

# An Investigation into the Stochasticity of Batch Whitening

Lei Huang<sup>†</sup> Lei Zhao Yi Zhou<sup>†</sup> Fan Zhu<sup>†</sup> Li Liu<sup>†</sup> Ling Shao<sup>†</sup>

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

{lei.huang, yi.zhou, fan.zhu, li.liu, ling.shao}@inceptioniai.org bhneo@126.com

## Abstract

Batch Normalization (BN) is extensively employed in various network architectures by performing standardization within mini-batches. A full understanding of the process has been a central target in the deep learning communities. Unlike existing works, which usually only analyze the standardization operation, this paper investigates the more general Batch Whitening (BW). Our work originates from the observation that while various whitening transformations equivalently improve the conditioning, they show significantly different behaviors in discriminative scenarios and training Generative Adversarial Networks (GANs). We attribute this phenomenon to the stochasticity that BW introduces. We quantitatively investigate the stochasticity of different whitening transformations and show that it correlates well with the optimization behaviors during training. We also investigate how stochasticity relates to the estimation of population statistics during inference. Based on our analysis, we provide a framework for designing and comparing BW algorithms in different scenarios. Our proposed BW algorithm improves the residual networks by a significant margin on ImageNet classification. Besides, we show that the stochasticity of BW can improve the GAN’s performance with, however, the sacrifice of the training stability.

## 1. Introduction

Normalization techniques have been extensively used for learning algorithms during data preprocessing [24, 22, 11]. It has been shown that centering, scaling and decorrelating the inputs speeds up the training [24]. Furthermore, whitening the input that combines all above operations, improves the conditioning of the Hessian, making the gradient descent updates similar to the Newton updates [24, 44, 15].

Batch Normalization (BN) [18] extends the idea of normalizing the input into the activations of intermediate layers of Deep Neural Networks (DNNs), and represents a milestone technique in the deep learning community [11, 41, 45]. BN standardizes the activations by executing centering and scaling within a mini-batch of data, facilitating the optimization [18, 21, 34] and generalization [3, 4]. Batch Whitening

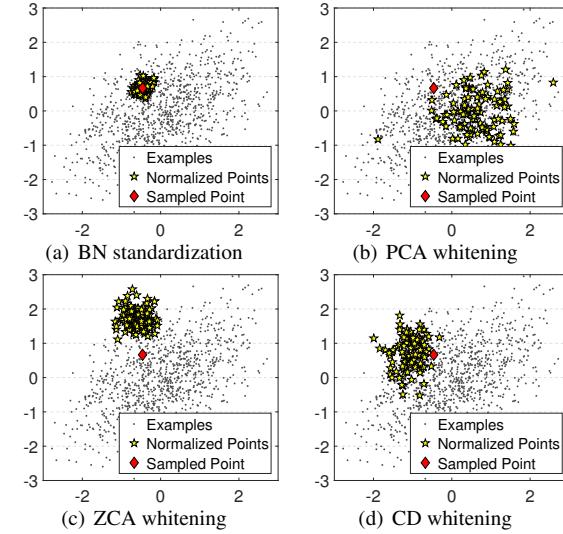


Figure 1. We sample 1000 examples (black points) from a Gaussian distribution in a 16-dimensional space, and show the examples in the two-dimension sub-space (the 6th and 16th dimension). Given an example  $\mathbf{x}$  (red diamond), when combining with 100 different mini-batches  $\mathbf{X}^B$  ( $B = 64$ ), we provide the normalized output  $\hat{\mathbf{x}}$  (yellow pentagram), where (a), (b), (c) and (d) show the results of BN standardization, PCA, ZCA and CD whitening, respectively.

(BW) further extends the scope of BN by removing the correlation of standardized activations [15]. It has been shown to improve the performance in discriminative scenarios [15] and Generative Adversarial Networks (GAN) [37].

Despite BW’s theoretical support in improving conditioning, there remains some intriguing observations relating to BW that are not well explored. Firstly, while various whitening transformations can equivalently improve the conditioning [19], they show significant differences in performance: 1) Principal Component Analysis (PCA) whitening hardly converges while Zero-phase Component Analysis (ZCA) whitening works well in discriminative scenarios [15]; 2) Cholesky Decomposition (CD) whitening achieves significantly better performance than ZCA in GAN training, while it has slightly worse performance in discriminative cases [37]. Secondly, while group based whitening—where features are divided into groups and whitening is performed within each one—is essential in discriminative scenarios

[15, 30], full feature whitening has been shown to achieve better performance in training GANs [37].

This paper focuses on explaining the above observations of BW. We find that the stochasticity introduced by normalization over batch data (Figure 1) can be key to uncovering the intriguing phenomena of BW. We quantitatively investigate the magnitude of stochasticity in different whitening transformations, using the evaluation method called Stochastic Normalization Disturbance (SND) [16]. By doing so, we demonstrate that PCA whitening has significantly larger stochasticity, which is difficult to be controlled by either increasing batch size or using group based whitening. On the other side, ZCA whitening has the smallest stochasticity, while CD whitening has a moderate value, and more importantly, their stochasticity can be well controlled. This suggests ZCA whitening should have better optimization behaviors in training, while PCA whitening has problems in converging, due to the increased stochasticity which slows down the progress of the optimization [40, 16].

We also investigate the stochasticity during inference, which is caused by the estimation of population statistics averaged over the mini-batch statistics during training. We show that in terms of estimating the population statistics of the whitening matrix, it is more stable to use the mini-batch covariance matrix indirectly (We calculate the whitening matrix after training) than the mini-batch whitening matrix directly. We further provide an empirical investigation to understand the reasons behind this observation, and find that the stochastic sequences of the mini-batch whitening matrix have a large diversity than the covariance matrix.

Based on the above analyses, we provide a general framework for designing and comparing BW algorithms in different scenarios. We design new BW algorithm and apply them to the residual networks [11] for ImageNet dataset [7], significantly improving the performance over the original one. We further conduct thorough experiments on training GANs. We show that full feature whitening, when combined with coloring, can improve the final score of evaluation. However, it reduces the training stability and is more sensitive to the hyper-parameters configurations. We attribute this phenomenon to two main effects caused by the introduced stochasticity of BW: 1) Strong stochasticity can increase the diversity of generated images and thus improve the GAN’s performance; 2) At the same time, high stochasticity harms optimization and thus is more sensitive to the hyper-parameters. We argue that controlling the magnitude of whitening (stochasticity) is also important in training GANs and we validate this argument with our experiments.

## 2. Related Work

The previous analyses on BN mainly focus on the optimization. One main argument is that BN can improve the conditioning of the optimization problem. This argument was initially introduced in the BN paper [18] and further

refined in [34], showing that BN leads to a smoother landscape of the optimization problem under certain assumptions. Ghorbani *et al.* [9] investigated this explanation by computing the spectrum of the Hessian for a large-scale dataset. It is believed that the improved conditioning enables large learning rates, thus improving the generalization, as shown in [4]. Another argument is that BN can adaptively adjust the learning rate [6, 14, 1] due to its scale invariance [18, 3]. This effect has been further discussed in combination with weight decay [47]. Other works have included an investigation into the signal propagation and gradient back-propagation [46]. Different from these approaches, our work focuses on analyzing the stochasticity of whitening over batch data.

The stochasticity introduced by normalization over batch data was first mentioned in the BN paper [18], and further explored in [2, 42] from the perspective of Bayesian optimization. This stochasticity results in differences between the training distribution (using mini-batch statistics) and the test distribution (using estimated population statistics) [17], which is believed to be the main cause of the small-batch-problem of BN [45]. To address this issue, a number of approaches have been proposed [45, 32, 28, 17, 43, 39]. Furthermore, it has been observed that BN also encounters difficulties in optimization during training [35, 16]. This phenomenon is explored by the stochastic analysis shown in [16]. Different from the above research which focuses on standardization, we analyze, for the first time, the stochasticity on batch whitening. We propose that analyzing whitening rather than standardization, has several advantages in understanding the behaviors of normalization over batch data: 1) There are an infinite number of whitening transformations and the main ones show significant differences as discussed in Section 1; 2) The extent of the whitening (stochasticity) can be well controlled by the batch and group size, which provides more information in designing experiments.

Our work is related to the previously proposed whitening methods regarding the activation of DNNs. One approach is to consider the whitening matrix as model parameters to be estimated using full data [8, 27]. This kind of whitening has also been exploited in image style transformation tasks [25, 36]. Another line of research is batch whitening, which is what this paper discusses. This approach treats the normalization as a function over a mini-batch input. The main works include PCA whitening, ZCA whitening [15] and its approximation ITN [16], and CD whitening [37]. Pan *et al.* [30] propose switchable whitening to learn different batch/instance whitening/standardization operations in DNNs. However, they used only the ZCA whitening transformation. Our work aims to understand different whitening transformation behaviors based on stochastic analysis.

## 3. Stochasticity Analysis of Batch Whitening

Let  $\mathbf{X} \in \mathbf{R}^{d \times m}$  be a data matrix denoting the mini-batch input of size  $m$  in a certain layer. For simplifying the discussion, we assume that the data is centered, by performing

$\mathbf{X} := \mathbf{X} - \mu \cdot \mathbf{1}^T$  where  $\mu = \frac{1}{m} \mathbf{X} \cdot \mathbf{1}$  is the mean of  $\mathbf{X}$ , and  $\mathbf{1}$  is a column vector of all ones. Whitening performs normalization over the mini-batch input as follows:

$$\hat{\mathbf{X}} = \mathbf{G}\mathbf{X}, \quad (1)$$

where  $\mathbf{G}$  is the mini-batch *whitening matrix* that is derived from the corresponding *covariance matrix*  $\Sigma = \frac{1}{m} \mathbf{X}\mathbf{X}^T$ .

The population statistics of the whitening matrix  $\hat{\mathbf{G}}$  used for inference, is usually calculated by running average over the mini-batches as follows:

$$\hat{\mathbf{G}} = (1 - \lambda)\hat{\mathbf{G}} + \lambda\mathbf{G}. \quad (2)$$

It is clear that both the whitened output  $\hat{\mathbf{X}}$  (Eqn. 1) and the population statistics  $\hat{\mathbf{G}}$  (Eqn. 2) can be viewed as stochastic variables, because they depend on the mini-batch inputs which are sampled over datasets. For illustration, we defer the analysis of the stochasticity to Sections 3.2 and 3.3, and first provide a review of the whitening transformations.

### 3.1. Whitening Transformations

There are an infinite number of possible whitening matrices, as shown in [19, 15], since any whitened data with a rotation is still whitened. This paper focuses on the whitening transformations based on PCA, ZCA and CD, since these three transformations have shown significant differences in performance when used in training DNNs [15, 37]. Note that BN [18] can be viewed as a special case of batch whitening, since it performs standardization without removing correlations, where  $\mathbf{G}_{BN} = (\text{diag}(\Sigma))^{-1/2}$  with  $\text{diag}(\cdot)$  setting the off-diagonal elements of a matrix to zeros. To simplify the description, this paper regards BN as a (reduced) whitening transformation, unless otherwise stated.

**PCA Whitening** uses  $\mathbf{G}_{PCA} = \Lambda^{-\frac{1}{2}}\mathbf{D}^T$ , where  $\Lambda = \text{diag}(\sigma_1, \dots, \sigma_d)$  and  $\mathbf{D} = [\mathbf{d}_1, \dots, \mathbf{d}_d]$  are the eigenvalues and associated eigenvectors of  $\Sigma$ , *i.e.*  $\Sigma = \mathbf{D}\Lambda\mathbf{D}^T$ . Under this transformation, the variables are first rotated by the eigen-matrix ( $\mathbf{D}$ ) of the covariance, then scaled by the square root inverse of the eigenvalues ( $\Lambda^{-\frac{1}{2}}$ ). PCA whitening over batch data suffers significant instability in training DNNs, and hardly converges, due to the so called Stochastic Axis Swapping (SAS), as explained in [15].

**ZCA Whitening** uses  $\mathbf{G}_{ZCA} = \mathbf{D}\Lambda^{-\frac{1}{2}}\mathbf{D}^T$ , where the PCA whitened input is rotated back by the corresponding rotation matrix  $\mathbf{D}$ . ZCA whitening works by stretching/squeezing the dimensions along the eigenvectors. It has been shown that ZCA whitening avoids the SAS and achieves better performance over standardization (used in BN) on discriminative classification tasks [15].

**CD Whitening** uses  $\mathbf{G}_{CD} = \mathbf{L}^{-1}$  where  $\mathbf{L}$  is a lower triangular matrix from the Cholesky decomposition, with  $\mathbf{L}\mathbf{L}^T = \Sigma$ . This kind of whitening works by recursively decorrelating the current dimension over the previous decorrelated ones, resulting in a triangular form of its whitening matrix. CD whitening has been shown to achieve the state-of-the-art performance in training GANs, while ZCA whitening has degenerated performance.

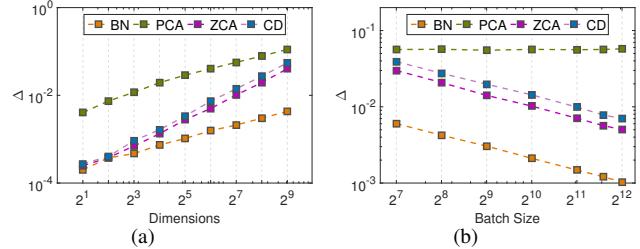


Figure 2. SND comparison of different batch whitening methods. We sample 60,000 examples from a Gaussian distribution as the training set. To calculate SND, we use  $s = 200$  and  $N = 20$ . We show (a) the SND with respect to the dimensions ranging from  $2^1$  to  $2^9$ , under a batch size of 1024; (b) the SND with respect to the batch size ranging from  $2^7$  to  $2^{12}$ , under a dimension of 128.

The primary motivation of this paper is to investigate the following problem: while all whitening methods equivalently improve the conditioning of the layer input, why do they show significantly different behaviors in training DNNs? In the following sections, we provide a unified analysis and demonstrate that the key is the stochasticity introduced by normalization over batch data.

### 3.2. Stochasticity During Training

Given a sample  $\mathbf{x} \in \mathbb{R}^d$  from a distribution  $P_\chi$ , we take a sample set  $\mathbf{X}^B = \{\mathbf{x}_1, \dots, \mathbf{x}_B, \mathbf{x}_i \sim P_\chi\}$  with a size of  $B$ . The whitened output  $\hat{\mathbf{x}}$  can be formulated as  $\hat{\mathbf{x}} = \mathbf{G}(\mathbf{X}^B; \mathbf{x})$ . For a certain  $\mathbf{x}$ ,  $\mathbf{X}^B$  can be viewed as a random variable [2, 42, 16].  $\hat{\mathbf{x}}$  is thus another random variable showing stochasticity. Here, we investigate the stochasticity effects of different whitening transformations.

To provide a more intuitive illustration, we conduct a toy experiment to show how the normalized output of one sample changes when combined with a different sample set  $\mathbf{X}^B$ , by performing different whitening methods. Figure 1 (a), (b), (c) and (d) show the results when performing BN, PCA, ZCA and CD whitening, respectively. It is clear that the distribution of the PCA whitened output of one sample is very sparse, which means that  $\hat{\mathbf{x}}$  has significant diversity. This suggests that PCA whitening provides large stochasticity. On the other side, the BN standardized output shows a tight Gaussian-style distribution, which suggests that BN has smaller stochasticity. Note that BN can not guarantee that the normalized output has an identity covariance matrix, while other whitening methods can. Similarly, we also observe that ZCA whitening provides reduced stochasticity, compared to CD. In fact, ZCA whitening has been shown to minimize the total squared distance between the original and whitened variables [19, 15]. We conjecture that such a property results in ZCA whitening having smaller stochasticity than CD.

#### 3.2.1 Quantitative Analysis

To provide a qualitative comparison, we exploit the evaluation called *Stochastic Normalization Disturbance* (SND), introduced in [16]. The empirical estimation of SND for

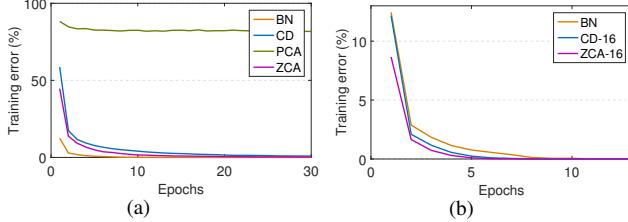


Figure 3. Experiments on a 4-layer MLP with 256 neurons in each layer, for MNIST classification. We use a batch size of 1024 and report the training errors. (a) The results of full whitening methods; (b) The results of group based whitening, where ‘ZCA-16’ indicates ZCA whitening with a group size of 16.

sample  $\mathbf{x}$  over the normalization  $\mathbf{G}(\cdot)$  is defined as:

$$\hat{\Delta}_{\mathbf{G}}(\mathbf{x}) = \frac{1}{s} \sum_{i=1}^s \|\mathbf{G}(\mathbf{X}_i^B; \mathbf{x}) - \frac{1}{s} \sum_{j=1}^s \mathbf{G}(\mathbf{X}_j^B; \mathbf{x})\|, \quad (3)$$

where  $s$  denotes the number of mini-batches  $\{\mathbf{X}_j^B\}_{j=1}^s$  that are randomly sampled from the dataset. SND can be used to evaluate the stochasticity of a sample after the normalization operation [16]. Further, a normalization operation  $G(\cdot)$ ’s empirical SND is defined as  $\hat{\Delta}_G = \frac{1}{N} \sum_{i=1}^N \hat{\Delta}(\mathbf{x}_i)$  given  $N$  samples.  $\hat{\Delta}_G$  describes the magnitudes of stochasticity for the corresponding normalization operations.

Here, we conduct experiments to quantitatively evaluate the effects of different normalization methods. Noticeably, the stochasticity is related to the batch size  $m$  and the dimension  $d$ . Figure 2 (a) shows the SND of different normalization methods with respect to the dimensions, when fixing the batch size to 1024. We find that PCA whitening shows the largest SND while BN the smallest, over all the dimensions, which is consistent with the observation shown in Figure 1. We notice that all whitening methods have an increased SND when the dimension increases. Besides, ZCA has a smaller SND than CD, over all the dimensions, which is also consistent with the data shown in Figure 1. Figure 2 (b) shows the SND of different normalization methods with respect to the batch size, when fixing the dimension to 128. An interesting observation is that PCA whitening has nearly the same large SND among different batch sizes. This suggests that the PCA whitening is extremely unstable, no matter how accurate the estimation of the mini-batch covariance matrix is. This effect is in accordance with the explanation of Stochastic Axis Swapping (SAS) shown in [15], where a small change over the examples (when performing PCA whitening) results in a large change of representation.

To further investigate how this stochasticity affects DNN training, we perform experiments on a four-layer Multilayer Perceptron (MLP) with 256 neurons in each layer. We evaluate the training loss with respect to the epochs, and show the results in Figure 3 (a). We find that, among all the whitening methods, ZCA works the best, while PCA is the worst. We argue that this correlates closely with the SND they produce. Apparently, the increased stochasticity can slow down training, even though all the whitening methods have equivalently

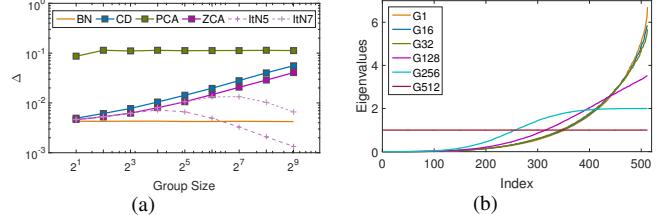


Figure 4. Group-based whitening experiments. (a) We show the SND of different normalization operations with respect to the group size. The experimental setup is the same as Figure 2 and the input dimension is  $d = 512$ . (b) We show the spectrum of covariance matrix of the ZCA whitened output (Note that CD/PCA whitening has the same spectrum as ZCA whitening.), where ‘G16’ indicates whitening with a group size of 16.

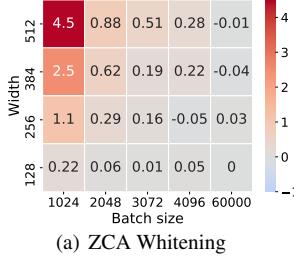
improved conditioning. An interesting observation is that, in this case, BN works better than ZCA whitening. This is surprising since ZCA has improved conditioning over BN by removing the correlation [15], and it should theoretically have a better optimization behavior. However, the amplified stochasticity of ZCA whitening mitigates this advantage in optimization, thus resulting in a degenerated performance. Therefore, from an optimization perspective, we should control the extent of the stochasticity.

### 3.2.2 Controlling the Stochasticity by Groups

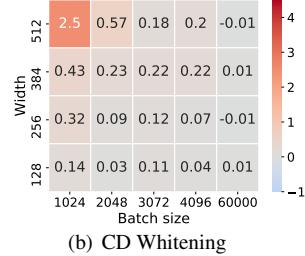
Huang *et al.* [15] proposed to use groups to control the extent of whitening. They argue that this method reduces the inaccuracy in estimating the full covariance matrix when the batch size is not sufficiently large. Here, we empirically show how group based whitening affects the SND, providing a good trade-off between introduced stochasticity and improved conditioning. This is essential for achieving a better optimization behavior.

We evaluate the SND of different whitening transformations by varying the group size ranging from 2 to 512, as shown in Figure 4 (a). We also display the spectrum of the covariance matrix of the whitened output (based on groups) in Figure 4 (b). We find that the group size effectively controls the SND of the ZCA/CD whitening. With decreasing group size, ZCA and CD show reduced stochasticity (Figure 4 (a)), while having also degenerated conditioning (Figure 4 (b)), since the output is only partially whitened. Besides, we observe that PCA whitening still has a large SND over all group sizes, and with no significant differences. This observation further corroborates the explanation of SAS given in [15], *i.e.*, that the PCA whitening is extremely unstable.

We also show the SND of the approximate ZCA whitening method (called ItN [16]) in Figure 4 (a), which uses Newton’s iteration to approximately calculate the whitening matrix. We denote ‘ItN5’ as the ItN method with an iteration number of 5. An interesting observation is that ItN has smaller SND than BN, when using a large group size (*e.g.*, 256) with a smaller iteration (*e.g.*, T=5). This suggests that we can further combine group size and iteration number to control the stochasticity for ItN, providing an efficient and stable solution to approximate ZCA whitening [16].



(a) ZCA Whitening



(b) CD Whitening

Figure 5. Comparison of estimating  $\hat{\mathbf{G}}$  using different estimation objects ( $\mathbf{G}/\Sigma$ ). We train MLPs with varying widths (number of neurons in each layer) and batch sizes, for MNIST classification. We evaluate the difference of test accuracy between  $\hat{\mathbf{G}}_\Sigma$  and  $\hat{\mathbf{G}}_\mathbf{G}$ :  $AC(\hat{\mathbf{G}}_\Sigma) - AC(\hat{\mathbf{G}}_\mathbf{G})$ . (a) and (b) show the results using ZCA and CD whitening, respectively, under a learning rate of 1. We also tried other learning rates and obtained similar observations (See *supplementary materials* for details).

The above observations furnish substantial insights into the application of whitening in neural networks (especially in scenarios that requires decorating the activation in a certain layer), and we will further elaborate on this in Section 4.

We also use group-based ZCA/CD whitening methods on the four-layer MLP experiments. The results are shown in Figure 3 (b). We observe that ZCA and CD whitening, with a group size of 16 to control the stochasticity, achieve better training behaviors than BN.

### 3.3. Stochasticity During Inference

In the previous section we have shown that performing different whitening transformations, by introducing different magnitudes of stochasticity, results in significantly different training behaviors from an optimization perspective. It's clear that such a stochasticity will also affect the final test performance during inference, because the population statistics  $\hat{\mathbf{G}}$  is estimated by running average (Eqn. 2) over the stochastic sequences  $\{\mathbf{G}_t\}_{t=1}^T$ , where  $T$  is the total number of training iterations. The more diverse of the stochastic sequence  $\{\mathbf{G}_t\}_{t=1}^T$ , the more difficult to accurately estimate  $\hat{\mathbf{G}}$ . Rather than estimating  $\hat{\mathbf{G}}$  directly, Siarohin *et al.* [37] proposed to first estimate the population statistic of the covariance matrix  $\hat{\Sigma}$ , then compute  $\hat{\mathbf{G}}$  based on  $\hat{\Sigma}$  after training. However, no further analysis was provided to explain why they do like this.

Here, we provide an empirical investigation on how the estimation object ( $\mathbf{G}/\Sigma$ ) in Eqn. 2 affects the test performance. Let  $\hat{\mathbf{G}}_\mathbf{G}$  and  $\hat{\mathbf{G}}_\Sigma$  denote the method to estimate  $\hat{\mathbf{G}}$  by  $\mathbf{G}$  and  $\Sigma$ , respectively. We conduct experiments on MLP with variations in width (the number of neurons in each layer) and batch size, for MNIST classification. Figure 5 (a) and (b) show the results of ZCA and CD whitening, respectively. We find that  $\hat{\mathbf{G}}_\Sigma$  has better performance than  $\hat{\mathbf{G}}_\mathbf{G}$ , especially under the scenarios with large width and small batch size (Intuitively, estimating in high-dimensional space with a small batch size will make the estimation noisier). This suggests that using  $\Sigma$  to estimate  $\hat{\mathbf{G}}$  indirectly is more stable than

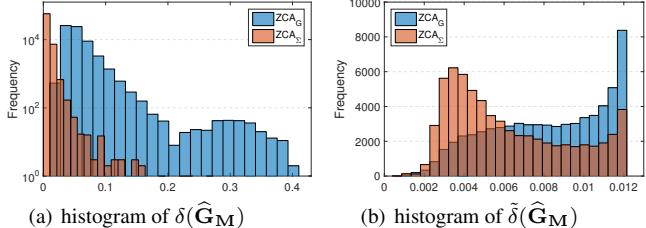


Figure 6. Analysis of the diversity of the stochastic sequences in estimating  $\hat{\mathbf{G}}$ . We report the histogram of  $\delta(\hat{\mathbf{G}}_\mathbf{M})$  and  $\tilde{\delta}(\hat{\mathbf{G}}_\mathbf{M})$ .

using  $\mathbf{G}$  directly. Another interesting observation is that, by comparing Figure 5 (a) and (b), the differences between the estimation methods using CD whitening is smaller than the ZCA whitening.

We further analyze how the diversity of stochastic sequences  $\{\mathbf{M}_t\}_{t=1}^T$  affect the estimation of  $\hat{\mathbf{G}}$ , where  $\mathbf{M} \in \{\mathbf{G}, \Sigma\}$ . Intuitively, if the stochastic sequence has high diversity, the effects of estimation will be worse. We view each element of  $\hat{\mathbf{G}}$  as an independent stochastic variable and calculate the standard deviation of each element during training, as follows:

$$\delta(\hat{\mathbf{G}}_{\mathbf{M}}^{ij}) = \sqrt{\frac{1}{T} \sum_{t=1}^T (\mathbf{M}_t^{ij} - \frac{1}{T} \sum_{t'=1}^T (\mathbf{M}_{t'}^{ij}))^2}, \quad (4)$$

where  $\hat{\mathbf{G}}^{ij}$  ( $\mathbf{M}_t^{ij}$ ) indicates the (i,j)-th element of  $\hat{\mathbf{G}}$  ( $\mathbf{M}_t$ ). Furthermore, we calculate the normalized standard deviation of each element  $\tilde{\delta}(\hat{\mathbf{G}}_{\mathbf{M}}^{ij})$ , which, as defined in Eqn. 4, is calculated over  $\tilde{\mathbf{M}}_t^{ij} = \mathbf{M}_t^{ij} / \sqrt{\sum_{t=1}^T (\mathbf{M}_t^{ij})^2}$ . Figure 6 (a) and (b) show the histogram of  $\delta(\hat{\mathbf{G}}_\mathbf{M})$  and  $\tilde{\delta}(\hat{\mathbf{G}}_\mathbf{M})$ , respectively, when using ZCA whitening. We clearly find that  $\hat{\mathbf{G}}_\mathbf{G}$  has a larger standard deviation on average, thus a large diversity in general, compared to  $\hat{\mathbf{G}}_\Sigma$ . This reveals why using  $\Sigma$  is more stable for estimating  $\hat{\mathbf{G}}_\mathbf{G}$  than using  $\mathbf{G}$ .

## 4. Evaluation on Vision Tasks

Based on the previous analysis, we can design new BW algorithms, and construct more effective DNNs by using BW. We investigate this in classification and training GANs. The code to reproduce the experiments is available at <https://github.com/huangleiBuaa/StochasticityBW>.

### 4.1. The Landscape of Batch Whitening Algorithms

We provide a general view of batch whitening algorithms in Algorithm 1 for vectorial inputs  $\mathbf{X} \in \mathbb{R}^{d \times m}$ . Backpropagation is necessary to pass through the whitening transformation, and we provide the details in *supplementary materials* for completeness. For feature map inputs  $\mathbf{X}_F \in \mathbb{R}^{h \times w \times d \times m}$ , where  $h$  and  $w$  indicate the height and width, the whitening transformation is performed over the unrolled  $\mathbf{X} \in \mathbb{R}^{d \times (mhw)}$ , since each spatial position of the feature map can be viewed as a sample [18].

Note that an extra step to recover the representation capacity of normalization is given in Line 11 of Algorithm 1, which is shown to be empirically effective in [18, 3, 15, 45].

---

**Algorithm 1** A general view of batch whitening algorithms.

---

- 1: **Input:** mini-batch inputs  $\mathbf{X} \in \mathbb{R}^{d \times m}$ .
- 2: **Output:**  $\mathbf{Y} \in \mathbb{R}^{d \times m}$ .
- 3: **if** Training **then**
- 4:   Calculate covariance matrix:  $\Sigma = \frac{1}{m} \mathbf{X} \mathbf{X}^T + \epsilon \mathbf{I}$ .
- 5:   Calculate whitening matrix:  $\mathbf{G} = \phi_1(\Sigma)$ .
- 6:   Calculate whitened output:  $\hat{\mathbf{X}} = \mathbf{G} \mathbf{X}$ .
- 7:   Update population statistics:  $\hat{\mathbf{G}} = \phi_2(\Sigma / \mathbf{G})$ .
- 8: **else**
- 9:   Calculate whitened output:  $\hat{\mathbf{X}} = \hat{\mathbf{G}} \mathbf{X}$ .
- 10: **end if**
- 11: Recover representation:  $\mathbf{Y} = \phi_3(\hat{\mathbf{X}})$ .

---

COMPONENT	VALUE
Whitening transformation	{‘ZCA’, ‘PCA’, ‘CD’, ‘ItN’}
Estimation object	{‘ $\Sigma$ ’, ‘ $\mathbf{G}$ ’}
Recovery operation	{‘Scale & Shift’, ‘Coloring’}

Table 1. The scope of value in Algorithm 1 for different components this paper discusses. The Cartesian product of the values considers the landscape of the batch whitening algorithms used in this study.

There are two alternatives for  $\phi_3$ : One is a dimension-wise scale and shift [18, 15]:  $\mathbf{Y}_k = \gamma_k \hat{\mathbf{X}}_k + \beta_k$ , ( $k = 1, \dots, d$ ); The other is a coloring transformation:  $\mathbf{Y} = \mathbf{W} \hat{\mathbf{X}} + \mathbf{b}$ , which was proposed in [37] to achieve better performance in training GANs. By combining whitening transformation  $\phi_1$ , estimation object  $\phi_2$  and recovery operation  $\phi_3$ , we can design different BW algorithms. Table 1 shows the scope of value for different components this paper discusses. Note that ItN [16] is the efficient and numerical stable approximation of ZCA whitening. We use the ‘Scale & Shift’ operation in  $\phi_3$  for all algorithm, unless otherwise stated.

## 4.2. Investigation on Discriminative Scenarios

In the section, we first investigate how different whitening transformations affect the training of VGG network [38] for CIFAR-10 datasets [22], from the perspective of optimization. We then show the performance on large-scale ImageNet classification [7], by applying the designed BW algorithm on residual networks [11].

### 4.2.1 VGG on CIFAR-10

We use the VGG networks [38] tailored for  $32 \times 32$  inputs (16 convolutional layers and 1 fully-connected layer). We add the normalization methods after each convolutional layer in the VGGs. We compare several methods including: 1) The full whitening methods ‘PCA’, ‘ZCA’ and ‘CD’; 2) The approximate whitening ‘ItN’ and the standardization ‘BN’; 3) Group-based whitening methods with the group size ranging in  $\{512, 256, 128, 64, 32, 16\}$ . We denote ‘ZCA-16’ as ZCA whitening with group size 16. We focus on comparing the training performance from an optimization perspective, and we use mini-batch covariance matrix ‘ $\Sigma$ ’ as the estimation object for all methods during inference. We use SGD with a batch size of 256 to optimize the model. We set the initial learning rate to 0.1, then divide it by 5 after 60

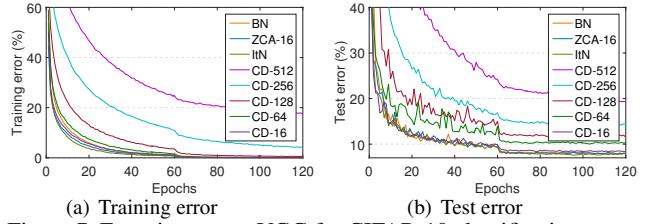


Figure 7. Experiments on VGG for CIFAR-10 classification.

epochs and finish the training at 120 epochs.

The main experimental observations include: 1) ‘PCA’ fails training under all the configurations, which means that either the training loss does not decrease or numerical instability appears. This observation is consistent with the previous MLP models. 2) ‘ZCA-16’ can train well, while other ‘ZCA’ related configurations fail training due to numerical instability. This is caused by the back-propagation through the eigen decomposition, which requires different eigenvalues for the mini-batch covariance matrix [15, 37]. 3) ‘CD’ has no numerical instability and can ensure full feature whitening of the models. Fully whitening features by ‘CD’ have significantly worse performance than the group-based ones shown in Figure 7. This again suggests that it is essential to control the extent of the whitening for discriminative models. 4) We find that ‘ItN’ (the approximates of the ZCA whitening) works the best.

### 4.2.2 Residual Network on ImageNet

We investigate the effectiveness of all kinds of whitening algorithms on residual networks for ImageNet classification with 1000 classes [7]. We use the given official 1.28M training images as a training set, and evaluate the top-1 accuracy on the validation set with 50k images.

**Ablation Study on ResNet-18** We first execute an ablation study on the 18-layer residual network (ResNet-18) to explore multiple positions for replacing BN with BW. We consider three architectures: 1) **ARC**<sub>A</sub>: where we only replace the first BN module of ResNet-18 proposed in [15]; 2) **ARC**<sub>B</sub>: where we further plug-in the BW layer after the last average pooling (before the last linear layer) to learn the decorrelated feature representations based on **ARC**<sub>A</sub>, as proposed in [16]; 3) **ARC**<sub>C</sub>: where we also replace the  $\{2n, n = 1, 2, \dots\}$ th BN modules (the  $\{3n\}$ th for ResNet-50) based on **ARC**<sub>B</sub>. We compare all the whitening transformations and estimation objects in Table 1. ‘ZCA’ and ‘ZCA $_{\Sigma}$ ’ denote the ZCA whitening that estimate the population statistics using  $\mathbf{G}$  and  $\Sigma$ , respectively. For ‘PCA’, ‘ZCA’ and ‘CD’, we use a group size in  $\{16, 64\}$  and report the best performance from these two configurations.

We follow the same experimental setup as described in [11], except that we use one GPU and train over 100 epochs. We apply SGD with a mini-batch size of 256, momentum of 0.9 and weight decay of 0.0001. The initial learning rate is set to 0.1 and divided by 10 at 30, 60 and 90 epochs.

The results are shown in Table 2. For **ARC**<sub>A</sub>, we find

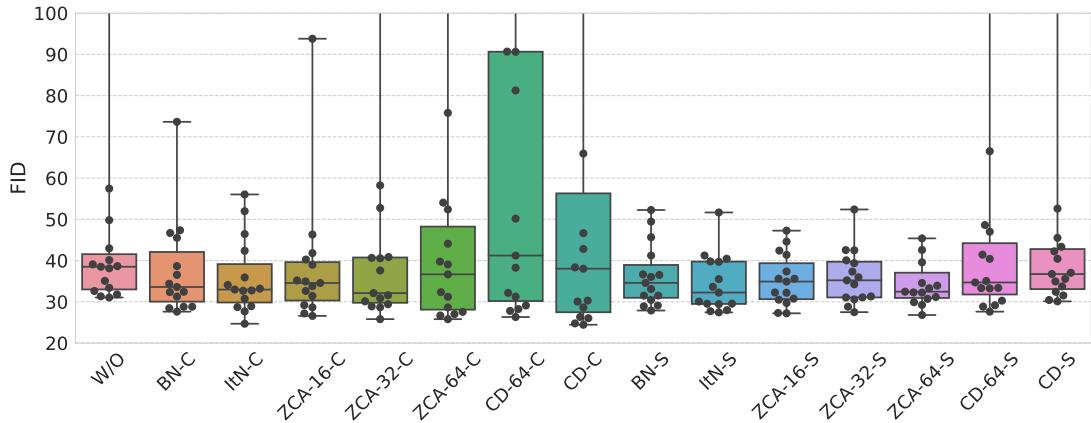


Figure 8. Stability experiments on GAN with hinge loss [29, 5] for unconditional image generation. The box shows the quartiles while the whiskers extend to show the rest of the distribution (we limit the FID ranging in (20,100) for better representation, and we show the full range of FID in the *supplementary materials*). Note that all methods use covariance matrix to estimate population statistics.

Method	ARC <sub>A</sub>	ARC <sub>B</sub>	ARC <sub>C</sub>
Baseline [11]	70.31	–	–
PCA [15]	59.93 ( $\downarrow$ 10.38)	–	–
PCA <sub><math>\Sigma</math></sub>	70.01 ( $\downarrow$ 0.30)	–	–
ZCA [15]	70.58 ( $\uparrow$ 0.27)	–	–
ZCA <sub><math>\Sigma</math></sub>	70.62 ( $\uparrow$ 0.31)	–	–
CD	70.46 ( $\uparrow$ 0.15)	70.80 ( $\uparrow$ 0.49)	68.15 ( $\downarrow$ 2.16)
CD <sub><math>\Sigma</math></sub> [37]	70.55 ( $\uparrow$ 0.24)	70.89 ( $\uparrow$ 0.58)	68.56 ( $\downarrow$ 1.75)
ItN [16]	70.62 ( $\uparrow$ 0.31)	71.14 ( $\uparrow$ 0.83)	71.26 ( $\uparrow$ 0.95)
ItN <sub><math>\Sigma</math></sub>	70.63 ( $\uparrow$ 0.32)	71.33 ( $\uparrow$ 1.02)	71.62 ( $\uparrow$ 1.31)

Table 2. Comparison of validation accuracy (%), single model and single-crop) on an 18-layer residual network for ImageNet.

that all whitening methods, except PCA related methods, improve the performance over the baselines. We observe that ZCA (ZCA <sub>$\Sigma$</sub> ) and its approximate ItN (ItN <sub>$\Sigma$</sub> ), achieve slightly better performance than CD (CD <sub>$\Sigma$</sub> ), and this observation is consistent with the results in [37]. This suggests ZCA whitening, which minimizes the distortion introduced by whitening under the L2 distance [15], usually works better than other whitening methods for discriminative classification tasks. Under ARC<sub>B</sub> and ARC<sub>C</sub>, ZCA/PCA related methods suffer from numerical instability. We also observe that CD-related methods have significantly degenerated performance under ARC<sub>C</sub>. This implies that the stochasticity introduced by decorrelating multiple layers with CD whitening harms the learning. We find that ItN-related methods can effectively control the stochasticity and achieve further performance improvement on ARC<sub>C</sub>. We try a ResNet-18 where all BN layers are replaced with ItN. However, the network has no performance improvement over ARC<sub>A</sub>, while the computations introduced are significant, as already observed in [16]. These results demonstrate that controlling the extent of whitening (stochasticity) is important for achieving performance improvements over the standardization.

From all the architectures and whitening methods, we observe that using  $\Sigma$  to estimate the population statistics is better than using  $\mathbf{G}$ , especially on ARC<sub>C</sub>.

**Results on ResNet-50** Based on the above observation, we further apply the ItN <sub>$\Sigma$</sub>  to ResNet-50. In addition to the standard step learning rate decay [11] used in the previous

Method	Step decay	Cosine decay
Baseline	76.20	76.62
ItN <sub><math>\Sigma</math></sub> -ARC <sub>B</sub>	77.18 ( $\uparrow$ 0.98)	77.68 ( $\uparrow$ 1.06)
ItN <sub><math>\Sigma</math></sub> -ARC <sub>C</sub>	<b>77.28</b> ( $\uparrow$ 1.08)	<b>77.92</b> ( $\uparrow$ 1.30)

Table 3. Results using ResNet-50 for ImageNet. We evaluate the top-1 validation accuracy (%), single model and single-crop).

setup, we also consider the cosine learning rate decay [26] that is also a basic setting [12] when training on ImageNet, and we want to address that the improvement of the proposed method can be obtained under different setups. For cosine decay, we start with a learning rate of 0.1 and decay it to 0.00001 over 100 epochs. We find that the proposed models significantly improves the performance over the original one, under all configurations. The additional time cost of ‘ItN <sub>$\Sigma$</sub> -ARC<sub>B</sub>’ and ‘ItN <sub>$\Sigma$</sub> -ARC<sub>C</sub>’ is 7.03% and 30.04% over the original. Note that ItN <sub>$\Sigma$</sub>  consistently provides a slight improvement (around 0.1 to 0.4 for all configurations in this experiment) over the original ItN proposed in [16], and please see *supplementary materials* for details. Moreover, ItN <sub>$\Sigma$</sub>  can bear a larger group size and smaller batch size than ItN, which is particularly important in the scenario of training GANs, as we will discuss.

### 4.3. Investigation on Training GANs

In this section, we provide an empirical analysis on BW algorithms for training GANs. We focus on investigating the effects of different BW algorithms in stabilizing the training and improving the performance. We evaluate on unconditional image generation for the CIFAR-10 dataset. In our qualitative evaluation of the generated samples, we use the two most common metrics: Fréchet Inception Distance (FID) [13] (the lower the better) and Inception Score (IS) [33] (the higher the better). Following the setup in [37], we only apply the BW on the generator. Unless otherwise noted, the batch whitening methods used in this section use the covariance matrix  $\Sigma$  to estimate the whitening matrix  $\widehat{\mathbf{G}}$ .

#### 4.3.1 Stability Experiments

Following the analysis in [23], we conduct experiments to demonstrate the effects of BW algorithms in stabilizing the

Method	Scale & Shift		Coloring	
	IS	FID	IS	FID
BN	7.243 $\pm$ 0.073	27.89	7.317 $\pm$ 0.093	27.61
CD	6.986 $\pm$ 0.065	30.11	<b>7.612 <math>\pm</math> 0.112</b>	<b>24.44</b>
ZCA-64	<b>7.265 <math>\pm</math> 0.069</b>	<b>26.8</b>	7.412 $\pm$ 0.118	25.79
ItN	7.246 $\pm$ 0.104	27.44	7.599 $\pm$ 0.089	24.68

Table 4. Results on stability experiments shown in Section 4.3.1. We report the best FID and the corresponding IS from all the configurations for different whitening methods.

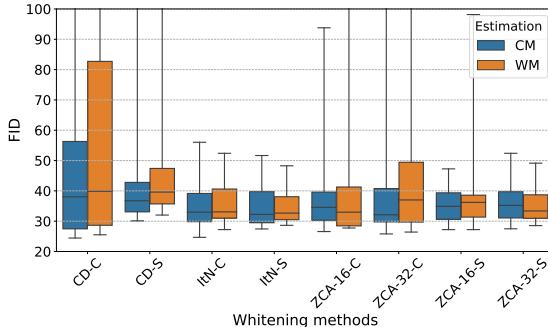


Figure 9. Comparison of BW methods between using Covariance Matrix (CM) and Whitening Matrix (WM) as estimation object.

training of GANs. We use DCGAN [31] architectures and use the hinge loss [29, 5]. We provide the implementation details in the *supplementary materials*.

We use the Adam optimizer [20] and train for 100 epochs. We consider 15 hyper-parameter settings (See *supplementary materials* for details) by varying the learning rate  $\alpha$ , first and second momentum ( $\beta_1, \beta_2$ ) of Adam, and number of discriminator updates per generator update  $n_{dis}$ . We use both the ‘Scale&Shift’ (denoted as ‘-S’) and ‘Coloring’ (denoted as ‘-C’) to recover the representation of BW. We calculate the FID distribution of the trained models under 15 configurations, and the results are shown in Figure 8. The lower the variance, the more stable the model from an optimization perspective. In Table 4, we also provide the best FID and the corresponding IS from all the configurations for different whitening methods. Note that full ZCA whitening also suffers from the numerical instability in this experiment, and we thus use the group-based ZCA whitening.

We observe that ‘CD’ whitening combined with coloring transformation (CD-C) achieves the best performances. However, it has the worst stability over all whitening methods. We argue that full feature whitening introduces strong stochasticity, which benefits the diversity of the generated examples. However, the strong stochasticity also harms the training and makes the model sensitive to the hyperparameter. We find that the coloring operation can improve the performances, especially for the whitening transformations (*e.g.*, BN benefits less). Nevertheless, it also makes the training unstable and sensitive to the hyperparameter. Besides, we observe that ‘ZCA-64’ achieves better performance than ‘CD-64’. This suggests that the advantage of ZCA whitening over CD in discriminative scenarios still exists when training GANS. We argue that controlling the magnitude of whiten-

Method	DCGAN		ResNet	
	IS	FID	IS	FID
CD	<b>7.820 <math>\pm</math> 0.099</b>	22.89	<b>8.522 <math>\pm</math> 0.113</b>	18.34
ZCA-64	7.672 $\pm$ 0.068	22.91	8.492 $\pm$ 0.106	17.86
ItN	7.652 $\pm$ 0.070	<b>22.5</b>	8.375 $\pm$ 0.093	<b>17.55</b>

Table 5. Performance comparison between different whitening methods on the DCGAN and ResNet architectures used in [37]. For fairness, we report all the results at the end of the training.

ing (stochasticity) is also still important in training GANs. For example, for ItN, the approximation of ZCA whitening has nearly the same performance as CD (Table 4), and shows better stability (Figure 8) due to its effective control of the stochasticity, as explained in [16].

We also have obtained similar results when using the non-saturated loss [10], as shown in the *supplementary material*.

**Comparison of Estimation Object** Considering that the CD whitening in [37] uses the covariance matrix to indirectly estimate the whitening matrix when training a GAN, we also compare the whitening methods with different estimation objects. The results are shown in Figure 9. We observe that the methods using the covariance matrix as the estimation object, achieve consistently better FID scores and are generally more stable than using the whitening matrix.

#### 4.3.2 Validation on Larger Architectures

Based on our previous analysis, we observe that ItN and ZCA whitening (with an appropriate group size) have similar performance to CD in training GANs, when using the covariance matrix as the estimation object. We further apply the ZCA whitening and ItN to the DCGAN and Resnet models, where CD whitening achieves nearly the state-of-the-art performance on CIFAR-10. We use the code provided in [37] and use the same setup as in [37]. We replace all the CD whitening in the models with ItN/ZCA whitening. The results are shown in Table 5. We observe that the CD whitening has no significant advantages over ItN/ZCA-G64. Ultimately, ItN achieves the best performances in terms of the FID evaluation, while CD whitening is the best in terms of IS. This suggests that there is still room to improve the performance of batch whitening in training GANs by further finely controlling the stochasticity.

## 5. Conclusions

In this paper, we provided a stochasticity analysis of batch whitening that thoroughly explains why different whitening transformations show significant differences in performance when training DNNs, despite their equivalent improvement of the conditioning. Our analysis provides insights for designing new normalization algorithms and constructing new network architectures. We believe that our analysis will open new avenues of research in better understanding normalization over batch data.

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

## References

- [1] Sanjeev Arora, Zhiyuan Li, and Kaifeng Lyu. Theoretical analysis of auto rate-tuning by batch normalization. In *ICLR*, 2019. 2
- [2] Andrei Atanov, Arsenii Ashukha, Dmitry Molchanov, Kirill Neklyudov, and Dmitry Vetrov. Uncertainty estimation via stochastic batch normalization. In *ICLR Workshop*, 2018. 2, 3
- [3] Lei Jimmy Ba, Ryan Kiros, and Geoffrey E. Hinton. Layer normalization. *CoRR*, abs/1607.06450, 2016. 1, 2, 5
- [4] Johan Bjorck, Carla Gomes, and Bart Selman. Understanding batch normalization. In *NIPS*, 2018. 1, 2
- [5] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *ICLR*, 2019. 7, 8
- [6] Minhyung Cho and Jaehyung Lee. Riemannian approach to batch normalization. In *NIPS*, 2017. 2
- [7] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR*, 2009. 2, 6
- [8] Guillaume Desjardins, Karen Simonyan, Razvan Pascanu, and koray kavukcuoglu. Natural neural networks. In *NIPS*, 2015. 2
- [9] Behrooz Ghorbani, Shankar Krishnan, and Ying Xiao. An investigation into neural net optimization via hessian eigenvalue density. In *ICML*, 2019. 2
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*. 2014. 8
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 1, 2, 6, 7
- [12] Tong He, Zhi Zhang, Hang Zhang, Zhongyue Zhang, Junyuan Xie, and Mu Li. Bag of tricks for image classification with convolutional neural networks. In *CVPR*, 2019. 7
- [13] 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 *NIPS*. 2017. 7
- [14] Elad Hoffer, Ron Banner, Itay Golan, and Daniel Soudry. Norm matters: efficient and accurate normalization schemes in deep networks. In *NIPS*, 2018. 2
- [15] Lei Huang, Dawei Yang, Bo Lang, and Jia Deng. Decorrelated batch normalization. In *CVPR*, 2018. 1, 2, 3, 4, 5, 6, 7
- [16] Lei Huang, Yi Zhou, Fan Zhu, Li Liu, and Ling Shao. Iterative normalization: Beyond standardization towards efficient whitening. In *CVPR*, 2019. 2, 3, 4, 6, 7, 8
- [17] Sergey Ioffe. Batch renormalization: Towards reducing mini-batch dependence in batch-normalized models. In *NIPS*, 2017. 2
- [18] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015. 1, 2, 3, 5, 6
- [19] Agnan Kessy, Alex Lewin, and Korbinian Strimmer. Optimal whitening and decorrelation. *The American Statistician*, 72(4):309–314, 2018. 1, 3
- [20] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014. 8
- [21] Jonas Kohler, Hadi Daneshmand, Aurelien Lucchi, Ming Zhou, Klaus Neymeyr, and Thomas Hofmann. Towards a theoretical understanding of batch normalization. *arXiv preprint arXiv:1805.10694*, 2018. 1
- [22] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009. 1, 6
- [23] Karol Kurach, Mario Lučić, Xiaohua Zhai, Marcin Michalski, and Sylvain Gelly. A large-scale study on regularization and normalization in GANs. In *ICML*, 2019. 7
- [24] Yann LeCun, Léon Bottou, Genevieve B. Orr, and Klaus-Robert Müller. Efficient backprop. In *Neural Networks: Tricks of the Trade*, pages 9–50, 1998. 1
- [25] Yijun Li, Chen Fang, Jimei Yang, Zhaowen Wang, Xin Lu, and Ming-Hsuan Yang. Universal style transfer via feature transforms. In *NIPS*, 2017. 2
- [26] Ilya Loshchilov and Frank Hutter. SGDR: stochastic gradient descent with restarts. In *ICLR*, 2017. 7
- [27] Ping Luo. Learning deep architectures via generalized whitened neural networks. In *ICML*, 2017. 2
- [28] Ping Luo, Jiamin Ren, and Zhanglin Peng. Differentiable learning-to-normalize via switchable normalization. *arXiv preprint arXiv:1806.10779*, 2018. 2
- [29] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *ICLR*, 2018. 7, 8
- [30] Xingang Pan, Xiaohang Zhan, Jianping Shi, Xiaou Tang, and Ping Luo. Switchable whitening for deep representation learning. In *ICCV*, 2019. 2
- [31] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *ICLR*, 2016. 8
- [32] Mengye Ren, Renjie Liao, Raquel Urtasun, Fabian H. Sinz, and Richard S. Zemel. Normalizing the normalizers: Comparing and extending network normalization schemes. In *ICLR*, 2017. 2
- [33] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, Xi Chen, and Xi Chen. Improved techniques for training gans. In *NIPS*, 2016. 7
- [34] Shibani Santurkar, Dimitris Tsipras, Andrew Ilyas, and Aleksander Madry. How does batch normalization help optimization? In *NIPS*, 2018. 1, 2
- [35] Alexander Shekhovtsov and Boris Flach. Stochastic normalizations as bayesian learning. In *ACCV*, 2018. 2
- [36] Lu Sheng, Ziyi Lin, Jing Shao, and Xiaogang Wang. Avatar-net: Multi-scale zero-shot style transfer by feature decoration. In *CVPR*, 2018. 2
- [37] Aliaksandr Siarohin, Enver Sangineto, and Nicu Sebe. Whitening and coloring batch transform for gans. In *ICLR*, 2019. 1, 2, 3, 5, 6, 7, 8
- [38] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 6
- [39] Saurabh Singh and Abhinav Shrivastava. Evalnorm: Estimating batch normalization statistics for evaluation. In *ICCV*, 2019. 2

- [40] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15(1):1929–1958, Jan. 2014. [2](#)
- [41] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, 2016. [1](#)
- [42] Mattias Teye, Hossein Azizpour, and Kevin Smith. Bayesian uncertainty estimation for batch normalized deep networks. In *ICML*, 2018. [2](#), [3](#)
- [43] Guangrun Wang, Jiefeng Peng, Ping Luo, Xinjiang Wang, and Liang Lin. Kalman normalization: Normalizing internal representations across network layers. In *NIPS*, 2018. [2](#)
- [44] Simon Wiesler and Hermann Ney. A convergence analysis of log-linear training. In *NIPS*, pages 657–665, 2011. [1](#)
- [45] Yuxin Wu and Kaiming He. Group normalization. In *ECCV*, 2018. [1](#), [2](#), [5](#)
- [46] Greg Yang, Jeffrey Pennington, Vinay Rao, Jascha Sohl-Dickstein, and Samuel S. Schoenholz. A mean field theory of batch normalization. In *ICLR*, 2019. [2](#)
- [47] Guodong Zhang, Chaoqi Wang, Bowen Xu, and Roger B. Grosse. Three mechanisms of weight decay regularization. In *ICLR*, 2019. [2](#)