

# Improving Test-Time Adaptation via Shift-agnostic Weight Regularization and Nearest Source Prototypes

Sungha Choi\*  Seunghan Yang  Seokeon Choi  Sungrack Yun 

Qualcomm AI Research<sup>†</sup>

{sunghac, seunghan, seokchoi, sungrack}@qti.qualcomm.com

**Abstract.** This paper proposes a novel test-time adaptation strategy that adjusts the model pre-trained on the source domain using only unlabeled online data from the target domain to alleviate the performance degradation due to the distribution shift between the source and target domains. Adapting the entire model parameters using the unlabeled online data may be detrimental due to the erroneous signals from an unsupervised objective. To mitigate this problem, we propose a shift-agnostic weight regularization that encourages largely updating the model parameters sensitive to distribution shift while slightly updating those insensitive to the shift, during test-time adaptation. This regularization enables the model to quickly adapt to the target domain without performance degradation by utilizing the benefit of a high learning rate. In addition, we present an auxiliary task based on nearest source prototypes to align the source and target features, which helps reduce the distribution shift and leads to further performance improvement. We show that our method exhibits state-of-the-art performance on various standard benchmarks and even outperforms its supervised counterpart.

**Keywords:** Test-time adaptation, Domain generalization, Source-free domain adaptation, On-Device AI

## 1 Introduction

After deep neural networks (DNNs) trained on a given dataset (*i.e.*, source domain) are deployed to a new environment (*i.e.*, target domain), the DNNs make predictions from the data in the target domain. However, in most cases, the distribution of the source and target domains varies significantly, which degrades the model’s performance in the target domain. If the deployed model does not remain stationary during test time but adapts to the new environment using clues about unlabeled target data, its performance can be improved [52,59,63,42,25,38,39,66].

Recently, several studies [59,63,42,25] have proposed test-time adaptation to update the model during test time after model deployment. However, it is extremely challenging to adapt the model to the target domain with only

---

\* Corresponding author.

<sup>†</sup> Qualcomm AI Research is an initiative of Qualcomm Technologies, Inc.

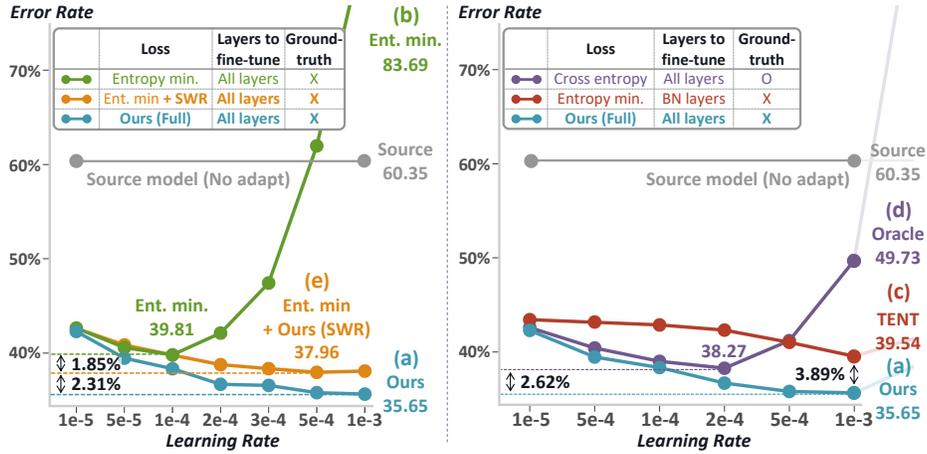


Fig. 1: Comparison of average error (%) between our approach and other methods with varying learning rates on CIFAR-100-C [20]. The x- and y-axes are the learning rate and average error rate, respectively. (a) Our method significantly outperforms the other three methods: (b) updating the entire parameters with only entropy minimization, (c) the state-of-the-art method, TENT [59], and (d) a supervised method. (e) Our proposed SWR keeps the performance stable with the combining of entropy minimization even at higher learning rates: [1e-3, 1e-4].

unlabeled online data. As shown in Fig. 1(b), the adaptation of the entire model parameters may be detrimental due to the erroneous signals from the unsupervised objective such as entropy minimization [15,26,40,57,59], and the performance may be highly dependent on the learning rate. In addition, since the test-time adaptation can access unlabeled target data only once, and the adaptation proceeds simultaneously with the evaluation, updating all network parameters may result in overfitting [62,17]. Thus, several approaches present the methods to update only some part of the network architecture [59,63,42,25] such as batch normalization [24] or classifier layers. Especially, T3A [25] proposes an optimization-free method to adapt only the classifier layers using unlabeled target data, and TENT [59] updates batch statistics and affine parameters in the batch normalization layers by entropy minimization on unlabeled target data. However, updating only partial parameters or layers of the model may only result in marginal performance improvement, as shown in Fig. 1(c). Furthermore, such methods cannot be applied to the model architecture without a specific layer such as batch normalization or classifier layers.

Other approaches [52,39] propose to jointly optimize a main task<sup>1</sup> and a self-supervised task, such as rotation prediction [13] or instance discrimination [18,6], during pre-training in the source domain, and then update the model using only the self-supervised task during test time. In contrast to the unsupervised objective for the main task that highly depends on the model’s prediction accuracy, the self-supervised task always obtains a proper supervisory signal. However, the self-

<sup>1</sup> refers to the ultimate objective of the model (*e.g.*, classification).

supervised task may interfere with the main task if both tasks are not properly aligned [39,50,64]. In addition, these approaches cannot be applied to adapt arbitrary pre-trained models to the target domain since they require specific pre-training methods in the source domain.

To resolve these issues, we present two novel approaches for the test-time adaptation. First, we consider a shift-agnostic weight regularization (SWR) that enables the model to quickly adapt to the target domain, which is beneficial when updating the entire model parameters with a high learning rate. In contrast to Fig. 1(b), the entropy minimization with the proposed SWR shows superior performance and less dependency on the learning rate choice, as shown in Fig. 1(e). In terms of distribution shift, the SWR identifies the entire model parameters into shift-agnostic and shift-biased parameters, updating the former less and the latter more. Second, we present an auxiliary task based on a non-parametric nearest source prototype (NSP) classifier, which pulls the target representation closer to its nearest source prototype. With the NSP classifier, both source and target representations can be well aligned, which significantly improves the performance of the main task. Our proposed method (Fig. 1(a)) outperforms the state-of-the-art method [59] (Fig. 1(c)) and even the supervised method using ground-truth labels (Fig. 1(d)).

Our method requires access to the source data to identify shift-agnostic and biased parameters and generate source prototypes before the model deployment, but it is applicable to any model regardless of its architecture or pre-training procedure. If a given model is pre-trained on open datasets, or if the source data owner deploys the model, source data is accessible before model deployment. In this case, our method significantly enhances the test-time adaptation capability by leveraging the source data without modifying the pre-trained model. Unlike TTT [52] and TTT++ [39], we do not change the pre-training method of a given model, so our method can take benefit from any pre-trained strong models, such as AugMix [21] (Table 1) or CORAL [51] (Table 6), as a good starting point for test-time adaptation. In these respects, we believe our method is practical.

The major contributions of this paper can be summarized as follows

- Two novel approaches for test-time adaptation are presented in this paper. The proposed SWR enables the model to quickly and reliably adapt to the target domain, and the NSP classifier aligns the source and target features to reduce the distribution shift, leading to further performance improvement.
- Our test-time adaptation method is model-agnostic and not dependent on the pre-training method in the source domain, and thus it can be applied to any pre-trained model. Therefore, our method can also complement other domain generalization approaches that mainly focus on the pre-training method in the source domain before model deployment.
- We show that our method achieves state-of-the-art performance through extensive experiments on CIFAR-10-C, CIFAR-100-C, ImageNet-C [20] and domain generalization benchmarks including PACS [31], OfficeHome [55], VLCS [10], and TerraIncognita [5]. Especially, our method even outperforms its supervised counterpart on CIFAR-100-C dataset.

## 2 Related Work

### 2.1 Source-Free Domain Adaptation

Unsupervised domain adaptation (UDA) [51,57,53,11,46,22,12] assumes simultaneous access to both the source and target domains. Data is often distributed across multiple devices. In such cases, UDA requires data sharing for simultaneous access to all data. However, it is often impossible due to data privacy concerns, limited bandwidth, and computational cost. Source-free domain adaptation [38,36,30,58,60,61,1] overcomes this challenge by adapting a source pre-trained model to the target domain using only unlabeled target data. These approaches focus on offline adaptation in which the same target sample is fed to the model multiple times during target adaptation, whereas our method concentrates on online adaptation.

### 2.2 Test-Time Adaptation and Training

Test-time adaptation focuses on online adaptation in which all target data can be accessed only once at test time, and adaptation is performed simultaneously with evaluation. More specifically, it forward propagates target samples through the model for evaluation and then backpropagates the error signal from the model’s output in an unsupervised manner for training [59]. Several studies adopt self-supervised learning, such as rotation prediction [52] or instance discrimination [39], to jointly optimize the main and self-supervised tasks on the source domain and then optimize only the self-supervised task on the target domain. However, these methods are not universally applicable to arbitrary pre-trained models as they require specific pre-training methods in the source domain. Recently, model-agnostic test-time adaptation methods independent of the pre-training method in the source domain have been proposed [59,25,42,63]. TENT [59] uses the batch statistics of the target domain and optimizes channel-wise affine parameters using entropy minimization loss. T3A [25] proposes an optimization-free method that adjusts a pre-trained linear classifier by updating the prototype for each class during test time. However, since these methods update only partial parameters or layers of the model, such as the batch normalization [59,42,63] or classifier layer [25], they may be suboptimal for target adaptation.

### 2.3 Domain Generalization

Since UDA aims to adapt the model to the predefined target domain before model deployment, it is not suitable to guarantee generalization performance to other arbitrary target domains. On the other hand, domain generalization (DG) differs from UDA in that it assumes that the model accesses only the source domain during training time before model deployment and aims to improve the generalization capability in arbitrary unseen target domains. Numerous DG approaches using meta-learning [32,3,33], normalization [44,47,8,7], adversarial training [34,37,45], and data augmentation [67,56,14,35] have been proposed to learn domain-agnostic feature representations for the target domain. However, these studies only focus on methods at training time before model deployment, whereas our method focuses on a test-time adaptation after model deployment.

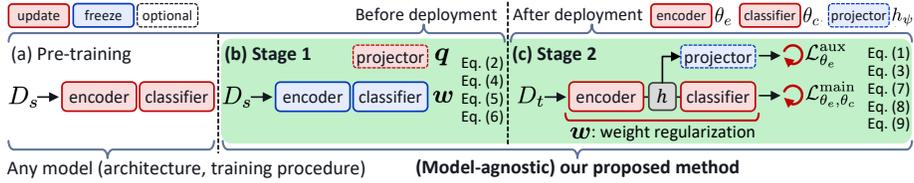


Fig. 2: Our method consists of two stages: (b) and (c). (a) our method takes the pre-trained model in an off-the-shelf manner and (b) generates penalty vector  $\mathbf{w}$  and source prototypes  $\mathbf{q}$  while keeping the model *frozen* before model deployment. After model deployment, (c) our method does not access labeled source data  $D_s$  other than unlabeled online target data  $D_t$  during test-time adaptation.

### 3 Proposed Method

Assume that the model parameters  $\theta$  trained on the source domain consist of an encoder part  $\theta_e$  and a classifier part  $\theta_c$ , as shown in Fig. 2(c). After being deployed to the target domain, the model infers the class probability distribution of the target sample and then optimizes our proposed test-time adaptation loss  $\mathcal{L}_{\theta_e, \theta_c}^{\text{target}}$ . The overall loss of our proposed method is defined as

$$\mathcal{L}_{\theta_e, \theta_c}^{\text{target}} = \mathcal{L}_{\theta_e, \theta_c}^{\text{main}} + \mathcal{L}_{\theta_e}^{\text{aux}} + \lambda_r \sum_l w_l \|\theta_l - \theta_l^*\|_2^2, \quad (1)$$

where  $w_l$  denotes the  $l$ -th element of the penalty vector  $\mathbf{w}$  used to control the update of the model parameters,  $\theta_l$  is the parameter vector of the  $l$ -th layer<sup>2</sup> of the model,  $\theta^*$  is the parameters from the previous update step,  $\lambda_r$  is the importance of the regularization term, and  $\mathcal{L}_{\theta_e, \theta_c}^{\text{main}}$  and  $\mathcal{L}_{\theta_e}^{\text{aux}}$  denote the main and auxiliary task losses, respectively. Optimizing the main task loss updates the entire model parameters  $\theta_e$  and  $\theta_c$ , whereas optimizing the auxiliary task loss updates only the encoder part  $\theta_e$ . We first present a shift-agnostic weight regularization (SWR) and then describe an entropy objective of the main task. Finally, we propose an auxiliary task based on a nearest source prototype (NSP) classifier, which directly benefits the main task.

#### 3.1 Shift-agnostic Weight Regularization

The main idea of the SWR is to impose different penalties for each parameter update during test-time adaptation, depending on the sensitivity of each model parameter to the distribution shift. Assuming that the distribution shift is mainly caused by color and blur shifts, we mimic the distribution shift using transformation techniques such as color distortion and Gaussian blur. Experiments on variations of the SWR, including the use of other transform functions, can be found in the supplementary Section B.

To obtain the penalty vector  $\mathbf{w}$  specified in Eq. (1), we first forward-propagate two input images (i.e., an original and its transformed image) through the pre-trained source model and then back-propagate the task loss (i.e., cross entropy)

<sup>2</sup> denotes a part divided into torch.nn.Module units defined in Pytorch. The gradient vector of each layer can be easily obtained using torch.nn.module.parameters().

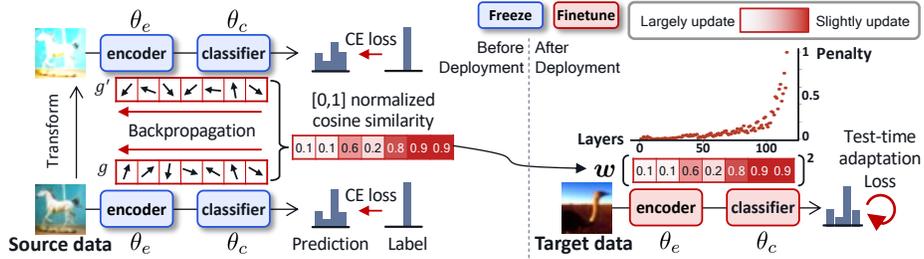


Fig. 3: Overall process of our proposed SWR. We first obtain the penalty vector  $\mathbf{w}$  before model deployment and then use it as layer-wise penalties to control the update of the model parameters at test-time adaptation after model deployment.

using the source labels to produce two sets  $\mathbf{g}$  and  $\mathbf{g}'$  of  $L$  gradient vectors, respectively. Note that  $L$  is the total number of layers in the model. Then the  $l$ -th element  $w_l$  of the penalty vector  $\mathbf{w}$  is calculated by employing the average cosine similarity  $s_l$  between two gradient vectors,  $\mathbf{g}_l$  and  $\mathbf{g}'_l$  from  $N$  source samples as

$$s_l = \frac{1}{N} \sum_{i=1}^N \frac{\mathbf{g}_l^i \cdot \mathbf{g}'_l{}^i}{\|\mathbf{g}_l^i\| \|\mathbf{g}'_l{}^i\|} \in \mathbb{R}, \quad (2)$$

$$\mathbf{w} = (\nu[s_1, \dots, s_l, \dots, s_L])^2 \in \mathbb{R}^L,$$

where  $\nu[\cdot]$  denotes min-max normalization with the range of  $[0,1]$ ,  $\mathbf{g}_l^i$  and  $\mathbf{g}'_l{}^i$  denote the  $l$ -th gradient vectors for  $i$ -th source sample and its transformed sample, respectively.  $N$  denotes the total number of samples. Note that the penalty vector  $\mathbf{w}$  is obtained from a frozen pre-trained source model before model deployment. Therefore, this process is independent of the source model's pre-training method and does not require source data after model deployment, as shown in Fig. 2.

As shown in Eq. (1) and Fig. 3, during test-time adaptation, we apply the layer-wise penalty value  $w_l$  to the difference between previous and current model parameters for each layer, and this controls the update of model parameters differently for each layer. Therefore, the model parameters belonging to the layers with high cosine similarity between the two gradient vectors are considered shift-agnostic, and we less update them by imposing high penalties. Section 4.6 experimentally shows that SWR takes advantage of using high learning rates to adapt the model to the target domain quickly.

### 3.2 Entropy Objective for the Main Task

The main task of the model  $f_\theta$  is defined as the task performed by the parameters  $\theta_e$  and  $\theta_c$ . The loss function for the main task during test time is built using the entropy of model predictions  $\tilde{y}$  on test samples from the target distribution. We adopt information maximization loss [27,48,23], validated in several test-time adaptation and source-free domain adaptation methods [38,58,42], as an unsupervised learning objective for the main task. This loss consists of entropy minimization [59,38,57,49] and mean entropy maximization [28,38,2,58] as

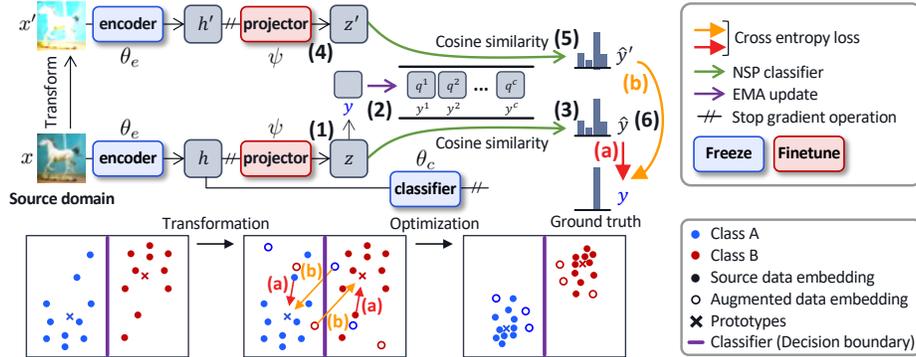


Fig. 4: Source prototype generation phase before model deployment. First, we repeat steps (1) and (2) until prototypes of all classes are generated, then train the projector and update the source prototype at the same time through an iterative process from (1) to (6) on the source data. (a) and (b) pull the original source projection and its transformed source projection, respectively, such that they become closer to the nearest source prototype from the original one.

$$\mathcal{L}_{\theta_e, \theta_c}^{\text{main}} = \lambda_{m_1} \frac{1}{N} \sum_{i=1}^N H(\tilde{y}_i) - \lambda_{m_2} H(\bar{y}), \quad (3)$$

where  $H(p) = -\sum_{k=1}^C p^k \log p^k$ ,  $\bar{y} = \frac{1}{N} \sum_i \tilde{y}_i$ ,  $\lambda_{m_1}$  and  $\lambda_{m_2}$  indicate the importance of each term. The number of classes and the batch size are denoted by  $C$  and  $N$ . Intuitively, entropy minimization makes individual predictions confident, and mean entropy maximization encourages average prediction within a batch to be close to the uniform distribution.

### 3.3 Auxiliary Task based on the Nearest Source Prototype

Due to the distribution shift between the source and target domains, the target features deviate from the source features at test time. To resolve this issue, we propose an auxiliary task based on the nearest source prototype (NSP) classifier, which pulls the target embeddings closer to their nearest source prototypes in the embedding space. Eventually, optimizing the auxiliary task improves performance significantly since it directly supports the main task by aligning the source and target representations. We first explain how to generate source prototypes and define the NSP classifier based on them.

**Source prototype generation** The source prototypes are defined as the averages over source embeddings for each class. As shown in Fig. 4, we freeze the model  $f_\theta$  trained on the source data and attach an additional projection layer  $h_\psi$  behind the encoder  $f_{\theta_e}$ . The encoder  $f_{\theta_e}$  infers the representation  $\mathbf{h}$  from the source sample  $x$ , and the projector  $h_\psi$  maps  $\mathbf{h}$  to the projection  $\mathbf{z}$  in another embedding space where the loss  $\mathcal{L}_\psi^{\text{emb}}$  is applied as  $\mathbf{z} = h_\psi(f_{\theta_e}(x))$ . The source prototype  $\mathbf{q}_t^k$  for class  $k$  is updated through exponential moving average (EMA) with the projection  $\mathbf{z}_t^k$  of the source sample  $(x, y^k)_{k \in [1, C]}$  at time  $t$  during the optimization trajectory as

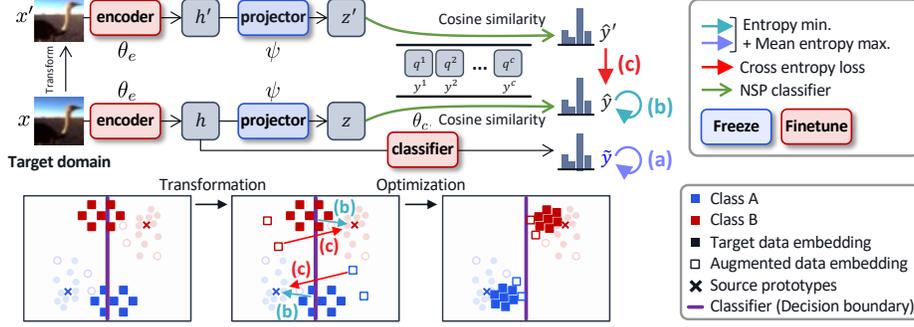


Fig. 5: Test-time adaptation phase after model deployment. (a) main task loss. (b),(c) auxiliary task loss. (b) and (c) pull the original target projection and its transformed target projection, respectively, such that they become closer to the nearest source prototype from the original one.

$$\mathbf{q}_t^k = \alpha \cdot \mathbf{q}_{t-1}^k + (1 - \alpha) \cdot \mathbf{z}_t^k, \quad (4)$$

where  $\alpha=0.99$  and  $\mathbf{q}_0^k = \mathbf{z}_0^k$ .

We define the NSP classifier as a non-parametric classifier. It measures the cosine similarity of a given target embedding to the source prototypes for all classes and then generates a class probability distribution  $\hat{y}$  as

$$\hat{y} = \sum_{k=1}^C \left( \frac{\exp(S(\mathbf{z}, \mathbf{q}^k)/\tau)}{\sum_{j=1}^C \exp(S(\mathbf{z}, \mathbf{q}^j)/\tau)} \right) y^k, \quad (5)$$

where  $S(\cdot, \cdot)$  is a cosine similarity function,  $S(\mathbf{a}, \mathbf{b}) = (\mathbf{a} \cdot \mathbf{b}) / \|\mathbf{a}\| \|\mathbf{b}\|$ ,  $\tau$  denotes a temperature that controls the sharpness of the distribution, and  $y^k$  is the one-hot ground-truth label vector of  $k$ -th class. In addition, inspired by recent self-supervised contrastive learning methods [6,18,4], we enable the projector  $h_\psi$  to learn transformation-invariant mapping. We obtain projection  $\mathbf{z}'$  of the transformed source sample by  $\mathbf{z}' = h_\psi(f_{\theta_e}(\mathcal{T}(x)))$ , where  $\mathcal{T}(\cdot)$  denotes an image transform function. The embedding loss  $\mathcal{L}_\psi^{\text{emb}}$  consisting of two cross entropy loss terms is applied to the embedding space to train the projector  $h_\psi$  as

$$\mathcal{L}_\psi^{\text{emb}} = \frac{1}{N} \sum_{i=1}^N (\text{CE}(y_i, \hat{y}_i) + \text{CE}(y_i, \hat{y}'_i)), \quad (6)$$

where  $\text{CE}(p, q) = -\sum_{k=1}^C p^k \log q^k$ , and  $y_i$  is the ground-truth label of  $i$ -th source sample. Here,  $\hat{y}$  and  $\hat{y}'$  denote the outputs of the NSP classifier for the projections  $\mathbf{z}$  and  $\mathbf{z}'$  of the source sample and its transformed one, respectively. As shown in Fig. 4, optimizing the embedding loss encourages the projector  $h_\psi$  to learn a mapping that pulls the projections belonging to the same class closer together and pushes source prototypes farther away from each other.

Note that this process is applied to a frozen pre-trained source model and completed before model deployment. Therefore, it is model-agnostic and does not require source data during test time.

**Auxiliary task loss at test time** Once the source prototypes are generated and the projection layer is trained, we can deploy the model and then jointly optimize both main and auxiliary tasks on unlabeled online data. The auxiliary task loss  $\mathcal{L}_{\theta_e}^{\text{aux}}$  consists of two objective functions: the entropy objective  $\mathcal{L}_{\theta_e}^{\text{aux.ent}}$  using the entropy of the NSP classifier’s prediction  $\hat{y}$ , and the self-supervised loss  $\mathcal{L}_{\theta_e}^{\text{aux.sel}}$  that encourages the model’s encoder  $f_{\theta_e}$  to learn transformation-invariant mappings as

$$\mathcal{L}_{\theta_e}^{\text{aux}} = \mathcal{L}_{\theta_e}^{\text{aux.ent}} + \lambda_s \mathcal{L}_{\theta_e}^{\text{aux.sel}}, \quad (7)$$

where  $\lambda_s$  denotes the importance of the self-supervised loss term. Similarly to Eq. (3), the entropy objective is built by using the entropy of the prediction  $\hat{y}$  of the NSP classifier on the target sample as

$$\mathcal{L}_{\theta_e}^{\text{aux.ent}} = \lambda_{a_1} \frac{1}{N} \sum_{i=1}^N H(\hat{y}_i) - \lambda_{a_2} H(\bar{y}), \quad (8)$$

where  $N$  is batch size,  $\lambda_{a_1}$  and  $\lambda_{a_2}$  indicate the importance of each term,  $H(p) = -\sum_{k=1}^C p^k \log p^k$ , and  $\bar{y} = \frac{1}{N} \sum_{i=1}^N \hat{y}_i$ . The self-supervised loss is applied to the prediction  $\hat{y}'$  of the NSP classifier on the transformed target sample as

$$\mathcal{L}_{\theta_e}^{\text{aux.sel}} = -\frac{1}{N} \sum_{i=1}^N \sum_{k=1}^C \hat{y}_i^k \log \hat{y}_i^k. \quad (9)$$

As shown in Fig. 5, the entropy objective function (Fig. 5(b)) pulls the projection  $\mathbf{z}$  of the target sample to move closer to its nearest source prototype, and the self-supervised objective (Fig. 5(c)) encourages the projection  $\mathbf{z}'$  of the transformed target sample to get closer to the same target as  $\mathbf{z}$ .

## 4 Experiments

This section describes the experimental setup, implementation details, and the experimental results of the comparisons with other state-of-the-art methods in test-time adaptation. We also show that generalization performance can be further improved by combining our proposed method with an existing domain generalization strategy that mainly focuses on training time in the source domain.

### 4.1 Experimental Setup

Following TENT [59] and T3A [25], all experiments in this paper are conducted on the online adaptation setting, where adaptation is performed concurrently with evaluation at test time without seeing the same data twice or more. After a prediction is obtained, the model is updated via back-propagation. We evaluate our proposed method on **CIFAR-10-C**, **CIFAR-100-C**, **ImageNet-C**<sup>3</sup> [20] and four domain generalization benchmarks such as **PACS** [31], **OfficeHome** [55], **VLCS** [10], and **TerraIncognita** [5]. Since our method can be used independently of the backbone networks and its pre-training method, we apply our method to publicly available pre-trained models for evaluation. We perform experiments

<sup>3</sup> Experiments on ImageNet-C are in the supplementary Section B.

on CIFAR datasets using WideResNet-28-10 [65] and WideResNet-40-2 [65] as backbone networks, based on RobustBench [9]. In the domain generalization setup, we use ResNet-50 [19] without the batch normalization layer, which is the default setting of DomainBed [16], DG benchmark framework. CIFAR-10/100 dataset [29] contains 50k images for training and 10k images for testing. Corruptions such as noise, blur, weather, and digital are applied to 10k images from CIFAR-10/100 test set to create CIFAR-10/100-C test images. For test-time adaptation, 50k images for CIFAR training set are defined as the source domain, and 10k images for CIFAR-C test set are defined as the target domain.

## 4.2 Implementation details

We integrate our proposed method within the frameworks officially provided by other state-of-the-art methods [59,39,25] for fair comparisons. Specifically, different frameworks are used for each experiment as follows: **TENT framework** [59] for all experiments with WRN-40-2 and WRN-28-10 backbone networks on CIFAR-10/100-C, **TTT++ framework** [39] for all experiments with ResNet-50 on CIFAR-10/100-C, and **T3A framework** [25] for all domain generalization benchmarks. For experiments on CIFAR, we follow the default values provided by each framework for experimental settings such as batch size and optimizer.

Color distortion, random grayscale and Gaussian blurring are used as the image transformations specified in Fig. 3 and Fig. 5, and random cropping and random horizontal flipping are additionally applied for the image transformations in Fig. 4. We use batch statistics on test data instead of using running estimates. The hyper-parameters are empirically set as  $\lambda_{m_1}=0.2$ ,  $\lambda_{a_1}=0.8$ ,  $\lambda_{m_2}=0.25$ ,  $\lambda_{a_2}=0.25$ ,  $\lambda_s=0.1$ ,  $\lambda_r=250$ , and softmax temperature  $\tau=0.1$ . The epoch for training the projector is 20, and  $N=1024$  in Eq. (2). Since these hyper-parameters are not sensitive to the backbone and datasets, they are fixed without individual tuning in most experiments in this paper unless noted otherwise. The projector as described in Section 3.3 can be configured as a single- or multi-layer perceptron (MLP). The MLP consists of a linear layer followed by batch normalization [24], ReLU [43], and a final linear layer with output dimension 512. The performance change according to the projector configuration is shown in Table 3, and the detailed architecture is described in the supplementary Section C.

## 4.3 Robustness against Image Corruptions

Table 1(a) shows a comparison of the robustness between our method and recent test-time adaptation methods for the most severe corruptions on CIFAR-100-C. TFA [39] and TTT++ [39] were originally implemented as offline adaptation methods that train a model by observing the same data multiple times across numerous epochs, so we change these methods to the online adaptation setting to reproduce the results. Our proposed method significantly outperforms other state-of-the-art methods with large margins of 3.89% for ResNet-50 and 2.59% for WRN-40-2. Table 1(b) shows the results on the most severe corruptions of CIFAR-10-C. Our method consistently outperforms other methods on CIFAR-10/100-C datasets across various backbone networks. In particular, WRN-40-2, which is trained with AugMix [21] for a data processing to increase the robustness

Table 1: Comparison with other methods. \* denotes the reported results from the original paper, and the others are reproduced values in our environment based on the official framework provided by TENT [59] and TTT++ [39]. Source denotes the source pre-trained model without test-time adaptation.

(a) Comparison of error rate (%) on CIFAR-100-C with severity level 5

Backbone	Methods	Avg. err	Gaus.	Shot	Impu.	Defo.	Glas.	Moti.	Zoom	Snow	Fros.	Fog	Brig.	Cont.	Elas.	Pixe.	Jpeg
WRN-40-2 (AugMix) [65,21]	Source	46.75	65.7	60.1	59.1	32.0	51.0	33.6	32.3	41.4	45.2	51.4	31.6	55.5	40.3	59.7	42.4
	TENT [59]	35.53	40.1	39.5	42.0	29.6	41.9	30.7	29.7	34.5	34.8	39.1	27.5	32.9	37.6	32.8	40.3
	Core* [63]	35.30	39.8	39.3	41.5	29.5	41.7	30.6	29.8	34.2	34.9	38.6	27.5	32.6	37.1	32.7	40.1
	<b>Ours</b>	<b>32.71</b>	<b>37.6</b>	<b>36.6</b>	<b>35.1</b>	<b>28.0</b>	<b>39.5</b>	<b>28.7</b>	<b>28.5</b>	<b>31.3</b>	<b>32.6</b>	<b>34.4</b>	<b>26.3</b>	<b>29.0</b>	<b>35.5</b>	<b>30.1</b>	<b>37.7</b>
ResNet-50 [19]	Source	60.35	80.8	77.8	87.8	39.6	82.3	54.2	38.4	54.6	60.2	68.1	28.9	50.9	59.5	72.3	50.0
	SHOT [38]	43.53	49.0	47.1	61.2	33.8	58.3	41.0	31.3	45.0	42.0	52.0	29.7	33.4	47.9	41.8	39.4
	TFA [39]	44.13	49.0	47.0	61.4	34.2	58.9	41.5	32.1	46.8	43.0	54.6	31.2	33.7	48.9	39.8	39.7
	TTT++ [39]	44.38	50.2	47.7	66.1	35.8	61.0	38.7	35.0	44.6	43.8	48.6	28.8	30.8	49.9	39.2	45.5
	TENT [59]	39.54	43.8	42.0	54.1	31.2	51.7	37.0	29.9	42.3	39.6	45.6	30.1	30.9	44.5	34.2	36.4
<b>Ours</b>	<b>35.65</b>	<b>40.0</b>	<b>38.4</b>	<b>46.3</b>	<b>29.3</b>	<b>46.0</b>	<b>32.5</b>	<b>27.9</b>	<b>37.3</b>	<b>36.6</b>	<b>37.3</b>	<b>27.5</b>	<b>28.8</b>	<b>41.0</b>	<b>31.4</b>	<b>34.7</b>	

(b) Comparison of error rate (%) on CIFAR-10-C with severity level 5

Backbone	Methods	Avg. err	Gaus.	Shot	Impu.	Defo.	Glas.	Moti.	Zoom	Snow	Fros.	Fog	Brig.	Cont.	Elas.	Pixe.	Jpeg
WRN-40-2 (AugMix) [65,21]	Source	18.27	28.8	23.0	26.2	9.5	20.6	10.6	9.3	14.2	15.3	17.5	7.6	20.9	14.7	41.3	14.7
	TENT [59]	12.08	15.6	13.2	18.8	7.9	18.2	9.0	8.0	10.4	10.9	12.4	6.7	10.0	14.0	11.4	14.8
	<b>Ours</b>	<b>10.37</b>	<b>13.1</b>	<b>11.4</b>	<b>14.7</b>	<b>7.4</b>	<b>15.8</b>	<b>8.3</b>	<b>7.4</b>	<b>9.2</b>	<b>9.4</b>	<b>9.5</b>	<b>6.3</b>	<b>7.8</b>	<b>13.1</b>	<b>9.4</b>	<b>12.9</b>
WRN-28-10 [65]	Source	43.51	72.3	65.7	72.9	49.9	54.3	34.8	42.0	25.1	41.3	26.0	9.3	46.7	26.6	58.5	30.3
	TENT [59]	18.58	24.8	23.5	33.1	11.9	31.8	13.7	10.8	15.9	16.2	13.7	7.9	12.0	22.0	17.3	24.2
	Core* [63]	16.80	22.5	20.3	29.8	11.0	29.2	12.3	<b>10.2</b>	14.4	14.8	12.4	7.7	10.6	20.4	15.3	21.4
<b>Ours</b>	<b>15.70</b>	<b>20.1</b>	<b>18.4</b>	<b>26.2</b>	<b>10.8</b>	<b>28.9</b>	<b>12.1</b>	<b>10.2</b>	<b>13.7</b>	<b>13.9</b>	<b>11.1</b>	<b>7.6</b>	<b>8.8</b>	<b>20.2</b>	<b>14.2</b>	<b>19.4</b>	
ResNet-50 [19]	Source	29.14	48.7	44.0	57.0	11.8	50.8	23.4	10.8	21.9	28.2	29.4	<b>7.0</b>	13.13	23.4	47.9	19.5
	SHOT [38]	16.19	20.0	18.8	29.6	9.9	27.1	15.0	8.5	15.4	14.5	19.8	7.3	8.5	18.7	15.8	14.0
	TFA [39]	15.97	18.8	17.9	29.2	9.8	27.3	14.6	8.0	16.0	14.0	20.3	7.8	8.6	19.4	14.1	13.9
	TTT++ [39]	15.82	18.0	17.1	30.8	10.4	29.9	13.0	9.9	14.8	14.1	15.8	<b>7.0</b>	7.8	19.3	12.7	16.4
	TENT [59]	14.02	16.0	14.5	24.7	9.1	23.5	12.6	7.6	14.3	13.1	16.8	8.2	8.0	18.1	10.8	13.3
<b>Ours</b>	<b>12.52</b>	<b>14.1</b>	<b>13.4</b>	<b>20.9</b>	<b>8.3</b>	<b>20.7</b>	<b>11.2</b>	<b>7.3</b>	<b>12.4</b>	<b>11.7</b>	<b>14.4</b>	<b>7.3</b>	<b>7.4</b>	<b>16.5</b>	<b>9.7</b>	<b>12.4</b>	

(c) Comparison of average error (%) on CIFAR-100-C with all severity levels

Methods	Backbone	Lv.5	Lv.4	Lv.3	Lv.2	Lv.1
TENT [59]	ResNet-50	39.54	36.09	33.35	31.30	29.62
<b>Ours</b>		<b>35.65</b>	<b>32.82</b>	<b>30.54</b>	<b>28.92</b>	<b>27.51</b>
TENT [59]	WRN-40-2	35.53	32.89	30.72	29.06	27.67
<b>Ours</b>		<b>32.71</b>	<b>30.49</b>	<b>28.85</b>	<b>27.53</b>	<b>26.38</b>

of the model, outperforms the other backbone networks, and our method further enhances the performance by complementing it. Table 1(c) shows the results on CIFAR-100-C with all severity levels. Because severity denotes the strength of the corruption, it shows how much the distribution shift presents, and our method outperforms TENT [59] at all levels with a large margin.

#### 4.4 Ablation Studies

Table 2 shows the effectiveness of our proposed shift-agnostic weight regularization (SWR) and nearest source prototype (NSP) classifier through ablation studies. At a high learning rate, optimizing only the main task loss based on the entropy of the model prediction results in poor performance, but adjusting the learning rate reduces the error rate to 39.44%. Adding the NSP to the main task loss leads to the performance improvement of 1.89%, and including the SWR improves the performance by 1.68% even at a high learning rate. Our method with both

Table 2: Ablation study on CIFAR-100-C. ResNet-50 is used.

Methods	Learning rate	Average err (%)
Main	1e-3	75.70
Main (optimal learning rate)	1e-4	39.44
Main + NSP ( $\mathcal{L}_{\theta_e}^{\text{aux.ent}} + \mathcal{L}_{\theta_e}^{\text{aux.sel}}$ )	1e-4	37.55
Main + SWR	1e-3	37.76
Main + SWR + NSP ( $\mathcal{L}_{\theta_e}^{\text{aux.ent}}$ )	1e-3	35.80
Main + SWR + NSP ( $\mathcal{L}_{\theta_e}^{\text{aux.ent}} + \mathcal{L}_{\theta_e}^{\text{aux.sel}}$ )	1e-3	<b>35.65</b>

Table 3: Comparison of error rate (%) according to changes in projector depth.

Datasets	Backbone	Projector depth			
		None	1	2	3
CIFAR-100-C	WRN-40-2	33.04	32.79	<b>32.71</b>	32.89
	ResNet-50	36.34	<b>35.43</b>	35.65	36.81
CIFAR-10-C	WRN-40-2	<b>10.37</b>	10.52	10.42	10.46
	WRN-28-10	<b>15.70</b>	16.45	16.09	16.39
	ResNet-50	<b>12.52</b>	12.91	12.95	12.87

SWR and NSP achieves 35.65% error rate with 3.79% performance enhancement compared to using only the main task loss.

#### 4.5 Projector Design and Hyper-parameter Impacts

Table 3 shows the performance impact of changing the projector depth (*i.e.*, number of projection layer). In addition, we conduct experiments to apply the auxiliary task loss  $\mathcal{L}_{\theta_e}^{\text{aux}}$  directly to the feature representation  $\mathbf{h}$ , the encoder’s output without using the projector. The model with the projector outperforms the one without the projector on CIFAR-100-C, and opposite results are obtained on CIFAR-10-C. Since the auxiliary task loss is applied to the embedding space based on the cosine similarity between the source prototypes and the target embeddings, its effect may be minimal if they are severely misaligned. To compensate for this issue, we attach and train the projector that minimizes the misalignment between the source and target embeddings by enabling transformation-invariant mapping and bringing the projections belonging to the same class closer together in the new embedding space. However, if the number of classes is small (*e.g.*, CIFAR-10-C), the source and target may already be relatively well aligned compared to the case with a large number of classes (*e.g.*, CIFAR-100-C). In this case, we conjecture that applying the auxiliary task loss directly to the encoder’s output  $\mathbf{h}$  rather than the new embedding space  $\mathbf{z}$ , the projector’s output, generates a better-aligned representation  $\mathbf{h}$  between the source and target, which can be more helpful to the classifier.

Table 4 shows the experimental results according to (a) the projector width (*i.e.*, output dimension of the last layer), (b) the transformation used for training the projector, and (c) whether to fine-tune or freeze the projector during test-time adaptation. Our default settings are marked with gray-colored cells, and these

Table 4: Hyper-parameter impacts on CIFAR-100-C. ResNet-50 is used.

(a)		(b)		(c)		
Width	Error (%)	Transformation	Error (%)	Backbone	Error (%)	
128	36.34	No transform	37.88	WRN-40-2	ResNet-50	<b>32.71</b>   <b>35.65</b>
256	35.85	Color distortion	35.79			
<b>512</b>	<b>35.65</b>	<b>+Crop. &amp; Blur.</b>	<b>35.65</b>	<b>Freeze.</b>		
1024	36.09			<b>Finetune.</b>	32.85	35.96

Table 5: Comparison of error rate (%) on CIFAR-100-C. Our method outperforms the supervised method in an online setting. LR denotes a learning rate.

Methods	GT	Optimal	Epoch		
	Label	LR	1 (online)	2 (offline)	3 (offline)
Entropy Minimization	No	1e-4	39.81	38.84	39.08
Cross Entropy (Supervised)	Yes	2e-4	38.27	<b>7.41</b>	<b>0.86</b>
<b>Ours</b>	No	1e-3	<b>35.65</b>	33.34	33.25

settings are also applied to the domain generalization benchmarks in the following section without additional tuning.

#### 4.6 Quick Adaptation

As shown in Table 5, it is natural that the supervised method performs perfectly when learning and evaluating the same test samples iteratively. However, interestingly our method outperforms the supervised one in an online setting where the test sample is seen only once. Unlike the other methods that require a low learning rate to train (Fig. 1(b),(d)), our method updates the entire parameters at a high learning rate. We conjecture that SWR enables quick convergence without performance degradation because only parameters sensitive to distribution shift (*i.e.*, parameters that need to quickly adapt to a new domain) are largely updated with a high learning rate.

#### 4.7 Domain Generalization Benchmarks

To evaluate our method on the DG benchmarks, we follow the protocol proposed by DomainBed [16] and T3A [25]. Our method is model-agnostic, so we apply it to the pre-trained models using empirical risk minimization (ERM) [54] or CORAL [51] on the source domain in order to adapt the models to the target domain at test time. We use the leave-one-domain-out validation [16] for model selection in all experiments in Table 6. Our methods show state-of-the-art performance on average over four datasets and especially outperform T3A [25] and the source pre-trained models with a large margin on PACS, OfficeHome, and TerraIncognita datasets. The detailed experimental setup can be found in the supplementary Section C.

#### 4.8 Qualitative Results

Fig. 6 visualizes the features on CIFAR-10-C using t-SNE [41]. The results in the first row are from WRN-40-2 as a source pre-trained model, and the results in the

Table 6: Comparison of accuracy (%) on four DG benchmarks.  $\dagger$  denotes the reported results from DomainBed [16], and the others are reproduced values.

Methods	VLCS	PACS	OfficeHome	Terra	Average
ERM $\dagger$	76.8 $\pm$ 1.0	83.3 $\pm$ 0.6	67.3 $\pm$ 0.3	46.2 $\pm$ 0.2	68.4
CORAL $\dagger$	<b>77.0</b> $\pm$ 0.5	<b>83.6</b> $\pm$ 0.6	<b>68.6</b> $\pm$ 0.2	<b>48.1</b> $\pm$ 1.3	69.3
ERM	77.4 $\pm$ 0.9	83.5 $\pm$ 0.7	65.6 $\pm$ 0.4	47.1 $\pm$ 1.1	68.4
+T3A	<b>79.4</b> $\pm$ 0.4	86.5 $\pm$ 0.3	67.8 $\pm$ 0.5	45.6 $\pm$ 0.7	69.8
<b>+Ours</b>	77.0 $\pm$ 0.5	<b>88.9</b> $\pm$ 0.1	<b>69.2</b> $\pm$ 0.1	<b>49.5</b> $\pm$ 0.8	<b>71.2</b>
CORAL	77.9 $\pm$ 0.9	85.3 $\pm$ 0.1	67.8 $\pm$ 0.3	44.1 $\pm$ 0.4	68.8
+T3A	<b>79.3</b> $\pm$ 0.3	86.3 $\pm$ 0.2	69.5 $\pm$ 0.2	45.4 $\pm$ 1.2	70.1
<b>+Ours</b>	78.7 $\pm$ 0.4	<b>89.9</b> $\pm$ 0.1	<b>71.0</b> $\pm$ 0.0	<b>47.5</b> $\pm$ 0.6	<b>71.8</b>

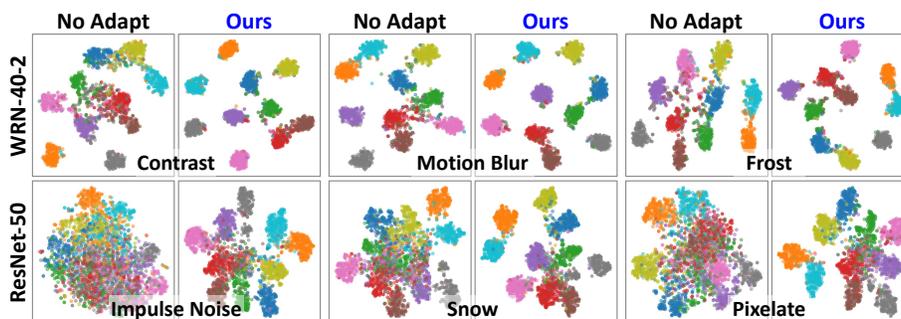


Fig. 6: t-SNE visualization of features from the target domain (CIFAR-10-C).

second row are from ResNet-50. Even without test-time adaptation, WRN-40-2 (AugMix) [21] is more robust against corruptions than ResNet-50, so better results can be obtained. Our method significantly improves the performance in terms of intra-class cohesion and inter-class separation in both backbones.

## 5 Conclusions

This paper proposed two novel approaches for model-agnostic test-time adaptation. Our proposed shift-agnostic weight regularization enables the model to reliably and quickly adapt to unlabeled online data from the target domain by controlling the update of the model parameters according to their sensitivity to the distribution shift. In addition, our proposed auxiliary task based on the nearest source prototype classifier boosts the performance by aligning the source and target representations. Test-time adaptation is a challenging but promising area in terms of allowing the model to evolve itself while adapting to a new environment without human intervention. In this regard, our efforts aim to promote the importance of this field and stimulate new research directions.

**Acknowledgement** We would like to thank Kyuwoong Hwang, Simyung Chang, Hyunsin Park, Juntae Lee, Janghoon Cho, Hyoungwoo Park, Byeonggeun Kim, and Hyesu Lim of the Qualcomm AI Research team for their valuable discussions.

## References

1. Agarwal, P., Paudel, D.P., Zaech, J.N., Van Gool, L.: Unsupervised robust domain adaptation without source data. In: IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) (2022) [4](#)
2. Assran, M., Caron, M., Misra, I., Bojanowski, P., Joulin, A., Ballas, N., Rabbat, M.: Semi-supervised learning of visual features by non-parametrically predicting view assignments with support samples. In: International Conference on Computer Vision (ICCV) (2021) [6](#)
3. Balaji, Y., Sankaranarayanan, S., Chellappa, R.: Metareg: Towards domain generalization using meta-regularization. In: Advances in Neural Information Processing Systems (NeurIPS) (2018) [4](#)
4. Bardes, A., Ponce, J., LeCun, Y.: Vicreg: Variance-invariance-covariance regularization for self-supervised learning. arXiv preprint arXiv:2105.04906 (2021) [8](#)
5. Beery, S., Van Horn, G., Perona, P.: Recognition in terra incognita. In: European Conference on Computer Vision (ECCV) (2018) [3](#), [9](#)
6. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International Conference on Machine Learning (ICML) (2020) [2](#), [8](#)
7. Choi, S., Kim, T., Jeong, M., Park, H., Kim, C.: Meta batch-instance normalization for generalizable person re-identification. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [4](#)
8. Choi, S., Jung, S., Yun, H., Kim, J.T., Kim, S., Choo, J.: Robustnet: Improving domain generalization in urban-scene segmentation via instance selective whitening. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [4](#)
9. Croce, F., Andriushchenko, M., Sehwag, V., Debenedetti, E., Flammarion, N., Chiang, M., Mittal, P., Hein, M.: Robustbench: a standardized adversarial robustness benchmark. arXiv preprint arXiv:2010.09670 (2020) [10](#)
10. Fang, C., Xu, Y., Rockmore, D.N.: Unbiased metric learning: On the utilization of multiple datasets and web images for softening bias. In: International Conference on Computer Vision (ICCV) (2013) [3](#), [9](#)
11. Ganin, Y., Lempitsky, V.: Unsupervised domain adaptation by backpropagation. In: International Conference on Machine Learning (ICML) (2015) [4](#)
12. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V.: Domain-adversarial training of neural networks. The journal of machine learning research (2016) [4](#)
13. Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. In: International Conference on Learning Representations (ICLR) (2018) [2](#)
14. Gong, R., Li, W., Chen, Y., Gool, L.V.: Dlow: Domain flow for adaptation and generalization. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019) [4](#)
15. Grandvalet, Y., Bengio, Y.: Semi-supervised learning by entropy minimization. In: Advances in Neural Information Processing Systems (NeurIPS) (2004) [2](#)
16. Gulrajani, I., Lopez-Paz, D.: In search of lost domain generalization. In: International Conference on Learning Representations (ICLR) (2020) [10](#), [13](#), [14](#)
17. Guo, Y., Shi, H., Kumar, A., Grauman, K., Rosing, T., Feris, R.: Spottune: transfer learning through adaptive fine-tuning. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019) [2](#)

18. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2020) [2](#), [8](#)
19. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016) [10](#), [11](#)
20. Hendrycks, D., Dietterich, T.: Benchmarking neural network robustness to common corruptions and perturbations. In: International Conference on Learning Representations (ICLR) (2018) [2](#), [3](#), [9](#)
21. Hendrycks, D., Mu, N., Cubuk, E.D., Zoph, B., Gilmer, J., Lakshminarayanan, B.: Augmix: A simple data processing method to improve robustness and uncertainty. In: International Conference on Learning Representations (ICLR) (2019) [3](#), [10](#), [11](#), [14](#)
22. Hoffman, J., Tzeng, E., Park, T., Zhu, J.Y., Isola, P., Saenko, K., Efros, A., Darrell, T.: Cycada: Cycle-consistent adversarial domain adaptation. In: International Conference on Machine Learning (ICML) (2018) [4](#)
23. Hu, W., Miyato, T., Tokui, S., Matsumoto, E., Sugiyama, M.: Learning discrete representations via information maximizing self-augmented training. In: International Conference on Machine Learning (ICML) (2017) [6](#)
24. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: International Conference on Machine Learning (ICML) (2015) [2](#), [10](#)
25. Iwasawa, Y., Matsuo, Y.: Test-time classifier adjustment module for model-agnostic domain generalization. *Advances in Neural Information Processing Systems (NeurIPS)* (2021) [1](#), [2](#), [4](#), [9](#), [10](#), [13](#)
26. Jain, H., Zepeda, J., Pérez, P., Gribonval, R.: Learning a complete image indexing pipeline. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018) [2](#)
27. Krause, A., Perona, P., Gomes, R.: Discriminative clustering by regularized information maximization. *Advances in Neural Information Processing Systems (NeurIPS)* (2010) [6](#)
28. Krause, A., Perona, P., Gomes, R.: Discriminative clustering by regularized information maximization. In: Lafferty, J., Williams, C., Shawe-Taylor, J., Zemel, R., Culotta, A. (eds.) *Advances in Neural Information Processing Systems (NeurIPS)* (2010) [6](#)
29. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009) [10](#)
30. Kundu, J.N., Venkat, N., Babu, R.V., et al.: Universal source-free domain adaptation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2020) [4](#)
31. Li, D., Yang, Y., Song, Y.Z., Hospedales, T.M.: Deeper, broader and artier domain generalization. In: International Conference on Computer Vision (ICCV) (2017) [3](#), [9](#)
32. Li, D., Yang, Y., Song, Y.Z., Hospedales, T.M.: Learning to generalize: Meta-learning for domain generalization. *arXiv preprint arXiv:1710.03463* (2017) [4](#)
33. Li, D., Zhang, J., Yang, Y., Liu, C., Song, Y.Z., Hospedales, T.M.: Episodic training for domain generalization. In: International Conference on Computer Vision (ICCV) (2019) [4](#)
34. Li, H., Jialin Pan, S., Wang, S., Kot, A.C.: Domain generalization with adversarial feature learning. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018) [4](#)

35. Li, L., Gao, K., Cao, J., Huang, Z., Weng, Y., Mi, X., Yu, Z., Li, X., Xia, B.: Progressive domain expansion network for single domain generalization. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [4](#)
36. Li, R., Jiao, Q., Cao, W., Wong, H.S., Wu, S.: Model adaptation: Unsupervised domain adaptation without source data. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2020) [4](#)
37. Li, Y., Tian, X., Gong, M., Liu, Y., Liu, T., Zhang, K., Tao, D.: Deep domain generalization via conditional invariant adversarial networks. In: European Conference on Computer Vision (ECCV) (2018) [4](#)
38. Liang, J., Hu, D., Feng, J.: Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In: International Conference on Machine Learning (ICML) (2020) [1](#), [4](#), [6](#), [11](#)
39. Liu, Y., Kothari, P., van Delft, B., Bellot-Gurlet, B., Mordan, T., Alahi, A.: Ttt++: When does self-supervised test-time training fail or thrive? Advances in Neural Information Processing Systems (NeurIPS) (2021) [1](#), [2](#), [3](#), [4](#), [10](#), [11](#)
40. Long, M., Zhu, H., Wang, J., Jordan, M.I.: Unsupervised domain adaptation with residual transfer networks. In: Advances in Neural Information Processing Systems (NeurIPS) (2016) [2](#)
41. Van der Maaten, L., Hinton, G.: Visualizing data using t-sne. *Journal of machine learning research* (2008) [13](#)
42. Mummadi, C.K., Hutmacher, R., Rambach, K., Levinkov, E., Brox, T., Metzen, J.H.: Test-time adaptation to distribution shift by confidence maximization and input transformation. *arXiv preprint arXiv:2106.14999* (2021) [1](#), [2](#), [4](#), [6](#)
43. Nair, V., Hinton, G.E.: Rectified linear units improve restricted boltzmann machines. In: International Conference on Machine Learning (ICML) (2010) [10](#)
44. Pan, X., Luo, P., Shi, J., Tang, X.: Two at once: Enhancing learning and generalization capacities via ibn-net. In: European Conference on Computer Vision (ECCV) (2018) [4](#)
45. Rahman, M.M., Fookes, C., Baktashmotlagh, M., Sridharan, S.: Correlation-aware adversarial domain adaptation and generalization. *Pattern Recognition* (2020) [4](#)
46. Saito, K., Watanabe, K., Ushiku, Y., Harada, T.: Maximum classifier discrepancy for unsupervised domain adaptation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018) [4](#)
47. Seo, S., Suh, Y., Kim, D., Han, J., Han, B.: Learning to optimize domain specific normalization for domain generalization. *arXiv preprint arXiv:1907.04275* (2019) [4](#)
48. Shi, Y., Sha, F.: Information-theoretical learning of discriminative clusters for unsupervised domain adaptation. In: International Conference on Machine Learning (ICML) (2012) [6](#)
49. Springenberg, J.T.: Unsupervised and semi-supervised learning with categorical generative adversarial networks. *International Conference on Learning Representations (ICLR)* (2016) [6](#)
50. Su, J.C., Maji, S., Hariharan, B.: When does self-supervision improve few-shot learning? In: European Conference on Computer Vision (ECCV) (2020) [3](#)
51. Sun, B., Saenko, K.: Deep coral: Correlation alignment for deep domain adaptation. In: European Conference on Computer Vision (ECCV) (2016) [3](#), [4](#), [13](#)
52. Sun, Y., Wang, X., Liu, Z., Miller, J., Efros, A., Hardt, M.: Test-time training with self-supervision for generalization under distribution shifts. In: International Conference on Machine Learning (ICML) (2020) [1](#), [2](#), [3](#), [4](#)
53. Tzeng, E., Hoffman, J., Saenko, K., Darrell, T.: Adversarial discriminative domain adaptation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017) [4](#)

54. Vapnik, V.N.: An overview of statistical learning theory. *IEEE transactions on neural networks* (1999) [13](#)
55. Venkateswara, H., Eusebio, J., Chakraborty, S., Panchanathan, S.: Deep hashing network for unsupervised domain adaptation. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017) [3](#), [9](#)
56. Volpi, R., Namkoong, H., Sener, O., Duchi, J.C., Murino, V., Savarese, S.: Generalizing to unseen domains via adversarial data augmentation. *Advances in neural information processing systems* **31** (2018) [4](#)
57. Vu, T.H., Jain, H., Bucher, M., Cord, M., Pérez, P.: Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2019) [2](#), [4](#), [6](#)
58. Wang, D., Liu, S., Ebrahimi, S., Shelhamer, E., Darrell, T.: On-target adaptation. *arXiv preprint arXiv:2109.01087* (2021) [4](#), [6](#)
59. Wang, D., Shelhamer, E., Liu, S., Olshausen, B., Darrell, T.: Tent: Fully test-time adaptation by entropy minimization. In: *International Conference on Learning Representations (ICLR)* (2020) [1](#), [2](#), [3](#), [4](#), [6](#), [9](#), [10](#), [11](#)
60. Yang, S., Wang, Y., van de Weijer, J., Herranz, L., Jui, S.: Generalized source-free domain adaptation. In: *International Conference on Computer Vision (ICCV)* (2021) [4](#)
61. Yeh, H.W., Yang, B., Yuen, P.C., Harada, T.: Sofa: Source-data-free feature alignment for unsupervised domain adaptation. In: *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)* (2021) [4](#)
62. Yosinski, J., Clune, J., Bengio, Y., Lipson, H.: How transferable are features in deep neural networks? *Advances in Neural Information Processing Systems (NeurIPS)* (2014) [2](#)
63. You, F., Li, J., Zhao, Z.: Test-time batch statistics calibration for covariate shift. *arXiv preprint arXiv:2110.04065* (2021) [1](#), [2](#), [4](#), [11](#)
64. Yu, T., Kumar, S., Gupta, A., Levine, S., Hausman, K., Finn, C.: Gradient surgery for multi-task learning. *Advances in Neural Information Processing Systems (NeurIPS)* (2020) [3](#)
65. Zagoruyko, S., Komodakis, N.: Wide residual networks. *British Machine Vision Conference (BMVC)* (2016) [10](#), [11](#)
66. Zhang, Y., Borse, S., Cai, H., Porikli, F.: Auxadapt: Stable and efficient test-time adaptation for temporally consistent video semantic segmentation. In: *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)* (2022) [1](#)
67. Zhou, K., Yang, Y., Hospedales, T., Xiang, T.: Learning to generate novel domains for domain generalization. In: *European Conference on Computer Vision (ECCV)* (2020) [4](#)