

Feature Alignment and Uniformity for Test Time Adaptation

Shuai Wang¹ Daoan Zhang² Zipei Yan³ Jianguo Zhang² Rui Li¹
¹Tsinghua University ²Southern University of Science and Technology
³Hongkong Polytechnic University

s-wang20@mails.tsinghua.edu.cn, dawn_change@126.com, zipei.yan@connect.polyu.hk
 zhangjg@sustech.edu.cn, leerui@tsinghua.edu.cn

Abstract

*Test time adaptation (TTA) aims to adapt deep neural networks when receiving out of distribution test domain samples. In this setting, the model can only access online unlabeled test samples and pre-trained models on the training domains. We first address TTA as a **feature revision** problem due to the domain gap between source domains and target domains. After that, we follow the two measurements **alignment** and **uniformity** to discuss the test time feature revision. For test time feature uniformity, we propose a **test time self-distillation** strategy to guarantee the consistency of uniformity between representations of the current batch and all the previous batches. For test time feature alignment, we propose a **memorized spatial local clustering** strategy to align the representations among the neighborhood samples for the upcoming batch. To deal with the common noisy label problem, we propound the entropy and consistency filters to select and drop the possible noisy labels. To prove the scalability and efficacy of our method, we conduct experiments on four domain generalization benchmarks and four medical image segmentation tasks with various backbones. Experiment results show that our method not only improves baseline stably but also outperforms existing state-of-the-art test time adaptation methods.*

1. Introduction

Deep learning has achieved great success in computer vision tasks when training and test data are sampled from the same distribution [17, 26, 37]. However, in real-world applications, performance degradation usually occurs when training (source) data and test (target) data are collected from different distributions, *i.e.* *domain shift*. In practice, test samples may encounter different types of variations or corruptions. Deep learning models will be sensitive to these variations or corruptions, which can cause performance degradation.

To tackle this challenging but practical problem, various

works have been proposed to adapt the model at *test* time [5, 8, 21, 36, 43, 55, 64, 65, 74]. Test time training (TTT) [36, 55] adapts model using self-supervised tasks, such as rotation classification [15] during both training and test phases. This paradigm relies on additional model modifications in both training and test phases, which is not feasible and scalable in the real world.

Similar to TTT, test time adaptation [64] (TTA) also adapts the model *i.e.*, updates the parameters in the test phase. But TTA does not require any specific modifications in training and requires only the pre-trained source model and unlabeled target data during the test phase, which is more practical and generalizable. In TTA, the model could be adapted with only online unlabeled data. Hence, the model trained on source data is incompatible with the target data due to the possible domain gap between source data and target data.

To deal with the above-mentioned problem, we address TTA as a *representation revision* problem in this paper. In the test phase of TTA, the accessed model has already learned the feature representations specialized to source domains and may generate inaccurate representations for the target domain due to the large domain gap. It is necessary to rectify the feature representations for the target domain. To achieve better representations for the target domain, we utilize the commonly used measurements for representation quality that can be summarized to feature *alignment* and *uniformity* [66, 71]. Alignment refers that the similar images should have the similar representations, while uniformity means that images of different classes should be distributed as uniform as possible in the latent space. Hence, we propose to address the TTA problem from the above-mentioned properties.

Most of the previous works about TTA can be inducted from the proposed representation revision perspective. Some methods adapt the source model by conducting a feature alignment process, such as feature matching [25, 36] and predictions adjustment [5]. One of the representative methods is LAME [5], which encourages neigh-

neighborhood samples in the feature space to have similar predictions using Laplacian adjusted maximum-likelihood estimation. Moreover, other methods aim to make target feature more uniform in feature space, including entropy minimization [43, 64, 74], prototype adjustment [21], information maximization [34] and batch normalization statistics alignment [33, 42, 51]. One representative method is T3A [21], which adjusts prototypes (class-wise centroids) to have a more uniform representation by building a support set. However, none of the methods address the TTA problem from the representation alignment and uniformity simultaneously. In this paper, we identify this limitation and propose a novel method that rectifies feature representation from both two properties. We formulate the two properties in TTA as *test time feature uniformity* and *test time feature alignment*.

Test Time Feature Uniformity. Following the feature uniformity perspective, we hope that representations of test images from different classes should be distributed as uniform as possible. However, only limited test samples can be accessed in an *online* manner in the TTA setting.

To better deal with all of the samples in the target domain, we propose to introduce historical temporal information for every arriving test sample. A memory bank is built to store feature representations and logits for all arriving samples to maintain the useful information from the previous data. We then calculate the pseudo-prototypes for every class by using the logits and features in the memory bank. After that, to guarantee the uniformity for the current batch of samples, the prediction distribution of prototype-based classification and model prediction (outputs of linear classifier) should be similar, *i.e.*, the feature distribution of the current images of one class should be consistent with the feature distribution of all the previous images of the same class. This can reduce the bias of misclassified outlier samples to form a more uniform latent space.

Motivated by this, we minimize the distance between the outputs of linear and prototype-based classifiers. This pattern is similar to self-distillation [23, 73] that transfers knowledge between different layers of the same network architecture. However, unlike typical self-distillation, our method does not require any ground truth supervision. We refer to this method as *Test Time Self-Distillation* (TSD).

Test Time Feature Alignment. Feature alignment encourages the images from the same class to have similar feature representations in the latent space. As for TTA, pseudo labels generated by the source model may be noisy due to the domain gap. Thus, instead of aligning all the positive pairs, we propose a *K-nearest feature alignment* to encourage features from the same class to be closed or features from different classes to be far away from each other. This can reduce the negative impact imposed by the noisy labels and maintain the alignment of images with the same

semantics. Specifically, we retrieve K -nearest features in the memory bank for the upcoming images and add consistency regularization between the representations and logits of the images. We refer to this as *Memorized Spatial Local Clustering* (MSLC). The ablation of hyperparameter K is shown in Table 5 and Fig. 3.

Entropy Filter and Consistency Filter. During the adaption, we use stored pseudo features and logits to compute the pseudo-prototypes. However, the noisy label problem cannot be completely alleviated despite the proposed efforts. To further reduce the impact, we adopt both *entropy* and *consistency* filters to filter noisy labels to boost performance. As for the entropy filter, we filter noisy features with high entropy when we compute prototypes because unreliable samples usually produce high entropy.

In addition, the predictions of prototype-based and linear classifiers of the network for reliable samples should be consistent ideally. We use this property to filter unreliable samples and back-propagate the gradient using only reliable samples. We refer to this filter as the consistency filter. The ablation study on the two proposed filters is presented in Table 5.

Finally, we demonstrate the effectiveness of our proposed approach on commonly used domain generalization benchmarks, including PACS [30], VLCS [58], OfficeHome [62] and DomainNet [47]. Furthermore, to prove the efficacy of our method, we conduct more experiments on four cross-domain medical image segmentation benchmarks, including prostate segmentation [1, 35], cardiac structure segmentation [4, 76, 77] and optic cup/disc segmentation [13, 45, 53]. Our method achieves the state-of-the-art performance on the above benchmarks.

We summarize our contributions as follows:

- We propose a new perspective for test time adaptation from the view of feature alignment and uniformity. The proposed test time feature uniformity encourages the representations of the current batch of samples along with the uniformity of all the previous samples. The test time feature alignment manipulates the representation of the test sample according to its neighbors in the latent space to align the representations based on the pseudo label.
- Specifically, to meet the online setting and noisy label problem in TTA, we propose two complementary strategies: unsupervised self-distillation for test time feature uniformity and memorized spatial local clustering for test time feature alignment. We also propose the entropy filter and consistency filter to further mitigate the effect of the noisy labels.
- The experiments demonstrate that our proposed method outperforms existing test time adaptation ap-

proaches on both the domain generalization benchmarks and medical image segmentation benchmarks.

2. Related Work

2.1. Domain Generalization

Domain generalization (DG) aims to learn knowledge from multi-source domains and generalize to the unseen target domain. A primary strategy of DG is domain-invariant feature learning which aims to reduce domain gaps, including aligning distributions among multiple domains with contrastive learning [22, 41, 68, 70], learning useful representations with self-supervised learning [6], matching statistics of feature distributions across domains [54, 59], domain adversarial learning [14, 31] and causality inference [38, 39]. Meta-learning based methods [2, 12, 29] simulate domain shift by dividing meta-train and meta-test domains from the original source domains. Data augmentation-based methods effectively solve the DG problem by diversifying training data [18, 32, 44, 63, 69, 72, 75].

Unlike the DG paradigm that focuses on training a generalizable model during the training phase with labeled source data, test time adaptation aims to adapt a pre-trained model in the test phase by using the online unlabeled target data.

2.2. Test Time Adaptation

Our work is mostly related to test time adaptation [5, 8, 21, 43, 64] or test time training [36, 55]. Test Time Training (TTT) [36, 55] adapts the model during the test phase via the self-supervised task, such as rotation classification [15]. However, TTT needs to optimize the source model by jointly training supervised loss and self-supervised loss, which may not be feasible in the real world. Different from this, test time adaptation [64] aims to adapt the model during the test phase without changing the training process. Tent [64] minimizes entropy to update the trainable parameters in Batch Normalization [20] layer. SHOT [34] combines information maximization and pseudo labeling. There are some works [43, 65] combining test time adaptation and continual learning to maintain the performance on the source domain. LAME [5] adapts the output of the model rather than parameters with a laplacian adjusted maximum-likelihood estimation. There are also some works using test time Batch Normalization statistics [33, 42, 51], self-training [8] and feature alignment [25, 36]. The above methods could be viewed as either feature uniformity or feature alignment as discussed in the Sec. 1, while our method benefits from both properties.

3. Method

Fig. 1 illustrates the overall pipeline of our method. We describe the details of the problem settings and our method in this section.

3.1. Preliminaries

In test time adaptation (TTA), we can only achieve target domain unlabeled images in an online manner and pre-trained model on the source domain. The source model is trained with standard empirical risk minimization on source domains, *e.g.* cross entropy loss for the image classification task. Given the model trained on \mathcal{D}_s , we aim to adapt the model using unlabeled target data $\{x_i\} \in \mathcal{D}_t, i \in \{1 \dots N\}$, where x_i denotes the i th image of target domain \mathcal{D}_t and N denotes the number of target images, \mathcal{D}_s denotes the source domain. During testing, we initialize the model $g = f \circ h$ with source model parameters trained on source domain \mathcal{D}_s , where f denotes backbone and h denotes linear classification head. The output of model g for image x_i is denoted as $p_i = g(x_i) \in \mathbb{R}^C$ where C is the number of class.

3.2. Test Time Self-Distillation

During adaption, given a batch of unlabeled test samples, we can generate image embeddings $z_i = f(x_i)$, logits $p_i = h(z_i)$ and pseudo labels $\hat{y}_i = \arg \max p_i$ via the pre-trained model.

We then maintain a memory bank $\mathcal{B} = \{(z_i, p_i)\}$ to store image embeddings z_i and logits p_i . Following T3A [21], the memory bank is initialized with the weights of the linear classifier. When the target sample x_i comes, for every image, we add the image embedding z_i and logits p_i into the memory bank. To build up the relations between the current samples and all of the previous samples, the pseudo-prototypes shall be generated for each class. The prototype of class k could be formulated as

$$c_k = \frac{\sum_i z_i \mathbb{1}[\hat{y}_i = k]}{\sum_i \mathbb{1}[\hat{y}_i = k]}, \quad (1)$$

where $\mathbb{1}(\cdot)$ is an indicator function, outputting value 1 if the argument is true or 0 otherwise. However, some pseudo labels may be assigned to the wrong class, leading to incorrect prototype computation. We use Shannon entropy [52] filter to filter noisy labels. For prediction p_i , its entropy could be computed as $H(p_i) = -\sum \sigma(p_i) \log \sigma(p_i)$ where σ denotes the softmax operation. We aim to filter unreliable features or predictions with high entropy because lower entropy usually means higher accuracy [64]. Specifically, for every class, image embeddings with the top- M highest entropy in the memory bank would be ignored. After that, we use filtered embeddings to calculate prototypes as Eq. 1 and define prototype-based classification output as the softmax over the feature similarity to prototypes for class k :

$$y_i^k = \frac{\exp(\text{sim}(z_i, c_k))}{\sum_{k'=1}^C \exp(\text{sim}(z_i, c_{k'}))}, \quad (2)$$

where $\text{sim}(z_i, c_k)$ denotes cosine similarity between z_i and c_k .

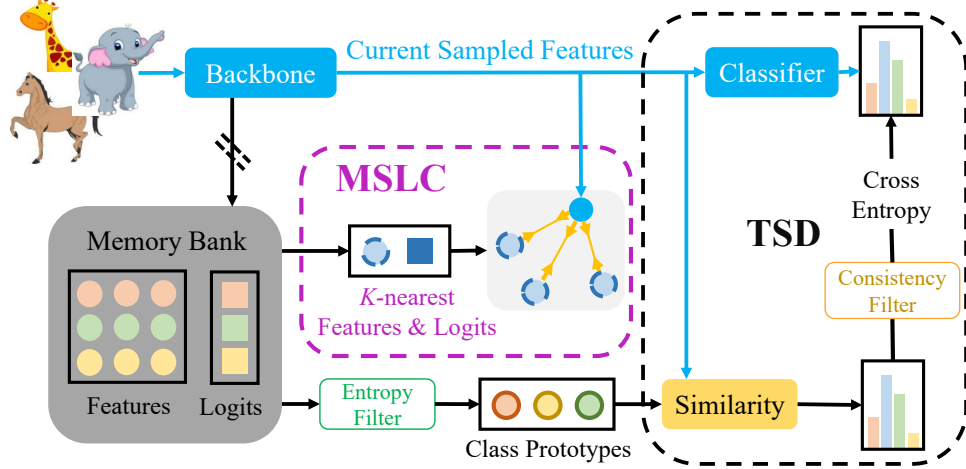


Figure 1. **Overview of our proposed method.** Blue lines denote forward and backward and black lines denote only forward (*i.e.* without gradient backpropagation). Different colours of features, logits and prototypes mean different classes. MSLC: Memorized Spatial Local Clustering. TSD: Test Time Self-Distillation.

The prototype-based classification results y_i and outputs p_i of the network g should share a similar distribution for the same input. Thus, the loss to maintain uniformity is proposed as

$$\mathcal{L}_i(p_i, y_i) = -\sigma(p_i) \log y_i. \quad (3)$$

Note that p_i is a soft pseudo label rather than a hard pseudo label. The reason for using a soft label is that a soft label usually provides more information [19]. By using the proposed test time self-distillation, the network could map the uniformity for the current samples to improve the quality of representations.

Although we use an entropy filter to drop noisy labels when computing prototypes, there are still some mistake predictions inevitable. We propose that, for a reliable sample, the outputs of the linear fully connected layer and prototype-based classifier should be similar. Hence, we adopt the consistency filter to identify the mistake predictions. Specially, if the linear classifier and prototype-based classifier produce the same predictions, *i.e.* the same result after executing argmax to the logits, we assume that this sample is reliable. This strategy could be implemented using a filter mask for image x_i as follows

$$\mathcal{M}_i = \mathbb{1}[\text{arg max } p_i = \text{arg max } y_i]. \quad (4)$$

By conducting a consistency filter, we further filter unreliable samples and the unsupervised self-distillation loss could be formulated as follows

$$\mathcal{L}_{tsd} = \frac{\sum_i \mathcal{L}_i * \mathcal{M}_i}{\sum_i \mathcal{M}_i}. \quad (5)$$

3.3. Memorized Spatial Local Clustering

As mentioned before, features belonging to the same class should be aligned in the latent space. However, this

situation may differ in TTA due to the domain gap between the target and source domains. We encourage K -nearest neighborhood features instead of all the features to be close to reducing the noisy label impact. A simple strategy is to add consistency regularization within a batch of samples. However, historical temporal information is ignored and the alignment is less effective. Moreover, there is a trivial solution that the model can easily map all images to a certain class if we use only one batch sample for alignment [3].

To handle these problems, we concatenate spatial local clustering with the memory bank. we begin from retrieving K -nearest features in the memory bank for image x . Based on our assumption, the logits of image x should be aligned with the logits of its nearest neighbors in the latent space. To achieve this, we align the two varieties of logits according to the distance between the image embeddings of image x and its neighbors. The formulation is presented as follows

$$\mathcal{L}_{mslc} = \frac{1}{K} \sum_{j=1}^K \text{sim}(z, z_j) (\sigma(p) - \sigma(p_j))^2, \quad (6)$$

where $\text{sim}(z, z_j)$ denotes cosine similarity between z and z_j . $\{z_j\}_{j=1}^K$ denotes the K -nearest image embeddings of z in the memory bank \mathcal{B} and p_j denotes the corresponding logits. If z_j and z are close in feature space, *i.e.* $\text{sim}(z, z_j)$ is large, this objective function will push p_j and p to get close. We detach the gradient of $\text{sim}(z, z_j)$, *i.e.* $\text{sim}(z, z_j)$ would be viewed as constant, to avoid the trivial solution that the model will output a constant result regardless of different samples.

3.4. Training Objective Function

Combing Eq. 5 and Eq. 6, we formulate the final objective function as

$$\mathcal{L} = \mathcal{L}_{tsd} + \lambda \mathcal{L}_{mslc}, \quad (7)$$

where λ is the trade-off parameter to balance different loss functions. In our implementation, we use cosine similarity as the similarity metric. Specifically, we define $\text{sim}(x, y) = x^\top y / \|x\| \|y\|$.

During testing phase, the adaptation is performed in an *online* manner. Specifically, when receiving image x_T at time point T , the model state is initialized with the parameters update from the last image x_{T-1} . The model produces the prediction $p_T = g(x_T)$ after receiving new samples x_T and updates the model using Eq. 7 with only *one* step gradient descent. It is noticed that the adaption could continue as long as there exists test data.

4. Experiments

4.1. Experimental Setup

Datasets. **PACS** [30] contains 9,991 examples and 7 classes that are collected from 4 domains: art, cartoons, photos, and sketches. **OfficeHome** [62] is consisted of 4 domains: art, clipart, product, and real, which includes 15,588 images and 65 classes. **VLCS** [58] comprises four domains: Caltech101, LabelMe, SUN09 and VOC2007 and includes 10,729 images and 5 classes. **DomainNet** [47], a large-scale dataset has six domains $d \in \{\text{clipart, infograph, painting, quickdraw, real, sketch}\}$ with 586,575 images and 345 classes.

Models. In the main experiments, we evaluated different methods on ResNet-18/50 [17] equipped with Batch Normalization [20], which is widely used in domain adaptation and generalization literature. Furthermore, we tested our algorithm on different backbones, including Vision Transformer (ViT-B/16) [11], ResNeXt-50 (32×4d) [67], EfficientNet (B4) [56] and MLP-Mixer (Mixer-L16) [57].

Implementation. For source training, we choose one domain as the target domain and the other domains as the source domains. We split all images from the source domains to 80%/20% for training and validation. We use the Adam optimizer [24] with $5e^{-5}$ as the learning rate. All models are initialized with ImageNet-1K [49] pre-trained weights except for ViT-B/16 and MLP-Mixer when we use ImageNet-21K pre-trained weights. We use `torchvision` implementations of all models except for ViT-B/16 and MLP-Mixer; instead, we use the implementations in the `timm` library.

For test time adaptation, we use the Adam optimizer [24] and set the batch size as 128. We empirically set the trade-off parameter $\lambda = 0.1$ (*cf.* Eq. 7). Unlike [43, 64], all the trainable layers are updated and no special selection is required in our method. We use PyTorch [46] for all implementations. We report the accuracy of the whole target domain for evaluation. For all experiments, we report the average of three repetitions with different weight initialization, random seed and data split. Please refer to our supple-

Table 1. Comparison of our method and existing test time adaptation methods with ResNet18 backbone. We highlight the **best** and the second results.

Method	PACS	OfficeHome	VLCS	DomainNet	Avg.
ERM [61]	82.07	63.12	72.75	38.95	64.22
BN [51]	82.82	62.30	64.31	37.80	61.81
Tent [64]	<u>84.92</u>	63.75	67.36	38.95	63.75
PL [28]	84.64	60.22	68.93	35.23	62.26
SHOT-IM [34]	82.55	63.42	64.90	39.50	62.59
T3A [21]	83.50	<u>64.25</u>	<u>73.03</u>	<u>39.61</u>	<u>65.10</u>
ETA [43]	82.70	62.46	64.35	39.43	62.24
LAME [5]	84.58	62.20	72.88	37.49	64.29
Ours	87.32	64.83	73.61	40.19	66.49

Table 2. Comparison of our method and existing test time adaptation methods with ResNet50 backbone. We highlight the **best** and the second results.

Method	PACS	OfficeHome	VLCS	DomainNet	Avg.
ERM [61]	84.59	67.37	<u>74.01</u>	45.20	67.74
BN [51]	85.03	66.10	64.78	43.38	64.82
Tent [64]	<u>87.48</u>	67.96	69.20	44.71	67.34
PL [28]	85.23	67.13	68.52	41.18	65.52
SHOT-IM [34]	85.50	67.39	65.23	<u>46.30</u>	66.11
T3A [21]	86.04	<u>68.29</u>	73.98	46.16	<u>68.62</u>
ETA [43]	85.04	66.21	64.79	46.13	65.54
LAME [5]	86.62	66.19	73.94	43.20	67.49
Ours	89.41	68.67	74.52	47.73	70.08

mentary material for more details, including detailed results and error bars.

Hyperparameter search and model selection. We use training domain validation [16] for source model training. We select the model with the highest accuracy on the validation set. For hyperparameter search, we search learning rate lr in $\{1e^{-3}, 1e^{-4}, 1e^{-5}, 1e^{-6}\}$, the hyperparameter for feature filter $M \in \{1, 5, 20, 50, 100, \text{NA}\}$ where NA denotes no entropy filter. We emphasize that all hyperparameters in the TTA setting should be selected before accessing test samples. We do hyperparameter search on *training domain validation set*.

Baselines. We compared our method with Empirical Risk Minimization (ERM) [61], Tent [64], T3A [21], SHOT-IM [34], ETA [43], Test Time Batch Normalization (BN) [51], Laplacian Adjusted Maximumlikelihood Estimation (LAME) [5] and PseudoLabeling (PL) [28]. We use the implementation from the released code of T3A library¹, except for LAME and ETA, we use the source code of authors.

4.2. Comparative Study

Comparison with TTA methods. Table 1 and Table 2 present the results on ResNet-18/50 of four different datasets. From Table 1 and 2, we can see that our method

¹<https://github.com/matsuolab/T3A>

Table 3. Compare with domain generalization methods on PACS dataset with ResNet50. †: numbers are from the original literature. ‡ denotes our implementation results. We highlight the **best** and the second results in each column.

Method	A	C	P	S	Avg.
ERM [61]	82.5	80.8	94.1	81.0	84.6
DNA† [10]	89.8	83.4	<u>97.7</u>	82.6	88.4
PCL† [68]	<u>90.2</u>	83.9	98.1	82.6	88.7
SWAD‡ [7]	89.3	83.2	96.9	83.4	88.2
Ours	87.7	<u>88.8</u>	96.2	<u>85.0</u>	89.4
SWAD + Ours	92.2	89.2	97.1	85.6	91.0

generally achieves the state-of-the-art performance. Specifically, the proposed method improves the ERM [61] baseline with 4.8%, 1.3%, 0.5%, 2.53% for each dataset, respectively. Other test time adaptation methods do not improve the baseline as stably as ours.

We also visualized the change in accuracy of different methods throughout the adaptation process which is shown in Fig. 2. The “Batch Numbers” in Fig. 2 indicates how many batches of images that the model has been updated with. We can see that our method could adapt the data much faster and achieve higher accuracy on the target domain.

Comparison with DG/SFDA methods. The above experiments mainly focus on test time adaptation that aims to adapt the model during *test* phase. It is natural to ask: *how about our method compared with domain generalization or source free domain adaptation methods?* To answer this question, we first compared our method with some recent domain generalization or source free domain adaptation methods, such as SWAD [7], PCL [68], DNA [10], and F-mix [27]² on PACS and DomainNet dataset. From Table 3, we can see that our method outperforms the state-of-the-art methods in domain generalization. Furthermore, combining SWAD [7], we achieve the impressive 91% accuracy with ResNet50 backbone.

We also report the result of the challenging DomainNet dataset. The results are listed in Table 4. It can be seen that our method outperforms the state-of-the-art methods of domain generalization, such as SWAD [7] and DNA [10]. Also, our method could significantly improve SWAD [7], which outperforms the current state-of-the-art SFDA method [27]. Note that online test time adaptation is more flexible in the real world because SFDA adapts the test data in an offline manner which requires more training loops and resources than TTA.

To summarize, our model can outperform all the domain generalization and test time adaptation methods on various datasets.

²“F-mix” denotes *feature-mixup* in [27].




Figure 2. Accuracy visualizations of different methods during the adaptation process on the PACS dataset (target domain: art). The “Batch Numbers” indicates how many batches of images the model has been updated with.

Table 4. Compare with domain generalization methods and source free domain adaptation methods on DomainNet dataset with ResNet50. †: numbers are from the original literature. ‡ denotes our implement results. We highlight the **best** and the second results in each column.

Method	clip	info	paint	quick	real	sketch	Avg.
Domain generalization methods							
ERM [61]	64.8	22.1	51.8	13.8	64.7	54.0	45.2
PCL† [68]	67.9	24.3	55.3	15.7	66.6	56.4	47.7
DNA† [10]	66.1	23.0	54.6	16.7	65.8	56.8	47.2
SWAD‡ [7]	66.1	22.4	53.6	16.3	65.5	56.2	46.7
Source free domain adaptation methods							
F-mix† [27]	75.4	<u>24.6</u>	<u>57.8</u>	<u>23.6</u>	65.8	<u>58.5</u>	<u>51.0</u>
Ours	66.1	24.1	52.8	18.2	68.5	56.7	47.7
SWAD + Ours	<u>69.2</u>	28.4	58.2	26.2	<u>68.1</u>	59.6	51.6

4.3. Ablation Study

In this section, we mainly conduct various ablation studies on the PACS dataset with ResNet50 backbone.

Effectiveness of key components. Table 5 shows the contribution of each term mentioned in our method. Compared with the ERM baseline, self-distillation improves the baseline 3.2%. This shows that our method could efficiently capture historical temporal features to maintain uniformity. The entropy filter and consistency filter improve the performance with 0.6% and 0.5%, respectively, which suggests that they mitigate the problem of noise pseudo labels. Finally, our proposed memorized spatial local clustering brings a 0.5% gain which suggests that the module also helps with the revision of representations.

Sensitivity to hyperparameter. Our method has three hyperparameters: top- M features for entropy filtering to compute prototypes, the number of nearest neighbors K in memorized spatial local clustering and the trade-off pa-

Table 5. Ablation Study. SD means self-training using self-distillation (*cf.* Eq 3) without any filter. EF and CF mean entropy filter and consistency filter. MSLC means memorized spatial local clustering loss (*cf.* Eq. 6).

#	SD	EF	CF	MSLC	Acc(% \uparrow)
0					84.59
1	✓				87.80 (+3.21)
2	✓	✓			88.40 (+3.81)
3	✓	✓	✓		88.92 (+4.33)
4	✓	✓	✓	✓	89.41 (+4.82)




Figure 3. Sensitivity analysis about the number of nearest neighbors K in memorized spatial local clustering, entropy filter hyperparameter M and trade-off parameter λ (*cf.* Eq. 7).

parameter λ in Eq. 7. Fig. 3 presents the ablation study on hyperparameters. First, for the number of nearest neighbors K , when $K \in \{1, 3, 5\}$ achieves good performance. However, though it still improves the ERM baseline, $K \in \{10, 15, 20\}$ leads to performance degradation. We conjecture that using too many neighbors may increase the probability of introducing misclassified features. Second, for the entropy filter hyperparameter M , we found that the proposed method is robust to the choice of M . The performance improves stably with the M increasing from 1 to 100. If we mute the entropy filter, our method still achieves a competitive performance. Third, for trade-off parameter λ , $\lambda \in [0.1, 0.5]$ achieves good performance and $\lambda = 0.1$ achieves the best performance.

4.4. Analysis

Scalability on different models. We validated our method on various backbones to verify that our method does not rely on any specific network structure design. Table 6 shows that the proposed method generally improves baseline performance regardless of different backbones. For example, for ViT-B/16, the proposed method improves ERM [61] 3.1%, 2.7%, 1.2% for different datasets, respectively. This shows that our method is “plug and play” for different deep neural networks which is important for applications in the real world.

Qualitative analysis by tSNE. We provide tSNE [60] visualizations of the ERM baseline and our method, as shown in Fig. 4. The learned features of the ERM base-

Table 6. Results with different backbones. **Bold** type indicates performance improvement compared with baseline. All baseline models are trained in a standard ERM manner.

Backbones	PACS	OfficeHome	VLCS
ResNet18 [17]	82.07	63.12	72.75
+Ours	87.32	64.83	73.61
ResNet50 [17]	84.59	67.37	74.01
+Ours	89.41	68.67	74.52
ResNeXt-50 [67]	86.67	72.66	78.50
+Ours	91.33	74.18	79.38
ViT-B/16 [11]	87.13	79.06	78.70
+Ours	90.20	81.80	79.90
EfficientNet-B4 [56]	85.11	74.65	77.14
+Ours	85.41	72.24	79.42
Mixer-L16 [57]	84.59	71.36	76.53
+Ours	88.47	74.82	79.75




Figure 4. tSNE [60] visualization of learned feature embeddings. (a) ERM baseline without adaptation. (b) After adaptation with our method. Different colours denote different classes.

line on the target domain are not well separated due to the large domain gap, as shown in Fig. 4a. After adaptation to the target domain, our method could generate more uniform and alignment features, as shown in Fig. 4b.

Efficiency analysis. Our method adapts a pre-trained model and needs to maintain a memory bank, which may lead to extra computational costs. The usage of GPU memory is mainly influenced by the batch size at test time. We present computational complexity analysis in Fig. 5. The default batch size of our method is 128 and the reported accuracy is 89.4% with an almost 14 GB GPU memory cost. While our method still achieves good performance (89.1% accuracy) when we set the batch size to 64, the GPU memory cost decreases to 7.85 GB and the running time is the fastest for all. When we apply our method to the real world applications, there is a trade-off between accuracy, GPU memory, and running time and we need to balance the choices due to different requirements.

Another strategy for reducing computational complexity




Figure 5. Computational complexity with different batch sizes at test time adaptation. From left to right: accuracy, GPU memory and running time on four domains of PACS [30] dataset.

is that we only update affine (*i.e.*, scale and shift) parameters when we adapt the model, like [64]. We conducted the experiments of updating affine parameters in batch normalization layers, and the result is 89% accuracy on the PACS dataset. Compared with 89.4% which is achieved by updating all parameters, the accuracy only decreases by 0.4% and the GPU memory cost drops by 2 GB. This shows that, with our method, only adapting affine parameters in batch normalization layers can achieve competitive performance with a salience GPU memory decrease.

4.5. Scalability to Medical Image Segmentation

Furthermore, we evaluated our algorithm on several medical image segmentation datasets. We conduct the experiments on three tasks: 1) prostate segmentation. We choose Promise12 [35] dataset, including 50 cases as source data and the MSD05 dataset (32 cases) from Medical Segmentation Decathlon [1] as test data. 2) cardiac structure segmentation. We choose ACDC [4] dataset (200 cases) as source data and LGE images (45 cases) from Multi-sequence Cardiac MR Segmentation Challenge [76, 77] as test data. The task is to segment the left ventricular blood pool, right ventricular blood pool, and myocardium. 3) optic disc and cup segmentation. We choose the training set of REFUGE challenge (400 images) [45] as the source domain, and RIM-ONE-r3 (159 images) [13] and Drishti-GS (101 images) [53] as two different target domains. We aim to segment the optic disc and optic cup. For RIM-ONE-r3 and Drishti-GS, we split them into 99/60 and 51/50 for training and test. For other datasets, we randomly split them to 80% for training and 20% for the test.

Implementation. We adopt DeepLabv3+ [9] with MobileNetV2 [50] backbone as segmentation models. In training and test phases, we use Adam [24] optimizer with fixed learning rate as $1e^{-4}$ without specific choice. We enlarge K (*cf.* Eq. 6) from 3 to 64 in the segmentation task. We use Dice Score [40] as the evaluation metric and we reported the average results of all the target structures in Table 7.

Results. We utilize ERM [61] and Tent [64] as the comparison methods. Table 7 shows that our method generally improves the baseline on different tasks. Specifically, we improve ERM baseline 3.9%, 4.44%, 3.57%, and 12.72%

Table 7. Results on medical image segmentation. We highlight the **best** result for each column.

Source	MSD05	ACDC	REFUGE	
Target	Promise12	LGE	DrishtiGS	RIM-ONE-r3
ERM [61]	83.00	82.34	81.05	68.17
Tent [64]	85.64	81.11	83.96	77.87
Ours	86.90	86.78	84.62	80.89

for the four tasks. Experiments prove that our method can also work well in image segmentation tasks.

5. Limitations

Our work has some limitations. First, even though we have already validated our method on several classification and segmentation tasks, we have not extended our method to low-level tasks, such as image dehazing, denoising, and super-resolution tasks. As for low-level tasks, there is no conception of semantic prototype or entropy which can lead to the failure of our method. The second limitation is *hyperparameter search*. In this paper, we do hyperparameter search on the training domain validation set. However, the best choice on the training domains does not mean the best performance on the test domain [48]. We may need to try more strategies to conduct experiments in the TTA setting like test domain validation.

6. Conclusion

In this work, we focus on the test time adaptation and propose a new perspective that test time adaptation could be viewed as a feature revision problem. Furthermore, feature revision includes two parts: test time feature uniformity and test time feature alignment. As for the feature uniformity perspective, we propose *Test Time Self-Distillation* to make the target feature as uniform as possible in the adaptation. To align features from the same class, we propose *Memorized Spatial Local Clustering* to encourage that the distance between feature representations in the latent space should align with the pseudo logits. A mass of experiments proves that our method not only generally improves the ERM baseline but also outperforms existing TTA or SFDA methods on four domain generalization benchmarks. Furthermore, our method could be adopted to different backbones. To broaden our model for real world applications, we then validate our method on four cross-domain medical image segmentation tasks. Experiment results show that our method is effect, flexible and scalable.

Acknowledgment

This work was supported by the National Natural Science Foundation of China, 81971604 and the Grant from the Tsinghua Precision Medicine Foundation, 10001020104.

References

- [1] Michela Antonelli, Annika Reinke, Spyridon Bakas, Keyvan Farahani, et al. The medical segmentation decathlon. *Nature Communications*, 13(1), July 2022. 2, 8
- [2] Yogesh Balaji, Swami Sankaranarayanan, and Rama Chelappa. Metareg: Towards domain generalization using meta-regularization. In *NeurIPS*, 2018. 3
- [3] Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15(6):1373–1396, June 2003. 4
- [4] Olivier Bernard, Alain Lalande, Clement Zotti, et al. Deep learning techniques for automatic MRI cardiac multi-structures segmentation and diagnosis: Is the problem solved? *IEEE Transactions on Medical Imaging*, 37(11):2514–2525, Nov. 2018. 2, 8
- [5] Malik Boudiaf, Romain Müller, Ismail Ben Ayed, and Luca Bertinetto. Parameter-free online test-time adaptation. In *CVPR*, 2022. 1, 3, 5
- [6] Fabio Maria Carlucci, Antonio D’Innocente, Silvia Bucci, Barbara Caputo, and Tatiana Tommasi. Domain generalization by solving jigsaw puzzles. In *CVPR*, 2019. 3
- [7] Junbum Cha, Sanghyuk Chun, Kyungjae Lee, Han-Cheol Cho, Seunghyun Park, Yunsung Lee, and Sungrae Park. SWAD: domain generalization by seeking flat minima. In *NeurIPS*, 2021. 6
- [8] Dian Chen, Dequan Wang, Trevor Darrell, and Sayna Ebrahimi. Contrastive test-time adaptation. In *CVPR*, 2022. 1, 3
- [9] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, 2018. 8
- [10] Xu Chu, Yujie Jin, Wenwu Zhu, Yasha Wang, Xin Wang, Shanghang Zhang, and Hong Mei. DNA: domain generalization with diversified neural averaging. In *ICML*, 2022. 6
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 5, 7
- [12] Qi Dou, Daniel Coelho de Castro, Konstantinos Kamnitsas, and Ben Glocker. Domain generalization via model-agnostic learning of semantic features. In *NeurIPS*, 2019. 3
- [13] F. Fumero, S. Alayon, J. L. Sanchez, J. Sigut, and M. Gonzalez-Hernandez. RIM-ONE: An open retinal image database for optic nerve evaluation. In *2011 24th International Symposium on Computer-Based Medical Systems (CBMS)*. IEEE, June 2011. 2, 8
- [14] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario March, and Victor Lempitsky. Domain-adversarial training of neural networks. *Journal of Machine Learning Research*, 17(59):1–35, 2016. 3
- [15] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Un-supervised representation learning by predicting image rotations. In *ICLR*, 2018. 1, 3
- [16] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. In *ICLR*, 2021. 5
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 1, 5, 7
- [18] Dan Hendrycks, Norman Mu, Ekin Dogus Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. Augmix: A simple data processing method to improve robustness and uncertainty. In *ICLR*, 2020. 3
- [19] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 4
- [20] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015. 3, 5
- [21] Yusuke Iwasawa and Yutaka Matsuo. Test-time classifier adjustment module for model-agnostic domain generalization. In *NeurIPS*, 2021. 1, 2, 3, 5
- [22] Daehee Kim, Youngjun Yoo, Seunghyun Park, Jinkyu Kim, and Jaekoo Lee. Selfreg: Self-supervised contrastive regularization for domain generalization. In *ICCV*, 2021. 3
- [23] Kyungyul Kim, Byeongmoon Ji, Doyoung Yoon, and Sangheum Hwang. Self-knowledge distillation with progressive refinement of targets. In *ICCV*, 2021. 2
- [24] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 5, 8
- [25] Takeshi Kojima, Yutaka Matsuo, and Yusuke Iwasawa. Robustifying vision transformer without retraining from scratch by test-time class-conditional feature alignment. In *IJCAI*, 2022. 1, 3
- [26] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, 2012. 1
- [27] Jogendra Nath Kundu, Akshay R. Kulkarni, Suvaansh Bhambri, Deepesh Mehta, Shreyas Anand Kulkarni, Varun Jampani, and Venkatesh Babu Radhakrishnan. Balancing discriminability and transferability for source-free domain adaptation. In *ICML*, 2022. 6
- [28] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, 2013. 5
- [29] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M. Hospedales. Learning to generalize: Meta-learning for domain generalization. In *AAAI*, 2018. 3
- [30] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generalization. In *ICCV*, 2017. 2, 5, 8
- [31] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C. Kot. Domain generalization with adversarial feature learning. In *CVPR*, 2018. 3
- [32] Xiaotong Li, Yongxing Dai, Yixiao Ge, Jun Liu, Ying Shan, and Lingyu Duan. Uncertainty modeling for out-of-distribution generalization. In *ICLR*, 2022. 3

- [33] Yanghao Li, Naiyan Wang, Jianping Shi, Xiaodi Hou, and Jiaying Liu. Adaptive batch normalization for practical domain adaptation. *Pattern Recognition*, 80:109–117, Aug. 2018. **2, 3**
- [34] Jian Liang, Dapeng Hu, and Jiashi Feng. Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In *ICML*, 2020. **2, 3, 5**
- [35] Geert Litjens, Robert Toth, Wendy van de Ven, Caroline Hoeks, Sjoerd Kerkstra, Bram van Ginneken, Graham Vincent, Gwenael Guillard, Neil Birbeck, Jindang Zhang, et al. Evaluation of prostate segmentation algorithms for mri: the promise12 challenge. *Medical image analysis*, 18(2):359–373, 2014. **2, 8**
- [36] Yuejiang Liu, Parth Kothari, Bastien van Delft, Baptiste Bellot-Gurlet, Taylor Mordan, and Alexandre Alahi. TTT++: when does self-supervised test-time training fail or thrive? In *NeurIPS*, 2021. **1, 3**
- [37] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. **1**
- [38] Fangrui Lv, Jian Liang, Shuang Li, Bin Zang, Chi Harold Liu, Ziteng Wang, and Di Liu. Causality inspired representation learning for domain generalization. In *CVPR*, 2022. **3**
- [39] Divyat Mahajan, Shruti Tople, and Amit Sharma. Domain generalization using causal matching. In *ICML*, 2021. **3**
- [40] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *3DV*, 2016. **8**
- [41] Saeid Motiian, Marco Piccirilli, Donald A. Adjeroh, and Gianfranco Doretto. Unified deep supervised domain adaptation and generalization. In *ICCV*, 2017. **3**
- [42] Zachary Nado, Shreyas Padhy, D Sculley, Alexander D’Amour, Balaji Lakshminarayanan, and Jasper Snoek. Evaluating prediction-time batch normalization for robustness under covariate shift. *arXiv preprint arXiv:2006.10963*, 2020. **2, 3**
- [43] Shuaicheng Niu, Jiayang Wu, Yifan Zhang, Yaofu Chen, Shijian Zheng, Peilin Zhao, and Mingkui Tan. Efficient test-time model adaptation without forgetting. In *ICML*, 2022. **1, 2, 3, 5**
- [44] Oren Nuriel, Sagie Benaim, and Lior Wolf. Permuted adain: Reducing the bias towards global statistics in image classification. In *CVPR*, 2021. **3**
- [45] José Ignacio Orlando, Huazhu Fu, João Barbosa Breda, et al. REFUGE challenge: A unified framework for evaluating automated methods for glaucoma assessment from fundus photographs. *Medical Image Analysis*, 59:101570, Jan. 2020. **2, 8**
- [46] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Z. Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, 2019. **5**
- [47] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *ICCV*, 2019. **2, 5**
- [48] Alexandre Rame, Corentin Dancette, and Matthieu Cord. Fishr: Invariant gradient variances for out-of-distribution generalization. In *International Conference on Machine Learning*, pages 18347–18377. PMLR, 2022. **8**
- [49] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, Apr. 2015. **5**
- [50] Mark Sandler, Andrew G. Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *CVPR*, 2018. **8**
- [51] Steffen Schneider, Evgenia Rusak, Luisa Eck, Oliver Bringmann, Wieland Brendel, and Matthias Bethge. Improving robustness against common corruptions by covariate shift adaptation. In *NeurIPS*, 2020. **2, 3, 5**
- [52] Claude E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948. **3**
- [53] Jayanthi Sivaswamy, S Krishnadas, Arunava Chakravarty, G Joshi, A Syed Tabish, et al. A comprehensive retinal image dataset for the assessment of glaucoma from the optic nerve head analysis. *JSM Biomedical Imaging Data Papers*, 2(1):1004, 2015. **2, 8**
- [54] Baochen Sun and Kate Saenko. Deep CORAL: correlation alignment for deep domain adaptation. In *ECCV*, 2016. **3**
- [55] 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 *ICML*, 2020. **1, 3**
- [56] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, 2019. **5, 7**
- [57] Ilya O. Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Steiner, Daniel Keysers, Jakob Uszkoreit, Mario Lucic, and Alexey Dosovitskiy. Mlp-mixer: An all-mlp architecture for vision. In *NeurIPS*, 2021. **5, 7**
- [58] Antonio Torralba and Alexei A Efros. Unbiased look at dataset bias. In *CVPR*, 2011. **2, 5**
- [59] Eric Tzeng, Judy Hoffman, Ning Zhang, Kate Saenko, and Trevor Darrell. Deep domain confusion: Maximizing for domain invariance. *arXiv preprint arXiv:1412.3474*, 2014. **3**
- [60] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(86):2579–2605, 2008. **7**
- [61] Vladimir Vapnik. *Statistical learning theory*. Wiley, 1998. **5, 6, 7, 8**
- [62] Hemant Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *CVPR*, 2017. **2, 5**
- [63] Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, David Lopez-Paz, and Yoshua Ben-

- gio. Manifold mixup: Better representations by interpolating hidden states. In *ICML*, 2019. 3
- [64] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In *ICLR*, 2021. 1, 2, 3, 5, 8
- [65] Qin Wang, Olga Fink, Luc Van Gool, and Dengxin Dai. Continual test-time domain adaptation. In *CVPR*, 2022. 1, 3
- [66] Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *ICML*, 2020. 1
- [67] Saining Xie, Ross B. Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *CVPR*, 2017. 5, 7
- [68] Xufeng Yao, Yang Bai, Xinyun Zhang, Yuechen Zhang, Qi Sun, Ran Chen, Ruiyu Li, and Bei Yu. PCL: proxy-based contrastive learning for domain generalization. In *CVPR*, 2022. 3, 6
- [69] Sangdoon Yun, Dongyoon Han, Sanghyuk Chun, Seong Joon Oh, Youngjoon Yoo, and Junsuk Choe. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *ICCV*, 2019. 3
- [70] Daoan Zhang, Mingkai Chen, Chenming Li, Lingyun Huang, and Jianguo Zhang. Aggregation of disentanglement: Reconsidering domain variations in domain generalization. *arXiv preprint arXiv:2302.02350*, 2023. 3
- [71] Daoan Zhang, Chenming Li, Haoquan Li, Wenjian Huang, Lingyun Huang, and Jianguo Zhang. Rethinking alignment and uniformity in unsupervised image semantic segmentation. *arXiv preprint arXiv:2211.14513*, 2022. 1
- [72] Hongyi Zhang, Moustapha Cissé, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *ICLR*, 2018. 3
- [73] Linfeng Zhang, Jiebo Song, Anni Gao, Jingwei Chen, Chenglong Bao, and Kaisheng Ma. Be your own teacher: Improve the performance of convolutional neural networks via self distillation. In *ICCV*, 2019. 2
- [74] Marvin Mengxin Zhang, Sergey Levine, and Chelsea Finn. MEMO: Test time robustness via adaptation and augmentation. In *Advances in Neural Information Processing Systems*, 2022. 1, 2
- [75] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain generalization with mixstyle. In *ICLR*, 2021. 3
- [76] Xiahai Zhuang. Multivariate mixture model for cardiac segmentation from multi-sequence MRI. In *MICCAI*. 2016. 2, 8
- [77] Xiahai Zhuang. Multivariate mixture model for myocardial segmentation combining multi-source images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(12):2933–2946, Dec. 2019. 2, 8