

Decorrelated Batch Normalization

Lei Huang^{†‡*} Dawei Yang[†] Bo Lang[†] Jia Deng[‡]

[†]State Key Laboratory of Software Development Environment, Beihang University, P.R.China

[‡]University of Michigan, Ann Arbor

Abstract

Batch Normalization (BN) is capable of accelerating the training of deep models by centering and scaling activations within mini-batches. In this work, we propose Decorrelated Batch Normalization (DBN), which not just centers and scales activations but whitens them. We explore multiple whitening techniques, and find that PCA whitening causes a problem we call stochastic axis swapping, which is detrimental to learning. We show that ZCA whitening does not suffer from this problem, permitting successful learning. DBN retains the desirable qualities of BN and further improves BN’s optimization efficiency and generalization ability. We design comprehensive experiments to show that DBN can improve the performance of BN on multilayer perceptrons and convolutional neural networks. Furthermore, we consistently improve the accuracy of residual networks on CIFAR-10, CIFAR-100, and ImageNet.

1. Introduction

Batch Normalization [25] is a technique for accelerating deep network training. Introduced by Ioffe and Szegedy, it has been widely used in a variety of state-of-the-art systems [17, 49, 19, 54, 48, 20]. Batch Normalization works by standardizing the activations of a deep network within a mini-batch—transforming the output of a layer, or equivalently the input to the next layer, to have a zero mean and unit variance. Specifically, let $\{x_i \in \mathbb{R}, i = 1, 2, \dots, m\}$ be the original outputs of a single neuron on m examples in a mini-batch. Batch Normalization produces the transformed outputs

$$\hat{x}_i = \gamma \frac{x_i - \mu}{\sqrt{\sigma^2 + \epsilon}} + \beta, \quad (1)$$

where $\mu = \frac{1}{m} \sum_{j=1}^m x_j$ and $\sigma^2 = \frac{1}{m} \sum_{j=1}^m (x_j - \mu)^2$ are the mean and variance of the mini-batch, $\epsilon > 0$ is a small number to prevent numerical instability, and γ, β are extra learnable parameters. Crucially, during training, Batch Normalization

is part of both the inference computation (forward pass) as well as the gradient computation (backward pass). Batch Normalization can be inserted extensively into a network, typically between a linear mapping and a nonlinearity.

Batch Normalization was motivated by the well-known fact that whitening inputs (*i.e.* centering, decorrelating, and scaling) speeds up training [32]. It has been shown that better conditioning of the covariance matrix of the input leads to better conditioning of the Hessian in updating the weights, making the gradient descent updates closer to Newton updates [32, 50]. Batch Normalization exploits this fact further by seeking to whiten not only the input to the first layer of the network, but also the inputs to each internal layer in the network. But instead of whitening, Batch Normalization only performs standardization. That is, the activations are centered and scaled, but not decorrelated. Such a choice was justified in [25] by citing the cost and differentiability of whitening, but no actual attempts were made to derive or experiment with a whitening operation.

While standardization has proven effective for Batch Normalization, it remains an interesting question whether full whitening—adding decorrelation to Batch Normalization—can help further. Conceptually, there are clear cases where whitening is beneficial. For example, when the activations are close to being perfectly¹ correlated, standardization barely improves the conditioning of the covariance matrix, whereas whitening remains effective. In addition, prior work has shown that decorrelated activations result in better features [3, 43, 5] and better generalization [9, 53], suggesting room for further improving Batch Normalization.

In this paper, we propose Decorrelated Batch Normalization, in which we whiten the activations of each layer within a mini-batch. Let $\mathbf{x}_i \in \mathbb{R}^d$ be the input to a layer for the i -th example in a mini-batch of size m . The whitened input is given by

$$\hat{\mathbf{x}}_i = \Sigma^{-\frac{1}{2}}(\mathbf{x}_i - \mu), \quad (2)$$

where $\mu = \frac{1}{m} \sum_{j=1}^m \mathbf{x}_j$ is the mini-batch mean and $\Sigma = \frac{1}{m} \sum_{j=1}^m (\mathbf{x}_j - \mu)(\mathbf{x}_j - \mu)^T$ is the mini-batch covariance matrix.

*This work was mainly done while Lei Huang was a visiting student at the University of Michigan.

¹For example, in 2D, this means all points lie close to the line $y = x$ and Batch Normalization does not change the shape of the distribution.

Several questions arise for implementing Decorrelated Batch Normalization. One is how to perform back-propagation, in particular, how to back-propagate through the inverse square root of a matrix (*i.e.* $\partial\Sigma^{-\frac{1}{2}}/\partial\Sigma$), whose key step is an eigen decomposition. The differentiability of this matrix transform was one of the reasons that whitening was not pursued in the Batch Normalization paper [25]. Desjardins *et al.* [12] whiten the activations but avoid back-propagation through it by treating the mean μ and the whitening matrix $\Sigma^{-\frac{1}{2}}$ as model parameters, rather than as functions of the input examples. However, as has been pointed out [25, 24], doing so may lead to instability in training.

In this work, we decorrelate the activations and perform proper back-propagation during training. We achieve this by using the fact that eigen-decomposition is differentiable, and its derivatives can be obtained using matrix differential calculus, as shown by prior work [15, 26]. We build upon these existing results and derive the back-propagation updates for Decorrelated Batch Normalization.

Another question is, perhaps surprisingly, the choice of how to compute the whitening matrix $\Sigma^{-\frac{1}{2}}$. The whitening matrix is not unique because a whitened input stays whitened after an arbitrary rotation [27]. It turns out that PCA whitening, a standard choice [13], does not speed up training at all and in fact inflicts significant harm. The reason is that PCA whitening works by performing rotation followed by scaling, but the rotation can cause a problem we call *stochastic axis swapping*, which, as will be discussed in Section 3.1, in effect randomly permutes the neurons of a layer for each batch. Such permutation can drastically change the data representation from one batch to another to the extent that training never converges.

To address this stochastic axis swapping issue, we discover that it is critical to use ZCA whitening [4, 27], which rotates the PCA-whitened activations back such that the distortion of the original activations is minimal. We show through experiments that the benefits of decorrelation are only observed with the additional rotation of ZCA whitening.

A third question is the amount of whitening to perform. Given a particular batch size, DBN may not have enough samples to obtain a suitable estimate for the full covariance matrix. We thus control the extent of whitening by decorrelating smaller groups of activations instead of all activations together. That is, for an output of dimension d , we divide it into groups of size $k_G < d$ and apply whitening within each group. This strategy has the added benefit of reducing the computational cost of whitening from $O(d^2 \max(m, d))$ to $O(mdk_G)$, where m is the mini-batch size.

We conduct extensive experiments on multilayer perceptrons and convolutional neural networks, and show that Decorrelated Batch Normalization (DBN) improves upon the original Batch Normalization (BN) in terms of training speed and generalization performance. In particular, experiments

demonstrate that using DBN can consistently improve the performance of residual networks [17, 19, 54] on CIFAR-10, CIFAR-100 [29] and ILSVRC-2012 [11].

2. Related Work

Normalized activations [35, 51] and gradients [44, 38] have long been known to be beneficial for training neural networks. Batch Normalization [25] was the first to perform normalization per mini-batch in a way that supports back-propagation. One drawback of Batch Normalization, however, is that it requires a reasonable batch size to estimate the mean and variance, and is not applicable when the batch size is very small. To address this issue, Ba *et al.* [2] proposed *Layer Normalization*, which performs normalization on a single example using the mean and variance of the activations from the same layer. Batch Normalization and Layer Normalization were later unified by Ren *et al.* under the Division Normalization framework [39]. Other attempts to improve Batch Normalization for small batch sizes, include Batch Renormalization [24] and Stream Normalization [33]. There have also been efforts to adapt Batch Normalization to Recurrent Neural Networks [31, 10]. Our work extends Batch Normalization by decorrelating the activations, which is a direction orthogonal to all these prior works.

Our work is closely related to Natural Neural Networks [12, 34], which whiten activations by periodically estimating and updating a whitening matrix. Our work differs from Natural Neural Networks in two important ways. First, Natural Neural Networks perform whitening by treating the mean and the whitening matrix as model parameters as opposed to functions of the input examples, which, as pointed out by Ioffe & Szegedy [25, 24], can cause instability in training, with symptoms such as divergence or gradient explosion. Second, during training, a Natural Neural Network uses a running estimate of the mean and whitening matrix to perform whitening for each mini-batch; as a result, it cannot ensure that the transformed activations within each batch are in fact whitened, whereas in our case the activations with a mini-batch are guaranteed to be whitened. Natural Neural Networks thus may suffer instability in training very deep neural networks [12].

Another way to obtain decorrelated activations is to introduce additional regularization in the loss function [9, 53]. Cogswell *et al.* [9] introduced the DeCov loss on the activations as a regularizer to encourage non-redundant representations. Xiong *et al.* [53] extends [9] to learn group-wise decorrelated representations. Note that these methods are not designed for speeding up training. In fact, empirically they often slow down training [9], probably because decorrelated activations are part of the learning objective and thus may not be achieved until later in training.

Our approach is also related to work that implicitly normalizes activations by either normalizing the network

weights—e.g. through re-parameterization techniques [41, 23, 22], Riemannian optimization methods [21, 7], or additional weight regularization [30, 37, 40]—or by designing special scaling coefficients and bias values that can induce normalized activations under certain assumptions [1]. Our work bears some similarity to that of Huang *et al.* [22], which also back-propagates gradients through a ZCA-like normalization transform that involves eigen-decomposition. But the work by Huang *et al.* normalizes weights instead of activations, which leads to significantly different derivations especially with regards to convolutional layers; in addition, unlike ours, it does not involve a separately estimated whitening matrix during inference, nor does it discuss the stochastic axis swapping issue. Finally, all of these works including [22] are orthogonal to ours in the sense that their normalization is data independent, whereas ours is data dependent. In fact, as shown in [23, 7, 22, 21], data-dependent and data-independent normalization can be combined to achieve even greater improvement.

3. Decorrelated Batch Normalization

Let $\mathbf{X} \in \mathbb{R}^{d \times m}$ be a data matrix that represents inputs to a layer in a mini-batch of size m . Let $\mathbf{x}_i \in \mathbb{R}^d$ be the i -th column vector of \mathbf{X} , *i.e.* the d -dimensional input from the i -th example. The whitening transformation $\phi : \mathbb{R}^{d \times m} \rightarrow \mathbb{R}^{d \times m}$ is defined as

$$\phi(\mathbf{X}) = \Sigma^{-1/2}(\mathbf{X} - \mu \cdot \mathbf{1}^T), \quad (3)$$

where $\mu = \frac{1}{m}\mathbf{X} \cdot \mathbf{1}$ is the mean of \mathbf{X} , $\Sigma = \frac{1}{m}(\mathbf{X} - \mu \cdot \mathbf{1}^T)(\mathbf{X} - \mu \cdot \mathbf{1}^T)^T + \epsilon \mathbf{I}$ is the covariance matrix of the centered \mathbf{X} , $\mathbf{1}$ is a column vector of all ones, and $\epsilon > 0$ is a small positive number for numerical stability (preventing a singular Σ). The whitening transformation $\phi(\mathbf{X})$ ensures that for the transformed data $\hat{\mathbf{X}} = \phi(\mathbf{X})$ is whitened, *i.e.*, $\hat{\mathbf{X}}\hat{\mathbf{X}}^T = \mathbf{I}$.

Although Eqn. 3 gives an analytical form of the whitening transformation, this transformation is in fact not unique. The reason is that $\Sigma^{-1/2}$, the inverse square root of the covariance matrix, is defined only up to rotation, and as a result there exist infinitely many whitening transformations. Thus, a natural question is whether the specific choice of $\Sigma^{-1/2}$ matters, and if so, which choice to use.

To answer this question, we first discuss a phenomenon we call *stochastic axis swapping* and show that not all whitening transformations are equally desirable.

3.1. Stochastic Axis Swapping

Given a data point represented as a vector $\mathbf{x} \in \mathbb{R}^d$ under the standard basis, its representation under another orthogonal basis $\{\mathbf{d}_1, \dots, \mathbf{d}_d\}$ is $\hat{\mathbf{x}} = \mathbf{D}^T \mathbf{x}$, where $\mathbf{D} = [\mathbf{d}_1, \dots, \mathbf{d}_d]$ is an orthogonal matrix. We define *stochastic axis swapping* as follows:

Definition 3.1 Assume a training algorithm that iteratively update weights using a batch of randomly sampled data points per iteration. **Stochastic axis swapping** occurs when a data point \mathbf{x} is transformed to be $\hat{\mathbf{x}}_1 = \mathbf{D}_1^T \mathbf{x}$ in one iteration and $\hat{\mathbf{x}}_2 = \mathbf{D}_2^T \mathbf{x}$ in another iteration such that $\mathbf{D}_1 = \mathbf{P}\mathbf{D}_2$ where $\mathbf{P} \neq \mathbf{I}$ is a permutation matrix solely determined by the statistics of a batch.

Stochastic axis swapping makes training difficult, because the random permutation of the input dimensions can greatly confuse the learning algorithm—in the extreme case where the permutation is completely random, what remains is only a bag of activation values (similar to scrambling all pixels in an image), potentially resulting in an extreme loss of information and discriminative power.

Here, we demonstrate that the whitening of activations, if not done properly, can cause *stochastic axis swapping* in training neural networks. We start with standard PCA whitening [13], which computes $\Sigma^{-1/2}$ through eigen decomposition: $\Sigma_{pca}^{-1/2} = \Lambda^{-1/2}\mathbf{D}^T$, where $\Lambda = \text{diag}(\sigma_1, \dots, \sigma_d)$ and $\mathbf{D} = [\mathbf{d}_1, \dots, \mathbf{d}_d]$ are the eigenvalues and eigenvectors of Σ , *i.e.* $\Sigma = \mathbf{D}\Lambda\mathbf{D}^T$. That is, the original data point (after centering) is rotated by \mathbf{D}^T and then scaled by $\Lambda^{-1/2}$. Without loss of generalization, we assume that \mathbf{d}_i is unique by fixing the sign of its first element. A first opportunity for stochastic axis swapping is that the columns (or rows) of Λ and \mathbf{D} can be permuted while still giving a valid whitening transformation. But this is easy to fix—we can commit to a unique Λ and \mathbf{D} by ordering the eigenvalues non-increasingly.

But it turns out that ensuring a unique Λ and \mathbf{D} is insufficient to avoid stochastic axis swapping. Fig. 1 illustrates an example. Given a mini-batch of data points in one iteration as shown in Fig. 1(a), PCA whitening rotates them by $\mathbf{D}^T = [\mathbf{d}_1^T, \mathbf{d}_2^T]^T$ and stretches them along the new axis system by $\Lambda^{-1/2} = \text{diag}(1/\sqrt{\sigma_1}, 1/\sqrt{\sigma_2})$, where $\sigma_1 > \sigma_2$. Considering another iteration shown in Figure 1(b), where all data points except the red points are the same, it has the same eigenvectors with different eigenvalues, where $\sigma_1 < \sigma_2$. In this case, the new rotation matrix is $(\mathbf{D}')^T = [\mathbf{d}_2^T, \mathbf{d}_1^T]^T$ because we always order the eigenvalues non-increasingly. The blue data points thus have two different representations with the axes swapped.

To further justify our conjecture, we perform an experiment on multilayer perceptrons (MLPs) over the MNIST dataset as shown in Figure 2. We refer to the network without whitening activations as ‘plain’ and the network with PCA whitening as DBN-PCA. We find that DBN-PCA has significantly inferior performance to ‘plain’. Particularly, on the 4-layer MLP, DBN-PCA behaves similarly to random guessing, which implies that it causes severe *stochastic axis swapping*.

The stochastic axis swapping caused by PCA whitening exists because the rotation operation is executed over varied activations. Such variation is a result of two factors: (1) the




Figure 1. Illustration that PCA whitening suffers from stochastic axis swapping. (a) The axis alignment of PCA whitening in the initial iteration; (b) The axis alignment in another iteration.




Figure 2. Illustration of different whitening methods in training an MLP on MNIST. We use full batch gradient descent and report the best results with respect to the training loss among learning rates $\{0.1, 0.5, 1, 5\}$. (a) and (b) show the training loss of the 2-layer and 4-layer MLP, respectively. The number of neurons in each hidden layer is 100. We refer to the network without whitening activation as ‘plain’, with PCA whitening activation as DBN-PCA, and with ZCA whitening as DBN-ZCA.

activations can change due to weight updates during training, following the *internal covariate shift* described in [25]; (2) the optimization is based on random mini-batches, which means that each batch will contain a different random set of examples in each training epoch.

A similar phenomenon is also observed in [22]. In this work, PCA-style orthogonalization failed to learn orthogonal filters effectively in neural networks. However, no further analysis was provided to explain why this is the case.

3.2. ZCA Whitening

To address the stochastic axis swapping problem, one straightforward idea is to rotate the transformed input back using the same rotation matrix \mathbf{D} :

$$\Sigma^{-1/2} = \mathbf{D}\Lambda^{-1/2}\mathbf{D}^T. \quad (4)$$

In other words, we scale along the eigenvectors to get the whitened activations under the original axis system. Such whitening is known as ZCA whitening [4], and has been shown to minimize the distortion introduced by whitening under L2 distance [4, 27, 22]. We perform the same experiments with ZCA whitening as we did with PCA whitening: with MLPs on MNIST. Shown in Figure 2, ZCA whitening (referred to as DBN-ZCA) improves training performance significantly compared to no whitening (‘plain’) and DBN-PCA. This shows that ZCA whitening is critical to addressing

the stochastic axis swapping problem.

Back-propagation It is important to note that the back-propagation through ZCA whitening is non-trivial. In our DBN, the mean μ and the covariance Σ are not parameters of the whitening transform ϕ , but are functions of the mini-batch data \mathbf{X} . We need to back-propagate the gradients through ϕ as in [25, 22]. Here, we use the results from [26] to derive the back-propagation formulations of whitening:

$$\frac{\partial L}{\partial \Sigma} = \mathbf{D}\{(\mathbf{K}^T \odot (\mathbf{D}^T \frac{\partial L}{\partial \mathbf{D}})) + (\frac{\partial L}{\partial \Lambda})_{diag}\}\mathbf{D}^T, \quad (5)$$

where L is the loss function, $\mathbf{K} \in \mathbb{R}^{d \times d}$ is 0-diagonal with $\mathbf{K}_{ij} = \frac{1}{\sigma_i \sigma_j} [i \neq j]$, the \odot operator is element-wise matrix multiplication, and $(\frac{\partial L}{\partial \Lambda})_{diag}$ sets the off-diagonal elements of $\frac{\partial L}{\partial \Lambda}$ as zero. Detailed derivation can be found in the supplementary materials. Here we only show the simplified formulation:

$$\frac{\partial L}{\partial \tilde{\mathbf{x}}_i} = \left(\frac{\partial L}{\partial \tilde{\mathbf{x}}_i} - \mathbf{f} + \tilde{\mathbf{x}}_i^T \mathbf{S} - \tilde{\mathbf{x}}_i^T \mathbf{M} \right) \Lambda^{-1/2} \mathbf{D}^T, \quad (6)$$

where $\mathbf{S} = 2(\mathbf{K}^T \odot (\Lambda \mathbf{F}_c^T + \Lambda^{1/2} \mathbf{F}_c \Lambda^{1/2}))_{sym}$, $\mathbf{M} = (\mathbf{F}_c)_{diag}$, $\mathbf{F}_c = \frac{1}{m} (\sum_{i=1}^m \frac{\partial L}{\partial \tilde{\mathbf{x}}_i} \tilde{\mathbf{x}}_i^T)$, and $\mathbf{f} = \frac{1}{m} \sum_{i=1}^m \frac{\partial L}{\partial \tilde{\mathbf{x}}_i}$. The notation $(\cdot)_{sym}$ represents symmetrizing the corresponding matrix.

3.3. Training and Inference

Decorrelated Batch Normalization (DBN) is a data-dependent whitening transformation with back-propagation formulations. Like Batch Normalization [25], it can be inserted extensively into a network. Algorithms 1 and 2 describe the forward pass and the backward pass of our proposed DBN respectively. During training, the mean μ and the whitening matrix $\Sigma^{-1/2}$ are calculated within each mini-batch to ensure that the activations are whitened for each mini-batch. We also maintain the expected mean μ_E and the expected whitening matrix $\Sigma_E^{-1/2}$ for use during inference. Specifically, during training, we initialize μ_E as $\mathbf{0}$ and $\Sigma_E^{-1/2}$ as \mathbf{I} and update them by running average as described in Line 10 and 11 of Algorithm 1.

Normalizing the activations constrains the model’s capacity for representation. To remedy this, Ioffe and Szegedy [25] introduce extra learnable parameters γ and β in Eqn. 1. These learnable parameters often marginally improve the performance in our observation. For DBN, we also recommend to use learnable parameters. Specifically, the learnable parameters can be merged into the following ReLU activation [36], resulting in the Translated ReLU (TReLU) [52].

For a convolutional neural network, the input to the DBN transformation is $\mathbf{X}_C \in \mathbb{R}^{h \times w \times d \times m}$ where h and w indicate the height and width of feature maps, and d and m are the numbers of feature maps and examples respectively. Following [25], we view each spatial position of the feature map as

Algorithm 1 Forward pass of DBN for each iteration.

- 1: **Input:** mini-batch inputs $\{\mathbf{x}_i, i = 1, 2, \dots, m\}$, expected mean μ_E and expected projection matrix $\Sigma_E^{-1/2}$.
 - 2: **Hyperparameters:** ϵ , running average momentum λ .
 - 3: **Output:** the ZCA-whitened activations $\{\tilde{\mathbf{x}}_i, i = 1, 2, \dots, m\}$.
 - 4: calculate: $\mu = \frac{1}{m} \sum_{j=1}^m \mathbf{x}_j$.
 - 5: calculate: $\Sigma = \frac{1}{m} \sum_{j=1}^m (\mathbf{x}_j - \mu)(\mathbf{x}_j - \mu)^T + \epsilon \mathbf{I}$.
 - 6: execute eigenvalue decomposition: $\Sigma = \mathbf{D} \Lambda \mathbf{D}^T$.
 - 7: calculate PCA-whitening matrix: $\mathbf{U} = \Lambda^{-1/2} \mathbf{D}^T$.
 - 8: calculate PCA-whitened activation: $\tilde{\mathbf{x}}_i = \mathbf{U}(\mathbf{x}_i - \mu)$.
 - 9: calculate ZCA-whitened output: $\hat{\mathbf{x}}_i = \mathbf{D} \tilde{\mathbf{x}}_i$.
 - 10: update: $\mu_E \leftarrow (1 - \lambda) \mu_E + \lambda \mu$.
 - 11: update: $\Sigma_E^{-1/2} \leftarrow (1 - \lambda) \Sigma_E^{-1/2} + \lambda \mathbf{D} \mathbf{U}$.
-

Algorithm 2 Backward pass of DBN for each iteration.

- 1: **Input:** mini-batch gradients respect to whitened outputs $\{\frac{\partial L}{\partial \tilde{\mathbf{x}}_i}, i = 1, 2, \dots, m\}$. Other auxiliary data from respective forward pass: (1) eigenvalues; (2) $\tilde{\mathbf{x}}$; (3) \mathbf{D} .
 - 2: **Output:** the gradients respect to the inputs $\{\frac{\partial L}{\partial \mathbf{x}_i}, i = 1, 2, \dots, m\}$.
 - 3: calculate the gradients respect to $\tilde{\mathbf{x}}$: $\frac{\partial L}{\partial \tilde{\mathbf{x}}_i} = \frac{\partial L}{\partial \tilde{\mathbf{x}}_i} \mathbf{D}$.
 - 4: calculate $\mathbf{f} = \frac{1}{m} \sum_{i=1}^m \frac{\partial L}{\partial \tilde{\mathbf{x}}_i}^T$.
 - 5: calculate 0-diagonal \mathbf{K} matrix by $\mathbf{K}_{ij} = \frac{1}{\sigma_i - \sigma_j} [i \neq j]$.
 - 6: generate diagonal matrix Λ from eigenvalues.
 - 7: calculate $\mathbf{F}_c = \frac{1}{m} (\sum_{i=1}^m \frac{\partial L}{\partial \tilde{\mathbf{x}}_i}^T \tilde{\mathbf{x}}_i^T)$ and $\mathbf{M} = (\mathbf{F}_c)_{diag}$.
 - 8: calculate $\mathbf{S} = 2(\mathbf{K}^T \odot (\Lambda \mathbf{F}_c^T + \Lambda^{\frac{1}{2}} \mathbf{F}_c \Lambda^{\frac{1}{2}}))_{sym}$.
 - 9: calculate $\frac{\partial L}{\partial \mathbf{x}_i}$ by formula 6.
-

a sample. We thus unroll \mathbf{X}_C as $\mathbf{X} \in \mathbb{R}^{d \times (mhw)}$ with mhw examples and d feature maps. The whitening operation is performed over the unrolled \mathbf{X} .

3.4. Group Whitening

As discussed in Section 1, it is necessary to control the extent of whitening such that there are sufficient examples in a batch for estimating the whitening matrix. To do this we use “group whitening”, specifically, we divide the activations along the feature dimension with size d into smaller groups of size k_G ($k_G < d$) and perform whitening within each group. The extent of whitening is controlled by the hyperparameter k_G . In the case $k_G = 1$, Decorrelated Batch Normalization reduces to the original Batch Normalization.

In addition to controlling the extent of whitening, group whitening reduces the computational complexity [22]. Full whitening costs $O(d^2 \max(m, d))$ for a batch of size m . When using group whitening, the cost is reduced to $O(\frac{d}{k_G} (k_G^2 (\max(m, k_G))))$. Typically, we choose $k_G < m$, therefore the cost of group whitening is $O(mdk_G)$.

3.5. Analysis and Discussion

DBN extends BN such that the activations are decorrelated over mini-batch data. DBN thus inherits the beneficial properties of BN, such as the ability to perform efficient training with large learning rates and very deep networks. Here, we further highlight the benefits of DBN over BN, in particular achieving better dynamical isometry [42] and improved conditioning.

Approximate Dynamical Isometry Saxe *et al.* [42] introduce dynamical isometry—the desirable property that occurs when the singular values of the product of Jacobians lie within a small range around 1. Enforcing this property, even approximately, is beneficial to training because it preserves the gradient magnitudes during back-propagation and alleviates the vanishing and exploding gradient problems [42]. Ioffe and Szegedy [25] find that Batch Normalization achieves approximate dynamical isometry under the assumption that (1) the transformation between two consecutive layers is approximately linear, and (2) the activations in each layer are Gaussian and uncorrelated. Our DBN inherently satisfies the second assumption, and therefore is more likely to achieve dynamical isometry than BN.

Improved Conditioning [12] demonstrated that whitening activations results in a block diagonal Fisher Information Matrix (FIM) for each layer under certain assumptions [16]. Their experiments show that such a block diagonal structure in the FIM can improve the conditioning. The proposed method in [12], however, cannot whiten the activations effectively, as shown in [34] and also discussed in Section 2. DBN, on the other hand, does this directly. Therefore, we conjecture that DBN can further improve the conditioning of the FIM, and we justify this experimentally in Section 4.1.

4. Experiments

We start with experiments to highlight the effectiveness of Decorrelated Batch Normalization (DBN) in improving the conditioning and speeding up convergence on multilayer perceptrons (MLP). We then conduct comprehensive experiments to compare DBN and BN on convolutional neural networks (CNNs). In the last section, we apply our DBN to residual networks on CIFAR-10, CIFAR-100 [29] and ILSVRC-2012 to show its power to improve modern network architectures. The code to reproduce the experiments is available at <https://github.com/huangleiBuaa/DecorrelatedBN>.

We focus on classification tasks and the loss function is the negative log-likelihood: $-\log P(\mathbf{y}|\mathbf{x})$. Unless otherwise stated, we use random weight initialization as described in [32] and ReLU activations [36].

4.1. Ablation Studies on MLPs

In this section, we verify the effectiveness of our proposed method in improving conditioning and speeding up conver-




Figure 3. Conditioning analysis with MLPs trained on the Yale-B dataset. (a) Condition number (log-scale) of relative FIM as a function of updates in the last layer; (b) training loss with respect to wall clock time.




Figure 4. Experiments on MLP architecture over PIE dataset. (a) The effects of group size of DBN, where ‘Gn’ indicates $k_G = n$; (b) Comparison of training loss with respect to wall clock time.

gence on MLPs. We also discuss the effect of the group size on the tradeoff between the performance and computation cost. We compare against several baselines, including the original network without any normalization (referred to as ‘plain’), Natural Neural Networks (NNN) [12], Layer Normalization (LN) [2], and Batch Normalization (BN) [25]. All results are averaged over 5 runs.

Conditioning Analysis We perform conditioning analysis on the Yale-B dataset [14], specifically, the subset [6] with 2,414 images and 38 classes. We resize the images to 32×32 and reshape them as 1024-dimensional vectors. We then convert the images to grayscale in the range $[0, 1]$ and subtract the per-pixel mean.

For each method, we train a 5-layer MLP with the numbers of neurons in each hidden layer $= \{128, 64, 48, 48\}$ and use full batch gradient descent. Hyper-parameters are selected by grid search based on the training loss. For all methods, the learning rate is chosen from $\{0.1, 0.2, 0.5, 1, 2\}$. For NNN, the revised term ε is one of $\{0.001, 0.01, 0.1, 1\}$ and the natural re-parameterization interval T is one of $\{20, 50, 100, 200, 500\}$.

We evaluate the condition number of the relative Fisher Information Matrix (FIM) [47] with respect to the last layer. Figure 3 (a) shows the evolution of the condition number over training iterations. Figure 3 (b) shows the training loss over the wall clock time. Note that the experiments are performed on CPUs and the model with DBN is $2 \times$ slower than the model with BN per iteration. From both figures, we see that NNN, LN, BN and DBN converge faster, and achieve better conditioning compared to ‘plain’. This shows

that normalization is able to make the optimization problem easier. Also, DBN achieves the best conditioning compared to other normalization methods, and speeds up convergence significantly.

Effects of Group Size As discussed in Section 3.4, the group size k_G controls the extent of whitening. Here we show the effects of the hyperparameter k_G on the performance of DBN. We use a subset [6] of the PIE face recognition [45] dataset with 68 classes with 11,554 images. We adopt the same pre-processing strategy as with Yale-B.

We trained a 6-layer MLP with the numbers of neurons in each hidden layer $= \{128, 128, 128, 128, 128\}$. We use Stochastic Gradient Descent (SGD) with a batch size of 256. Other configurations are chosen in the same way as the previous experiment. Additionally, we explore group sizes in $\{1, 8, 16, 32, 64, 128\}$ for DBN. Note that when $k_G = 1$, DBN is reduced to the original BN without the extra learnable parameters.

Figure 4 (a) shows the training loss of DBN with different group sizes. We find that the largest (G128) and smallest group sizes (G1) both have noticeably slower convergence compared to the ones with intermediate group sizes such as G16. These results show that (1) decorrelating activations over a mini-batch can improve optimization, and (2) controlling the extent of whitening is necessary, as the estimate of the full whitening matrix might be poor over mini-batch samples. Also, the eigendecomposition with small group sizes (e.g. 16) is less computationally expensive. We thus recommend using group whitening in training deep models.

We also compared DBN with group whitening ($k_G = 16$) to other baselines and the results are shown in Figure 4 (b). We find that DBN converges significantly faster than other normalization methods.

4.2. Experiments on CNNs

We design comprehensive experiments to evaluate the performance of DBN with CNNs against BN, the state-of-the-art normalization technique. For these experiments we use the CIFAR-10 dataset [29], which contains 10 classes, 50k training images, and 10k test images.

4.2.1 Comparison of DBN and BN

We compare DBN to BN over different experimental configurations, including the choice of optimization method, non-linearities, and the position of DBN/BN in the network. We adopt the VGG-A architecture [46] for all experiments, and pre-process the data by subtracting the per-pixel mean and dividing by the variance.

We use SGD with a batchsize of 256, momentum of 0.9 and weight decay of 0.0005. We decay the learning rate by half every T iterations. The hyper-parameters are chosen by grid search over a random validation set of 5k examples




Figure 5. Comprehensive performance comparison between DBN and BN with the VGG-A architecture on CIFAR-10. We show the training accuracy (solid line) and test accuracy (line marked with plus) for each epoch.




Figure 6. DBN can make optimization easier and benefits from a higher learning rate. Results are reported on the *S-plain* architecture over the CIFAR-10 dataset. (a) Comparison by varying the depth of network. ‘ nL ’ means the network has n layers; (b) Comparison by using a higher learning rate.

taken from the training set. The grid search includes the initial learning rate $lr = \{1, 2, 4, 8\}$ and the decay interval $T = \{1000, 2000, 4000, 8000\}$. We set the group size of DBN as $k_G = 16$. Figure 5 (a) compares the performance of BN and DBN under this configuration.

We also experiment with other configurations, including using (1) Adam [28] as the optimization method, (2) replacing ReLU with another widely used non-linearity called Exponential Linear Units (ELU) [8], and (3) inserting BN/DBN after the non-linearity. All the experimental setups are otherwise the same, except that Adam [28] is used with an initial learning rate in $\{0.001, 0.005, 0.01, 0.05\}$. The respective results are shown in Figure 5 (b), (c) and (d).

In all configurations, DBN converges faster with respect to the epochs and generalizes better, compared to BN. Particularly, in the four experiments above, DBN reduces the absolute test error by 0.61%, 1.34%, 1.44% and 0.38% respectively. The results demonstrate that our Decorrelated Batch Normalization outperforms Batch Normalization in terms of optimization quality and regularization ability.

4.2.2 Analyzing the Properties of DBN

We conduct experiments to support the conclusions from Section 3.5, specifically that DBN has better stability and converges faster than BN with high learning rates in very deep networks. The experiments were conducted on the *S-plain* network, which follows the design of the residual network [18] but removes the identity maps and uses the same feature maps for simplicity.

Method	Res-20	Res-32	Res-44	Res-56
Baseline*	8.75	7.51	7.17	6.97
Baseline	7.94	7.31	7.17	7.21
DBN-L1	7.94	7.28	6.87	6.63
DBN-scale-L1	7.77	6.94	6.83	6.49

Table 1. Comparison of test errors (%) with residual networks on CIFAR-10. ‘Res- L ’ indicates residual network with L layers, and ‘Baseline*’ indicates the results reported in [17] with only one run. Our results are averaged over 5 runs.

Going Deeper He *et al.* [18] addressed the degradation problem for the network without identity mappings: that is, when the network depth increases, the training accuracy degrades rapidly, even when Batch Normalization is used. In our experiments, we demonstrate that DBN will relieve this problem to some extent. In other words, a model with DBN is easier to optimize. We validate this on the *S-plain* architecture with feature maps of dimension $d = 48$ and number of layers 20, 32 and 44. The models are trained with a mini-batch size of 128, momentum of 0.9 and weight decay of 0.0005. We set the initial learning rate to be 0.1, dividing it by 5 at 80 and 120 epochs, and end training at 160 epochs. The results in Figure 6 (a) show that, with increased depth, the model with BN was more difficult to optimize than with DBN. We conjecture that the approximate dynamical isometry of DBN alleviates this problem.

Higher Learning Rate A network with Batch Normalization can benefit from high learning rates, and thus faster training, because it reduces internal covariate shift [25]. Here, we show that DBN can help even more. We train the networks with BN and DBN with a $4\times$ higher learning rate — 0.4. We use the *S-plain* architecture with feature maps of dimensions $d = 42$ and divide the learning rate by 5 at 60, 100, 140, and 180 epochs. The results in Figure 6 (b) show that DBN has significantly better training accuracy than BN. We argue that DBN benefits from higher learning rates because of its property of improved conditioning.

4.3. Applying DBN to Residual Network in Practice

Due to our current un-optimized implementation of DBN, it would incur a high computational cost to replace all BN

Method	CIFAR-10			CIFAR-100		
	Baseline* [54]	Baseline	DBN-scale-L1	Baseline* [54]	Baseline	DBN-scale-L1
WRN-28-10	3.89	3.99 ± 0.13	3.79 ± 0.09	18.85	18.75 ± 0.28	18.36 ± 0.17
WRN-40-10	3.80	3.80 ± 0.11	3.74 ± 0.11	18.3	18.7 ± 0.22	18.27 ± 0.19

Table 2. Test errors (%) on wide residual networks over CIFAR-10 and CIFAR-100. ‘Baseline’ and ‘DBN-scale-L1’ refer to the results we perform, based on the released code of paper [54], and the results are shown in the format of ‘mean ± std’ computed over 5 random seeds. ‘Baseline*’ refers to the results reported by authors of [54] on their Github. They report the median of 5 runs on WRN-28-10 and only perform one run on WRN-40-10.

modules of a network with DBN². We instead only decorrelate the activations among a subset of layers. We find that this in practice is already effective for residual networks [17], because the information in previous layers can pass directly to the later layers through the identity connections. We also show that we can improve upon residual networks [17, 19, 54] by using only one DBN module before the first residual block, which introduces negligible computation cost. In principle, an optimized implementation of DBN will be much faster, and could be injected in multiple places in the network with little overhead. However, optimizing the implementation of DBN is beyond the scope of this work.

Residual Network on CIFAR-10 We apply our method on residual networks [17] by using only one DBN module before the first residual block (denoted as DBN-L1). We also consider DBN with adjustable scale (denoted as DBN-scale-L1) as discussed in Section 3.3. We adopt the Torch implementation of residual networks³ and follow the same experimental protocol as described in [17]. We train the residual networks with depth 20, 32, 44 and 56 on CIFAR-10. Table 1 shows the test errors of these networks. Our methods obtain lower test errors compared to BN over all 4 networks, and the improvement is more dramatic for deeper networks. Also, we see that DBN-scale-L1 marginally outperforms DBN-L1 in all cases. Therefore, we focus on comparing DBN-scale-L1 to BN in later experiments.

Wide Residual Network on CIFAR We apply DBN to Wide Residual Network (WRN) [54] to improve the performance on CIFAR-10 and CIFAR-100. Following the convention set in [54], we use the abbreviation WRN-d-k to indicate a WRN with depth d and width k . We again adopt the publicly available Torch implementation⁴ and follow the same setup as in [54]. The results in Table 2 show that DBN improves the original WRN on both datasets and both networks. In particular, we reduce the test error by 3.74% and 18.27% on CIFAR-10 and CIFAR-100, respectively.

Residual Network on ILSVRC-2012 We further validate the scalability of our method on ILSVRC-2012 with 1000 classes [11]. We use the given official 1.28M training images

²See the supplementary material for more details on computational cost.

³<https://github.com/facebook/fb.resnet.torch>

⁴<https://github.com/szagoruyko/wide-residual-networks>

Method	Res-50		Res-101	
	Top-1	Top-5	Top-1	Top-5
Baseline*	24.70	7.80	23.60	7.10
Baseline	24.87	7.58	22.54	6.38
DBN-scale-L1	24.29	7.08	22.17	6.09

Table 3. Comparison of test errors (% , single model and single-crop) on 50 and 101-layer residual networks on ILSVRC-2012. ‘Baseline*’ indicates that the results are obtained from the website: <https://github.com/KaimingHe/deep-residual-networks>

as a training set, and evaluate the top-1 and top-5 classification errors on the validation set with 50k images. We use the 50 and 101-layer residual network (Res-50 and Res-101) and perform single model and single-crop testing. We follow the same experimental setup as described in [17], except that we used 4 GPUs instead of 8 GPUs for training Res-50: we apply SGD with a mini-batch size of 256 over 4 GPUs for Res-50 and 8 GPUs for Res-101, momentum of 0.9 and weight decay of 0.0001; we set the initial learning rate of 0.1, dividing it by 10 at 30 and 60 epochs, and end the training at 90 epochs. The results are shown in Table 3. We can see that the DBN-scale-L1 achieves lower test errors compared to the original residual networks.

Conclusions

In this paper, we propose Decorrelated Batch Normalization (DBN), which extends Batch Normalization to include whitening over mini-batch data. We find that PCA whitening can sometimes be detrimental to training because it causes *stochastic axis swapping*, and demonstrate that it is critical to use ZCA whitening, which avoids this issue. DBN retains the advantages of Batch Normalization while using decorrelated representations to further improve models’ optimization efficiency and generalization abilities. This is because DBN can maintain approximate dynamical isometry and improve the conditioning of the Fisher Information Matrix. These properties are experimentally validated, suggesting DBN has great potential to be used in designing DNN architectures.

Acknowledgement This work was partially supported by China Scholarship Council, NSFC-61370125 and SKLSDE-2017ZX-03. We also thank Jonathan Stroud and Lanlan Liu for their help with proofreading and editing.

References

- [1] D. Arpit, Y. Zhou, B. U. Kota, and V. Govindaraju. Normalization propagation: A parametric technique for removing internal covariate shift in deep networks. In *ICML*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 1168–1176. JMLR.org, 2016. 3
- [2] L. J. Ba, R. Kiros, and G. E. Hinton. Layer normalization. *CoRR*, abs/1607.06450, 2016. 2, 6
- [3] H. B. Barlow. *Possible principles underlying the transformations of sensory messages*, pages 217–234. MIT Press, Cambridge, MA, 1961. 1
- [4] A. J. Bell and T. J. Sejnowski. The “independent components” of natural scenes are edge filters. *Vision research*, 37(23):3327–3338, Dec. 1997. 2, 4
- [5] Y. Bengio and J. S. Bergstra. Slow, decorrelated features for pretraining complex cell-like networks. In *Advances in Neural Information Processing Systems 22*, pages 99–107. 2009. 1
- [6] D. Cai, X. He, Y. Hu, J. Han, and T. Huang. Learning a spatially smooth subspace for face recognition. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition Machine Learning (CVPR’07)*, 2007. 6
- [7] M. Cho and J. Lee. Riemannian approach to batch normalization. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5225–5235. Curran Associates, Inc., 2017. 3
- [8] D. Clevert, T. Unterthiner, and S. Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). 2016. 7
- [9] M. Cogswell, F. Ahmed, R. B. Girshick, L. Zitnick, and D. Batra. Reducing overfitting in deep networks by decorrelating representations. In *ICLR*, 2016. 1, 2
- [10] T. Cozijmans, N. Ballas, C. Laurent, and A. C. Courville. Recurrent batch normalization. In *ICLR*, 2017. 2
- [11] 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, 8
- [12] G. Desjardins, K. Simonyan, R. Pascanu, and k. kavukcuoglu. Natural neural networks. In *Advances in Neural Information Processing Systems 28*, pages 2071–2079, 2015. 2, 5, 6
- [13] G. Desjardins, K. Simonyan, R. Pascanu, and K. Kavukcuoglu. Natural neural networks. In *Proceedings of the 28th International Conference on Neural Information Processing Systems, NIPS’15*, pages 2071–2079, 2015. 2, 3
- [14] A. Georghiades, P. Belhumeur, and D. Kriegman. From few to many: Illumination cone models for face recognition under variable lighting and pose. *IEEE Trans. Pattern Anal. Mach. Intelligence*, 23(6):643–660, 2001. 6
- [15] M. B. Giles. *Collected Matrix Derivative Results for Forward and Reverse Mode Algorithmic Differentiation*. 2008. 2
- [16] R. B. Grosse and J. Martens. A kronecker-factored approximate fisher matrix for convolution layers. In *ICML*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 573–582. JMLR.org, 2016. 5
- [17] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *arXiv preprint arXiv:1506.01497*, 2015. 1, 2, 7, 8
- [18] K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *ICCV*. IEEE Computer Society, 2015. 7
- [19] K. He, X. Zhang, S. Ren, and J. Sun. Identity mappings in deep residual networks. In B. Leibe, J. Matas, N. Sebe, and M. Welling, editors, *Computer Vision – ECCV 2016*, pages 630–645, Cham, 2016. Springer International Publishing. 1, 2, 8
- [20] G. Huang, Z. Liu, and K. Q. Weinberger. Densely connected convolutional networks. *CoRR*, abs/1608.06993, 2016. 1
- [21] L. Huang, X. Liu, B. Lang, and B. Li. Projection based weight normalization for deep neural networks. *CoRR*, abs/1710.02338, 2017. 3
- [22] L. Huang, X. Liu, B. Lang, A. W. Yu, Y. Wang, and B. Li. Orthogonal weight normalization: Solution to optimization over multiple dependent stiefel manifolds in deep neural networks. In *AAAI*, 2018. 3, 4, 5
- [23] L. Huang, X. Liu, Y. Liu, B. Lang, and D. Tao. Centered weight normalization in accelerating training of deep neural networks. In *ICCV*, 2017. 3
- [24] S. Ioffe. Batch renormalization: Towards reducing mini-batch dependence in batch-normalized models. *CoRR*, abs/1702.03275, 2017. 2
- [25] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015*, 2015. 1, 2, 4, 5, 6, 7
- [26] C. Ionescu, O. Vantzos, and C. Sminchisescu. Training deep networks with structured layers by matrix backpropagation. In *Proceedings of International Conference on Computer Vision, ICCV 2015*, 2015. 2, 4
- [27] A. Kessy, A. Lewin, and K. Strimmer. Optimal whitening and decorrelation. *The American Statistician*, 0(ja):0–0, 2017. 2, 4
- [28] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014. 7
- [29] A. Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009. 2, 5, 6
- [30] A. Krogh and J. A. Hertz. A simple weight decay can improve generalization. In *NIPS*. 1992. 3
- [31] C. Laurent, G. Pereyra, P. Brakel, Y. Zhang, and Y. Bengio. Batch normalized recurrent neural networks. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2016*, pages 2657–2661, 2016. 2
- [32] Y. LeCun, L. Bottou, G. B. Orr, and K.-R. Müller. Efficient backprop. In *Neural Networks: Tricks of the Trade, This Book is an Outgrowth of a 1996 NIPS Workshop*, pages 9–50, 1998. 1, 5
- [33] Q. Liao, K. Kawaguchi, and T. Poggio. Streaming normalization: Towards simpler and more biologically-plausible normalizations for online and recurrent learning. *arXiv preprint arXiv:1610.06160*, 2016. 2
- [34] P. Luo. Learning deep architectures via generalized whitened neural networks. In *Proceedings of the 34th International*

- Conference on Machine Learning*, pages 2238–2246, 2017. [2](#), [5](#)
- [35] G. Montavon and K.-R. Müller. *Deep Boltzmann Machines and the Centering Trick*, volume 7700 of *LNCS*. Springer, 2nd edn edition, 2012. [2](#)
- [36] V. Nair and G. E. Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on Machine Learning ICML 2010*, 2010. [4](#), [5](#)
- [37] B. Neyshabur, R. Salakhutdinov, and N. Srebro. Path-sgd: Path-normalized optimization in deep neural networks. In *Annual Conference on Neural Information Processing Systems NIPS 2015*, pages 2422–2430, 2015. [3](#)
- [38] T. Raiko, H. Valpola, and Y. LeCun. Deep learning made easier by linear transformations in perceptrons. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 924–932, 2012. [2](#)
- [39] M. Ren, R. Liao, R. Urtasun, F. H. Sinz, and R. S. Zemel. Normalizing the normalizers: Comparing and extending network normalization schemes. 2017. [2](#)
- [40] P. Rodríguez, J. González, G. Cucurull, J. M. Gonfaus, and F. X. Roca. Regularizing cnns with locally constrained decorrelations. 2017. [3](#)
- [41] T. Salimans and D. P. Kingma. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. *CoRR*, abs/1602.07868, 2016. [3](#)
- [42] A. M. Saxe, J. L. McClelland, and S. Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *CoRR*, abs/1312.6120, 2013. [5](#)
- [43] J. Schmidhuber. Learning factorial codes by predictability minimization. *Neural Computation*, 4(6):863–879, 1992. [1](#)
- [44] N. N. Schraudolph. Accelerated gradient descent by factor-centering decomposition. Technical report, 1998. [2](#)
- [45] T. Sim, S. Baker, and M. Bsat. The cmu pose, illumination, and expression (pie) database. In *Proceedings of the Fifth IEEE International Conference on Automatic Face and Gesture Recognition, FGR '02*, pages 53–, 2002. [6](#)
- [46] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014. [6](#)
- [47] K. Sun and F. Nielsen. Relative natural gradient for learning large complex models. *CoRR*, abs/1606.06069, 2016. [6](#)
- [48] C. Szegedy, S. Ioffe, and V. Vanhoucke. Inception-v4, inception-resnet and the impact of residual connections on learning. *CoRR*, abs/1602.07261, 2016. [1](#)
- [49] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the inception architecture for computer vision. *CoRR*, abs/1512.00567, 2015. [1](#)
- [50] S. Wiesler and H. Ney. A convergence analysis of log-linear training. In *NIPS*, pages 657–665, 2011. [1](#)
- [51] S. Wiesler, A. Richard, R. Schlüter, and H. Ney. Mean-normalized stochastic gradient for large-scale deep learning. In *ICASSP*, pages 180–184. IEEE, 2014. [2](#)
- [52] S. Xiang and H. Li. On the effects of batch and weight normalization in generative adversarial networks. *CoRR*, abs/1704.03971, 2017. [4](#)
- [53] W. Xiong, B. Du, L. Zhang, R. Hu, and D. Tao. Regularizing deep convolutional neural networks with a structured decorrelation constraint. In *IEEE 16th International Conference on Data Mining, ICDM 2016*, 2016. [1](#), [2](#)
- [54] S. Zagoruyko and N. Komodakis. Wide residual networks. *CoRR*, abs/1605.07146, 2016. [1](#), [2](#), [8](#)