

MIXNORM: TEST-TIME ADAPTATION THROUGH ONLINE NORMALIZATION ESTIMATION

Anonymous authors

Paper under double-blind review

ABSTRACT

We present a simple and effective way to estimate the batch-norm statistics during test time, to fast adapt a source model to target test samples. Known as Test-Time Adaptation, most prior works studying this task follow two assumptions in their evaluation where (1) test samples come together as a large batch, and (2) all from a single test distribution. However, in practice, these two assumptions may not stand, the reasons for which we propose two new evaluation settings where batch sizes are arbitrary and multiple distributions are considered. Unlike the previous methods that require a large batch of single distribution during test time to calculate stable batch-norm statistics, our method avoid any dependency on large online batches and is able to estimate accurate batch-norm statistics with a single sample. The proposed method significantly outperforms the State-Of-The-Art in the newly proposed settings in Test-Time Adaptation Task, and also demonstrates improvements in various other settings such as Source-Free Unsupervised Domain Adaptation and Zero-Shot Classification.

1 INTRODUCTION

Deep learning models have brought tremendous performance improvements over many different domains (Krizhevsky et al. (2012); He et al. (2016); Deng et al. (2009)) in recent years. In spite of this, we often find that a well-trained model can perform unexpectedly poorly in real-world settings. Many factors play a role in this challenging problem; but one of the most probable causes stems from a domain shift from the training data (Quiñero-Candela et al. (2009)). To study the domain shift in production environment and to evaluate the generalization ability of deep learning models, many benchmarks have been proposed (Saenko et al. (2010); Peng et al. (2017); Hendrycks & Dietterich (2019)). Figure 1 demonstrates the samples from one of the benchmarks, ImageNet-C (Hendrycks & Dietterich (2019)), in which synthesized corruptions are added to ImageNet (Deng et al. (2009)) test images in an attempt to simulate domain shift.

To increase deep learning models’ robustness when there is domain shift, various approaches have been developed in different settings, such as Domain Adaptation (Patel et al. (2015)), Unsupervised Domain Adaptation (Wilson & Cook (2020)), Source-Free Domain Adaptation (Liang et al. (2020)) and Fully Test-Time Adaptation (Wang et al. (2020)). Among all those different settings, the task of Fully Test-Time Adaptation is one of the most challenging settings. With only access to a model that is pre-trained offline, the goal of Test-Time Adaptation (TTA) is to rapidly adapt the model to the test samples while making predictions. The task is essentially Domain Adaptation plus three more challenging conditions:

- **Online:** The model gets updated during inference, and has access to each test sample only once.
- **Source-Free:** The model does not have access to training data after the offline training is done. It will only have access to the pre-trained weights.
- **Unsupervised:** There is no label for the test data.

Several works have studied the Test-Time Adaptation problem in recent years (Li et al. (2016); Sun et al. (2020); Wang et al. (2020)). Existing methods focus on adapting the Batch Normalization Layers (Ioffe & Szegedy (2015)) in the deep learning models, and have significantly improved deep learning model’s performance against test-time domain shift.

Such Batch-Norm Layer based methods make two underlying assumptions to estimate stable normalization statistics during evaluation: 1) test data come together in large batches; 2) all the test samples are from the same shifted distribution (single corruption types, single domain, etc). However, such assumptions might not be practical in the real world as the samples might come from an arbitrary distribution such as: varying sample sizes, different ways of collection, or different types of corruptions. Further, during inference, the AI system might not have the freedom to postpone the prediction in order to collect enough data to apply test-time adaptation algorithm on the incoming test samples.

To eliminate the dependency on large batch size and sole distribution, we propose to replace the Batch-Norm Layers with a novel MixNorm operation that performs adaptive prediction of the normalization statistics from a single input. The MixNorm Layer considers not only global statistics calculated from all historical samples, but also local statistics calculated from incoming test sample and its spatial augmentations. Figure 2 illustrates our overall approach. The novel MixNorm method demonstrate significant improvement over the State-of-the-art methods TENT (Wang et al. (2020)), especially in the two newly proposed evaluation protocols, where the all Test-Time Adaptation methods are evaluated at various batch sizes, and against mixture of test samples from different shifted domains. In addition to the Test-Time Adaptation task, the proposed MixNorm Layer demonstrates improvement over two other tasks, source-free Domain Adaptation and Zero-Shot Image Classification.

Our contributions are:

- We propose two new evaluation conditions for the task of Test-Time Adaptation, which avoid the impractical assumptions made by previous protocols.
- We propose a novel and simple way to adaptively estimate the test-time normalization statistics from single sample.
- We performed extensive experiments that demonstrate the effectiveness of the new MixNorm Layer in various tasks, such as Test-Time Adaptation, Source-Free Unsupervised Domain Adaptation and Zero-Shot Image Classification.

2 RELATED WORK

Distribution shift between training domain and evaluation domain often causes a drop in the performance of deep learning models. In recent years, many proposals have been made to increase deep learning model’s robustness against distribution shifts, a field of research known as Domain Adaptation. Domain Adaptation methods can be classified into different settings depending on factors like access to source (training) samples, existence of new or missing categories, availability of supporting samples in evaluation domain, etc. In this paper, we focus on settings where there is no access to the source (training) samples. In particular, we compare our new method with existing methods on Test-Time Adaptation, Source-Free Unsupervised Domain Adaptation and Zero-Shot Classification settings.

2.1 TEST-TIME ADAPTATION

The task of Test-Time Adaptation (TTA) is one of the most challenging settings in the field of Domain Adaptation, as it requires rapid adaptation while making predictions, without access to

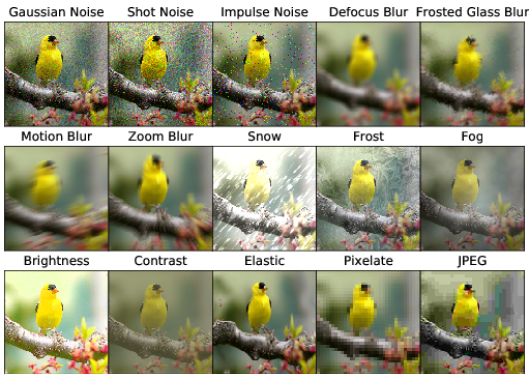


Figure 1: Examples from ImageNet-C dataset Hendrycks & Dietterich (2019). ImageNet-C contains 75 corrupted copies of the original ImageNet-1K test set, with 5 severity levels and 15 different corruption types. This dataset is used to simulate the distribution shift between training domain (original) and test domains (corrupted).

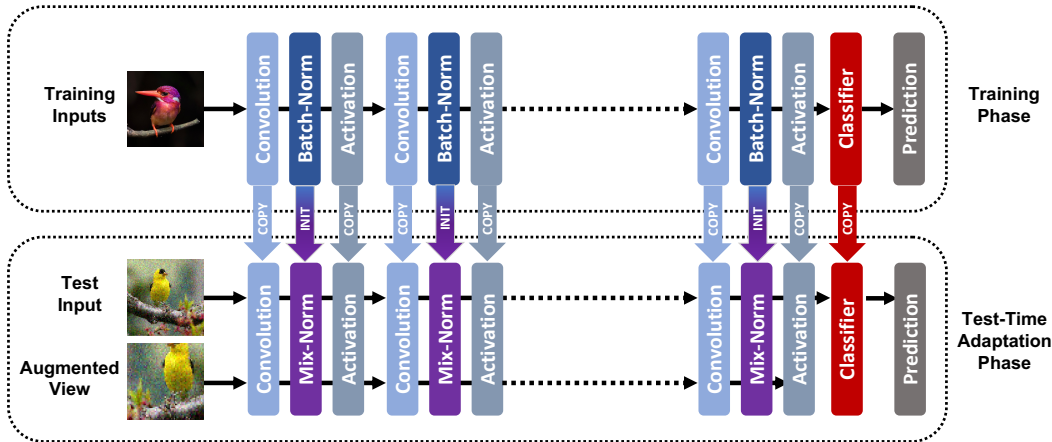


Figure 2: Method overview for Test-Time Adaptation with MixNorm layers. Before the inference, we replace all the Batch-Norm layers in the pre-trained model with the new proposed MixNorm layers. The MixNorm layers are initialized by the training set (source) statistics (means and variances) and the corresponding weights and biases from the pre-trained Batch-Norm layers. During inference, each input is paired with another augmented view for calculation in the MixNorm layers and the final prediction is made only on the original input.

anything but the pre-trained model and testing samples. BN (Li et al. (2016)) was one of the earliest work that explores the TTA task. It focuses on adapting the Batch-Norm Layers in deep learning models. A traditional Batch-Norm module collects running means and variances at training time from the training data, and uses the collected means and variance at test time. BN argued that when there are distribution shifts between training and testing data, it should not keep using the same training statistics. Instead, the model should collect statistic directly from the online batch. BN demonstrates that this simple approach of using online statistics is effective in handling the distribution shifts caused by corruptions. However, one clear drawback of the method is that it requires relative large batches to obtain stable online statistics. Following the work of BN (Li et al. (2016)), TTT (Sun et al. (2020)) proposed to update not only the batch-norm statistics, but also the affine weights and biases, which are learnable weights in Batch-Norm Layers that further re-scale the normalized output. TTT performs a single step gradient update over the batch-norm parameters after each prediction, with an unsupervised loss defined by rotation prediction task. TENT (Wang et al. (2020)) further proposed to use entropy loss instead of rotation prediction loss and achieved state-of-the-art performance on CIFAR-10C, CIFAR-100C and ImageNet-C datasets.

Most of the existing works in TTA follow the same evaluation protocol proposed by Li et al. (2016), where the models are adapted using large batches of test samples coming from a single type of corruption. However, such an assumption is not realistic in practice, because the test samples in the real world could come in any batch size (in particular, real time high performance inference platforms often process one input at a time, Anderson et al. (2021)), or from any shifted distribution.

2.2 SOURCE-FREE UNSUPERVISED DOMAIN ADAPTATION

Similar to the task of TTA, in Source-Free Unsupervised Domain Adaptation (UDA), the models only have access to pre-trained models and test samples. Different than the TTA setting, Source-Free UDA models have access to the evaluation samples as many times as needed during adaptation.

SHOT Liang et al. (2020) is currently the state of the art in Source-Free UDA, and also one of the first to propose UDA without needing any source data. In SHOT, the parameters are fine-tuned with a self-supervised loss that minimize the prediction entropy in the target samples, while keeping the classification head intact. SHOT demonstrates that it is possible to get reasonable performance on standard benchmarks such as Office (Saenko et al. (2010)), Office-Home (Venkateswara et al. (2017)) and VisDA (Peng et al. (2017)).

Another method called TENT (Wang et al. (2020)) also has shown potential for the Source-Free UDA task with promising performance on digit datasets (e.g. adapting SVHN model to MNIST/MNIST-M/USPS). TENT cannot reach SHOT’s performance when scaled to much larger scale datasets such as VisDA (Peng et al. (2017)), but TENT adapts to target domain inputs with much cheaper operations and sees each test sample only once during testing.

Algorithm 1: Mix-Norm Update

Input:

Input Feature: $F \in R^{D \times H \times W}$
 Augmentation Features of the same sample:
 $F' \in R^{N \times D \times H \times W}$
 Training Set Mean μ^0 and Variance σ^0
 Exponential Moving Speed $\tau = 0.001$
 Mixing Scale $m = 0.05$
 Weights, Bias: α, β

Procedure:

```
// Update  $\mu$  and  $\sigma$  with most recent sample
 $\mu^{t+1} \leftarrow (1 - \tau)\mu^t + \tau F.mean(1, 2)$ 
 $\sigma^{t+1} \leftarrow (1 - \tau)\sigma^t + \tau(F - \mu).var(1, 2)$ 
// Get global statistics
 $\mu_{global} \leftarrow \mu^t$ 
 $\sigma_{global} \leftarrow \sigma^t$ 
// Get local statistics from Augmentations
 $\mu_{local} \leftarrow F'.mean(0, 2, 3)$ 
 $\sigma_{local} \leftarrow (F' - \mu_{local}).var(0, 2, 3)$ 
// Mix Global and Local
 $\mu_{mixed} \leftarrow (1 - m)\mu_{global} + m\mu_{local}$ 
 $\sigma_{mixed} \leftarrow (1 - m)\sigma_{global} + m\sigma_{local}$ 
// Calculate shifted f
 $F \leftarrow \alpha \frac{F - \mu_{mixed}}{\sigma_{mixed}} + \beta$ 
return F
```

Algorithm 2: MixNormBN: Combines MixNorm and BatchNorm

Input:

Batch of Features: $F \in R^{B \times D \times H \times W}$
 Augmented Features of the same batch:
 $F' \in R^{B \times N \times D \times H \times W}$
 Training Set Mean μ^0 and Variance σ^0
 Max Moving Speed $\tau_{max} = 0.9$
 Mixing Scale $m = 0.05$
 Weights, Bias: α, β

Procedure:

```
// Adjust  $\tau$  based on Batch Size B
 $\tau \leftarrow \tau_{max} 10^{-\frac{3}{B}}$ 
// Update  $\mu$  and  $\sigma$  with most recent sample
 $\mu^{t+1} \leftarrow (1 - \tau)\mu^t + \tau F.mean(0, 1, 2)$ 
 $\sigma^{t+1} \leftarrow (1 - \tau)\sigma^t + \tau(F - \mu).var(0, 1, 2)$ 
// Update global statistics with current batch
 $\mu_{global} \leftarrow \mu^t$ 
 $\sigma_{global} \leftarrow \sigma^t$ 
// Get local statistics from Augmentations
 $\mu_{local} \leftarrow F'.mean(0, 1, 3, 4)$ 
 $\sigma_{local} \leftarrow (F' - \mu_{local}).var(0, 1, 3, 4)$ 
// Mix Global and Local
 $\mu_{mixed} \leftarrow (1 - m)\mu_{global} + m\mu_{local}$ 
 $\sigma_{mixed} \leftarrow (1 - m)\sigma_{global} + m\sigma_{local}$ 
// Calculate shifted f
 $F \leftarrow \alpha \frac{F - \mu_{mixed}}{\sigma_{mixed}} + \beta$ 
return F
```

2.3 ZERO-SHOT CLASSIFICATION

Unlike Test-Time Adaptation or Source-Free UDA, where methods are developed to close the domain gap during evaluation stage, the goal of Zero-Shot Classification (Xian et al. (2018)) is to increase the generalization ability during the training stage so that it can be applied to any test domain without adaptation. Without access to any support samples from the new categories, Zero-Shot Classification methods typically take advantage of cues from syntax information such as label embedding (Frome et al. (2013)) or attributes (Lampert et al. (2013)) to establish knowledge about the new categories.

Recently, along with the success in large-scale pre-training (Brown et al. (2020)), self-supervised learning and contrastive learning (Chen et al. (2020)), CLIP (Radford et al. (2021)) has brought a huge improvement over almost all of the common benchmark datasets. CLIP collects 400 millions image-caption pairs and learns a joint representation space of visual and language representations. With the joint representation space, CLIP has the ability to do Zero-Shot Classification by directly looking at the potential category names, and its Zero-Shot Classification performance is close to that of fully supervised models on many datasets including ImageNet (Deng et al. (2009)).

3 TEST-TIME ADAPTATION WITH MIX-NORM LAYER

In this section, we define our new Mix-Norm layer and provide details on how it can be utilized for closing domain gap during test time on Deep Neural Networks (DNN). Given a pre-trained

DNN, we replace its existing Batch-Norm layers (Ioffe & Szegedy (2015)). During test time, for arbitrary sized batches our new MixNorm layer calculates empirical normalization statistics from two sources:

- **Global Statistics:** The Batch-Norm layers from the pre-trained model store the statistics from the training distribution. Our MixNorm layer uses this training statistics as a global anchor point. The global statistics, initialized by the training statistics, is updated by an exponential moving average of the online test sample.
- **Local Statistics:** We create additional augmented views of the test sample to form a small augmentation batch. We calculate the empirical statistics from this augmentation batch to capture the local distribution shift of the test sample. The augmentations are spatial, including random-resized-crop and random-flipping. In practice, we only use one additional augmentation as it has the best performance.

In a typical deep learning model such as ResNet (He et al. (2016)) or WideResNet (Zagoruyko & Komodakis (2016)), there will be multiple batch-norm layers across the whole backbone network. Figure 3 below, illustrates our proposed operation in only one such layer. As shown earlier in Figure 2, each sample X comes along with its augmented view X' . After the encoding layers and right before our MixNorm layer, input images $X, X' \in \mathbb{R}^{3 \times 224 \times 224}$ become feature tensor $F, F' \in \mathbb{R}^{D \times H \times W}$. For simplicity, we denote the i -th row, j -th column of F, F' to be $F_{i,j}, F'_{i,j} \in \mathbb{R}^D$.

As depicted in Figure 3, our global statistics μ^t, σ^t are initialized with train-time statistics, copied from the batch-norm layers in the pre-trained network:

$$\mu^0 = \mu^{training} \qquad \sigma^0 = \sigma^{training},$$

and they will be slowly updated by the statistics on each new sample. At $(t + 1)$ -th test example:

$$\mu^{t+1} = (1 - \tau)\mu^t + \tau\mu \qquad \sigma^{t+1} = (1 - \tau)\sigma^t + \tau\sigma,$$

where τ is the exponential moving speed, and μ and σ are the means and variances from F :

$$\mu = \frac{\sum_{i=0}^H \sum_{j=0}^W F_{ij}}{HW} \qquad \sigma = \frac{\sum_{i=0}^H \sum_{j=0}^W (F_{ij} - \mu)^2}{HW}.$$

On the other hand, we calculate the local statistics $\mu_{local}, \sigma_{local}$ by examining both F, F' :

$$\mu_{local} = \frac{\sum_{F^* \in \{F, F'\}} \sum_{i=0}^H \sum_{j=0}^W F_{ij}^*}{2HW} \qquad \sigma_{local} = \frac{\sum_{F^* \in \{F, F'\}} \sum_{i=0}^H \sum_{j=0}^W (F_{ij}^* - \mu_{local})^2}{2HW}.$$

Renaming μ^{t+1}, σ^{t+1} to $\mu_{global}, \sigma_{global}$, we obtain our final statistics:

$$\mu_{mixed} = (1 - m)\mu_{global} + m\mu_{local} \qquad \sigma_{mixed} = (1 - m)\sigma_{global} + m\sigma_{local}.$$

Finally, with $\mu_{mixed}, \sigma_{mixed}$, we can perform the normalization operation, and forward the resulting F, F' to next layers:

$$F = \alpha \frac{F - \mu_{mixed}}{\epsilon + \sqrt{\sigma_{mixed}}} + \beta \qquad F' = \alpha \frac{F' - \mu_{mixed}}{\epsilon + \sqrt{\sigma_{mixed}}} + \beta,$$

where α, β, ϵ are the parameters copied from the pre-trained network batch-norm layers.

While the above process describes the procedure in Algorithm 1, we also present an variant of MixNorm in Algorithm 2. Different from the vanilla MixNorm which only consider single sample, Algorithm 2 updates the global statistics with batch statistics, and have an adaptive moving speed τ which approximates Batch-Norm behavior for large batch size, and MixNorm behavior for small batch size.

4 EXPERIMENTS AND RESULTS

Datasets. For the Test-Time Adaptation task, we follow TENT (Wang et al. (2020)) and evaluate our method over the standard benchmarks CIFAR-10C and ImageNet-C (Hendrycks & Dietterich

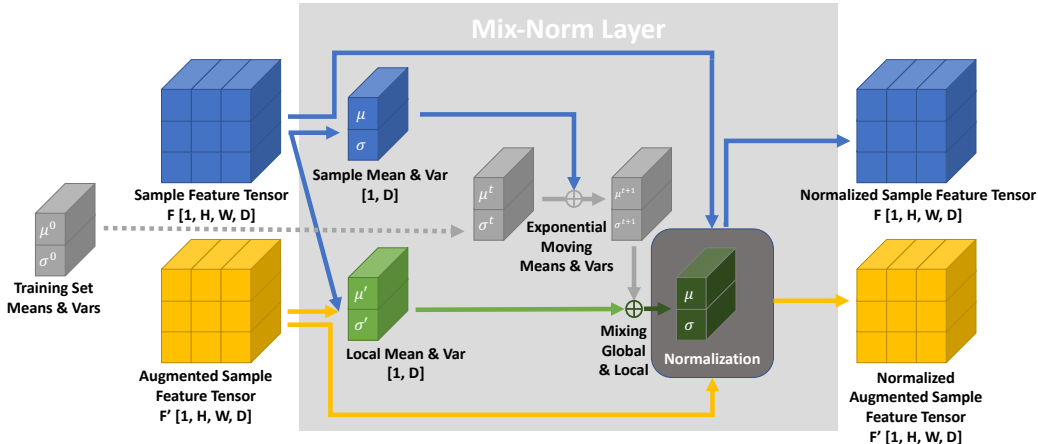


Figure 3: Illustration of the MixNorm Layer. Inside the MixNorm layer we estimate the normalization statistics by looking at both global statistics from historical data and local statistics from augmented views.

(2019)). CIFAR-10C includes 75 subsets that are composed of 15 types of corruptions at 5 severity levels. It comprises the 10,000 images, each of size 32x32 pixels, from the 10 classes in the CIFAR10 test set, but with each modified by the corresponding corruption type at certain severity level. Similarly, ImageNet-C is composed of a set of 75 common visual corruptions, applied to the original ImageNet validation set, with 50,000 images of 224x224 pixels from 1000 classes.

For Source-Free Unsupervised Domain Adaptation task, we follow the experiment setup from Wang et al. (2020), and use SVHN (Netzer et al. (2011)) as the source domain and MNIST (LeCun et al. (1998))/MNIST-M (Ganin & Lempitsky (2015))/USPS (Netzer et al. (2011)) as the target domains. SVHN, MNIST, MNIST-M and USPS are all 10-class digits datasets, with training/testing data sizes of 73257/26032, 60000/10000, 60000/10000 and 7291/2007 respectively.

For Zero-Shot Image Classification, we follow the setup from CLIP (Radford et al. (2021)) and evaluate our method on CIFAR-10/100 (Krizhevsky et al. (2009)), STL-10 (Coates et al. (2011)), Stanford Cars (Krause et al. (2013)) and Food101 (Bossard et al. (2014)). Details of these datasets are provided in the Appendix C.

Implementation Details. For Test-Time Adaptation experiments, we follow the optimization setting from TENT (Wang et al. (2020)), which uses Adam optimizer on CIFAR-10C with learning rate 0.001 and SGD optimizer on ImageNet-C with learning rate 0.00025, both optimized by the entropy minimization loss on output logits as described in Wang et al. (2020). For CIFAR-10C we compare with TENT on both Wide-ResNet-28 and Wide-ResNet-40 (Zagoruyko & Komodakis (2016)) and for ImageNet-C we compare with TENT on ResNet-50 (He et al. (2016)) following the official public model provided by the TENT repository¹.

For Source-Free Unsupervised Domain Adaptation experiments, we replicated TENT (Wang et al. (2020)) and BN (Li et al. (2016)) methods using the pre-trained SVNH model repository² since the original ResNet-26 model reported in TENT (Wang et al. (2020)) was not public. For hyper-parameter selection, on CIFAR-10C/ImageNet-C we use the “Gaussian Noise at Severity 5” set to choose the hyper parameters, and then use the same hyper-parameters during all experiments. For Source-Free UDA and Zero-Shot Classification experiments, as there is no trivial way to split a validation set, we just report the optimal hyper-parameters for each dataset. The details of the selected hyper-parameters will be provided in Appendix B. For all the experiments, we use RandomResizedCrop and RandomFlip as the augmentation methods.

¹<https://github.com/DequanWang/tent>

²<https://github.com/aaron-xichen/pytorch-playground>

4.1 TEST-TIME ADAPTATION

To understand the merit of our new method in real world scenarios, we introduce two new conditions during evaluations. In TENT, error rates for each of the 15 corruption types at severity 5 are reported with a fixed batch size of 200 on CIFAR-10C and 64 on ImageNet-C. In addition to this standard protocol, we compare our method with TENT under two new conditions:

1) **Different Batch Sizes.** In addition to the fixed 200/64 batch size in the previous protocol, we also compare our method with TENT at batch sizes of 1, 5, 8, 16, 32, 64, 100, 200;

2) **Mixed Distribution.** In addition to the standard protocol where the model’s final error rate is averaged from 15 separate evaluations, each of which contains only one type of corruption, we further evaluate the models with all 15 corruptions types shuffled and mixed together.

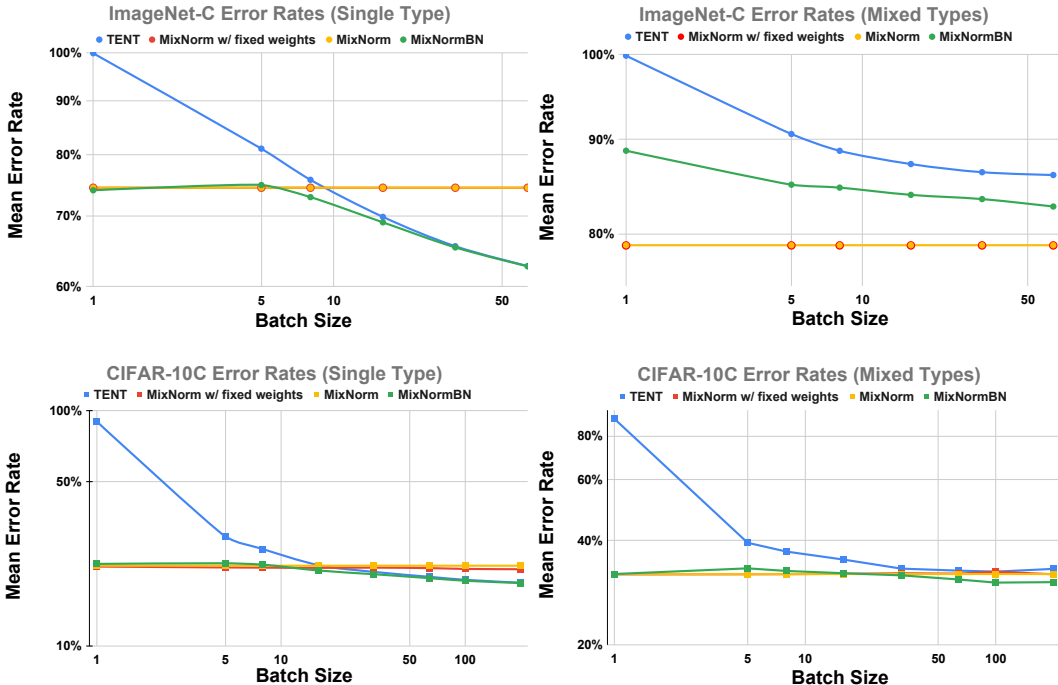


Figure 4: Comparison of MixNorm, with and without learnable affine weights (Algorithm 1), and MixNormBN (Algorithm 2) methods to the baseline method TENT in two different test time settings i) test samples come only from a single type of corruption; ii) test samples come from mixed types of corruptions. For the Single Type experiments, we report average error rates over all 15 corruption types at severity 5. Numerical results are reported in the Appendix A.

Figure 4 shows our main results. Blue curves represent the SOTA model TENT, the red and orange curves represent MixNorm model with and without learnable affine weights α and β . Green curves represent the performance of MixNormBN as described in Algorithm 2. Tent’s performance drops as the batch size decreases, while MixNorm’s performance stays stable for any batch size. Additional to the improvement to small batch size, Figure 4 also shows the advantages of MixNorm when test samples come from different distributions. In TENT’s testing protocol, all test samples are collected from the same corruption distribution, which makes it easier to obtain reliable normalization statistics especially when the batch size is large. However, when tested with batches of mixed inputs sampled from different corruption types, such assumption does not hold and Tent’s performance drops significantly compared to MixNorm. This result holds even when tested with larger batch size (up to 64 on ImageNet-C and up to 200 on CIFAR-10C). We also observe that when batch size is large enough and all samples come from a similar distribution, TENT might have an advantage over MixNorm; because TENT can get to estimate stable normalization statistics from empirical samples. In order to take advantage of large batch size, when we test our MixNormBN method (green curves), again we report consistently better scores than TENT at all batch sizes.

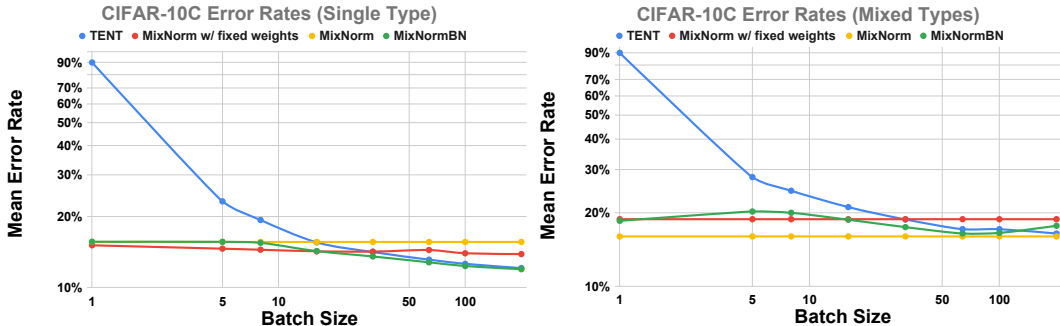


Figure 5: Comparison of our MixNorm and MixNormBN methods to the baseline method TENT, with a different backbone architecture (WideResNet40). Numerical results are reported in the Appendix A.

4.2 ABLATION STUDIES

In this section, we present the results of several ablations on architecture selection, ways of collecting normalization statistics and number of augmentations.

Architectures: Figure 5 provides the results on CIFAR-10C where both TENT and MixNorm adopt a more powerful backbone, WideResNet40, as opposed to the WideResNet28 used in Figure 4. Along with the ResNet-50 used in the ImageNet-C experiments, we demonstrate the effectiveness of MixNorm on architecture with different depth and width.

Normalization Statistics Estimation: In Table 1, we present an ablation study on different ways to collect normalization parameters. They include: 1) *Instance Norm*: each feature is normalized by itself along the spatial dimensions; 2) *Augmentation (Local) Norm*: features are normalized within the local batch of original samples and corresponding augmented views, i.e., only local statistics are used; 3) *Fixed Global Norm*: each feature is normalized by the fixed statistics stored in the pre-trained weights learned during training; 4) *Moving Global Norm*: each feature is normalized by the exponential moving statistics, i.e., only global statistics are used.

As shown in the table row 1 and 2, Augmentation Norm improves on Instance Norm by having augmented samples in addition to the original samples, but both of them perform poorly due to the absence of global statistics. Rows 3 and 4 show that Moving Global Norm works better than Fixed Global Norm, which is expected due to the existence of domain shift. Rows 5, 6, 7 indicate that combining the global and local statistics boosted performance by a significant scale. Row 8 indicates that learnable affine weights, α and β , do not seem to provide additional boost to the performance observed in TENT, possibly because MixNorm can already make accurate prediction so it does not need to be adjusted by the affine weights.

Augmentations: Table 2 shows that one augmentation is enough to get good performance, and the number of augmentations in local statistics calculation does not have significant influence on performance. Therefore, we just use one augmentation for faster computation.

4.3 TEST-TIME ADAPTATION IN ZERO-SHOT LEARNING SETTING

For Zero-Shot Classification, samples are evaluated separately at batch size of 1. Therefore, we use the regular MixNorm with fixed affine weights for comparison. We replace the Batch-Norm layers from the pre-trained CLIP models, with backbone architectures of ResNet-50, and then use the updated ResNet-50 model to calculate the visual features in Zero-Shot Classification. The results in Table 3 shows MixNorm improves the CLIP’s zero-shot learning performance on CIFAR-100, Stanford-Cars and CIFAR-10 and is on par with CLIP on Food-101 and STL10. This experiment aims to show MixNorm can bring reasonable improvements even over extremely large-scale pre-training models. Interestingly, the results show that our new method can bring larger improvements when the CLIP has a lower performance, which implicitly could be indication of the presence of larger domain shifts.

No	Experiment Name	Way of Collecting Means & Vars			Error Rate
		Moving Speed	Global Scale	Local Scale	
1	Instance Norm	0	0	1	78.21%
2	Augmentation (Local) Norm	0	0	1	68.73%
3	Fixed Global Norm	0	1	0	50.50%
4	Moving Global Norm	0.001	1	0	38.45%
5	Mixed Norm	0.001	0.95	0.95	34.45%
6	Mixed Norm	0.001	0.9	0.1	32.66%
7	Mixed Norm	0.001	0.8	0.2	31.96%
8	Mixed Norm w/ fixed params	0.001	0.8	0.2	32.01%

Table 1: Ablation studies on different ways to collect normalization statistics, tested with CIFAR-10C, with a mixed set of corruptions.

Number of Augmentations	1	3	5	7	9	15
ImageNet-C Error Rate	58.16%	58.33%	58.35%	58.44%	58.49%	58.56%

Table 2: Ablation studies on the effect of number of augmentations to MixNormBN method when collecting normalization statistics, tested with ImageNet-C and ResNet-50, with a mixed set of corruptions.

Model	CIFAR-100	Stanford-Cars	CIFAR-10	Food-101	STL10
CLIP (RN50)	40.97%	53.80%	72.26%	78.99%	94.84%
MixNorm	45.13%	54.01%	75.36%	78.78%	94.67%

Table 3: Zero-Shot Image Classification Accuracy. We replaced the Batch-Norm layers in the pre-trained CLIP ResNet-50 weights and tested it under the zero-shot setting on the standard benchmarks. Result shows MixNorm improved the CLIP performance

4.4 SOURCE-FREE UNSUPERVISED DOMAIN ADAPTATION

Model	SVHN \rightarrow MNIST	SVHN \rightarrow USPS	SVHN \rightarrow MNIST-M
Source	23.84%	37.02%	45.45%
BN (Li et al. (2016))	22.70%	36.87%	44.04%
TENT (Wang et al. (2020))	8.17%	34.63%	42.28%
MixNormBN	7.95%	25.71%	42.19%

Table 4: Error rate for Source-Free Unsupervised Domain Adaptation experiment on Digits Sets from SVHN to MNIST/MNIST-M/USPS.

For Source-Free UDA experiments, as we compare to TENT (Wang et al. (2020)) and BN (Li et al. (2016)), both of which use large batch sizes (200), we use MixNormBN for a fair comparison. As shown in Table 4, MixNormBN outperforms both TENT and BN by a significant margin on all three datasets, indicating the effectiveness of MixNormBN in the task of Source-Free UDA, for other types of domain shift that are not just corruptions.

5 CONCLUSION

We have developed a novel way to estimate the batch-norm statistics during test time, to fast adapt a source model to target test samples. Unlike the previous methods that require a large batch from a single distribution to calculate stable batch-norm statistics, our proposed method eliminates any dependency on online batches and is able to estimate accurate batch-norm statistics with single samples. The newly proposed method significantly outperformed the State-Of-The-Art models in the newly proposed real-world settings, and also demonstrated potential to improve zero-shot learning performance.

REFERENCES

- Michael J. Anderson, Benny Chen, Stephen Chen, Summer Deng, Jordan Fix, Michael K., and et al. First-generation inference accelerator deployment at facebook. *ArXiv*, abs/2107.04140, 2021.
- Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101—mining discriminative components with random forests. In *European conference on computer vision*, pp. 446–461. Springer, 2014.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PMLR, 2020.
- Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp. 215–223. JMLR Workshop and Conference Proceedings, 2011.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.
- Andrea Frome, Greg Corrado, Jonathon Shlens, Samy Bengio, Jeffrey Dean, Marc’ Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. 2013.
- Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pp. 1180–1189. PMLR, 2015.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261*, 2019.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pp. 448–456. PMLR, 2015.
- Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13)*, Sydney, Australia, 2013.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012.
- Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. Attribute-based classification for zero-shot visual object categorization. *IEEE transactions on pattern analysis and machine intelligence*, 36(3):453–465, 2013.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Yanghao Li, Naiyan Wang, Jianping Shi, Jiaying Liu, and Xiaodi Hou. Revisiting batch normalization for practical domain adaptation. *arXiv preprint arXiv:1603.04779*, 2016.

- Jian Liang, Dapeng Hu, and Jiashi Feng. Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In *International Conference on Machine Learning*, pp. 6028–6039. PMLR, 2020.
- Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011.
- Vishal M Patel, Raghuraman Gopalan, Ruonan Li, and Rama Chellappa. Visual domain adaptation: A survey of recent advances. *IEEE signal processing magazine*, 32(3):53–69, 2015.
- Xingchao Peng, Ben Usman, Neela Kaushik, Judy Hoffman, Dequan Wang, and Kate Saenko. Visda: The visual domain adaptation challenge. *arXiv preprint arXiv:1710.06924*, 2017.
- Joaquin Quiñero-Candela, Masashi Sugiyama, Neil D Lawrence, and Anton Schwaighofer. *Dataset shift in machine learning*. Mit Press, 2009.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020*, 2021.
- Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *European conference on computer vision*, pp. 213–226. Springer, 2010.
- Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei Efros, and Moritz Hardt. Test-time training with self-supervision for generalization under distribution shifts. In *International Conference on Machine Learning*, pp. 9229–9248. PMLR, 2020.
- Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5018–5027, 2017.
- Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. *arXiv preprint arXiv:2006.10726*, 2020.
- Garrett Wilson and Diane J Cook. A survey of unsupervised deep domain adaptation. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 11(5):1–46, 2020.
- Yongqin Xian, Christoph H Lampert, Bernt Schiele, and Zeynep Akata. Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly. *IEEE transactions on pattern analysis and machine intelligence*, 41(9):2251–2265, 2018.
- Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.

A FULL RESULTS ON CIFAR-10C, IMAGENET-C

In Table 9,10, 5,6 we present the exact error rates shown in Figure 4. In Table 7,8 we present the error rates shown in Figure 5.

Batch Size	TENT	MixNorm w/ Fixed Parameters	MixNorm	MixNormBN
1	90.00%	21.76%	21.94%	22.33%
5	29.21%	21.57%	21.94%	22.47%
8	25.81%	21.53%	21.94%	22.16%
16	22.07%	21.48%	21.94%	20.94%
32	20.65%	21.53%	21.94%	20.19%
64	19.69%	21.43%	21.94%	19.43%
100	19.11%	21.26%	21.94%	18.94%
200	18.58%	21.20%	21.94%	18.52%

Table 5: **Single Distribution** Error Rates of TENT, MixNorm (Algorithm 1) with and without fixed affine parameters, MixNormBN (Algorithm 2) on **CIFAR-10C** datasets. All scores are average error rates from all 15 corruption datasets at severity level 5. All methods adopt **Wide-ResNet-28** as backbone architectures.

Batch Size	TENT	MixNorm w/ Fixed Parameters	MixNorm	MixNormBN
1	89.69%	31.94%	32.01%	32.01%
5	39.39%	31.96%	32.01%	33.24%
8	37.15%	31.98%	32.01%	32.67%
16	35.22%	32.06%	32.01%	32.18%
32	33.18%	32.17%	32.01%	31.73%
64	32.74%	32.09%	32.01%	30.87%
100	32.52%	32.51%	32.01%	30.24%
200	33.12%	31.96%	32.01%	30.33%

Table 6: **Mixed Distribution** Error Rates of TENT, MixNorm (Algorithm 1) with and without fixed affine parameters, MixNormBN (Algorithm 2) on **CIFAR-10C** datasets. All methods are tested with 10,000 randomly ordered samples composed of corrupted images from 15 different corruptions types at severity level 5. All methods adopt **Wide-ResNet-28** as backbone architectures.

Batch Size	TENT	MixNorm w/ Fixed Parameters	MixNorm	MixNormBN
1	90.00%	15.09%	15.59%	15.62%
5	23.18%	14.60%	15.59%	15.61%
8	19.32%	14.45%	15.59%	15.46%
16	15.53%	14.23%	15.59%	14.26%
32	14.13%	14.19%	15.59%	13.53%
64	13.12%	14.41%	15.59%	12.78%
100	12.59%	13.96%	15.59%	12.32%
200	12.08%	13.85%	15.59%	11.95%

Table 7: **Single Distribution** Error Rates of TENT, MixNorm (Algorithm 1) with and without fixed affine parameters, MixNormBN (Algorithm 2) on **CIFAR-10C** datasets. All scores are average error rates from all 15 corruption datasets at severity level 5. All methods adopt **Wide-ResNet-40** as backbone architectures.

B HYPER PARAMETER SELECTION

In this section we report the hyper parameters we used for each experiments. As we mentioned in Section 4, we use the “Gaussian Noise at Severity 5” set to choose the hyper parameters on CIFAR-10C and ImageNet-C experiments. For other dataset, because of the absence of proper validation set, we just pick the hyper-parameters that give the optimal performance. All learning rate stay the

Batch Size	TENT	MixNorm w/ Fixed Parameters	MixNorm	MixNormBN
1	89.81%	18.80%	15.99%	18.52%
5	27.92%	18.80%	15.99%	20.21%
8	24.57%	18.80%	15.99%	19.97%
16	21.07%	18.80%	15.99%	18.71%
32	18.76%	18.80%	15.99%	17.45%
64	17.12%	18.80%	15.99%	16.44%
100	17.12%	18.80%	15.99%	16.53%
200	16.43%	18.80%	15.99%	17.67%

Table 8: **Mixed Distribution** Error Rates of TENT, MixNorm (Algorithm 1) with and without fixed affine parameters, MixNormBN (Algorithm 2) on **CIFAR-10C** datasets. All methods are tested with 10,000 randomly ordered samples composed of corrupted images from 15 different corruptions types at severity level 5. All methods adopt **Wide-ResNet-40** as backbone architectures.

Batch Size	TENT	MixNorm w/ Fixed Parameters	MixNorm	MixNormBN
1	99.87%	74.44%	74.44%	74.03%
5	81.08%	74.43%	74.44%	74.89%
8	75.74%	74.44%	74.44%	72.93%
16	69.85%	74.41%	74.44%	69.02%
32	65.51%	74.41%	74.44%	65.34%
64	62.74%	74.43%	74.44%	62.69%

Table 9: **Single Distribution** Error Rates of TENT, MixNorm (Algorithm 1) with and without fixed affine parameters, MixNormBN (Algorithm 2) on **ImageNet-C** datasets. All scores are average error rates from all 15 corruption datasets at severity level 5. All methods adopt **ResNet-50** as backbone architectures.

same as the learning rate used by other baselines. Finding optimal hyper-parameters remains to be a challenging problem for methods in Test-Time Adaptation.

C ZERO-SHOT CLASSIFICATION DATASETS

In Table 13 we report the details of datasets we used for Zero-Shot Classification Experiments.

Batch Size	TENT	MixNorm w/ Fixed Parameters	MixNorm	MixNormBN
1	99.86%	78.88%	78.89%	88.73%
5	90.59%	78.87%	78.89%	85.06%
8	88.72%	78.87%	78.89%	84.75%
16	87.27%	78.85%	78.89%	83.99%
32	86.39%	78.87%	78.89%	83.55%
64	86.09%	78.86%	78.89%	82.78%

Table 10: **Mixed Distribution** Error Rates of TENT, MixNorm (Algorithm 1) with and without fixed affine parameters, MixNormBN (Algorithm 2) on **ImageNet-C** datasets. All methods are tested with 50,000 randomly ordered samples composed of corrupted images from 15 different corruptions types at severity level 5. All methods adopt **ResNet-50** as backbone architectures.

	Scale m	Moving Speed τ	Learning Rate
CIFAR10-C Single	0.05	1.00E-03	1.00E-03
CIFAR10-C Mixed	0.2	1.00E-03	1.00E-03
ImageNet-C Single	0.01	1.00E-03	2.50E-04
ImageNet-C Mixed	0.05	1.00E-06	2.50E-04
CIFAR-10	0.01	1.00E-06	-
CIFAR-100	0.01	1.00E-06	-
STL10	0.01	1.00E-06	-
Stanford-Cars	0.01	1.00E-07	-
Food-101	0.005	1.00E-07	-

Table 11: Hyper-Parameters of MixNorm Experiments. In Figure 4 and 5 we reported scores for both MixNorm with and without fixed affine parameters. The MixNorm uses the same hyper-parameter in the both cases.

	Scale m	Max Moving Speed τ_{max}	Learning Rate
CIFAR10-C Single	0.01	0.75	1.00E-03
CIFAR10-C Mixed	0.2	0.75	1.00E-03
ImageNet-C Single	0.01	1.1	2.50E-04
ImageNet-C Mixed	0.05	1.1	2.50E-04
MNIST	0.1	0.75	0.75
MNIST-M	0.01	0.75	0.05
USPS	0.01	0.5	0.75

Table 12: Hyper-Parameters of MixNormBN Experiments.

Datasets	Classes	Train Size	Test Size
CIFAR-10	10	50000	10000
CIFAR-100	100	50000	10000
Stanford-Cars	196	8144	8041
Food-101	101	75750	25250
STL10	10	1000	8000

Table 13: Dataset Size and Class Number