This ICCV paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Switchable Whitening for Deep Representation Learning

Xingang Pan¹, Xiaohang Zhan¹, Jianping Shi², Xiaoou Tang¹, and Ping Luo^{1,3}

¹CUHK-SenseTime Joint Lab, The Chinese University of Hong Kong ²SenseTime Group Limited ³The University of Hong Kong ⁷ Strong stand standard and standard s

{px117, zx017, xtang, pluo}@ie.cuhk.edu.hk, shijianping@sensetime.com

Abstract

Normalization methods are essential components in convolutional neural networks (CNNs). They either standardize or whiten data using statistics estimated in predefined sets of pixels. Unlike existing works that design normalization techniques for specific tasks, we propose Switchable Whitening (SW), which provides a general form unifying different whitening methods as well as standardization methods. SW learns to switch among these operations in an end-to-end manner. It has several advantages. First, SW adaptively selects appropriate whitening or standardization statistics for different tasks (see Fig.1), making it well suited for a wide range of tasks without manual design. Second, by integrating the benefits of different normalizers, SW shows consistent improvements over its counterparts in various challenging benchmarks. Third, SW serves as a useful tool for understanding the characteristics of whitening and standardization techniques.

We show that SW outperforms other alternatives on image classification (CIFAR-10/100, ImageNet), semantic segmentation (ADE20K, Cityscapes), domain adaptation (GTA5, Cityscapes), and image style transfer (COCO). For example, without bells and whistles, we achieve state-ofthe-art performance with 45.33% mIoU on the ADE20K dataset.

1. Introduction

Normalization methods have been widely used as a basic module in convolutional neural networks (CNNs). In various applications, different normalization techniques like Batch Normalization (BN) [10], Instance Normalization (IN) [27] and Layer Normalization (LN) [1] are proposed. These normalization techniques generally perform *standardization* that *centers* and *scales* features. Nevertheless, the features are not *decorrelated*, hence their correlation still exists.

Another type of normalization methods is whitening,



Figure 1. (a) SW outperforms its counterparts in a variety of benchmarks. (b) SW learns to select appropriate whitening or standardization methods in different tasks and datasets. The CNNs are ResNet50 for ImageNet and ADE20K, ResNet44 for CIFAR-10, and VGG16 for GTA5 \rightarrow Cityscapes. GTA5 \rightarrow Cityscapes indicates adapting from GTA5 to Cityscapes using domain adaptation.

which not only standardizes but also *decorrelates* features. For example, Decorrelated Batch Normalization (DBN) [13], or namely Batch Whitening (BW), whitens a mini-batch using its covariance matrix, which gives rise to better optimization efficiency than BN in image classification. Moreover, whitening features of an individual image is used in image style transfer [15] to filter out information of image appearance. Here we refer to this operation as instance whitening (IW). Despite their successes, existing works applied these whitening techniques separately to different tasks, preventing them from benefiting each other. Besides, whitening and standardization methods are typically employed in different layers of a CNN, which complicates model design.

To address the above issues, we propose *Switchable Whitening (SW)*. SW provides a general form that integrates different whitening techniques (*e.g.* BW, IW), as well as standardization techniques (*e.g.* BN, IN and LN). SW controls the ratio of each technique by learning their importance weights. It is able to select appropriate normalizers with respect to various vision tasks, as shown in Fig.1(b). For example, semantic segmentation prefers BW and BN, while IW is mainly chosen to address image diversity in

image classification. Compared to semantic segmentation, domain adaptation selects more IW and IN, which alleviates domain discrepancy in CNN features. In image style transfer, IW dominates to handle image style variance.

SW can be inserted into advanced CNN architectures and effectively boosts their performances. Owing to the rich statistics and selectivity of SW, models trained with SW consistently outperform other counterparts in a number of popular benchmarks, such as CIFAR-10/100 [12] and ImageNet [4] for image classification, ADE20K [35] and Cityscapes [3] for semantic segmentation, domain adaptation between GTA5 [23] and Cityscapes, and image style transfer on COCO [18]. For example, when using ResNet50 [7] for ImageNet, ADE20K, and Cityscapes, as well as using VGG16 [25] for domain adaptation, SW significantly outperforms the BN-based baselines by 1.51%, 3.2%, 4.1%, and 3.0% respectively.

SW serves as a useful tool for analyzing the characteristics of these whitening or standardization techniques. This work answers two questions: (1) *Is IW beneficial for highlevel vision tasks like classification and domain adaptation?* (2) *Is standardization still necessary when whitening is presented?* Our experiments suggest that (1) IW works extremely well for handling image appearance diversity and reducing domain gap, giving rise to better performance in high-level vision tasks; (2) Using BW+IW in SW performs comparably well compared to using all the normalizers mentioned above in SW, indicating that full whitening generally works well, and the requirement for standardization is marginal when whitening is presented.

Overall, our **contributions** are summarized as follows. (1) We propose Switchable Whitening (SW), which unifies existing whitening and standardization methods in a general form and learns to switch among them during training. (2) SW adapts to various tasks and is used as a new building block in advanced CNNs. We show that SW outperforms its counterparts in multiple challenging benchmarks. (3) SW could be used as a tool to analyze the effects and characteristics of different normalization methods, and the interactions between whitening and standardization. We will make the code of SW available and hope it would deepen our understanding on various normalization methods.

2. Related Work

Normalization. Existing normalization techniques generally performs *standardization*. For example, Batch Normalization (BN) [10] *centers* and *scales* activations using the mean and variance estimated over a mini-batch, accelerating training and enhancing generalization. In contrast, Instance Normalization (IN) [27] and Layer Normalization (LN) [1] standardize activations with statistics computed over each individual channel and all channels of a layer respectively. IN is mainly used in image generation [8, 27]

while LN has been proved beneficial for training recurrent neural networks [1]. The above three normalizers are combined in Switchable Normalization (SN) [20] that learns the ratio of each one. The combination of BN and IN is also explored in IBN-Net [22] and Batch-Instance Normalization [21]. Besides, there have been other attempts to improve BN for small batch sizes such as Group Normalization [29], Batch Renormalization [9], and Batch Kalman Normalization [28]. All these normalization methods perform centering and scaling to the activations, whereas the correlation between activations remains, leading to suboptimal optimization efficiency. Our work provides a general form that integrates both whitening and standardization techniques, having SN as a special case.

Whitening. Another paradigm towards improving optimization is whitening. Desjardins et al. [5] proposes Natural Neural Network, which implicitly whitens the activations to improve conditioning of the Fisher Information Matrix. This improves optimization efficiency of deep neural networks. Decorrelated Batch Normalization (DBN) [13] whitens features using covariance matrix computed over a mini-batch. It extends BN by decorrelating features. In this paper, we refer to DBN as Batch Whitening (BW) for consistency. Moreover, in the field of image style transfer. whitening and coloring operations are used to manipulate the image appearance [15, 24]. This is because the appearance of an individual image is well encoded in the covariance matrix of its features. We call whitening of an individual image as instance whitening (IW). In this work, we make the first attempt to apply IW in high-level vision tasks like image classification and semantic segmentation.

3. Switchable Whitening (SW)

We first present a general form of whitening as well as standardization operations, and then introduce SW.

3.1. A General Form

Our discussion is mainly based on CNNs, where the data have four dimensions. Let $\mathbf{X} \in \mathbb{R}^{C \times NHW}$ be the data matrix of a mini-batch, where N, C, H, W indicate the number of samples, number of channels, height, and width respectively. Here N, H and W are viewed as a single dimension for convenience. Let matrix $\mathbf{X}_n \in \mathbb{R}^{C \times HW}$ be the *n*th sample in the mini-batch, where $n \in \{1, 2, ..., N\}$. Then the whitening transformation $\phi : \mathbb{R}^{C \times HW} \to \mathbb{R}^{C \times HW}$ for a sample \mathbf{X}_n could be formulated as

$$\phi(\mathbf{X}_n) = \mathbf{\Sigma}^{-1/2} (\mathbf{X}_n - \boldsymbol{\mu} \cdot \mathbf{1}^T)$$
(1)

where μ and Σ are the mean vector and the covariance matrix calculated from the data, and 1 is a column vector of all ones. Note that different whitening methods could be achieved by calculating μ and Σ using different sets of pixels. We discuss them in detail as below.

Batch Whitening (BW). In BW [13], the statistics are calculated in a mini-batch. Thus

$$\boldsymbol{\mu}_{bw} = \frac{1}{NHW} \mathbf{X} \cdot \mathbf{1}$$
$$\boldsymbol{\Sigma}_{bw} = \frac{1}{NHW} (\mathbf{X} - \boldsymbol{\mu} \cdot \mathbf{1}^T) (\mathbf{X} - \boldsymbol{\mu} \cdot \mathbf{1}^T)^T + \epsilon \mathbf{I} \quad (2)$$

where $\epsilon > 0$ is a small positive number to prevent a singular Σ_{bw} . In this way, the whitening transformation ϕ whitens the data of the entire mini-batch, *i.e.*, $\phi(\mathbf{X})\phi(\mathbf{X})^T = \mathbf{I}$.

Instance Whitening (IW). In contrast, for IW [15], μ and Σ are calculated within each individual sample,

$$\boldsymbol{\mu}_{iw} = \frac{1}{HW} \mathbf{X}_n \cdot \mathbf{1}$$
$$\boldsymbol{\Sigma}_{iw} = \frac{1}{HW} (\mathbf{X}_n - \boldsymbol{\mu} \cdot \mathbf{1}^T) (\mathbf{X}_n - \boldsymbol{\mu} \cdot \mathbf{1}^T)^T + \epsilon \mathbf{I} \quad (3)$$

for *n* in {1, 2, ..., N}. IW whitens each samples separately, *i.e.*, $\phi(\mathbf{X}_n)\phi(\mathbf{X}_n)^T = \mathbf{I}$.

Note that Eq.(1) also naturally incorporates standardization operations as its special cases. In the covariance matrix Σ , the diagonal elements are the variance for each channel, while the off-diagonal elements are the correlation between channels. Therefore, by simply setting the off-diagonal elements to zeros, the left multiplication of $\Sigma^{-1/2}$ equals to dividing the standard variance, so that Eq.(1) becomes standardization.

Batch Normalization (BN). BN[10] centers and scales data using the mean and standard deviation of a mini-batch. Hence its mean is the same as in BW *i.e.*, $\mu_{bn} = \mu_{bw}$. As discussed above, since BN does not decorrelate data, the covariance matrix becomes $\Sigma_{bn} = diag(\Sigma_{bw})$, which is a diagonal matrix that only preserves the diagonal of Σ_{bw} .

Instance Normalization (IN). Similarly, in IN [27] we have $\mu_{in} = \mu_{iw}$ and $\Sigma_{in} = diag(\Sigma_{iw})$.

Layer Normalization (LN). LN [1] uses the mean and variance of all channels in a sample to normalize. Let μ_{ln} and σ_{ln} denote the mean and the variance, then $\mu_{ln} = \mu_{ln} \mathbf{1}$ and $\Sigma_{ln} = \sigma_{ln} \mathbf{I}$. In practice μ_{ln} and σ_{ln} could be calculated efficiently from μ_{in} and Σ_{in} using the results in [20].

In Eq.(1), the inverse square root of the covariance matrix is typically calculated by using ZCA whitening,

$$\boldsymbol{\Sigma}^{-1/2} = \mathbf{D}\boldsymbol{\Lambda}^{-1/2}\mathbf{D}^T \tag{4}$$

where $\mathbf{\Lambda} = diag(\sigma_1, ..., \sigma_c)$ and $\mathbf{D} = [\mathbf{d}_1, ..., \mathbf{d}_c]$ are the eigenvalues and the eigenvectors of $\boldsymbol{\Sigma}$, *i.e.*, $\boldsymbol{\Sigma} = \mathbf{D}\mathbf{\Lambda}\mathbf{D}^T$, which is obtained via eigen decomposition.

So far we have formulated different whitening and normalization transforms in a general form. In the next section, we introduce switchable whitening based on this formulation.

3.2. Formulation of SW

For a data sample X_n , a natural way to unify the aforementioned whitening and standardization transforms is to combine the mean and covariance statistics of those methods, and perform whitening using this unified statistics, giving rise to

$$SW(\mathbf{X}_n) = \hat{\boldsymbol{\Sigma}}^{-1/2} (\mathbf{X}_n - \hat{\boldsymbol{\mu}} \cdot \mathbf{1}^T)$$
(5)

where
$$\hat{\boldsymbol{\mu}} = \sum_{k \in \Omega} \omega_k \boldsymbol{\mu}_k, \quad \hat{\boldsymbol{\Sigma}} = \sum_{k \in \Omega} \omega'_k \boldsymbol{\Sigma}_k$$
 (6)

Here Ω is a set of statistics estimated in different ways. In this work, we mainly focus on two cases, *i.e.*, $\Omega = \{bw, iw\}$ and $\Omega = \{bw, iw, bn, in, ln\}$, where the former switches between two whitening methods, while the later incorporates both whitening and standardization methods. ω_k are importance ratios to switch among different statistics. In practice, ω_k are generated by the corresponding control parameters λ_k via softmax function, *i.e.*, $\omega_k = \frac{e^{\lambda_k}}{\sum_{z \in \Omega} e^{\lambda_z}}$. And ω'_k are defined similarly using another group of control parameters λ'_k . This relieves the constraint of consistency between mean and covariance, which is a more general form.

Note that the above formulation incorporates SN [20] as its special case by letting $\Omega = \{bn, in, ln\}$. Our formulation is more flexible and general in that it takes into account the whole covariance matrix rather than only the diagonal. This provides the possibility of producing decorrelated features, giving rise to either better optimization conditioning or style invariance. SW could be easily extended to incorporate some other normalization methods like Batch Renormalization [9] or Group Normalization [29], which is out of the scope of this work.

3.3. Training and Inference

Switchable Whitening could be inserted extensively into a convolutional neural network (CNN). Let Θ be a set of parameters of a CNN, and Φ be a set of importance weights in SW. The importance weights are initialized uniformly, *e.g.* $\lambda_k = 1$. During training, Θ and Φ are optimized jointly by minimizing a loss function $\mathcal{L}(\Theta, \Phi)$ using backpropagation. The forward and backward calculation of our proposed SW is presented in Algorithm 1 and 2. For clearance, we use $\Omega = \{bw, iw\}$ as an illustrative example.

In the training phase, μ_{bw} and Σ_{bw} are calculated within each mini-batch and used to update the running mean and running covariance as in Line 7 and 8 of Algorithm 1. During inference, the running mean and the running covariance are used as μ_{bw} and Σ_{bw} , while μ_{iw} and Σ_{iw} are calculated independently for each sample.

In practice, the scale and shift operations are usually added right after the normalization or whitening transform

Algorithm 1 Forward pass of SW for each iteration.

- 1: **Input:** mini-batch inputs $\mathbf{X} \in \mathbb{R}^{C \times NHW}$, where the *n*th sample in the batch is $\mathbf{X}_n \in \mathbb{R}^{C \times HW}$, $n \in \{1, 2, ..., N\}$; importance weights λ_k and λ'_k , $k \in \{bw, iw\}$; expected mean μ_E and expected covariance Σ_E .
- 2: Hyperparameters: ϵ , running average momentum α .
- 3: **Output:** the whitened activations $\{\hat{\mathbf{X}}_n, n = 1, 2, ..., N\}$.
- 4: calculate: $\omega_{bw}, \omega_{iw} = Softmax(\lambda_{bw}, \lambda_{iw}), \omega'_{bw}, \omega'_{iw} = Softmax(\lambda'_{bw}, \lambda'_{iw}).$
- 5: calculate: $\boldsymbol{\mu}_{bw} = \frac{1}{NHW} \mathbf{X} \cdot \mathbf{1}$. 6: calculate: $\Sigma_{bw} = \frac{1}{1} \frac{NHW}{NHW} (\mathbf{X} - \boldsymbol{\mu} \cdot \mathbf{1}^T) (\mathbf{X} - \boldsymbol{\mu} \cdot \mathbf{1}^T)^T + \epsilon \mathbf{I}.$ 7: update: $\boldsymbol{\mu}_E \leftarrow (1 - \alpha) \boldsymbol{\mu}_E + \alpha \boldsymbol{\mu}_{bw}.$ 8: update: $\Sigma_E \leftarrow (1 - \alpha) \boldsymbol{\Sigma}_E + \alpha \boldsymbol{\Sigma}_{bw}.$ 9: for m = 1 to N d = 0

- 9: for n = 1 to N do
- 10:
- $n = 1 \text{ to } i_{V} \text{ to }$ calculate: $\mu_{iw}^{(n)} = \frac{1}{HW} \mathbf{X}_{n} \cdot \mathbf{1}.$ calculate: $\Sigma_{iw}^{(n)} = \frac{1}{HW} (\mathbf{X}_{n} \boldsymbol{\mu} \cdot \mathbf{1}^{T}) (\mathbf{X}_{n} \boldsymbol{\mu} \cdot \mathbf{1}^{T})^{T} + \epsilon \mathbf{I}.$ calculate: $\hat{\boldsymbol{\mu}}_{n} = \sum_{k} \omega_{k} \boldsymbol{\mu}_{k}^{(n)}, \hat{\boldsymbol{\Sigma}}_{n} = \sum_{k} \omega_{k}' \boldsymbol{\Sigma}_{k}^{(n)}, k \in \{bw, iw\}.$ 11:
- 12:
- execute eigenvalue decomposition: $\hat{\Sigma}_n = \mathbf{D} \mathbf{\Lambda} \mathbf{D}^T$ 13:
- calculate ZCA-whitening matrix: $\mathbf{U}_n = \mathbf{D} \mathbf{\Lambda}^{-1/2} \mathbf{D}^T$. 14:
- calculate ZCA-whitened output: $\hat{\mathbf{X}}_n = \mathbf{U}_n (\mathbf{X}_n \hat{\boldsymbol{\mu}}_n \cdot \mathbf{1}^T).$ 15:
- 16: end for

to enhance the model's representation capacity. For SW, we follow this design to introduce scale and shift parameters γ and β as in BN.

3.4. Analysis and Discussion

We have introduced the formulation and training of SW. Here we discuss some of its important properties and analyze its complexity.

Instance Whitening for Appearance Invariance. In style transfer, researchers have found that image appearance information (i.e. color, contrast, style etc.) is well encoded in the covariance matrix of features produced by CNNs [15]. In this work, we take the first attempt to induce appearance invariance by leveraging IW, which is beneficial for domain adaptation or high-level vision tasks like classification or semantic segmentation. Although IN also introduces invariance by standardizing each sample separately, the difference in correlation could be easily enlarged in highly non-linear deep neural networks. In IW, features of different samples are not only standardized but also whitened individually, giving rise to the same covariance matrix, *i.e.*, identity matrix. Therefore, IW has better invariance property than IN.

Switching between Whitening and Standardization. Our formulation of SW makes it possible to switch between whitening and standardization. For example, considering $\Omega = \{bw, bn\}, i.e., \Sigma = \omega_{bw} \Sigma_{bw} + \omega_{bn} \Sigma_{bn}, (\omega_{bw} + \omega_{bn} = \omega_{bw} \Sigma_{bw})$ 1). As ω_{bn} grows larger, the diagonal of $\hat{\Sigma}$ would remain the same, while the off-diagonal would be weaken. This would make the features less decorrelated after whitening. This is beneficial when the extent of whitening requires careful adjustment, which is an important issue of BW as pointed out in [13].

Group SW. Huang et al. [13] uses group whitening to re-

Algorithm 2 Backward pass of SW for each iteration.

- 1: Input: mini-batch gradients respect to whitened outputs $\{\frac{\partial L}{\partial \hat{\mathbf{x}}_n}, n = 1, 2, ..., N\}$. Other auxiliary data from respective forward pass.
- the gradients respect to the inputs 2: Output: $\{\frac{\partial L}{\partial \mathbf{X}_n}, n = 1, 2, ..., N\}$; the gradients respect to the importance weights $\frac{\partial L}{\partial \lambda_k}$ and $\frac{\partial L}{\partial \lambda'_k}$, $k \in \{bw, iw\}$.
- 3: for n = 1 to N do 4: calculate $\frac{\partial L}{\partial \hat{\Sigma}_n}$ using results in Appendix A.
- calculate $\frac{\partial L}{\partial \hat{\boldsymbol{\mu}}_n}$ using results in Appendix A. 5:
- 6: **end for**
- 7: for n = 1 to N do $\begin{aligned} & \text{calculate} \quad \frac{\partial L}{\partial \mathbf{X}_n} = \frac{\partial L}{\partial \mathbf{\hat{X}}_n} \mathbf{U}_n + \left(\frac{\omega_{bw}}{NHW} \sum_{i=1}^N \frac{\partial L}{\partial \mathbf{\hat{\mu}}_i} + \frac{\omega_{iw}}{HW} \frac{\partial L}{\partial \mathbf{\hat{\mu}}_n}\right) \\ &+ \left[\frac{2\omega'_{bw}(\mathbf{X}_n - \boldsymbol{\mu}_{bw})^T}{NHW} \sum_{i=1}^N \left(\frac{\partial L}{\partial \mathbf{\hat{\Sigma}}_i}\right) sym + \frac{2\omega'_{iw}(\mathbf{X}_n - \boldsymbol{\mu}_{iw})^T}{HW} \left(\frac{\partial L}{\partial \mathbf{\hat{\Sigma}}_n}\right) sym\right] \end{aligned}$ 8.
- 9: end for
- 10: calculate:

<i>.</i>	calculate.
	$\frac{\partial L}{\partial \lambda_{bw}} = \omega_{bw} (1 - \omega_{bw}) \sum_{n=1}^{N} \left(\frac{\partial L}{\partial \hat{\mu}_n} \mu_{bw} \right) - \omega_{iw} \omega_{bw} \sum_{n=1}^{N} \left(\frac{\partial L}{\partial \hat{\mu}_n} \mu_{iw}^{(n)} \right)$
	$\frac{\partial L}{\partial \lambda_{iw}} = \omega_{iw} (1 - \omega_{iw}) \sum_{n=1}^{N} (\frac{\partial L}{\partial \hat{\mu}_n} \boldsymbol{\mu}_{iw}^{(n)}) - \omega_{bw} \omega_{iw} \sum_{n=1}^{N} (\frac{\partial L}{\partial \hat{\mu}_n} \boldsymbol{\mu}_{bw})$
	$\frac{\partial L}{\partial \lambda'_{bw}} = \omega'_{bw} (1 - \omega'_{bw}) \sum_{n=1}^{N} \langle \frac{\partial L}{\partial \hat{\mathbf{\Sigma}}_n}, \mathbf{\Sigma}_{bw} \rangle_F - \omega'_{iw} \omega'_{bw} \sum_{n=1}^{N} \langle \frac{\partial L}{\partial \hat{\mathbf{\Sigma}}_n}, \mathbf{\Sigma}_{iw}^{(n)} \rangle_F$
	$\frac{\partial L}{\partial \lambda'_{iw}} = \omega'_{iw} (1 - \omega'_{iw}) \sum_{n=1}^{N} \left(\frac{\partial L}{\partial \hat{\Sigma}_n}, \hat{\Sigma}_{iw}^{(n)} \right)_F - \omega'_{bw} \omega'_{iw} \sum_{n=1}^{N} \left(\frac{\partial L}{\partial \hat{\Sigma}_n}, \hat{\Sigma}_{bw} \right)_F$

duce complexity and to address the inaccurate estimation of large covariance matrices. In SW we follow the same design, *i.e.*, the features are divided into groups along the channel dimension and SW is performed for each group. The importance weights λ_k could be shared or independent for each group. In this work we let groups of a layer share the same λ_k to simplify discussion.

Table 1. Comparisons of computational complexity. N, C, H, Ware the number of samples, number of channels, height, and width of the input tensor respectively. G denotes the number of channels for each group in group whitening.

Method	Computational complexity							
Methou	w/o group	w/ group						
BN,IN,LN,SN	O(NCHW)	O(NCHW)						
BW	$O(C^2 \max(NHW, C))$	$O(CG\max(NHW,G))$						
IW	$O(NC^2 \max(HW, C))$	$O(NCG\max(HW,G))$						
SW	$O(NC^2 \max(HW, C))$	$O(NCG\max(HW,G))$						

Complexity Analysis. The computational complexities for different normalization methods are compared in Table 1. The flop of SW is comparable with IW. And applying group whitening could reduce the computation by C/G times. Usually we have HW > G, thus the computation cost of SW and BW would be roughly the same (i.e., O(CGNHW)).

In practice, the running time bottleneck of SW lies in the eigen decomposition. To make training efficient, we implement a distributed version for the 'for' loop in Line 9 of Algorithm 1 and Line 3, 7 of Algorithm 2, making SW have roughly the same time complexity as BW. More analysis on time complexity is provided in supplementary material.

 $^{{}^{1}\}langle\rangle_{F}$ denotes Frobenius inner product.



Figure 2. Training and validation error curve on CIFAR-10 and ImageNet. Models with different normalization methods are reported. Here SW has $\Omega = \{bw, iw\}$.

Table 2. Test errors (%) on CIFAR-10/100 and ImageNet validation sets [12]. For each model, we evaluate different normalization or whitening methods. SW^{*a*} and SW^{*b*} correspond to $\Omega = \{$ bw, iw $\}$ and $\Omega = \{$ bw, iw, bn, in, ln $\}$ respectively. Results on CIFAR are averaged over 5 runs.

Method	BN	SN	BW	SW^a	SW^b
ResNet20	8.45	8.34	8.28	7.64	7.75
ResNet44	7.01	6.75	6.83	6.27	6.35
ResNet56	6.88	6.57	6.62	6.07	6.25
ResNet110	6.21	5.97	5.99	5.69	5.78
ResNet20	32.09	32.28	32.44	31.00	30.87
ResNet110	27.32	27.25	27.76	26.64	26.48
ResNet50 (top1)	23.58	23.10	23.31	22.10	22.07
ResNet50 (top5)	7.00	6.55	6.72	5.96	5.91
	Method ResNet20 ResNet44 ResNet56 ResNet110 ResNet20 ResNet110 ResNet50 (top1) ResNet50 (top5)	Method BN ResNet20 8.45 ResNet44 7.01 ResNet56 6.88 ResNet110 6.21 ResNet20 32.09 ResNet110 27.32 ResNet50 (top1) 23.58 ResNet50 (top5) 7.00	Method BN SN ResNet20 8.45 8.34 ResNet44 7.01 6.75 ResNet56 6.88 6.57 ResNet110 6.21 5.97 ResNet20 32.09 32.28 ResNet110 27.32 27.25 ResNet50 (top1) 23.58 23.10 ResNet50 (top5) 7.00 6.55	MethodBNSNBWResNet208.458.348.28ResNet447.016.756.83ResNet566.886.576.62ResNet1106.215.975.99ResNet2032.0932.2832.44ResNet11027.3227.2527.76ResNet50 (top1)23.5823.1023.31ResNet50 (top5)7.006.556.72	MethodBNSNBWSW ^a ResNet208.458.348.28 7.64 ResNet447.016.756.83 6.27 ResNet566.886.576.62 6.07 ResNet1106.215.975.99 5.69 ResNet2032.0932.2832.4431.00ResNet11027.3227.2527.7626.64ResNet50 (top1)23.5823.1023.3122.10ResNet50 (top5)7.006.556.725.96

4. Experiments

We evaluate SW on image classification (CIFAR-10/100, ImageNet), semantic segmentation (ADE20K, Cityscapes), domain adaptation (GTA5, Cityscapes), and image style transfer (COCO). For each task, SW is compared with previous normalization methods.

4.1. Classification

CIFAR-10, CIFAR-100 [12] and ImageNet [4] are standard image classification benchmarks. Our training policies and settings are the same as in [7].

Implementation. We evaluate different normalization methods based on standard ResNet [7]. Note that introducing whitening after all convolution layers of ResNet is redundant and would incur a high computational cost, as also pointed out in [13]. Hence we replace part of the BN layers in ResNet to the desired normalization layers. For CIFAR, we apply SW or other counterparts after the 1*st* and the $\{4n\}th$ (n = 1,2,3,...) convolution layers. And for ImageNet, the normalization layers considered are those at the 1*st* and the $\{6n\}th$ (n = 1,2,3,...) layers. The residual blocks

with 2048 channels are not considered to save computation. More discussions for such choices could be found in supplementary material.

The normalization layers studied here are BN, SN, BW, and SW. For SW, we consider two cases: $\Omega = \{bw, iw\}$ and $\Omega = \{$ bw, iw, bn, in, ln $\}$, which are denoted as SW^a and SW^b respectively. In all experiments, we adopt group whitening with group size G = 16 for SW and BW. Since [19] shows that applying early stop to the training of SN reduces overfitting, we stop the training of SN and SW at the 80th epoch for CIFAR and the 30th epoch for ImageNet. **Results**. The results are given in Table. 2 and the training curves are shown in Fig. 2. In both datasets, SW^a and SW^b show better results and faster convergence than BN, SN, and BW over various network depth. Specifically, with only 7 SW^b layers, the top1 and top5 error of ResNet50 on ImageNet is significantly reduced by 1.51% and 1.09%. This performance is comparable with the original ResNet152 which has 5.94% top5 error.

Our results reveal that combining different normalization methods in a suitable manner surpasses every single normalizer. For example, the superiority of SW^b over SN attributes to the better optimization conditioning brought out by whitening. And the better performance of SW^a over BW shows that instance whitening is beneficial as it introduces style invariance. Moreover, SW^a and SW^b perform comparably well, which indicates that full whitening generally performs well, and the need for standardization is marginal while whitening is presented.

Discussions. SW has two groups of importance weights λ_k and λ'_k . We observe that allowing λ_k and λ'_k to share weight produces slightly worse results. For example, ResNet20 has 8.17% test error when using SW with shared importance weights. We conjecture that mean and covariance have different impacts in training, and recommend to maintain independent importance weights for mean and covariance.

Note that IW is not reported here because it generally produces worse results due to diminished feature discrimi-

Mathod	ADI	E20K	Cityscapes				
Method	mIoU _{ss}	mIoU _{ms}	mIoU _{ss}	mIoU _{ms}			
ResNet50-BN	36.6	37.9	72.1	73.4			
ResNet50-SN	37.8	38.8	75.0	76.2			
ResNet50-BW	35.9	37.8	72.5	73.7			
ResNet50-SW ^a	39.8	40.8	76.2	77.1			
ResNet50-SW ^b	39.8	40.7	76.0	77.0			

Table 3. Results on Cityscapes and ADE20K datasets. 'ss' and 'ms' indicate single-scale and multi-scale test respectively.

Table 4. Comparison with advanced methods on the ADE20K validation set. * indicates our implementation.

Method	mIoU(%)	Pixel Acc.(%)
DilatedNet [31]	32.31	73.55
CascadeNet [36]	34.90	74.52
RefineNet [17]	40.70	-
PSPNet101 [33]	43.29	81.39
SDDPN [16]	43.68	81.13
WiderNet [30]	43.73	81.17
PSANet101 [34]	43.77	81.51
EncNet [32]	44.65	81.69
PSPNet101*	43.59	81.41
PSPNet101-SW ^a	45.33	82.05

nation. For example, ResNet20 with IW gives 12.57% test error on CIFAR-10, which is worse than other normalization methods. This also implies that SW borrows the benefits of different normalizers so that it could outperform any individual of them.

4.2. Semantic Segmentation

We further verify the scalability of our method on ADE20K [35] and Cityscapes [3], which are standard and challenging semantic segmentation benchmarks. We evaluate SW based on ResNet and PSPNet [33].

Implementation. We adopt the same ResNet architecture, training setting, and data augmentation scheme as in [33]. The normalization layers considered are the 1*st* and the $\{3n\}$ *th* (n = 1,2,3,...) layers except those with 2048 channels. Since overfitting is not observed in these two benchmarks, early stop is not used here. The BN and BW involved are synchronized across multiple GPUs.

Results. Table.3 reports mIoU on the validation sets of the two benchmarks. For ResNet50, simply replacing part of BN with SW would significantly boost mIoU_{ss} by 3.2% and 4.1% for ADE20K and Cityscapes respectively. SW also notably outperforms SN and BW, which is consistent with the results of classification.

Furthermore, we show that SW could improve even the most advanced models for semantic segmentation. We apply SW to PSPNet101 [33], and compare with other methods on the ADE20K dataset. The results are shown in Table.4. Simply using some SW layers could improve the



Figure 3. MMD distance between Cityscapes and GTA5.

strong PSPNet by 1.74% on mIoU. And our final score, 45.33%, outperforms other more advanced semantic segmentation methods like PSANet [34] and EncNet [32].

4.3. Domain Adaptation

The adaptive style invariance of SW making it suitable for handling appearance discrepancy between two image domains. To verify this, we evaluate SW on domain adaptation task. The datasets employed are the widely used GTA5 [23] and Cityscapes [3] datasets. GTA5 is a street view dataset generated semi-automatically from the computer game Grand Theft Auto V (GTA5), while Cityscapes contains traffic scene images collected from the real world. **Implementation.** We conduct our experiments based on the AdaptSegNet [26] framework, which is a recent stateof-the-art domain adaptation approach. It adopts adversarial learning to shorten the discrepancy between two domains with a discriminator. The segmentation network is DeepLab-v2 [2] model with VGG16 [25] backbone. The training setting is the same as in [26].

Note that the VGG16 model has five convolutional groups, where the number of convolution layers for these groups are $\{2,2,3,3,3\}$. We add SW or its counterparts after the first convolution layer of each group, and report the results using different normalization layers.

Results. Table.5 reports the results of adapting GTA5 to Cityscapes. The models with SW achieve higher performance when evaluated on a different image domain. Particularly, compared with BN, and SN, SW^{*a*} improves the mIoU by 3.0%, and 1.6% respectively.

To understand how SW performs better under crossdomain evaluation, we analyze the maximum mean discrepancy (MMD)[6] of deep features between the two datasets. MMD is a commonly used metric for evaluating domain discrepancy. Specifically, we use the MMD with Gaussian kernels as in [14]. We calculate the MMD for features of the first 13 layers in VGG16 with different normalization layers. The results are shown in Fig.3. Compared with BN and SN, SW significantly reduces MMD for both shallow and deep features. This shows that the IW introduced effectively reduces domain discrepancy in the CNN features, making the model easier to generalize to other data domains.

Table 5. Results of adapting GTA5 to Cityscapes. mIoU of models with different normalization layers are reported.

Method	road	sidewalk	building	wall	fence	pole	light	sign	veg	terrain	sky	person	rider	car	truck	pus	train	mbike	bike	mIoU
AdaptSetNet-BN	88.3	42.7	74.9	22.0	14.0	16.5	17.8	4.2	83.5	34.3	72.1	44.8	1.7	76.9	18.0	6.7	0.0	3.0	0.1	32.7
AdaptSetNet-SN	87.0	41.6	77.5	21.2	20.0	18.3	20.9	8.3	82.4	35.4	72.6	48.4	1.4	81.1	18.7	5.2	0.0	8.4	0.0	34.1
AdaptSetNet-SW ^a	91.8	50.2	78.1	25.3	17.5	17.5	21.4	6.2	83.4	36.6	74.0	50.7	7.4	83.4	16.7	6.3	0.0	10.4	0.8	35.7
AdaptSetNet-SW ^b	91.8	50.5	78.4	23.5	16.5	17.2	19.8	5.5	83.6	38.4	74.6	48.9	5.3	83.6	17.6	3.9	0.1	7.7	0.7	35.1



Figure 4. Training loss in style transfer and the learned importance ratios of SW^b . The importance ratios are averaged over all SW^b layers in the image stylizing network.



Figure 5. Visualization of style transfer using different normalization layers.

4.4. Image Style Transfer

Thanks to the rich statistics, SW could work not only in high-level vision tasks, but also in low-level vision tasks like image style transfer. To show this, we employ a popular style transfer algorithm [11]. It has an image stylizing network trained with the content loss and style loss calculated by a loss network. The MS-COCO dataset [18] is used as content images while the style images selected are *candy* and *starry night*. We follow the same training policy as in [11], and adopt different normalization layers for the image stylizing network.

Results. The training loss curve is shown in Fig.4. As revealed in former works, IW and IN perform better than BN. Besides, we observe that IW has smaller content loss and style loss than IN, which verifies that IW works better in manipulating image style. Although SW converges slower than IW at the beginning, it soon catches up with IW as SW learns to select IW as the normalizer. Moreover, SW has smaller content loss than IW when the training converges,

as BW preserves important content information.

Qualitative examples of style transfer using different normalization layers are shown in Fig.5. BN produces poor stylization images, while IW gives satisfactory results. SW works comparably well with IW, showing that SW is able to select appropriate normalizer according to the task. More examples are provided in supplementary material.

4.5. Analysis on SW

In order to understand the behavior of SW, in this section we study its learning dynamics and the learned importance ratios.

Learning Dynamics. The importance ratios of SW is initialized to have uniform values, i.e. 0.5 for $\Omega = \{bw, iw\}$ and 0.2 for $\Omega = \{$ bw, iw, bn, in, ln $\}$. To see how the ratios of SW in different layers change during training, we plot the learning curves of ω_k and ω'_k in Fig.6 and Fig.7. It can be seen that the importance ratios shift quickly at the beginning and gradually become stable. There are several interesting observations. (1) The learning dynamics vary across different tasks. In CIFAR-10, SW mostly selects IW and occasionally selects BW, while in Cityscapes BW or BN is mostly chosen. (2) The learning behaviours of SW across different layers tend to be distinct rather than homogeneous. For example, in Fig.7 (a), SW selects IW for layer {15, 21, 39, and BW for the rest except for layer $\{6, 9\}$ where the ratios keep uniform. (3) The behaviors of ω_k and ω'_k are mostly coherent and sometimes divergent. For instance, in layer {15, 21} of Fig.7, ω_k chooses IW while ω'_k chooses BW or BN. This implies that μ and Σ are not necessarily have to be consistent, as they might have different impacts in training.



Figure 6. Learning curve of importance weights in ResNet56 on CIFAR-10. (a) and (b) show ω_k and ω'_k in SW with $\Omega = \{$ bw, iw $\}$. (c) and (d) correspond to ω_k and ω'_k in SW with $\Omega = \{$ bw, iw, bn, in, ln $\}$.



Figure 7. Learning curve of importance weights in ResNet50 on Cityscapes. (a)(b)(c)(d) have the same meanings as in Fig.6.



Figure 8. Learned importance ratios of SW in various tasks. Above and below correspond to $\Omega = \{bw, iw\}$ and $\Omega = \{bw, iw, bn, in, ln\}$ respectively.

Importance Ratios for Various Tasks.

We further analyze the learned importance ratios in SW for various tasks, as shown in Fig.8. The results are obtained by taking average over the importance ratios ω'_k of all SW layers in a CNN. The models are ResNet50 for ImageNet, ADE20K, and Cityscapes, ResNet44 for CIFAR-10, VGG16 for GTA5 \rightarrow Cityscapes, and a ResNet-alike network for style transfer as in [11]. Both $\Omega = \{bw, iw\}$ and $\Omega = \{bw, iw, bn, in, ln\}$ are reported.

We make the following remarks: (1) For semantic segmentation, SW chooses mainly BW and BN, and partially the rest, while in classification more IW are selected. This is because the diversity between images is higher in classification datasets than in segmentation datasets. Thus more IW is required to alleviate the intra-dataset variance. (2) Semantic segmentation on Cityscapes tends to produce more IW and IN under domain adaptation setting than in the normal setting. Since domain adaptation introduces a domain discrepancy loss, more IW and IN would be beneficial for reducing the feature discrepancy between the two domains, *i.e.*, GTA5 and Cityscapes. (3) In image style transfer, SW switches to IW aggressively. This phenomenon is consistent with the common knowledge that IW is well suited for style transfer, as image level appearance information is well encoded in the covariance of CNN features. Our experiments also verify that IW is a better choice than IN in this task.

5. Conclusion

In this paper, we propose Switchable Whitening, which integrates various whitening and standardization techniques in a general form. SW adapts to various tasks by learning to select appropriate normalizers in different layers of a CNN. Our experiments show that SW achieves consistent improvements over previous normalization methods in a number of computer vision tasks, including classification, segmentation, domain adaptation, and image style transfer. Investigation of SW reveals the importance of leveraging different whitening methods in CNNs. We hope that our findings in this work would benefit other research fields and tasks that employ deep learning.

References

- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 1, 2, 3
- [2] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *TPAMI*, 2018. 6
- [3] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. *CVPR*, 2016. 2, 6
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. *CVPR*, 2009. 2, 5
- [5] Guillaume Desjardins, Karen Simonyan, Razvan Pascanu, et al. Natural neural networks. *NIPS*, 2015. 2
- [6] Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 2012. 6
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CVPR*, 2016.
 2, 5
- [8] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. *ICCV*, 2017.
 2
- [9] Sergey Ioffe. Batch renormalization: Towards reducing minibatch dependence in batch-normalized models. *NIPS*, 2017. 2, 3
- [10] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *ICML*, 2015. 1, 2, 3
- [11] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. *ECCV*, 2016. 7, 8
- [12] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical report, Technical report, 2009. 2, 5
- [13] Huang Lei, Yang Dawei, Lang Bo, and Deng Jia. Decorrelated batch normalization. CVPR, 2018. 1, 2, 3, 4, 5
- [14] Chun-Liang Li, Wei-Cheng Chang, Yu Cheng, Yiming Yang, and Barnabás Póczos. Mmd gan: Towards deeper understanding of moment matching network. *NIPS*, 2017. 6
- [15] Yijun Li, Chen Fang, Jimei Yang, Zhaowen Wang, Xin Lu, and Ming-Hsuan Yang. Universal style transfer via feature transforms. *NIPS*, 2017. 1, 2, 3, 4
- [16] Xiaodan Liang, Hongfei Zhou, and Eric Xing. Dynamicstructured semantic propagation network. CVPR, 2018. 6
- [17] Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. Refinenet: Multi-path refinement networks for highresolution semantic segmentation. *CVPR*, 2017. 6
- [18] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. ECCV, 2014. 2, 7

- [19] Ping Luo, Zhanglin Peng, Jiamin Ren, and Ruimao Zhang. Do normalization layers in a deep convnet really need to be distinct? *arXiv preprint arXiv:1811.07727*, 2018. 5
- [20] Ping Luo, Jiamin Ren, and Zhanglin Peng. Differentiable learning-to-normalize via switchable normalization. *ICLR*, 2019. 2, 3
- [21] Hyeonseob Nam and Hyo-Eun Kim. Batch-instance normalization for adaptively style-invariant neural networks. *NIPS*, 2018. 2
- [22] Xingang Pan, Ping Luo, Jianping Shi, and Xiaoou Tang. Two at once: Enhancing learning and generalization capacities via ibn-net. *ECCV*, 2018. 2
- [23] Stephan R. Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. ECCV, 2016. 2, 6
- [24] Aliaksandr Siarohin, Enver Sangineto, and Nicu Sebe. Whitening and coloring transform for gans. arXiv preprint arXiv:1806.00420, 2018. 2
- [25] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556, 2014. 2, 6
- [26] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schulter, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. *CVPR*, 2018. 6
- [27] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis. *CVPR*, 2017. 1, 2, 3
- [28] Guangrun Wang, Jiefeng Peng, Ping Luo, Xinjiang Wang, and Liang Lin. Batch kalman normalization: Towards training deep neural networks with micro-batches. *NIPS*, 2018.
- [29] Yuxin Wu and Kaiming He. Group normalization. ECCV, 2018. 2, 3
- [30] Zifeng Wu, Chunhua Shen, and Anton Van Den Hengel. Wider or deeper: Revisiting the resnet model for visual recognition. *Pattern Recognition*, 2019. 6
- [31] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *ICLR*, 2016. 6
- [32] Hang Zhang, Kristin Dana, Jianping Shi, Zhongyue Zhang, Xiaogang Wang, Ambrish Tyagi, and Amit Agrawal. Context encoding for semantic segmentation. *CVPR*, 2018. 6
- [33] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. *CVPR*, 2017. 6
- [34] Hengshuang Zhao, Yi Zhang, Shu Liu, Jianping Shi, Chen Change Loy, Dahua Lin, and Jiaya Jia. Psanet: Point-wise spatial attention network for scene parsing. *ECCV*, 2018. 6
- [35] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. CVPR, 2017. 2, 6
- [36] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ade20k dataset. *IJCV*, 2016.
 6