

# Controllable Orthogonalization in Training DNNs

Lei Huang<sup>1</sup> Li Liu<sup>1</sup> Fan Zhu<sup>1</sup> Diwen Wan<sup>1,2</sup> Zehuan Yuan<sup>3</sup> Bo Li<sup>4</sup> Ling Shao<sup>1</sup>

<sup>1</sup>Inception Institute of Artificial Intelligence (IIAI), Abu Dhabi, UAE

<sup>2</sup>University of Electronic Science and Technology of China, Chengdu, China

<sup>3</sup>ByteDance AI Lab, Beijing, China

<sup>4</sup>University of Illinois at Urbana-Champaign Illinois, USA

## Abstract

Orthogonality is widely used for training deep neural networks (DNNs) due to its ability to maintain all singular values of the Jacobian close to 1 and reduce redundancy in representation. This paper proposes a computationally efficient and numerically stable orthogonalization method using Newton's iteration (ONI), to learn a layer-wise orthogonal weight matrix in DNNs. ONI works by iteratively stretching the singular values of a weight matrix towards 1. This property enables it to control the orthogonality of a weight matrix by its number of iterations. We show that our method improves the performance of image classification networks by effectively controlling the orthogonality to provide an optimal tradeoff between optimization benefits and representational capacity reduction. We also show that ONI stabilizes the training of generative adversarial networks (GANs) by maintaining the Lipschitz continuity of a network, similar to spectral normalization (SN), and further outperforms SN by providing controllable orthogonality.

## 1. Introduction

Training deep neural networks (DNNs) is often difficult due to the occurrence of vanishing/exploding gradients [7, 14, 47]. Preliminary research [36, 14] has suggested that weight initialization techniques are essential for avoiding these issues. As such, various works have tried to tackle the problem by designing weight matrices that can provide nearly equal variance to activations from different layers [14, 18]. Such a property can be further amplified by orthogonal weight initialization [54, 42, 58], which shows excellent theoretical results in convergence due to its ability to obtain a DNN's *dynamical isometry* [54, 49, 67], i.e. all singular values of the input-output Jacobian are concentrated near 1. The improved performance of orthogonal initialization is empirically observed in [54, 42, 49, 65] and it makes training even 10,000-layer DNNs possible [65]. However, the initial orthogonality can be broken down and is not necessarily sustained throughout training [66].

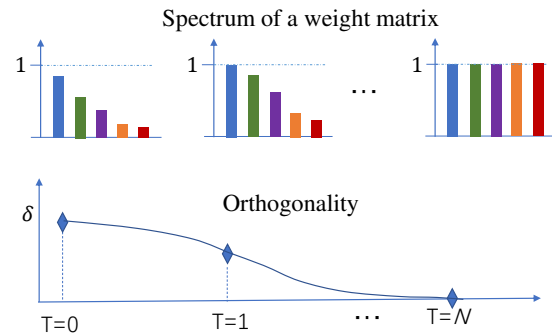


Figure 1. ONI controls a weight matrix's magnitude of orthogonality (measured as  $\delta = \|\mathbf{W}\mathbf{W}^T - \mathbf{I}\|_F$ ), by iteratively stretching its singular values towards 1.

Previous works have tried to maintain the orthogonal weight matrix by imposing an additional orthogonality penalty on the objective function, which can be viewed as a 'soft orthogonal constraint' [47, 61, 66, 5, 3]. These methods show improved performance in image classification [66, 71, 37, 5], resisting attacks from adversarial examples [12], neural photo editing [11] and training generative adversarial networks (GAN) [10, 43]. However, the introduced penalty works like a pure regularization, and whether or not the orthogonality is truly maintained or training benefited is unclear. Other methods have been developed to directly solve the 'hard orthogonal constraint' [61, 5], either by Riemannian optimization [46, 17] or by orthogonal weight normalization [62, 24]. However, Riemannian optimization often suffers from training instability [17, 24], while orthogonal weight normalization [24] requires computationally expensive eigen decomposition, and the necessary back-propagation through this eigen decomposition may suffer from numerical instability, as shown in [30, 40].

We propose to perform orthogonalization by Newton's iteration (ONI) [41, 8, 26] to learn an exact orthogonal weight matrix, which is computationally efficient and numerically stable. To further speed up the convergence of Newton's iteration, we propose two techniques: 1) we perform centering to improve the conditioning of the proxy matrix; 2) we explore a more compact spectral bounding method to make the initial singular value of the proxy matrix closer to 1.

We provide an insightful analysis and show that ONI works by iteratively stretching the singular values of the weight matrix towards 1 (Figure 1). This property makes ONI work well even if the weight matrix is singular (with multiple zero singular values), under which the eigen-decomposition based method [24] often suffers from numerical instability [30, 40]. Moreover, we show that controlling orthogonality is necessary to balance the increase in optimization and reduction in representational capacity, and ONI can elegantly achieve this through its iteration number (Figure 1). Besides, ONI provides a unified solution for the row/column orthogonality, regardless of whether the weight matrix’s output dimension is smaller or larger than the input.

We also address practical strategies for effectively learning orthogonal weight matrices in DNNs. We introduce a constant of  $\sqrt{2}$  to initially scale the orthonormal weight matrix so that the *dynamical isometry* [54] can be well maintained for deep ReLU networks [45]. We conduct extensive experiments on multilayer perceptrons (MLPs) and convolutional neural networks (CNNs). Our proposed method benefits the training and improves the test performance over multiple datasets, including ImageNet [51]. We also show that our method stabilizes the training of GANs and achieves improved performance on unsupervised image generation, compared to the widely used spectral normalization [43].

## 2. Related Work

Orthogonal filters have been extensively explored in signal processing since they are capable of preserving activation energy and reducing redundancy in representation [72]. Saxe *et al.* [54] introduced an orthogonal weight matrix for DNNs and showed that it achieves approximate *dynamical isometry* [54] for deep linear neural networks, therefore significantly improving the optimization efficiency [42, 58]. Pennington *et al.* [49] further found that the nonlinear sigmoid network can also obtain *dynamical isometry* when combined with orthogonal weight initialization [49, 65, 67].

Research has also been conducted into using *orthogonal matrices* to avoid the gradient vanishing/explosion problems in recurrent neural networks (RNNs). These methods mainly focus on constructing square orthogonal matrices/unitary matrices for the hidden-to-hidden transformations in RNNs [4, 62, 13, 61, 27, 32, 21]. This is done by either constructing a decomposed unitary weight matrix with a restricted [4] or full representational capability [62, 21], or by using soft constraints [61]. Different from these methods requiring a square weight matrix and limited to hidden-to-hidden transformations in RNNs, our method is more general and can adapt to situations where the weight matrix is not square.

Our method is related to the methods that impose orthogonal penalties on the loss functions [47, 61, 5]. Most works propose to use soft orthogonality regularization under the standard Frobenius norm [47, 61, 5], though other alternative orthogonal penalties were explored in [5]. There are also methods that propose to bound the singular values with

periodical projection [31]. Our method targets at solving the ‘hard constraint’ and providing controllable orthogonality.

One way to obtain exact orthogonality is through Riemannian optimization methods [46, 17]. These methods usually require a retract operation [2] to project the updated weight back to the Stiefel manifold [46, 17], which may result in training instability for DNNs [17, 24]. Our method avoids this by employing re-parameterization to construct the orthogonal matrix [24]. Our work is closely related to orthogonal weight normalization [24], which also uses re-parameterization to design an orthogonal transformation. However, [24] solves the problem by computationally expensive eigen decomposition and may result in numeric instability [30, 40]. We use Newton’s iteration [41, 8], which is more computationally efficient and numerically stable. We further argue that fully orthogonalizing the weight matrix limits the network’s learning capacity, which may result in degenerated performance [43, 10]. Another related work is spectral normalization [43], which uses reparametrization to bound only the maximum eigenvalue as 1. Our method can effectively interpolate between spectral normalization and full orthogonalization, by altering the iteration number.

Newton’s iteration has also been employed in DNNs for constructing bilinear/second-order pooling [40, 38], or whitening the activations [26]. [40] and [38] focused on calculating the square root of the covariance matrix, while our method computes the square root inverse of the covariance matrix, like the work in [26]. However, our work has several main differences from [26]: 1) In [26], they aimed to whiten the activation [25] over batch data using Newton’s iteration, while our work seeks to learn the orthogonal weight matrix, which is an entirely different research problem [29, 52, 25]; 2) We further improve the convergence speed compared to the Newton’s iteration proposed in [26] by providing more compact bounds; 3) Our method can maintain the Lipschitz continuity of the network and thus has potential in stabilizing the training of GANs [43, 10]. It is unclear whether or not the work in [26] has such a property, since it is data-dependent normalization [29, 43, 10].

## 3. Proposed Method

Given the dataset  $D = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^M$  composed of an input  $\mathbf{x}_i \in \mathbb{R}^d$  and its corresponding labels  $\mathbf{y}_i \in \mathbb{R}^c$ , we represent a standard feed-forward neural network as a function  $f(\mathbf{x}; \theta)$  parameterized by  $\theta$ .  $f(\mathbf{x}; \theta)$  is a composition of  $L$  simple nonlinear functions. Each of these consists of a linear transformation  $\hat{\mathbf{h}}^l = \mathbf{W}^l \mathbf{h}^{l-1} + \mathbf{b}^l$  with learnable weights  $\mathbf{W}^l \in \mathbb{R}^{n_l \times d_l}$  and biases  $\mathbf{b}^l \in \mathbb{R}^{n_l}$ , followed by an element-wise nonlinearity:  $\mathbf{h}^l = \varphi(\hat{\mathbf{h}}^l)$ . Here  $l \in \{1, 2, \dots, L\}$  indexes the layers. We denote the learnable parameters as  $\theta = \{\mathbf{W}^l, \mathbf{b}^l | l = 1, 2, \dots, L\}$ . Training neural networks involves minimizing the discrepancy between the desired output  $\mathbf{y}$  and the predicted output  $f(\mathbf{x}; \theta)$ , described by a loss function  $\mathcal{L}(\mathbf{y}, f(\mathbf{x}; \theta))$ . Thus, the optimization objective is:  $\theta^* = \arg \min_{\theta} \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \in D} [\mathcal{L}(\mathbf{y}, f(\mathbf{x}; \theta))]$ .

---

**Algorithm 1** Orthogonalization by Newton’s Iteration (ONI).

---

- 1: **Input:** proxy parameters  $\mathbf{Z} \in \mathbb{R}^{n \times d}$  and iteration numbers  $T$ .
  - 2: Bounding  $\mathbf{Z}$ ’s singular values:  $\mathbf{V} = \frac{\mathbf{Z}}{\|\mathbf{Z}\|_F}$ .
  - 3: Calculate covariance matrix:  $\mathbf{S} = \mathbf{V}\mathbf{V}^T$ .
  - 4:  $\mathbf{B}_0 = \mathbf{I}$ .
  - 5: **for**  $t = 1$  to  $T$  **do**
  - 6:    $\mathbf{B}_t = \frac{3}{2}\mathbf{B}_{t-1} - \frac{1}{2}\mathbf{B}_{t-1}^3\mathbf{S}$ .
  - 7: **end for**
  - 8:  $\mathbf{W} = \mathbf{B}_T\mathbf{V}$
  - 9: **Output:** orthogonalized weight matrix:  $\mathbf{W} \in \mathbb{R}^{n \times d}$ .
- 




Figure 2. Convergence behaviors of the proposed Orthogonalization by Newton’s Iteration. The entries of proxy matrix  $\mathbf{Z} \in \mathbb{R}^{64 \times 256}$  are sampled from the Gaussian distribution  $N(3, 1)$ . We show (a) the magnitude of the orthogonality, measured as  $\delta = \|\mathbf{W}\mathbf{W}^T - \mathbf{I}\|_F$ , with respect to the iterations and (b) the distribution (log scale) of the eigenvalues of  $\mathbf{W}\mathbf{W}^T$  with different iterations.

### 3.1. Preliminaries

This paper starts with learning orthogonal filter banks (row orthogonalization of a weight matrix) for deep neural networks (DNNs). We assume  $n \leq d$  for simplicity, and will discuss the situation where  $n > d$  in Section 3.4. This problem is formulated in [24] as an optimization with layer-wise orthogonal constraints, as follows:

$$\begin{aligned} \theta^* &= \arg \min_{\theta} \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \in D} [\mathcal{L}(\mathbf{y}, f(\mathbf{x}; \theta))] \\ \text{s.t. } &\mathbf{W}^l (\mathbf{W}^l)^T = \mathbf{I}, \quad l = 1, 2, \dots, L. \end{aligned} \quad (1)$$

To solve this problem directly, Huang *et al.* [24] proposed to use the proxy parameters  $\mathbf{V}$  and construct the orthogonal weight matrix  $\mathbf{W}$  by minimizing them in a Frobenius norm over the feasible transformation sets, where the objective is:

$$\begin{aligned} \min_{\phi(\mathbf{V})} & \text{tr} \left( (\mathbf{W} - \mathbf{V})(\mathbf{W} - \mathbf{V})^T \right) \\ \text{s.t. } & \mathbf{W} = \phi(\mathbf{V}) \text{ and } \mathbf{W}\mathbf{W}^T = \mathbf{I}. \end{aligned} \quad (2)$$

They solved this in a closed-form, with the orthogonal transformation as:

$$\mathbf{W} = \phi(\mathbf{V}) = \mathbf{D}\mathbf{\Lambda}^{-1/2}\mathbf{D}^T\mathbf{V}, \quad (3)$$

where  $\mathbf{\Lambda} = \{\lambda_1, \dots, \lambda_n\}$  and  $\mathbf{D}$  are the eigenvalues and corresponding eigenvectors of the covariance matrix  $\mathbf{S} = \mathbf{V}\mathbf{V}^T$ . Given the gradient  $\frac{\partial \mathcal{L}}{\partial \mathbf{W}}$ , back-propagation must pass through the orthogonal transformation to calculate  $\frac{\partial \mathcal{L}}{\partial \mathbf{V}}$  for updating  $\mathbf{V}$ . The closed formulation is concise; however, it encounters the following problems in practice: 1) Eigen decomposition is required, which is computationally expensive, especially on GPU devices [40]; 2) The back-propagation through the eigen decomposition requires the element-wise multiplication of a matrix  $\mathbf{K}$  [24], whose elements are given

---

**Algorithm 2** ONI with Acceleration.

---

- 1: **Input:** proxy parameters  $\mathbf{Z} \in \mathbb{R}^{n \times d}$  and iteration numbers  $N$ .
  - 2: Centering:  $\mathbf{Z}_c = \mathbf{Z} - \frac{1}{d}\mathbf{Z}\mathbf{1}\mathbf{1}^T$ .
  - 3: Bounding  $\mathbf{Z}$ ’s singular values:  $\mathbf{V} = \frac{\mathbf{Z}_c}{\sqrt{\|\mathbf{Z}_c\mathbf{Z}_c^T\|_F}}$ .
  - 4: Execute Step. 3 to 8 in Algorithm 1.
  - 5: **Output:** orthogonalized weight matrix:  $\mathbf{W} \in \mathbb{R}^{n \times d}$ .
- 

by  $\mathbf{K}_{i,j} = \frac{1}{(\lambda_i - \lambda_j)}$ , where  $i \neq j$ . This may cause numerical instability when there exists equal eigenvalues of  $\mathbf{S}$ , which is discussed in [30, 40] and observed in our preliminary experiments, especially for high-dimensional space.

We observe that the solution of Eqn. 2 can be represented as  $\mathbf{W} = \mathbf{S}^{-\frac{1}{2}}\mathbf{V}$ , where  $\mathbf{S}^{-\frac{1}{2}}$  can be computed by Newton’s iteration [41, 8, 26], which avoids eigen decomposition in the forward pass and potential numerical instability during the back-propagation.

### 3.2. Orthogonalization by Newton’s Iteration

Newton’s iteration calculates  $\mathbf{S}^{-\frac{1}{2}}$  as follows:

$$\begin{cases} \mathbf{B}_0 = \mathbf{I} \\ \mathbf{B}_t = \frac{1}{2}(3\mathbf{B}_{t-1} - \mathbf{B}_{t-1}^3\mathbf{S}), \quad t = 1, 2, \dots, T, \end{cases} \quad (4)$$

where  $T$  is the number of iterations. Under the condition that  $\|\mathbf{I} - \mathbf{S}\|_2 < 1$ ,  $\mathbf{B}_T$  will converge to  $\mathbf{S}^{-\frac{1}{2}}$  [8, 26].

$\mathbf{V}$  in Eqn. 2 can be initialized to ensure that  $\mathbf{S} = \mathbf{V}\mathbf{V}^T$  initially satisfies the convergence condition, e.g. ensuring  $0 \leq \sigma(\mathbf{V}) \leq 1$ , where  $\sigma(\mathbf{V})$  are the singular values of  $\mathbf{V}$ . However, the condition is very likely to be violated when training DNNs, since  $\mathbf{V}$  varies.

To address this problem, we propose to maintain another proxy parameter  $\mathbf{Z}$  and conduct a transformation  $\mathbf{V} = \phi_N(\mathbf{Z})$ , such that  $0 \leq \sigma(\mathbf{V}) \leq 1$ , inspired by the re-parameterization method [53, 24]. One straightforward way to ensure  $0 \leq \sigma(\mathbf{V}) \leq 1$  is to divide the spectral norm of  $\mathbf{Z}$ , like the spectral normalization method does [43]. However, it is computationally expensive to accurately calculate the spectral norm, since singular value decomposition is required. We thus propose to divide the Frobenius norm of  $\mathbf{Z}$  to perform spectral bounding:

$$\mathbf{V} = \phi_N(\mathbf{Z}) = \frac{\mathbf{Z}}{\|\mathbf{Z}\|_F}. \quad (5)$$

It’s easy to demonstrate that Eqn. 5 satisfies the convergence condition of Newton’s iteration and we show that this method is equivalent to the Newton’s iteration proposed in [26] (See the [supplementary materials \(SM\)](#) for details). Algorithm 1 describes the proposed method, referred to as Orthogonalization by Newton’s Iteration (ONI), and its corresponding back-propagation is shown in the [SM](#). We find that Algorithm 1 converges well (Figure 2). However, the concern is the speed of convergence, since 10 iterations are required in order to obtain a good orthogonalization. We thus further explore methods to speed up the convergence of ONI.




Figure 3. Analysis of speeding up Newton’s iteration. The entries of proxy matrix  $\mathbf{Z} \in \mathbb{R}^{64 \times 256}$  are sampled from the Gaussian distribution  $N(3, 1)$ . (a) Comparison of convergence; (b) Comparison of the distribution of the eigenvalues of  $\mathbf{W}\mathbf{W}^T$  at iteration  $t = 1$ .

### 3.3. Speeding up Convergence of Newton’s Iteration

Our Newton’s iteration proposed for obtaining orthogonal matrix  $\mathbf{W}$  works by iteratively stretching the singular values of  $\mathbf{V}$  towards 1, as shown in Figure 2 (b). The speed of convergence depends on how close the singular values of  $\mathbf{V}$  initially are to 1 [8]. We observe that the following factors benefit the convergence of Newton’s iteration: 1) The singular values of  $\mathbf{Z}$  have a balanced distribution, which can be evaluated by the condition number of the matrix  $\mathbf{Z}$ ; 2) The singular values of  $\mathbf{V}$  should be as close to 1 as possible after spectral bounding (Eqn. 5).

**Centering** To achieve more balanced distributions for the eigenvalues of  $\mathbf{Z}$ , we perform a centering operation over the proxy parameters  $\mathbf{Z}$ , as follows

$$\mathbf{Z}_c = \mathbf{Z} - \frac{1}{d} \mathbf{Z} \mathbf{1} \mathbf{1}^T. \quad (6)$$

The orthogonal transformation is then performed over the centered parameters  $\mathbf{Z}_c$ . As shown in [36, 55], the covariance matrix of centered matrix  $\mathbf{Z}_c$  is better conditioned than  $\mathbf{Z}$ . We also experimentally observe that orthogonalization over centered parameters  $\mathbf{Z}_c$  (indicated as ‘ONI+Center’) produces larger singular values on average at the initial stage (Figure 3 (b)), and thus converges faster than the original ONI (Figure 3 (a)).

**Compact Spectral Bounding** To achieve larger singular values of  $\mathbf{V}$  after spectral bounding, we seek a more compact spectral bounding factor  $f(\mathbf{Z})$  such that  $f(\mathbf{Z}) > \|\mathbf{Z}\|_F$  and  $\mathbf{V}$  satisfies the convergence condition. We find that  $f(\mathbf{Z}) = \sqrt{\|\mathbf{Z}\mathbf{Z}^T\|_F}$  satisfies the requirements, which is demonstrated in the SM. We thus perform spectral bounding based on the following formulation:

$$\mathbf{V} = \phi_N(\mathbf{Z}) = \frac{\mathbf{Z}}{\sqrt{\|\mathbf{Z}\mathbf{Z}^T\|_F}}. \quad (7)$$

More compact spectral bounding (CSB) is achieved using Eqn. 7, compared to Eqn. 5. For example, assuming that  $\mathbf{Z}$  has  $g$  equivalent singular values, the initial singular values of  $\mathbf{V}$  after spectral bounding will be  $\frac{1}{\sqrt[4]{n}}$  when using Eqn. 7, while  $\frac{1}{\sqrt{n}}$  when using Eqn. 5. We also experimentally observe that using Eqn. 7 (denoted with ‘+CSB’ in Figure 3) results in a significantly faster convergence.

Algorithm 2 describes the accelerated ONI method with centering and more compact spectral bounding (Eqn. 7).




Figure 4. Unified row and column orthogonalization. The entries of proxy matrix  $\mathbf{Z} \in \mathbb{R}^{64 \times 32}$  are sampled from the Gaussian distribution  $N(0, 1)$ . (a) Orthogonalization comparison between  $\delta_{Row} = \|\mathbf{W}\mathbf{W}^T - \mathbf{I}\|_F$  and  $\delta_{Column} = \|\mathbf{W}^T \mathbf{W} - \mathbf{I}\|_F$ ; (b) The distribution of the eigenvalues of  $\mathbf{W}\mathbf{W}^T$  with different iterations.

### 3.4. Unified Row and Column Orthogonalization

In previous sections, we assume  $n \leq d$ , and obtain an orthogonalization solution. One question that remains is how to handle the situation when  $n > d$ . When  $n > d$ , the rows of  $\mathbf{W}$  cannot be orthogonal, because the rank of  $\mathbf{W}$  is less than/equal to  $d$ . Under this situation, full orthogonalization using the eigenvalue decomposition based solution (Eqn. 3) may cause numerical instability, since there exists at least  $n - d$  zero eigenvalues for the covariance matrix. These zero eigenvalues specifically lead to numerical instability during back-propagation (when element-wisely multiplying the scaling matrix  $\mathbf{K}$ , as discussed in Section 3.1).

Our orthogonalization solution by Newton’s iteration can avoid such problems, since there are no operations relating to dividing the eigenvalues of the covariance matrix. Therefore, our ONI can solve Eqn. 2 under the situation  $n > d$ . More interestingly, our method can achieve column orthogonality for the weight matrix  $\mathbf{W}$  (that is,  $\mathbf{W}^T \mathbf{W} = \mathbf{I}$ ) by solving Eqn. 2 directly under  $n > d$ . Figure 4 shows the convergence behaviors of the row and column orthogonalizations. We observe ONI stretches the non-zero eigenvalues of the covariance matrix  $\mathbf{S}$  towards 1 in an iterative manner, and thus equivalently stretches the singular values of the weight matrix  $\mathbf{V}$  towards 1. Therefore, it ensures column orthogonality under the situation  $n > d$ . Our method unifies the row and column orthogonalizations, and we further show in Section 3.5 that they both benefit in preserving the norm/distribution of the activation/gradient when training DNNs.

Note that, for  $n > d$ , Huang *et al.* [24] proposed the group based methods by dividing the weights  $\{w_i\}_{i=1}^n$  into groups of size  $N_G \leq d$  and performing orthogonalization over each group, such that the weights in each group are row orthogonal. However, such a method cannot ensure the whole matrix  $\mathbf{W}$  to be either row or column orthogonal (See the SM for details).

### 3.5. Controlling Orthogonality

One remarkable property of the orthogonal matrix is that it can preserve the norm and distribution of the activation for a linear transformation, given appropriate assumptions. Such properties are described in the following theorem.

**Theorem 1.** Let  $\hat{\mathbf{h}} = \mathbf{W}\mathbf{x}$ , where  $\mathbf{W}\mathbf{W}^T = \mathbf{I}$  and  $\mathbf{W} \in \mathbb{R}^{n \times d}$ .



Assume: (1)  $\mathbb{E}_{\mathbf{x}}(\mathbf{x}) = \mathbf{0}$ ,  $\text{cov}(\mathbf{x}) = \sigma_1^2 \mathbf{I}$ , and (2)  $\mathbb{E}_{\frac{\partial \mathcal{L}}{\partial \mathbf{h}}}(\frac{\partial \mathcal{L}}{\partial \mathbf{h}}) = \mathbf{0}$ ,  $\text{cov}(\frac{\partial \mathcal{L}}{\partial \mathbf{h}}) = \sigma_2^2 \mathbf{I}$ . If  $n = d$ , we have the following properties: (1)  $\|\hat{\mathbf{h}}\| = \|\mathbf{x}\|$ ; (2)  $\mathbb{E}_{\hat{\mathbf{h}}}(\hat{\mathbf{h}}) = \mathbf{0}$ ,  $\text{cov}(\hat{\mathbf{h}}) = \sigma_1^2 \mathbf{I}$ ; (3)  $\|\frac{\partial \mathcal{L}}{\partial \mathbf{x}}\| = \|\frac{\partial \mathcal{L}}{\partial \hat{\mathbf{h}}}\|$ ; (4)  $\mathbb{E}_{\frac{\partial \mathcal{L}}{\partial \mathbf{x}}}(\frac{\partial \mathcal{L}}{\partial \mathbf{x}}) = \mathbf{0}$ ,  $\text{cov}(\frac{\partial \mathcal{L}}{\partial \mathbf{x}}) = \sigma_2^2 \mathbf{I}$ . In particular, if  $n < d$ , property (2) and (3) hold; if  $n > d$ , property (1) and (4) hold.

The proof is provided in the SM. Theorem 1 shows the benefits of orthogonality in preventing gradients from exploding/vanishing, from an optimization perspective. Besides, the orthonormal weight matrix can be viewed as the embedded Stiefel manifold  $\mathcal{O}^{n \times d}$  with a degree of freedom  $nd - n(n+1)/2$  [1, 24], which regularizes the neural networks and can improve the model’s generalization [1, 24].

However, this regularization may harm the representational capacity and result in degenerated performance, as shown in [24] and observed in our experiments. Therefore, controlling orthogonality is necessary to balance the increase in optimization benefit and reduction in representational capacity, when training DNNs. Our ONI can effectively control orthogonality using different numbers of iterations.

### 3.6. Learning Orthogonal Weight Matrices in DNNs

Based on Algorithm 2 and its corresponding backward pass, we can wrap our method in linear modules [53, 24], to learn filters/weights with orthogonality constraints for DNNs. After training, we calculate the weight matrix  $\mathbf{W}$  and save it for inference, as in the standard module.

**Layer-wise Dynamical Isometry** Theorem 1 shows that the orthogonal matrix has remarkable properties for preserving the norm/distributions of activations during the forward and backward passes, for linear transformations. However, in practice, we need to consider the nonlinearity function as well. Here, we show that we can use an additional constant to scale the magnitude of the weight matrix for ReLU nonlinearity [45], such that the output-input Jacobian matrix of each layer has dynamical isometry.

**Theorem 2.** Let  $\mathbf{h} = \max(0, \mathbf{W}\mathbf{x})$ , where  $\mathbf{W}\mathbf{W}^T = \sigma^2 \mathbf{I}$  and  $\mathbf{W} \in \mathbb{R}^{n \times d}$ . Assume  $\mathbf{x}$  is a normal distribution with  $\mathbb{E}_{\mathbf{x}}(\mathbf{x}) = \mathbf{0}$ ,  $\text{cov}(\mathbf{x}) = \mathbf{I}$ . Denote the Jacobian matrix as  $\mathbf{J} = \frac{\partial \mathbf{h}}{\partial \mathbf{x}}$ . If  $\sigma^2 = 2$ , we have  $\mathbb{E}_{\mathbf{x}}(\mathbf{J}\mathbf{J}^T) = \mathbf{I}$ .

The proof is shown in the SM. We propose to multiply the orthogonal weight matrix  $\mathbf{W}$  by a factor of  $\sqrt{2}$  for networks with ReLU activation. We experimentally show this improves the training efficiency in Section 4.1. Note that Theorems 1 and 2 are based on the assumption that the layer-wise input is Gaussian. Such a property can be approximately satisfied using batch normalization (BN) [29]. Besides, if we apply BN before the linear transformation, there is no need to apply it again after the linear module, since the normalized property of BN is preserved according to Theorem 1. We experimentally show that such a process improves performance in Section 4.1.3.

**Learnable Scalar** Following [24], we relax the constraint of orthonormal to orthogonal, with  $\mathbf{W}^T \mathbf{W} = \Lambda$ , where  $\Lambda$  is the diagonal matrix. This can be viewed as the orthogonal




Figure 5. Effects of maintaining orthogonality. Experiments are performed on a 10-layer MLP. (a) The training (solid lines) and testing (dashed lines) errors with respect to the training epochs; (b) The distribution of eigenvalues of the weight matrix  $\mathbf{W}$  of the 5th layer, at the 200 iteration.




Figure 6. Effects of scaling the orthogonal weights. ‘-NS’ indicates orthogonalization without scaling by  $\sqrt{2}$ . We evaluate the training (solid lines) and testing (dashed lines) errors on (a) a 6-layer MLP and (b) a 20-layer MLP.

filters having different contributions to the activations. To achieve this, we propose to use a learnable scalar parameter  $g$  to fine-tune the norm of each filter [53, 24].

**Convolutional Layer** With regards to the convolutional layer parameterized by weights  $\mathbf{W}^C \in \mathbb{R}^{n \times d \times F_h \times F_w}$ , where  $F_h$  and  $F_w$  are the height and width of the filter, we reshape  $\mathbf{W}^C$  as  $\mathbf{W} \in \mathbb{R}^{n \times p}$ , where  $p = d \cdot F_h \cdot F_w$ , and the orthogonalization is executed over the unrolled weight matrix  $\mathbf{W} \in \mathbb{R}^{n \times (d \cdot F_h \cdot F_w)}$ .

**Computational Complexity** Consider a convolutional layer with filters  $\mathbf{W} \in \mathbb{R}^{n \times d \times F_h \times F_w}$ , and  $m$  mini-batch data  $\{\mathbf{x}_i \in \mathbb{R}^{d \times h \times w}\}_{i=1}^m$ . The computational cost of our method, coming mainly from the Lines 3, 6 and 8 in Algorithm 1, is  $2n^2 d F_h F_w + 3N n^3$  for each iteration during training. The relative cost of ONI over the constitutional layer is  $\frac{2n}{mhw} + \frac{3N n^2}{mdhw F_h F_w}$ . During inference, we use the orthogonalized weight matrix  $\mathbf{W}$ , and thus do not introduce additional computational or memory costs. We provide the wall-clock times in the SM.

## 4. Experiments

### 4.1. Image Classification

We evaluate our ONI on the Fashion-MNIST [64], CIFAR-10 [34] and ImageNet [51] datasets. We provide an ablation study on the iteration number  $T$  of ONI in Section 4.1.4. Due to space limitations, we only provide essential components of the experimental setup; for more details, please refer to the SM. The code is available at <https://github.com/huangleiBuaa/ONI>.

| methods         | g=2,k=1     | g=2,k=2     | g=2,k=3     | g=3,k=1     | g=3,k=2     | g=3,k=3     | g=4,k=1     | g=4,k=2     | g=4,k=3     |
|-----------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| plain           | 11.34       | 9.84        | 9.47        | 10.32       | 8.73        | 8.55        | 10.66       | 9.00        | 8.43        |
| WN              | 11.19       | 9.55        | 9.49        | 10.26       | 9.26        | 8.19        | 9.90        | 9.33        | 8.90        |
| OrthInit        | 10.57       | 9.49        | 9.33        | 10.34       | 8.94        | 8.28        | 10.35       | 10.6        | 9.39        |
| OrthReg         | 12.01       | 10.33       | 10.31       | 9.78        | 8.69        | 8.61        | 9.39        | 7.92        | 7.24        |
| OLM-1 [24]      | 10.65       | 8.98        | 8.32        | 9.23        | 8.05        | 7.23        | 9.38        | 7.45        | 7.04        |
| OLM- $\sqrt{2}$ | 10.15       | 8.32        | 7.80        | 8.74        | 7.23        | 6.87        | <b>8.02</b> | 6.79        | 6.56        |
| ONI             | <b>9.95</b> | <b>8.20</b> | <b>7.73</b> | <b>8.64</b> | <b>7.16</b> | <b>6.70</b> | 8.27        | <b>6.72</b> | <b>6.52</b> |

Table 1. Test errors (%) on VGG-style networks for CIFAR-10 classification. The results are averaged over three independent runs.

#### 4.1.1 MLPs on Fashion-MNIST

We use an MLP with a ReLU activation [45], and vary the depth. The number of neurons in each layer is 256. We employ stochastic gradient descent (SGD) optimization with a batch size of 256, and the learning rates are selected based on the validation set (5,000 samples from the training set) from  $\{0.05, 0.1, 0.5, 1\}$ .

**Maintaining Orthogonality** We first show that maintaining orthogonality can improve the training performance. We compare two baselines: 1) ‘plain’, the original network; and 2) ‘OrthInit’, in which the orthogonal initialization [54] is used. The training performances are shown in Figure 5 (a). We observe orthogonal initialization can improve the training efficiency in the initial phase (comparing with ‘plain’), after which the benefits of orthogonality degenerate (comparing with ‘ONI’) due to the updating of weights (Figure 5 (b)).

**Effects of Scaling** We experimentally show the effects of initially scaling the orthogonal weight matrix by a factor of  $\sqrt{2}$ . We also apply this technique to the ‘OLM’ method [24], in which the orthogonalization is solved by eigen decomposition. We refer to ‘OLM-NS’/‘ONI-NS’ as the ‘OLM’/‘ONI’ without scaling by  $\sqrt{2}$ . The results are shown in Figure 6. We observe that the scaling technique has no significant effect on shallow neural networks, *e.g.*, the 6-layer MLP. However, for deeper neural networks, it produces significant performance boosts. For example, for the 20-layer MLP, neither ‘OLM’ nor ‘ONI’ can converge without the additional scaling factors, because the activation and gradient exponentially vanish (see the SM). Besides, our ‘ONI’ has a nearly identical performance compared to ‘OLM’, which indicates the effectiveness of our approximate orthogonalization with few iterations (*e.g.* 5).

#### 4.1.2 CNNs on CIFAR-10

**VGG-Style Networks** Here, we evaluate ONI on VGG-style neural networks with  $3 \times 3$  convolutional layers. The network starts with a convolutional layer of  $32k$  filters, where  $k$  is the varying width based on different configurations. We then sequentially stack three blocks, each of which has  $g$  convolutional layers with filter numbers of  $32k$ ,  $64k$  and  $128k$ , respectively. We vary the depth with  $g$  in  $\{2, 3, 4\}$  and the width with  $k$  in  $\{1, 2, 3\}$ . We use SGD with a momentum of 0.9 and batch size of 128. The best initial learning rate is chosen from  $\{0.01, 0.02, 0.05\}$  over the validation set of 5,000 samples from the training set, and we divide the

|                   | BatchSize=128 | BatchSize=2 |
|-------------------|---------------|-------------|
| w/BN* [68]        | 6.61          | –           |
| Xavier Init* [24] | 7.78          | –           |
| Fixup-init* [69]  | 7.24          | –           |
| w/BN              | 6.82          | 7.24        |
| Xavier Init       | 8.43          | 9.74        |
| GroupNorm         | 7.33          | 7.36        |
| ONI               | <b>6.56</b>   | <b>6.67</b> |

Table 2. Test errors (%) comparison on 110-layer residual network [19] without BN [29] under CIFAR-10. ‘w/BN’ indicates with BN. We report the median of five independent runs. The methods with ‘\*’ indicate the results reported in the cited paper.

learning rate by 5 at 80 and 120 epochs, ending the training at 160 epochs. We compare our ‘ONI’ to several baselines, including orthogonal initialization [54] (‘OrthInit’), using soft orthogonal constraints as the penalty term [66] (‘OrthReg’), weight normalization [53] (‘WN’), ‘OLM’ [24] and the ‘plain’ network. Note that OLM [24] originally uses a scale of 1 (indicated as ‘OLM-1’), and we also apply the proposed scaling by  $\sqrt{2}$  (indicated as ‘OLM- $\sqrt{2}$ ’).

Table 1 shows the results. ‘ONI’ and ‘OLM- $\sqrt{2}$ ’ have significantly better performance under all network configurations (different depths and widths), which demonstrates the beneficial effects of maintaining orthogonality during training. We also observe ‘ONI’ and ‘OLM- $\sqrt{2}$ ’ converge faster than other baselines, in terms of training epochs (See the SM). Besides, our proposed ‘ONI’ achieves slightly better performance than ‘OLM- $\sqrt{2}$ ’ on average, over all configurations. Note that we train ‘OLM- $\sqrt{2}$ ’ with a group size of  $G = 64$ , as suggested in [24]. We also try full orthogonalization for ‘OLM- $\sqrt{2}$ ’. However, we observe either performance degeneration or numerical instability (*e.g.*, the eigen decomposition cannot converge). We argue that the main reason for this is that full orthogonalization solved by OLM over-constrains the weight matrix, which harms the performance. Moreover, eigen decomposition based methods are more likely to result in numerical instability in high-dimensional space, due to the element-wise multiplication of a matrix  $\mathbf{K}$  during back-propagation [40], as discussed in Section 3.1.

**Residual Network without Batch Normalization** Batch normalization (BN) is essential for stabilizing and accelerating the training [29] of DNNs [19, 23, 20, 59]. It is a standard configuration in residual networks [19]. However, it sometimes suffers from the small batch size problem [28, 63] and introduces too much stochasticity [60] when debugging

|          | Top-1 (%)    | Top-5 (%)   | Time (min./epoch) |
|----------|--------------|-------------|-------------------|
| plain    | 27.47        | 9.08        | 97                |
| WN       | 27.33        | 9.07        | 98                |
| OrthInit | 27.75        | 9.21        | 97                |
| OrthReg  | 27.22        | 8.94        | 98                |
| ONI      | <b>26.31</b> | <b>8.38</b> | 104               |

Table 3. Test errors (%) on ImageNet validation set (single model and single crop test) evaluated with VGG-16 [57]. The time cost for each epoch is averaged over the training epochs.

| Method | ResNet w/o BN |              | ResNet       |              | ResNetVar    |              |
|--------|---------------|--------------|--------------|--------------|--------------|--------------|
|        | Train         | Test         | Train        | Test         | Train        | Test         |
| plain  | 31.76         | 33.84        | 29.33        | 29.64        | 28.82        | 29.56        |
| ONI    | <b>27.05</b>  | <b>31.17</b> | <b>29.28</b> | <b>29.57</b> | <b>28.12</b> | <b>28.92</b> |

Table 4. Ablation study on ImageNet with an 18-layer ResNet. We evaluate the top-1 training and test errors (%).

| Method          | Test error (%) |              | Time (min./epoch) |     |
|-----------------|----------------|--------------|-------------------|-----|
|                 | 50             | 101          | 50                | 101 |
| ResNet          | 23.85          | 22.40        | 66                | 78  |
| ResNet + ONI    | 23.55          | 22.17        | 74                | 92  |
| ResNetVar       | 23.94          | 22.76        | 66                | 78  |
| ResNetVar + ONI | <b>23.30</b>   | <b>21.89</b> | 74                | 92  |

Table 5. Results on ImageNet with the 50- and 101-layer ResNets.

neural networks. Several studies have tried to train deep residual networks without BN [56, 69]. Here, we show that, when using our ONI, the residual network without BN can also be well trained.

The experiments are executed on a 110-layer residual network (Res-110). We follow the same experimental setup as in [19], except that we run the experiments on one GPU. We also compare against the Xavier Init [14, 9], and group normalization (GN) [63]. ONI can be trained with a large learning rate of 0.1 and converge faster than BN, in terms of training epochs (see the SM). We observe that ONI has slightly better test performance than BN (Table 2). Finally, we also test the performance on a small batch size of 2. We find ONI continues to have better performance than BN in this case, and is not sensitive to the batch size, like GN [63].

### 4.1.3 Large-scale ImageNet Classification

To further validate the effectiveness of our ONI on a large-scale dataset, we evaluate it on the ImageNet-2012 dataset. We keep almost all the experimental settings the same as the publicly available PyTorch implementation [48]: We apply SGD with a momentum of 0.9, and a weight decay of 0.0001. We train for 100 epochs in total and set the initial learning rate to 0.1, lowering it by a factor of 10 at epochs 30, 60 and 90. For more details on the slight differences among different architectures and methods, see the SM.

**VGG Network** Table 3 shows the results on the 16-layer VGG [57]. Our ‘ONI’ outperforms ‘plain’, ‘WN’, ‘OrthInit’ and ‘OrthReg’ by a significant margin. Besides, ‘ONI’ can be trained with a large learning rate of 0.1, while the other methods cannot (the results are reported for an initial learn-




Figure 7. Effects of the iteration number  $T$  of the proposed ONI. (a) 6-layer MLP for Fashion-MNIST; (b) VGG-Style network with ( $g = 2, k = 3$ ) for CIFAR-10; (c) 18-layer ResNet for ImageNet.

ing rate of 0.01). We also provide the running times in Table 3. The additional cost introduced by ‘ONI’ compared to ‘plain’ is negligible (7.2%).

**Residual Network** We first perform an ablation study on an 18-layer residual network (ResNet) [19], applying our ONI. We use the original ResNet and the ResNet without BN [29]. We also consider the architecture with BN inserted after the nonlinearity, which we refer to as ‘ResNetVar’. We observe that our ONI improves the performance over all three architectures, as shown in Table 4. One interesting observation is that ONI achieves the lowest training error on the ResNet without BN, which demonstrates its ability to facilitate optimization for large-scale datasets. We also observe that ONI has no significant difference in performance compared to ‘plain’ on ResNet. One possible reason is that the BN module and residual connection are well-suited for information propagation, causing ONI to have a lower net gain for such a large-scale classification task. However, we observe that, on ResNetVar, ONI obtains obviously better performance than ‘plain’. We argue that this boost is attributed to the orthogonal matrix’s ability to achieve approximate dynamical isometry, as described in Theorem 2.

We also apply our ONI on a 50- and 101-layer residual network. The results are shown in Table 5. We again observe that ONI can improve the performance, without introducing significant computational cost.

#### 4.1.4 Ablation Study on Iteration Number

ONI controls the spectrum of the weight matrix by the iteration number  $T$ , as discussed before. Here, we explore the effect of  $T$  on the performance of ONI over different datasets and architectures. We consider three configurations: 1) the 6-layer MLP for Fashion-MNIST; 2) the VGG-Style network with ( $g = 2, k = 3$ ) for CIFAR-10; and 3) the 18-layer ResNet without BN for ImageNet. The corresponding experimental setups are the same as described before. We vary  $T$  and show the results in Figure 7. Our primary observation is that using either a small or large  $T$  degrades performance. This indicates that we need to control the magnitude of orthogonality to balance the increased optimization benefit and diminished representational capacity. Our empirical observation is that  $T = 5$  usually works the best for networks without residual connections, whereas  $T = 2$  usually works better for residual networks. We argue that the residual network itself already has good optimization [19], which reduces the optimization benefits of orthogonality.




Figure 8. Comparison of SN and ONI on DCGAN. (a) The FID with respect to training epochs. (b) Stability experiments on six configurations, described in [43].

Besides, we also observe that larger  $T$ 's have nearly equivalent performance for simple datasets, *e.g.* Fashion-MNIST, as shown in 7 (a). This suggests that amplifying the eigenbasis corresponding to a small singular value cannot help more, even though the network with a fully orthogonalized weight matrix can well fit the dataset. We further show the distributions of the singular values of the orthogonalized weight matrix in the SM.

## 4.2. Stabilizing Training of GANs

How to stabilize GAN training is an open research problem [15, 52, 16]. One pioneering work is spectral normalization (SN) [43], which can maintain the Lipschitz continuity of a network by bounding the maximum eigenvalue of its weight matrices as 1. This technique has been extensively used in current GAN architectures [44, 70, 10, 35]. As stated before, our method is not only capable of bounding the maximum eigenvalue as 1, but can also control the orthogonality to amplify other eigenbasis with increased iterations, meanwhile orthogonal regularization is also a good technique for training GANs [10]. Here, we conduct a series of experiments for unsupervised image generation on CIFAR-10, and compare our method against the widely used SN [43].

**Experimental Setup** We strictly follow the network architecture and training protocol reported in the SN paper [43]. We use both DCGAN [50] and ResNet [19, 16] architectures. We provide implementation details in the SM. We replace all the SN modules in the corresponding network with our ONI. Our main metric for evaluating the quality of generated samples is the Fréchet Inception Distance (FID) [22] (the lower the better). We also provide the corresponding Inception Score (IS) [52] in the SM.

**DCGAN** We use the standard non-saturating function as the adversarial loss [15, 35] in the DCGAN architecture, following [43]. For optimization, we use the Adam optimizer [33] with the default hyper-parameters, as in [43]: learning rate  $\alpha = 0.0002$ , first momentum  $\beta_1 = 0$ , second momentum  $\beta_2 = 0.9$ , and the number of discriminator updates per generator update  $n_{dis} = 5$ . We train the network over 200 epochs with a batch size of 64 (nearly 200k generator updates) to determine whether it suffers from training instability. Figure 8 (a) shows the FID of SN and ONI when varying Newton's iteration number  $T$  from 0 to 5. One interesting observation is that the ONI with only the initial spectral bounding described in Formula 7 ( $T = 0$ ) can also stabilize training, even though it has downgraded perfor-




Figure 9. Comparison of SN and ONI on ResNet GAN. We show the FID with respect to training epochs when using (a) the non-saturating loss and (b) the hinge loss.

mance compared to SN. When  $T = 1$ , ONI achieves better performance than SN. This is because, based on what we observed, ONI stretches the maximum eigenvalue to nearly 1, while simultaneously amplifying other eigenvalues. Finally, we find that ONI achieves the best performance when  $T = 2$ , yielding an  $FID = 20.75$ , compared to SN's  $FID = 23.31$ . Further increasing  $T$  harms the training, possibly because too strong an orthogonalization downgrades the capacity of a network, as discussed in [43, 10].

We also conduct experiments to validate the stability of our proposed ONI under different experimental configurations: we use six configurations, following [43], by varying  $\alpha, \beta_1, \beta_2$  and  $n_{dis}$  (denoted by A-F, for more details please see the SM). Figure 8 (b) shows the results of SN and ONI (with  $T=2$ ) under these six configurations. We observe that our ONI is consistently better than SN.

**ResNet GAN** For experiments on the ResNet architecture, we use the same setup as the DCGAN. Besides the standard non-saturating loss [15], we also evaluate the recently popularized hinge loss [39, 43, 10]. Figure 9 shows the results. We again observe that our ONI achieves better performance than SN under the ResNet architecture, both when using the non-saturating loss and hinge loss.

## 5. Conclusion

In this paper, we proposed an efficient and stable orthogonalization method by Newton's iteration (ONI) to learn layer-wise orthogonal weight matrices in DNNs. We provided insightful analysis for ONI and demonstrated its ability to control orthogonality, which is a desirable property in training DNNs. ONI can be implemented as a linear layer and used to learn an orthogonal weight matrix, by simply substituting it for the standard linear module.

ONI can effectively bound the spectrum of a weight matrix in  $(\sigma_{min}, \sigma_{max})$  during the course of training. This property makes ONI a potential tool for validating some theoretical results relating to DNN's generalization (*e.g.*, the margin bounds shown in [6]) and resisting attacks from adversarial examples [12]. Furthermore, the advantage of ONI in stabilizing training w/o BN (BN usually disturbs the theoretical analysis since it depends on the sampled mini-batch input with stochasticity [29, 26]) makes it possible to validate these theoretical arguments under real scenarios.

**Acknowledgement** We thank Anna Hennig and Ying Hu for their help with proofreading.



## References

- [1] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, NJ, 2008. 5
- [2] Pierre-Antoine Absil and Jerome Malick. Projection-like retractions on matrix manifolds. *SIAM Journal on Optimization*, 22(1):135–158, 2012. 2
- [3] Jaweria Amjad, Zhaoyan Lyu, and Miguel RD Rodrigues. Deep learning for inverse problems: Bounds and regularizers. *arXiv preprint arXiv:1901.11352*, 2019. 1
- [4] Martín Arjovsky, Amar Shah, and Yoshua Bengio. Unitary evolution recurrent neural networks. In *ICML*, 2016. 2
- [5] Nitin Bansal, Xiaohan Chen, and Zhangyang Wang. Can we gain more from orthogonality regularizations in training deep cnns? In *NeurIPS*, 2018. 1, 2
- [6] Peter L Bartlett, Dylan J Foster, and Matus J Telgarsky. Spectrally-normalized margin bounds for neural networks. In *NeurIPS*. 2017. 8
- [7] Y. Bengio, P. Simard, and P. Frasconi. Learning long-term dependencies with gradient descent is difficult. *Trans. Neur. Netw.*, 5(2):157–166, Mar. 1994. 1
- [8] Dario A. Bini, Nicholas J. Higham, and Beatrice Meini. Algorithms for the matrix pth root. *Numerical Algorithms*, 39(4):349–378, Aug 2005. 1, 2, 3, 4
- [9] Nils Bjorck, Carla P Gomes, Bart Selman, and Kilian Q Weinberger. Understanding batch normalization. In *NeurIPS*. 2018. 7
- [10] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *ICLR*, 2019. 1, 2, 8
- [11] Andrew Brock, Theodore Lim, James M. Ritchie, and Nick Weston. Neural photo editing with introspective adversarial networks. In *ICLR*, 2017. 1
- [12] Moustapha Cisse, Piotr Bojanowski, Edouard Grave, Yann Dauphin, and Nicolas Usunier. Parseval networks: Improving robustness to adversarial examples. In *ICML*, 2017. 1, 8
- [13] Victor Dorobantu, Per Andre Stromhaug, and Jess Renteria. Dizzyrnn: Reparameterizing recurrent neural networks for norm-preserving backpropagation. *CoRR*, abs/1612.04035, 2016. 2
- [14] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *AISTATS*, 2010. 1, 7
- [15] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*. 2014. 8
- [16] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *NeurIPS*. 2017. 8
- [17] Mehrtash Harandi and Basura Fernando. Generalized backpropagation, etude de cas: Orthogonality. *CoRR*, abs/1611.05927, 2016. 1, 2
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *ICCV*, 2015. 1
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 6, 7, 8
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *ECCV*, 2016. 6
- [21] Kyle Helfrich, Devin Willmott, and Qiang Ye. Orthogonal recurrent neural networks with scaled Cayley transform. In *ICML*, 2018. 2
- [22] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*. 2017. 8
- [23] Gao Huang, Zhuang Liu, and Kilian Q. Weinberger. Densely connected convolutional networks. In *CVPR*, 2017. 6
- [24] Lei Huang, Xianglong Liu, Bo Lang, Adams Wei Yu, Yongliang Wang, and Bo Li. Orthogonal weight normalization: Solution to optimization over multiple dependent stiefel manifolds in deep neural networks. In *AAAI*, 2018. 1, 2, 3, 4, 5, 6
- [25] Lei Huang, Dawei Yang, Bo Lang, and Jia Deng. Decorrelated batch normalization. In *CVPR*, 2018. 2
- [26] Lei Huang, Yi Zhou, Fan Zhu, Li Liu, and Ling Shao. Iterative normalization: Beyond standardization towards efficient whitening. In *CVPR*, 2019. 1, 2, 3, 8
- [27] Stephanie Hyland and Gunnar Rätsch. Learning unitary operators with help from u(n). In *AAAI*, 2017. 2
- [28] Sergey Ioffe. Batch renormalization: Towards reducing mini-batch dependence in batch-normalized models. In *NeurIPS*, 2017. 6
- [29] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015. 2, 5, 6, 7, 8
- [30] Catalin Ionescu, Orestis Vantzos, and Cristian Sminchisescu. Training deep networks with structured layers by matrix back-propagation. In *ICCV*, 2015. 1, 2, 3
- [31] Kui Jia. Improving training of deep neural networks via singular value bounding. In *CVPR*, 2017. 2
- [32] Li Jing, Çağlar Gülçehre, John Peurifoy, Yichen Shen, Max Tegmark, Marin Soljacic, and Yoshua Bengio. Gated orthogonal recurrent units: On learning to forget. *CoRR*, abs/1706.02761, 2017. 2
- [33] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 8
- [34] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009. 5
- [35] Karol Kurach, Mario Lučić, Xiaohua Zhai, Marcin Michalski, and Sylvain Gelly. A large-scale study on regularization and normalization in GANs. In *ICML*, 2019. 8
- [36] Yann LeCun, Léon Bottou, Genevieve B. Orr, and Klaus-Robert Müller. Efficient backprop. In *Neural Networks: Tricks of the Trade*, 1998. 1, 4
- [37] José Lezama, Qiang Qiu, Pablo Musé, and Guillermo Sapiro. OIÉ: Orthogonal low-rank embedding - a plug and play geometric loss for deep learning. In *CVPR*, June 2018. 1
- [38] Peihua Li, Jiangtao Xie, Qilong Wang, and Zilin Gao. Towards faster training of global covariance pooling networks by iterative matrix square root normalization. In *CVPR*, 2018. 2
- [39] Jae Hyun Lim and Jong Chul Ye. Geometric gan. *CoRR*, abs/1705.02894, 2017. 8

- [40] Tsung-Yu Lin and Subhransu Maji. Improved bilinear pooling with cnns. In *BMVC*, 2017. 1, 2, 3, 6
- [41] Per-Olov Löwdin. On the non-orthogonality problem connected with the use of atomic wave functions in the theory of molecules and crystals. *The Journal of Chemical Physics*, 18(3):365–375, 1950. 1, 2, 3
- [42] Dmytro Mishkin and Jiri Matas. All you need is a good init. In *ICLR*, 2016. 1, 2
- [43] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *ICLR*, 2018. 1, 2, 3, 8
- [44] Takeru Miyato and Masanori Koyama. cgans with projection discriminator. In *ICLR*, 2018. 8
- [45] Vinod Nair and Geoffrey E. Hinton. Rectified linear units improve restricted boltzmann machines. In *ICML*, 2010. 2, 5, 6
- [46] Mete Ozay and Takayuki Okatani. Optimization on submanifolds of convolution kernels in cnns. *CoRR*, abs/1610.07008, 2016. 1, 2
- [47] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks. In *ICML*, 2013. 1, 2
- [48] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. In *NeurIPS Autodiff Workshop*, 2017. 7
- [49] Jeffrey Pennington, Samuel Schoenholz, and Surya Ganguli. Resurrecting the sigmoid in deep learning through dynamical isometry: theory and practice. In *NeurIPS*, 2017. 1, 2
- [50] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *ICLR*, 2016. 8
- [51] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. 2, 5
- [52] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, Xi Chen, and Xi Chen. Improved techniques for training gans. In *NeurIPS*, 2016. 2, 8
- [53] Tim Salimans and Diederik P. Kingma. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. In *NeurIPS*, 2016. 3, 5, 6
- [54] Andrew M. Saxe, James L. McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *CoRR*, abs/1312.6120, 2013. 1, 2, 6
- [55] Nicol N. Schraudolph. Accelerated gradient descent by factor-centering decomposition. Technical report, 1998. 4
- [56] Wenling Shang, Justin Chiu, and Kihyuk Sohn. Exploring normalization in deep residual networks with concatenated rectified linear units. In *AAAI*, 2017. 7
- [57] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 7
- [58] Piotr A. Sokol and Il Memming Park. Information geometry of orthogonal initializations and training. *CoRR*, abs/1810.03785, 2018. 1, 2
- [59] Christian Szegedy, Sergey Ioffe, and Vincent Vanhoucke. Inception-v4, inception-resnet and the impact of residual connections on learning. *CoRR*, abs/1602.07261, 2016. 6
- [60] Mattias Teye, Hossein Azizpour, and Kevin Smith. Bayesian uncertainty estimation for batch normalized deep networks. In *ICML*, 2018. 6
- [61] Eugene Vorontsov, Chiheb Trabelsi, Samuel Kadoury, and Chris Pal. On orthogonality and learning recurrent networks with long term dependencies. In *ICML*, 2017. 1, 2
- [62] Scott Wisdom, Thomas Powers, John Hershey, Jonathan Le Roux, and Les Atlas. Full-capacity unitary recurrent neural networks. In *NeurIPS*, 2016. 1, 2
- [63] Yuxin Wu and Kaiming He. Group normalization. In *ECCV*, 2018. 6, 7
- [64] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *CoRR*, abs/1708.07747, 2017. 5
- [65] Lechao Xiao, Yasaman Bahri, Jascha Sohl-Dickstein, Samuel S. Schoenholz, and Jeffrey Pennington. Dynamical isometry and a mean field theory of cnns: How to train 10,000-layer vanilla convolutional neural networks. In *ICML*, 2018. 1, 2
- [66] Di Xie, Jiang Xiong, and Shiliang Pu. All you need is beyond a good init: Exploring better solution for training extremely deep convolutional neural networks with orthonormality and modulation. In *CVPR*, 2017. 1, 6
- [67] Greg Yang, Jeffrey Pennington, Vinay Rao, Jascha Sohl-Dickstein, and Samuel S. Schoenholz. A mean field theory of batch normalization. In *ICLR*, 2019. 1, 2
- [68] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *BMVC*, 2016. 6
- [69] Hongyi Zhang, Yann N. Dauphin, and Tengyu Ma. Fixup initialization: Residual learning without normalization. In *ICLR*, 2019. 6, 7
- [70] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *ICML*, 2019. 8
- [71] Liheng Zhang, Marzieh Edraki, and Guo-Jun Qi. Cappronet: Deep feature learning via orthogonal projections onto capsule subspaces. In *NeurIPS*, 2018. 1
- [72] Jianping Zhou, Minh N. Do, and Jelena Kovacevic. Special paraunitary matrices, cayley transform, and multidimensional orthogonal filter banks. *IEEE Trans. Image Processing*, 15(2):511–519, 2006. 2