

Learning to Optimize Domain Specific Normalization for Domain Generalization

Seonguk Seo¹, Yumin Suh², Dongwan Kim¹,
Geeho Kim¹, Jongwoo Han³, and Bohyung Han¹

¹ Seoul National University

² NEC Laboratories America

³ LG Electronics

Abstract. We propose a simple but effective multi-source domain generalization technique based on deep neural networks by incorporating optimized normalization layers that are specific to individual domains. Our approach employs multiple normalization methods while learning separate affine parameters per domain. For each domain, the activations are normalized by a weighted average of multiple normalization statistics. The normalization statistics are kept track of separately for each normalization type if necessary. Specifically, we employ batch and instance normalizations in our implementation to identify the best combination of these two normalization methods in each domain. The optimized normalization layers are effective to enhance the generalizability of the learned model. We demonstrate the state-of-the-art accuracy of our algorithm in the standard domain generalization benchmarks, as well as viability to further tasks such as multi-source domain adaptation and domain generalization in the presence of label noise.

Keywords: Domain generalization

1 Introduction

Domain generalization aims to learn generic feature representations agnostic to domains and make trained models perform well in completely new domains. To achieve this challenging goal, one needs to train models that can capture useful information observed commonly in multiple domains and recognize semantically related but visually inconsistent examples effectively. Many real-world problems have similar objectives so this task can be widely used in various practical applications. Domain generalization is closely related to unsupervised domain adaptation but there is a critical difference regarding the availability of target domain data; contrary to unsupervised domain adaptation, domain generalization cannot access any examples in target domain during training but is still required to capture transferable information across domains. Due to this constraint, the domain generalization problem is typically considered to be more challenging, so multiple source domains are usually involved to make the problem more feasible.

Domain generalization techniques are classified into several groups depending on their approaches. Some algorithms define novel loss functions to learn

domain-agnostic representations [21, 20, 15, 7] while others are more interested in designing deep neural network architectures to achieve similar goals [13, 6, 17]. The algorithms based on meta-learning have been proposed under the assumption that there exists a held-out validation set [2, 14, 12].

Our algorithm belongs to the second category, *i.e.* network architecture design methods. In particular, we are interested in exploiting normalization layers in deep neural networks to handle the domain generalization task. A naïve approach would be to train a single deep neural network with batch normalization using all training examples regardless of their domain memberships. This method works fairly well partly because batch normalization regularizes feature representations from heterogeneous domains and the trained model is often capable of adapting to unseen domains. However, the benefit of batch normalization is limited when domain shift is significant, and we are often required to remove domain-specific styles for better generalization. Instance normalization [27] turns out to be an effective scheme for the goal and incorporating both batch and instance normalization techniques further improves accuracy by a data-driven balancing of two normalization methods [22, 23]. Our approach also employs the two normalizations but proposes a more sophisticated algorithm designed for domain generalization.

We explore domain-specific normalizations to learn representations that are both domain-agnostic and semantically discriminative by discarding domain-specific ones. The goal of our algorithm is to optimize the combination of normalization techniques in each domain while different domains learn separate parameters for the mixture of normalizations. The intuition behind this approach is that we can learn domain-invariant representations by controlling types of normalization and parameters in normalization layers. Note that all other parameters, including the ones in convolutional layers, are shared across domains. Although our approach is somewhat similar to [16] in that the optimal mixing weights between normalization types are learned, we emphasize that the motivations are different; [16] aims for a differentiable normalization for universal tasks while we set our sights on how to remove style information without losing semantics to generalize on unseen domains. In addition, our domain-specific properties—learning normalization parameters, batch statistics and mixture weights for each domain separately—makes it unique and more effective to construct domain-agnostic representations, thereby outperforming all the established normalization techniques. We illustrate the main idea of our approach in Fig. 1.

Our contributions are as follows:

- Our approach leverages instance normalization to optimize the trade-off between cross-category variance and domain invariance, which is desirable for domain generalization in unseen domains.
- We propose a simple but effective domain generalization technique combining heterogeneous normalization methods specific to individual domains, which facilitates the extraction of domain-agnostic feature representations by removing domain-specific information effectively.

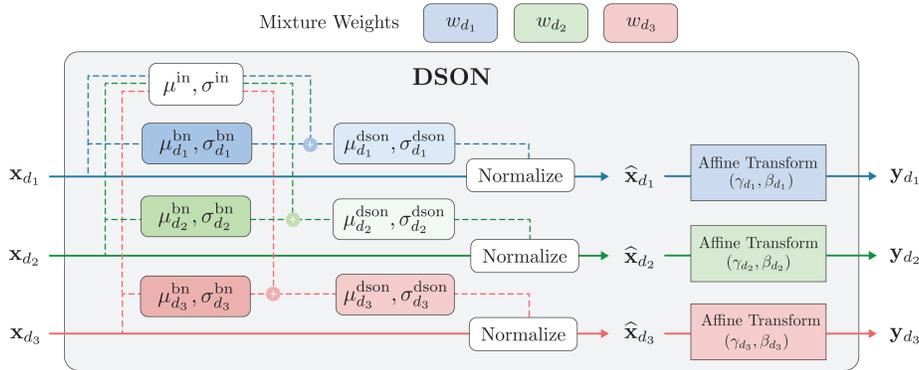


Fig. 1. Illustration of Domain Specific Optimized Normalization (DSON). Each domain maintains domain-specific batch normalization statistics and affine parameters, as well as mixture weights.

- The proposed algorithm achieves the state-of-the-art accuracy in multiple standard benchmark datasets and outperforms all established normalization methods.

2 Related Work

This section discusses existing domain generalization approaches and reviews two related problems, multi-source domain adaptation and normalization techniques in deep neural networks.

2.1 Domain Generalization

Domain generalization algorithms learn domain-invariant representations given input examples regardless of their domain memberships. Since target domain information is not available at training time, they typically rely on multiple source domains to extract knowledge applicable to any unseen domain. The existing domain generalization approaches can be roughly categorized into three classes. The first group of methods proposes novel loss functions that encourage learned representations to generalize well to new domains. Muandet *et al.* [21] propose domain-invariant component analysis, which generates invariant feature representation via dimensionality reduction. A few recent works [20, 15, 7] also attempt to learn a shared embedding space appropriate for semantic matching across domains. Another kind of approach tackles the domain generalization task by manipulating deep neural network architectures. Domain-specific information is handled by designated modules within deep neural networks [13, 6, 18] while [17] proposes a soft model selection technique to obtain generalized representations. Recently, meta-learning based techniques start to be used to solve domain

generalization problems. MLDG [12] extends MAML [8] to domain generalization task. Balaji *et al.* [2] points out the limitation of [12] and proposes a regularizer to address domain generalization in a meta-learning framework directly. Also, [14] presents an episodic training technique appropriate for domain generalization. Note that, to the best of our knowledge, none of the existing methods exploit normalization types and their optimization for domain generalization.

2.2 Multi-Source Domain Adaptation

Multi-source domain adaptation can be considered as the middle-ground between domain adaptation and generalization, where data from multiple source domains are used for training in addition to examples in an unlabeled target domain. Although unsupervised domain adaptation is a very popular problem, its multi-source version is relatively less investigated. Zhao *et al.* [31] propose to learn features that are invariant to multiple domain shifts through adversarial training, and Guo *et al.* [9] use a mixture-of-experts approach by modeling the inter-domain relationships between source and target domains. A recent work using domain-specific batch normalization (DSBN) [5] has shown competitive performance in multi-source domain adaptation by aligning the representations in heterogeneous domains to a single common feature space.

2.3 Normalization in Neural Networks

Normalization techniques in deep neural networks are originally designed for regularizing trained models and improving their generalization performance. Various normalization techniques [11, 27, 1, 19, 24, 29, 22, 16, 25, 5] have been studied actively in recent years. The most popular technique is batch normalization (BN) [11], which normalizes activations over individual channels using data in a mini-batch while instance normalization (IN) [27] performs the same operation per instance instead of mini-batch. In general, IN is effective to remove instance-specific characteristics (*e.g.* style in an image) and adding IN makes a trained model focus on instance-invariant information and increases generalization capability of the model to an unseen domain. Other normalizations such as layer normalization (LN) [1] and group normalization (GN) [29] have the same concept while weight normalization [24] and spectral normalization [19] normalize weights over parameter space.

Recently, batch-instance normalization (BIN) [22], switchable normalization (SN) [16], and sparse switchable normalization (SSN) [25] employ the combinations of multiple normalization types to maximize the benefit. Note that BIN considers batch and instance normalizations while SN uses LN additionally. On the other hand, DSBN [5] adopts separate batch normalization layers for each domain to deal with domain shift and generate domain-invariant representations.

3 Domain-Specific Optimized Normalization for Domain Generalization

This section describes our main algorithm called domain-specific optimized normalization (DSON) in details and also presents how the proposed method is employed to solve domain generalization problems.

3.1 Overview

Domain generalization aims to learn a domain-agnostic model that can be applied to an unseen domain by leveraging multiple source domains. Consider a set of training examples \mathcal{X}_s with its corresponding label set \mathcal{Y}_s in a source domain s . Our goal is to train a classifier using the data in multiple source domains $\{\mathcal{X}_s\}_{s=1}^S$ to correctly classify an image $\mathbf{x}_t \in \mathcal{X}_t$ in a target domain t , which is unavailable during training.

In our approach, we aim to learn a joint embedding space across all source domains, which is expected to be valid in target domains as well. To this end, we train domain-invariant classifiers from each of the source domains and ensemble their predictions. To embed each example onto a domain-invariant feature space, we employ domain-specific normalization, which is to be described in the following sections.

Our classification network consists of a set of feature extractors $\{F_s\}_{s=1}^S$ and a single fully connected layer D . Specifically, the feature extractors share all parameters across domains except for the ones in the normalization layers. For each source domain s , loss function is defined as

$$\mathcal{L}_C(\mathcal{X}_s, \mathcal{Y}_s) = \frac{1}{|\mathcal{X}_s|} \sum_{x \in \mathcal{X}_s, y \in \mathcal{Y}_s} \ell(y, D(F_s(x))), \quad (1)$$

where $\ell(\cdot, \cdot)$ is the cross-entropy loss. All the parameters are jointly optimized to minimize the sum of classification losses of source domains:

$$\mathcal{L} = \sum_{s=1}^S \mathcal{L}_C(\mathcal{X}_s, \mathcal{Y}_s). \quad (2)$$

Our domain-specific deep neural network model is obtained by minimizing the total loss \mathcal{L} . To facilitate generalization, in the validation phase, we follow the leave-one-domain-out validation strategy proposed in [6]; the label of a validation example from domain s is predicted by averaging predictions from all domain-specific classifiers, except for the one with domain s .

3.2 Instance Normalization for Domain Generalization

Normalization techniques [11, 27, 1, 29] are widely applied in recent network architectures for better optimization and regularization. Particularly, batch normalization (BN) [11] improves performance and generalization ability in training

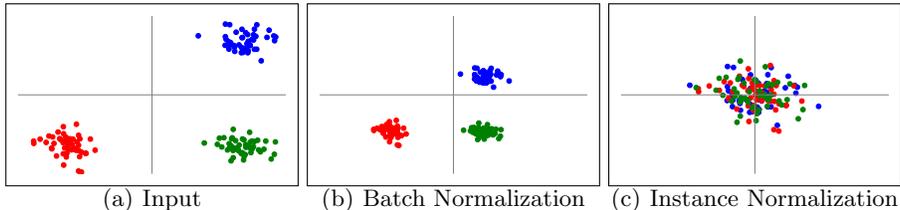


Fig. 2. Comparing feature distributions of three classes, where color represents the class label and each dot represents a feature map with two channels, where each axis corresponds to one channel. For given (a) input activations, (c) instance normalization makes the features less discriminative over classes when compared to (b) batch normalization. Although instance normalization loses discriminability, it makes the normalized representations less overfit to a particular domain and eventually improves the quality of features when combined with batch normalization. (Best viewed in color.)

Table 1. Effects of training batch normalization on the PACS dataset using a ResNet-18 architecture. Each column shows the performance on the target domain when a network is trained using the remaining domains as sources. Fine-tuning BN parameters degrades the generalization performance by overfitting to source domains.

	Art painting	Cartoon	Sketch	Photo	Avg.
BN fixed	79.25	74.61	71.52	95.99	80.34
BN finetuned	78.47	70.41	70.68	95.87	78.86

neural networks and has become indispensable in many deep neural networks. However, when BN is applied to cross-domain scenarios, it may not be optimal [3]. Table 1 empirically validates our claim about the effects of training BN in domain generalization. We employ a pretrained ResNet-18 on the ImageNet dataset as the backbone network, and fine-tune it on the PACS dataset with two different settings; fixing the BN parameters and statistics or fine-tuning them. Although BN generally works well in a variety of vision tasks, it consistently degrades performance when it is trained in the presence of a large domain divergence. This is because the batch statistics overfit to the particular training domains, resulting in poor generalization performance in unseen domains. This motivates us to construct a domain-agnostic feature extractor.

To achieve our goal, we combine BN with instance normalization (IN) to obtain domain-agnostic features. Our intuition about the two normalization methods is as follows. Instance normalization has been widely adopted in many works regarding style transfer due to its ability to perform *style normalization* [10]. Inspired by this, we employ IN as a means of reducing the inherent style information in each domain. In addition, IN does not depend on mini-batch construction or batch statistics, which can be helpful to extrapolate on unseen domains. These properties allow the network to learn feature representations that less overfit to

a particular domain. The downside of IN, however, is that it makes the features less discriminative with respect to object categories. This is illustrated in a simplified setting (Fig. 2), where we represent an instance by a cluster of data points and the corresponding classes by color. Unlike BN, which retains variation across the different classes, IN largely reduces the inter-class variance. To reap the benefits of IN while maintaining good classification performance, we utilize a mixture of IN and BN by optimizing the tradeoff between cross-category variance and domain invariance. More specifically, we fuse IN into all the BN layers of our network by linearly interpolating the means and variances of the two normalization statistics. The combination serves as a regularization method which results in a strong classifier that tends to focus on high-level semantic information but is much less vulnerable to domain shifts.

3.3 Optimization for Domain-Specific Normalization

Based on the intuitions above, we propose a domain-specific optimized normalization (DSON) for domain generalization. Given an example from domain d , the proposed domain-specific normalization layer transforms channel-wise whitened activations using affine parameters β_d and γ_d . Note that whitening is also performed for each domain. At each channel, the activations $\mathbf{x}_d \in \mathbb{R}^{H \times W \times N}$ are transformed as

$$\text{DSON}_d(\mathbf{x}_d[i, j, n]; \gamma_d, \beta_d) = \gamma_d \cdot \hat{\mathbf{x}}_d[i, j, n] + \beta_d, \quad (3)$$

where the whitening is performed using the domain-specific mean and variance, μ_{dn} and σ_{dn}^2 ,

$$\hat{\mathbf{x}}_d[i, j, n] = \frac{\mathbf{x}_d[i, j, n] - \mu_{dn}}{\sqrt{\sigma_{dn}^2 + \epsilon}}. \quad (4)$$

We combine batch normalization (BN) with instance normalization (IN) in a similar manner to [16] as

$$\mu_{dn} = w_d \mu_d^{\text{bn}} + (1 - w_d) \mu_n^{\text{in}}, \quad (5)$$

$$\sigma_{dn}^2 = w_d \sigma_d^{\text{bn}2} + (1 - w_d) \sigma_n^{\text{in}2}, \quad (6)$$

where both are calculated separately in each domain as

$$\mu_d^{\text{bn}} = \frac{\sum_n \sum_{i,j} \mathbf{x}_d[i, j, n]}{N \cdot H \cdot W} \quad \text{and} \quad \sigma_d^{\text{bn}2} = \frac{\sum_n \sum_{i,j} (\mathbf{x}_d[i, j, n] - \mu_d^{\text{bn}})^2}{N \cdot H \cdot W}, \quad (7)$$

and

$$\mu_n^{\text{in}} = \frac{\sum_{i,j} \mathbf{x}_d[i, j, n]}{H \cdot W} \quad \text{and} \quad \sigma_n^{\text{in}2} = \frac{\sum_{i,j} (\mathbf{x}_d[i, j, n] - \mu_n^{\text{in}})^2}{H \cdot W}. \quad (8)$$

The optimal mixture weight, w_d , between BN and IN are trained to minimize the loss in Equation 2. Note that our domain-specific mixture weights are shared across all layers for each domain, which facilitates to find the optimal point.

3.4 Inference

A test example x in a target domain is unknown during training. Hence, for inference, we feed the example to the feature extractors of all domains. The final label prediction is given by computing the logits using the fully connected layer D , averaging the logits, *i.e.* $\frac{1}{S} \sum_{s=1}^S D(F_s(x))$, and finally applying a softmax function.

One potential issue in the inference step is whether target domains can rely on the model trained only on source domains. This is the main challenge in domain generalization, which assumes that reasonably good representations of target domains can be obtained from the information in source domains only. In our algorithm, instance normalization in each domain has the capability to remove domain-specific styles and standardize the representation. Since each domain has different characteristics, we learn the relative weights of instance normalization in each domain separately. Thus, predictions in each domain should be accurate enough even for the data in target domains. Additionally, the accuracy given by aggregating the predictions of multiple networks trained on different source domains should further improve accuracy.

4 Experiments

To depict the effectiveness of domain-specific optimized normalization (DSON), we implement it on domain generalization benchmarks and provide an extensive ablation study of the algorithm.

4.1 Experimental Settings

Datasets We evaluate the proposed method on three domain generalization benchmarks. The PACS dataset [13] is commonly used in domain generalization and is favored due to its large inter-domain shift across four domains: Photo, Art Painting, Cartoon, and Sketch. It contains a total of 9,991 images in 7 categories, with an image resolution of 227×227 . We follow the experimental protocol in [13], where the model is trained on any three of the four domains (source domains), and then tested on the remaining domain (target domain). Office-Home [28] is a popular domain adaptation dataset, which consists of four distinct domains: Artistic Images, Clip Art, Product, and Real-world Images. Each domain contains 65 categories, with around 15,500 images in total. While the dataset is mostly used in the domain adaptation context, it can easily be repurposed for domain generalization by following the same protocol used in the PACS dataset. Finally, we employ five datasets—MNIST, MNIST-M, USPS, SVHN and Synthetic Digits—for digit recognition and split training and testing subsets following [30].

Implementation details For the fair comparison with prior arts [2, 6, 4, 14], we employ ResNet as the backbone network in all experiments. The convolutional and BN layers are initialized with ImageNet pretrained weights. We use

Table 2. Comparison with the state-of-the-art domain generalization methods (%) on the PACS dataset using ResNet-18 and ResNet-50 architectures. Column title indicates the target domain. *All experiments use the “train” split for the source domains, except MetaReg [2], which uses both “train” and “validation” splits.

Architecture	Method	Art painting	Cartoon	Sketch	Photo	Avg.
ResNet-18	Baseline	78.47	70.41	70.68	95.87	78.86
	JiGen [4]	79.42	75.25	71.35	96.03	80.51
	D-SAM [6]	77.33	72.43	77.83	95.30	80.72
	Epi-FCR [14]	82.10	77.00	73.00	93.90	81.50
	MetaReg* [2]	83.70	77.20	70.30	95.50	81.70
	MASF [7]	80.29	77.17	71.69	94.99	81.04
	MMLD [18]	81.28	77.16	72.29	96.09	81.83
	DSON (Ours)	84.67	77.65	82.23	95.87	85.11
ResNet-50	Baseline	80.22	78.52	76.10	95.09	82.48
	MetaReg* [2]	87.20	79.20	70.30	97.60	83.60
	MASF [7]	82.89	80.49	72.29	95.01	82.67
	DSON (Ours)	87.04	80.62	82.90	95.99	86.64

Table 3. Comparison with the state-of-the-art domain generalization methods (%) on the Office-Home dataset using a ResNet-18 architecture. Column title indicates the target domain.

	Art	Clipart	Product	Real-World	Avg.
Baseline	58.71	44.20	71.75	73.19	61.96
JiGen [4]	53.04	47.51	71.47	72.79	61.20
D-SAM [6]	58.03	44.37	69.22	71.45	60.77
DSON (Ours)	59.37	45.70	71.84	74.68	62.90

a batch size of 32 images per source domain, and optimize the network parameters over 10K iterations using SGD-M with a momentum 0.9 and an initial learning rate $\eta_0 = 0.02$. As suggested in [31], the learning rate is annealed by $\eta_p = \frac{\eta_0}{(1+\alpha p)^\beta}$, where $\alpha = 10$, $\beta = 0.75$, and p increases linearly from 0 to 1 as training progresses. We follow the domain generalization convention by training with the “train” split from each of the source domains, then testing on the combined “train” and “validation” splits of the target domain.

We made the mixture weights shared across all layers in our network to facilitate optimization. In our experiments, the convergence rates of local mixture weights in lower layers were significantly slower than higher layers. We sidestepped this issue by sharing the mixture weights across all the layers. This strategy improved accuracy substantially and consistently in all settings.

4.2 Comparison with Other Methods

In this section, we compare our method with other domain generalization methods on PACS and Office-Home datasets.

Table 4. Domain generalization accuracy (%) in ablation study. We compare DSON (ours) with its variants given by different implementations of normalization layers. ‘Art.’ denotes ‘Art painting’ domain in the PACS dataset.

	PACS					Office-Home				
	Art.	Cartoon	Sketch	Photo	Avg.	Art	Clipart	Product	Real-World	Avg.
Baseline	78.47	70.41	70.68	95.87	78.86	58.71	44.20	71.75	73.19	61.96
IBN [23]	75.29	72.95	77.42	92.04	79.43	55.41	44.82	68.28	71.95	60.09
DSBN [5]	78.61	66.17	70.15	95.51	77.61	59.04	45.02	72.67	71.98	62.18
SN [16]	82.50	76.80	80.77	93.47	83.38	54.10	44.97	64.54	71.40	58.75
DSON (Ours)	84.67	77.65	82.23	95.87	85.11	59.37	45.70	71.84	74.68	62.90

Table 5. Domain generalization accuracy (%) on Digits datasets using a ResNet-18 architecture. We compare DSON (ours) with its variants given by different implementations of normalization layers.

	Digits					
	MNIST	MNIST-M	USPS	SVHN	Synthetic	Avg.
Baseline	86.15	74.44	90.07	81.29	94.46	85.28
DSBN [5]	87.01	71.20	91.18	78.23	94.30	84.38
SN [16]	89.28	78.40	88.54	79.12	95.66	86.20
DSON (Ours)	89.62	79.00	91.63	81.02	95.34	87.32

PACS Table 2 portrays the domain generalization accuracy on the PACS dataset. The proposed algorithm is compared with several existing methods, which include JiGen [4], D-SAM [6], Epi-FCR [14], MetaReg [2], MASF [7], and MMLD [18]. Our method outperforms both the baseline and other state-of-the-art techniques by significant margins, which is particularly effective for *hard* domain (Sketch). When ResNet-50 is employed as a backbone architecture, our method still achieves better performance than other baselines; this result implies the scalability and stability of DSON when it is incorporated into a larger model.

Office-Home We also evaluate DSON on the Office-Home dataset, and the results are presented in Table 3. As in PACS, DSON outperforms the recently proposed JiGen [4] and D-SAM [6] as well as our baseline. We find that DSON achieves the best score on all target domains. Again, DSON is more advantageous in *hard* domain (Clipart).

4.3 Ablation Study

PACS and Office-Home Dataset We conduct an ablation study to assess the contribution of individual components within our full algorithm on the PACS and Office-Home datasets. Table 4 presents the results, where our complete method is denoted by DSON. It also presents accuracies of its variants given by different implementations of normalization layers. We first present results from the

Table 6. Domain generalization accuracy (%) using ResNet-18 in the presence of label noise on the PACS dataset. Note that Δ Avg. denotes the amount of accuracy drop with respect to the results from clean data.

Noise level	Method	Art painting	Cartoon	Sketch	Photo	Avg.	Δ Avg.
0.2	Baseline (BN)	75.16	70.41	68.17	92.13	76.47	-2.89
	IBN [23]	77.25	69.75	69.53	90.60	76.78	-2.65
	DSBN [5]	77.10	66.00	59.43	94.85	74.35	-3.27
	SN [16]	78.56	75.21	77.42	91.08	80.57	-2.57
	DSON (Ours)	83.11	79.07	80.15	94.79	84.03	-1.08
0.5	Baseline (BN)	73.49	64.68	58.95	89.22	71.59	-7.77
	IBN [23]	67.60	63.58	65.08	86.53	70.70	-8.73
	DSBN [5]	75.05	57.98	59.99	93.35	71.59	-6.02
	SN [16]	78.27	74.06	72.23	88.86	78.36	-4.78
	DSON (Ours)	80.22	73.85	77.37	94.91	81.59	-3.52

baseline method, where the model is trained naïvely with BN layers that are not specific to any single domain. Then, to examine the effects of domain-specific normalization layers, the BN layers are made specific to each of the source domains, which is denoted by DSBN [5]. We also examine the suitability of SN [16] by replacing BN layers with adaptive mixtures of BN, IN and LN. IBN-Net [23] concatenates instance normalization and batch normalization in a channel axis, but its improvement is marginal. We do not include batch-instance normalization (BIN) [22] in our experiment because it is hard to optimize and the results are unstable. The ablation study clearly illustrates the benefits of individual components in our algorithm: optimization of multiple normalization methods and domain-specific normalization. Note that other normalization methods can degrade the performance compared to baseline depending on the dataset, while DSON consistently displays superior results.

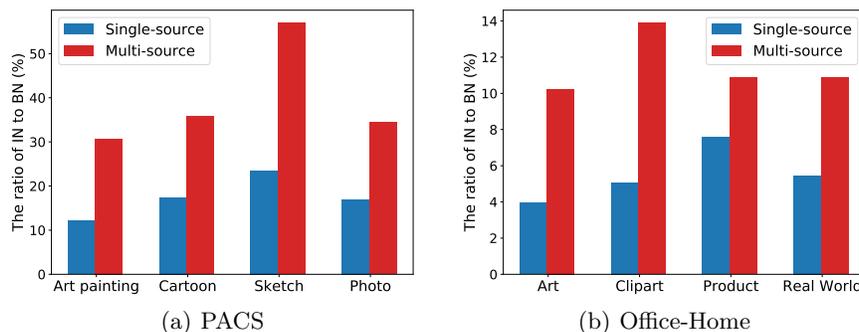
Digits Dataset The results on five digits datasets are shown in Table 5. Our model achieves 87.32% of average accuracy, outperforming all other baselines by large margins.

4.4 Additional Experiments

Robustness against Label Noise The performance of the proposed algorithm on the PACS dataset is tested in the presence of label noise, and the results are investigated against other approaches. Two different noise levels are tested (0.2 and 0.5), and the results are presented in Table 6. Although all algorithms undergo performance degradation, the amount of accuracy drops is marginal in general and DSON turns out to be more reliable with respect to label noise compared to other models. This is partly because DSON makes the network less overfit to class discrimination, reducing the effects of noise.

Table 7. Multi-source domain adaptation results on the PACS dataset using [26] as a backbone with a ResNet-18 architecture.

	Art painting	Cartoon	Sketch	Photo	Avg.
Baseline	78.85	76.79	78.14	99.42	83.30
DSBN [5]	88.94	83.54	77.39	99.42	87.32
SN [16]	79.00	77.66	79.11	97.78	83.39
DSON (Ours)	86.54	88.61	86.93	99.42	90.38

**Fig. 3.** Analysis of the mixture weights with single-source and multi-source scenarios on the PACS (left) and Office-Home (right) datasets. We present the weight ratio of IN to BN in our DSON module. (Best viewed in color.)

Multi-Source Domain Adaptation DSON can be extended to the multi-source domain adaptation task, where we gain access to unlabeled data from the target domain. To compare the effect of different normalization methods, we adopt the algorithm proposed by Shu *et al.* [26] as the baseline method and vary the normalization method only. The results are shown in Table 7, where we compare DSON with baseline, SN [16], and DSBN [5]. All compared methods illustrate a large improvement over the baseline. In direct contrast to the results from the ablation analysis in Table 4, DSBN is clearly superior to SN. This is unsurprising, given that DSBN is focused specifically on the domain adaptation task. Interestingly, we find that DSON outperforms not only the baseline but also DSBN, which demonstrates how effectively DSON can be extended to the domain adaptation task. Domain-specific models consistently outperform their domain-agnostic counterparts in this task.

4.5 Analysis

Mixture Weights For a trained network with DSON in domain generalization tasks, the mixture weights of IN and BN are 3:7 and 1:9 on average for PACS and Office-Home datasets, respectively. Additionally, we analyzed the effect that

Table 8. Effects of separating mixture weights for each domain in our DSON module on the PACS dataset with multi-source domain generalization scenario. Domain-agnostic denotes the mixture weights are shared across domains.

Mixture weights	Art painting	Cartoon	Sketch	Photo	Avg.
Domain-agnostic	82.13	74.10	80.02	95.03	82.82
Domain-specific (Ours)	84.67	77.65	82.23	95.87	85.11

Table 9. Single-source domain generalization accuracy on the PACS dataset. Rows and columns denote source and target domains, respectively. We compare DSON (ours) with BN in terms of the amount of change (%p).

	BN (%)				DSON (%p)			
	A	C	S	P	A	C	S	P
Art painting	89.42	60.71	48.74	94.25	- 1.20	+ 7.42	+ 20.36	- 0.44
Cartoon	69.68	94.51	71.60	81.56	+ 2.83	- 0.11	+ 0.88	+ 3.65
Sketch	44.34	63.65	93.50	49.28	+ 8.64	+ 0.64	- 0.67	+ 9.82
Photo	61.72	29.10	33.98	97.86	+ 3.56	+ 16.85	+ 1.19	- 0.52

the number of source domains has on the value of the mixture weights. To this end, we tested two scenarios; single-source and multi-source. For each domain, single-source denotes that only the specified domain was used for training, while multi-source indicates that other domains, along with the specified domain, were used for training. The graphs in Fig. 3 show the weight ratio of IN to BN for each domain in the two scenarios. As shown in the plot, training with multi-source domains led to a large and consistent increase in the usage of IN for all domains, because the multi-source scenario requires more domain-invariant representations than the single-source scenario. We also validated the effectiveness of separating mixture weights for each domain independently in Table 8, in which domain-specific mixture weights boosts the performance compared to the domain-agnostic ones.

Effects of Instance Normalization To analyze the effects of combining instance normalization, we tested single-source domain generalization on every source-target combination of domains, as shown in Table 9. Rows and columns denote source domains and target domains, respectively. For evaluation, we used the validation split of the target domain in case source and target domains are the same. Although DSON marginally sacrifices the accuracy compared to BN when source and target domains are the same, it brought a significant performance gain in most cross-domain scenarios. This presents a desirable trade-off in combining instance normalization; cross-category variance for domain invariance.

Table 10. Results from single source domain branch on the PACS dataset. Columns denote target domains and rows denote the single source domain branch of our model.

	Art painting	Cartoon	Sketch	Photo
Art painting	-	73.68	79.51	95.41
Cartoon	80.96	-	75.87	93.77
Sketch	78.71	71.72	-	91.98
Photo	79.49	75.71	76.38	-
DSON (all)	84.67	77.65	82.23	95.87

Results from Single Source Domain Branch We present the classification results using each single domain branch of our model on the PACS dataset in Table 10. Columns denote target domains and rows denote each single source domain branch of our model. It shows that the results from single domain branches differ slightly from each other in accuracy, and integrating them gives consistent performance gain.

5 Conclusion

We presented a simple but effective domain generalization algorithm based on domain-specific optimized normalization layers. The proposed algorithm uses multiple normalization methods while learning a separate affine parameter per domain. The mixing weights are employed to compute the weighted average of multiple normalization statistics for each domain separately. This strategy turns out to be helpful for learning domain-invariant representations since instance normalization removes domain-specific style while preserving semantic category information effectively. The proposed algorithm achieves the state-of-the-art accuracy consistently on multiple standard benchmarks even with substantial label noise. We showed that the algorithm is well-suited for unsupervised domain adaptation as well. Finally we analyzed the characteristics and effects of our method with diverse ablative study.

Acknowledgement

This work was supported by Institute for Information & Communications Technology Promotion (IITP) grant funded by the Korea government (MSIT) [2016-0-00563, 2017-0-01779].

References

1. Ba, L.J., Kiros, R., Hinton, G.E.: Layer normalization. arXiv preprint arXiv:1607.06450 (2016) 4, 5
2. Balaji, Y., Sankaranarayanan, S., Chellappa, R.: Metareg: Towards domain generalization using meta-regularization. In: NeurIPS (2018) 2, 4, 8, 9, 10

3. Bilen, H., Vedaldi, A.: Universal representations: The missing link between faces, text, planktons, and cat breeds. arXiv preprint arXiv:1701.07275 (2017) [6](#)
4. Carlucci, F.M., D’Innocente, A., Bucci, S., Caputo, B., Tommasi, T.: Domain generalization by solving jigsaw puzzles. In: CVPR (2019) [8, 9, 10](#)
5. Chang, W.G., You, T., Seo, S., Kwak, S., Han, B.: Domain specific batch normalization for unsupervised domain adaptation. In: CVPR (2019) [4, 10, 11, 12](#)
6. DaInnocente, A., Caputo, B.: Domain generalization with domain-specific aggregation modules. In: GCPR (2018) [2, 3, 5, 8, 9, 10](#)
7. Dou, Q., de Castro, D.C., Kamnitsas, K., Glocker, B.: Domain generalization via model-agnostic learning of semantic features. In: NeruIPS (2019) [2, 3, 9, 10](#)
8. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: ICML (2017) [4](#)
9. Guo, J., Shah, D., Barzilay, R.: Multi-source domain adaptation with mixture of experts. In: EMNLP (2018) [4](#)
10. Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: ICCV (2017) [6](#)
11. Ioffe, S., Szegedy, C.: Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In: ICML (2015) [4, 5](#)
12. Li, D., Yang, Y., Song, Y.Z., Hospedales, T.: Learning to generalize: Meta-learning for domain generalization. In: AAAI (2018) [2, 4](#)
13. Li, D., Yang, Y., Song, Y.Z., Hospedales, T.M.: Deeper, broader and artier domain generalization. In: ICCV (2017) [2, 3, 8](#)
14. Li, D., Zhang, J., Yang, Y., Liu, C., Song, Y.Z., Hospedales, T.M.: Episodic training for domain generalization (2019) [2, 4, 8, 9, 10](#)
15. Li, Y., Gong, M., Tian, X., Liu, T., Tao, D.: Domain generalization via conditional invariant representations. In: AAAI (2018) [2, 3](#)
16. Luo, P., Ren, J., Peng, Z.: Differentiable learning-to-normalize via switchable normalization. In: ICLR (2019) [2, 4, 7, 10, 11, 12](#)
17. Mancini, M., Bulò, S.R., Caputo, B., Ricci, E.: Best sources forward: domain generalization through source-specific nets. In: ICIP (2018) [2, 3](#)
18. Matsuura, T., Harada, T.: Domain generalization using a mixture of multiple latent domains. In: AAAI (2020) [3, 9, 10](#)
19. Miyato, T., Kataoka, T., Koyama, M., Yoshida, Y.: Spectral normalization for generative adversarial networks. In: ICLR (2018) [4](#)
20. Motiian, S., Piccirilli, M., Adjeroh, D.A., Doretto, G.: Unified deep supervised domain adaptation and generalization. In: ICCV (2017) [2, 3](#)
21. Muandet, K., Balduzzi, D., Schölkopf, B.: Domain generalization via invariant feature representation. In: ICML (2013) [2, 3](#)
22. Nam, H., Kim, H.E.: Batch-instance normalization for adaptively style-invariant neural networks. In: NeurIPS (2018) [2, 4, 11](#)
23. Pan, X., Luo, P., Shi, J., Tang, X.: Two at once: Enhancing learning and generalization capacities via ibn-net. In: ECCV (2018) [2, 10, 11](#)
24. Salimans, T., Kingma, D.P.: Weight normalization: A simple reparameterization to accelerate training of deep neural networks. In: NIPS (2016) [4](#)
25. Shao, W., Meng, T., Li, J., Zhang, R., Li, Y., Wang, X., Luo, P.: Ssn: Learning sparse switchable normalization via sparsestmax. In: CVPR (2019) [4](#)
26. Shu, Y., Cao, Z., Long, M., Wang, J.: Transferable curriculum for weakly-supervised domain adaptation. In: AAAI (2018) [12](#)
27. Ulyanov, D., Vedaldi, A., Lempitsky, V.: Improved Texture Networks: Maximizing Quality and Diversity in Feed-Forward Stylization and Texture Synthesis. In: CVPR (2017) [2, 4, 5](#)

28. Venkateswara, H., Eusebio, J., Chakraborty, S., Panchanathan, S.: Deep Hashing Network for Unsupervised Domain Adaptation. In: CVPR (2017) 8
29. Wu, Y., He, K.: Group Normalization. In: ECCV (2018) 4, 5
30. Xu, R., Chen, Z., Zuo, W., Yan, J., Lin, L.: Deep cocktail network: Multi-source unsupervised domain adaptation with category shift. In: CVPR (2018) 8
31. Zhao, H., Zhang, S., Wu, G., Moura, J.M.F., Costeira, J.P., Gordon, G.J.: Adversarial multiple source domain adaptation. In: NeurIPS (2018) 4, 9