more accommodating to gradient-based optimization than a ConvNet. Our results are the first to provide an exact quantification of the singular values/conditioning evolution under smooth transition from a ConvNet to a network with skip connections. Our results confirm the empirical observations from [20] depicted in Fig. 1.

Our results also imply that extreme values of the training data affect the optimization, an effect that is tempered when using mini-batches due to the inherent averaging that takes place. Finally, our results shed new light on the impact of different nonlinear activation functions on a DN's singular values, regardless of its architecture. In particular, absolute value activations (such as those employed in the Scattering Network [22]) do not impact the singular values governing the local loss landscape, as opposed to ReLU, which can decrease the singular values.

Our derivations have assumed both deterministic settings and random settings. In the latter, some of our results are asymptotic in the sense that the bounds hold with probability as close to 1 as desired provided that the parameters are set accordingly. We have also provided results in the pre-asymptotic regime that enable one to assess the probability that the bounds will hold.

Finally, while our application focus here has been on analyzing the loss landscape of an important family of DNs, our journey has enabled us to develop novel theoretical results that characterize the evolution of matrix singular values under a range of perturbation regimes that we hope will be of independent interest and broader use.

**Acknowledgements** Richard Baraniuk was supported by NSF Grants CCF-1911094, IIS-1838177, and IIS-1730574; ONR Grants N00014-18-12571, N00014-20-1-2534, and N00014-20-1-2787; AFOSR Grant

FA9550-22-1-0060; and a Vannevar Bush Faculty Fellowship, ONR Grant N00014-18-1-2047. Rudolf Riedi was supported by ONR Grant N00014-20-1-2787. Thanks to Micha Wasem for stimulating discussions leading up to Lemma 21 and to Stephen Wright for pointing out the classification connection in [16].

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Balestrieri, R., Baraniuk, R.G.: A spline theory of deep networks. *Int. Conf. Mach. Learn.* **80**, 374–383 (2018)
- Balestrieri, R., Baraniuk, R.G.: Mad max: affine spline insights into deep learning. *Proc. IEEE* **109**(5), 704–727 (2021)
- Benhamou, E., Atif, J., Laraki, R.: A short note on the operator norm upper bound for sub-Gaussian tailed random matrices. [arXiv:1812.09618](https://arxiv.org/abs/1812.09618) (2019)
- Bruna, J., Mallat, S.: Invariant scattering convolution networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**(8), 1872–1886 (2013)
- Chafai, D., Guedon, O., Lecue, G., Pajor, A.: Singular values of random matrices. In: Chafai, D. (ed.) *Lecture Notes*, chapter 6, pp. 147–184. <https://djalil.chafai.net/docs/sing.pdf> (2009)
- Daubechies, I., DeVore, R., Dym, N., Faigenbaum-Golovin, S., Kovalsky, S.Z., Lin, K.C., Park, J., Petrova, G., Sober, B.: Neural network approximation of refinable functions. [arXiv:2107.13191](https://arxiv.org/abs/2107.13191) (2021)
- Dumaz, L., Virag, B.: The right tail exponent of the Tracy-Widom  $\beta$  distribution. *Ann. l'I.H.P. Probab. Stat.* **49**(4), 915–933 (2013)
- Edelman, A.: Eigenvalues and condition numbers of random matrices. *SIAM J. Matrix Anal. Appl.* **9**(4), 543–560 (1988)
- Gal, Y., Ghahramani, Z.: Dropout as a Bayesian approximation: representing model uncertainty in deep learning. In: *International Conference on Machine Learning*, pp. 1050–1059 (2016)
- Glorot, X., Bordes, A., Bengio, Y.: Deep sparse rectifier neural networks. In: *International Conference on Artificial Intelligence and Statistics*, pp. 315–323 (2011)
- Goodfellow, I., Bengio, Y., Courville, A.: *Deep Learning*. MIT Press, Cambridge (2016)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778 (2016)
- He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In: *European Conference on Computer Vision*, pp. 630–645 (2016)
- Hoffman, A.J., Wielandt, H.W.: The variation of the spectrum of a normal matrix. *Duke Math. J.* **20**, 37–39 (1953)
- Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4700–4708 (2017)
- Hui, L., Belkin, M.: Evaluation of neural architectures trained with square loss vs. cross-entropy in classification tasks. In: *International Conference on Learning Representations* (2021)
- Laurent, B., Massart, P.: Adaptive estimation of a quadratic functional by model selection. *Ann. Stat.* **28**(5), 1302–1338 (2000)
- LeCun, Y., Bengio, Y.: Convolutional networks for images, speech, and time-series. In: *The Handbook of Brain Theory and Neural Networks*, vol. 3361, no. 10 (1995)
- LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *Nature* **521**(7553), 436–444 (2015)
- Li, H., Xu, Z., Taylor, G., Studer, C., Goldstein, T.: Visualizing the loss landscape of neural nets. [arXiv:1712.09913](https://arxiv.org/abs/1712.09913) (2017)
- Li, H., Xu, Z., Taylor, G., Studer, C., Goldstein, T.: Visualizing the loss landscape of neural nets. In: *Advances in Neural Information Processing Systems*, pp. 6389–6399 (2018)

22. Mallat, S.: Group invariant scattering. *Commun. Pure Appl. Math.* **65**(10), 1331–1398 (2012)
23. Tracy, C.A., Widom, H.: Level-spacing distributions and the airy kernel. *Commun. Math. Phys.* **159**(1), 151–174 (1994)
24. Tracy, C.A., Widom, H.: Distribution functions for largest eigenvalues and their applications. *Proc. Int. Congr. Math.* **I**, 587 (2002)
25. Xu, B., Wang, N., Chen T., Li, M.: Empirical evaluation of rectified activations in convolutional network. CoRR. abs/1505.00853 (2015)

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.