

Domain-Agnostic Feature Modulation for Semi-Supervised Domain Generalization

Venuri Amarasinghe, Kalinga Bandara, Isun Randila, Asini Jayakody
University of Moratuwa

{amarasingheamvm.20, bandarammkr.20, kandedegarapmirb.20, jayakodyjaau.20}@uom.lk

Chamuditha Jayanga Galappaththige
Queensland University of Technology

chamuditha.galappaththige@qut.edu.au

Ranga Rodrigo
University of Moratuwa

ranga@uom.lk

Abstract

Semi-supervised domain generalization (SSDG) leverages a small fraction of labeled data alongside unlabeled data to enhance model generalization. Most of the existing SSDG methods rely on pseudo-labeling (PL) for unlabeled data, often assuming access to domain labels—a privilege not always available. However, domain shifts introduce domain noise, leading to inconsistent PLs that degrade model performance. Methods derived from semi-supervised learning baselines suffer particularly from lower PL accuracy, reducing the effectiveness of unlabeled data. We address this in a more challenging setting—domain-label agnostic SSDG—where domain labels for unlabeled data are not available during training. First, we propose a feature modulation strategy that enhances class-discriminative features while suppressing domain-specific information. This modulation shifts features toward Similar Average Representations—a modified version of class prototypes—that are robust across domains, encouraging the classifier to distinguish between closely related classes and feature extractor to form tightly clustered, domain-invariant representations. Second, to mitigate domain noise and improve pseudo-label accuracy, we introduce a loss-scaling function that dynamically lowers the fixed confidence threshold for pseudo-labels, optimizing the use of unlabeled data. With these key innovations, our approach achieves significant improvements on four major domain generalization benchmarks—even without domain labels. We will make the code available.

1. Introduction

Domain shift—the discrepancy between the distributions of training and testing data, violating the independent and

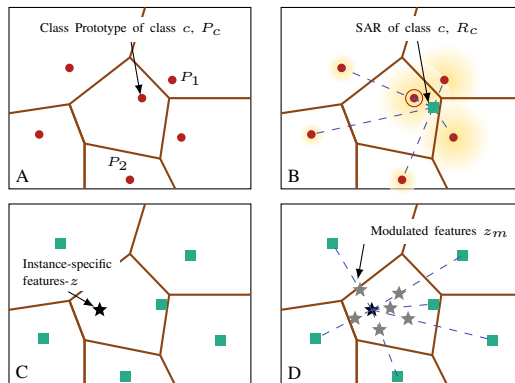


Figure 1. Feature modulation towards SARs in the classwise separated feature space. Polygon lines denote class boundaries. **A:** Class prototypes P_c in the feature space. **B:** SAR of class C (Sec. 3.1.1) generated by averaging class prototypes based on similarity. The glow intensity represents the similarity. **C:** Instance-specific features z in the feature space with SARs of all classes. **D:** Feature modulation shifts z more towards its true class.

identically distributed (i.i.d.) assumption—is a fundamental challenge in deep learning models [8, 16]. Domain Generalization (DG) [15, 20, 36, 41] and Domain Adaptation [13, 24, 39] have emerged as two key paradigms to address this issue. DG aims to develop models that generalize to unseen domains using multiple labeled source domains, whereas domain adaptation leverages knowledge from both source and target domains to improve test performance.

To enhance generalization across domains, DG methods typically focus on learning domain-invariant representations using techniques such as invariant risk minimization [2], meta-learning [9, 21, 26], data augmentation [31, 32, 36, 42], and adversarial feature learning [22, 23, 31]. However, these approaches rely heavily on labeled data from multiple source domains, which is often scarce in real-world settings. Semi-supervised learning (SSL) presents a viable




Figure 2. Average accuracy, keep rate, and pseudo-label (PL) accuracy. We achieve higher average accuracy as a result of the increased PL accuracy and higher keep rate. Note that different y -axis ranges were chosen to enhance the visualization.

solution to address labeled data scarcity by leveraging abundant unlabeled data alongside a limited set of labeled instances [4, 27, 28]. Recent SSL methods [5, 27, 34, 43] predominantly utilize pseudo-labeling and consistency regularization. Pseudo-labeling [19] assigns artificial labels to unlabeled data with high prediction confidence, while consistency regularization ensures that perturbed versions of the same data receive consistent predictions. Although DG methods struggle when a majority of training data is unlabeled, SSL methods naturally exploit unlabeled data. However, their effectiveness is hindered by noisy pseudo-labels and the inability to extract robust domain-invariant features, making them inferior to fully supervised DG methods.

Semi-supervised domain generalization (SSDG) [11, 33, 40, 45] aims to bridge this gap by combining DG and SSL to improve generalization to unseen domains while mitigating the domain shift, poor data utilization, and pseudo-label noise inherent in SSL. Recent SSDG methods introduce multi-view consistency learning, stochastic classifiers [45], feature-based conformity, and semantic alignment constraints [11]. However, these methods rely on domain labels, which are difficult to obtain in practice. Moreover, their utilization of unlabeled data is suboptimal in early training stages, as the number of valid pseudo-labels passing the fixed confidence threshold is low. Simply lowering this threshold to incorporate more unlabeled data results in increased pseudo-label noise and inaccurate predictions which result in inferior DG performances. This presents a fundamental *dilemma*—balancing accurate pseudo-labels with higher unlabeled data utilization.

To address this challenge, we propose a novel SSDG method that enhances pseudo-label accuracy and improves unlabeled data utilization under a more challenging setting where domain labels are unavailable during training. Our approach is motivated by the intuition that class representations in each domain comprise two signals: a class-specific signal, and a domain-related signal. We aim to isolate and emphasize the class-specific signal while minimizing domain information propagation to the classifier through a proposed **feature modulator**.

We also notice that some features are common to more

than one class. For example, dogs and cats have many similar features. Our method encourages the classifier to focus on class-specific features while minimizing dependence on features common to multiple classes. We adopt similar average representations SAR—a modified version of class prototypes—to reduce the effect of domain-specific features, ensuring the classifier focuses on class-specific information and improves domain generalization. The feature extractor and the classifier are trained while shifting extracted features toward SARs using the feature modulator (Fig. 1). This approach helps the classifier distinguish similar classes more accurately while guiding the feature extractor to form clearer class boundaries. As a result, it improves pseudo-label accuracy by **3.3%** over the baseline (Fig. 2).

The intervention of the feature modulator for better pseudo-labels allows us to incorporate more unlabeled data for model training when coupled with our proposed loss scaling function that dynamically adjusts the effect of pseudo-labels. Instead of naively lowering the threshold—introducing label noise—we scale loss gradients based on pseudo-label confidence and uncertainty scores. This strategy increases the portion of unlabeled data used for training by **4.1%** compared to the baseline (Fig. 2). The improved data utilization with accurate pseudo labels improves the DG capability of the learned model. In summary, our contributions are:

- We propose a feature modulator that enhances class-specific feature representations, leading to more accurate pseudo-labels without relying on domain labels.
- We introduce a loss scaling function that dynamically lowers the effect of low confidence pseudo-labels while reducing label noise, improving unlabeled data utilization.
- We address the SSDG problem under more challenging setting where domain labels are unavailable and demonstrate the effectiveness of our method on four major domain generalization benchmarks.

2. Related Work

Domain Generalization: With the objective of creating models that can generalize across unseen domains, the DG methods can be subdivided into data augmentation [31, 36], feature augmentation [46] and meta-learning [9, 21, 26] based approaches. Data augmentation [31, 36] based DG models generate diverse training samples to improve cross-domain generalization. Mixing instance-level feature statistics in [46], can be identified as a feature augmentation method for DG. The primary objective of these data and feature augmentation methods is to learn domain-invariant features. Subsequent works apply meta-learning to DG by designing tasks and losses to generalize across domains [9, 21, 26]. However, the performance of these fully super-

vised DG methods relies on the fully labeled data in multiple source domains, limiting their effectiveness in practical semi-supervised settings.

Semi-Supervised Learning: Semi-Supervised Learning (SSL) addresses the problem of scarcity of labeled data by using a large amount of unlabeled data with a few labeled data. In recent literature, consistency regularization [1, 18], entropy minimization [14] and pseudo-labeling-based approaches show the most promising results in SSL. Consistency regularization forces the model to give the same prediction to different perturbed versions [4, 27, 35] of the same input. Pseudo-labeling assigns labels to unlabeled data, selecting only high-confidence predictions. MixMatch [4] unifies existing data augmentations, pseudo-labeling, and mixup [38] to achieve both consistency regularization and entropy minimization. RemixMatch [3] shows that introducing strong augmentations boosts the performance. FixMatch [27] is a strong SSL baseline for SSDG which combines pseudo-labeling and consistency regularization. FreeMatch [34] extends FixMatch [27] showing that adaptive thresholds perform better than fixed thresholds for pseudo-labeling. FlexMatch [37] improves FixMatch to utilize more unlabeled data by introducing curriculum pseudo labeling, while FullMatch [6] uses entropy meaning loss (EML) and adaptive negative learning (ANL) for a better leveraging of unlabeled data. HyperMatch [43] separates pseudo-labels into noisy and clean pseudo-labels and leverages both of them to learn better-clustered feature representations. However, efforts to improve pseudo-label quality often ignore the number of low-confidence samples discarded, leading to poor use of unlabeled data [27, 34, 43]. Conversely, some SSL methods [6, 37] prioritize data utilization but compromise pseudo-label accuracy.

Semi-Supervised Domain Generalization: SSL and semi-supervised domain generalization (SSDG) leverage unlabeled data during model training while SSDG faces a significant challenge because the data comes from different source domains. SSDG is a relatively under-explored research area, yet it presents a more realistic real-world scenario. StyleMatch [45] extends FixMatch [27] for SSDG by incorporating stochastic modeling to reduce over-fitting in scarce labels and multi-view consistency learning to enhance DG. The feature-based conformity and semantic alignment loss functions introduced in [11] improve pseudo-label accuracy, leveraging the domain labels as well. MultiMatch [25] further extends FixMatch [27] to a multi-task learning framework, enhancing the pseudo-label quality by treating individual domains as separate tasks and introducing a global task to generalize to unseen domains. In [40], the authors study SSDG for a new paradigm, known and unknown classes, while authors in [33] address a different SSDG setting where only one source domain is fully labeled. Their approach involves a joint domain-aware label

and dual-classifier system, which improves pseudo-labels by using a separate classifier for pseudo-labeling and combining its output probabilities with similarity to its class representation from a domain-aware class representation bank. The authors in [12] propose a method to generate accurate pseudo-labels by domain-level information vectors and utilize them to learn a soft domain-aware mask to modulate the domain-shared classifier weights. However, the methods in [11, 12] require domain labels to generate domain-aware class prototypes and domain-level information vectors respectively, imposing an additional constraint on the unlabeled data. Existing works train the model using multiple domain data, where the classifier learns some domain information of source domains as well. In our work, we introduce our novel method, feature modulation which improves generalization by de-emphasizing the features containing domain information while keeping the class-specific information.

3. Methodology

Problem Setting: We adopt notation from [11]. Let \mathcal{X} and \mathcal{Y} denote the input and label space, respectively, where $y \in \mathcal{Y} = \{1, 2, \dots, C\}$ is the class label. A joint distribution $\mathcal{P}(\mathcal{X}, \mathcal{Y})$ over $\mathcal{X} \times \mathcal{Y}$ is defined as a domain d . We only consider the distribution shift in $\mathcal{P}(\mathcal{X})$ while $\mathcal{P}(\mathcal{Y})$ remains the same. Following standard SSDG protocol [11, 12, 45], we use multiple source domains during training and an unseen domain as the target. We denote $d^L = \{(x_i^d, y_i^d)\}_{i=1}^{n_L}$ as the labeled part and $d^U = \{(u_i^d)\}_{i=1}^{n_U}$ as the unlabeled part of data for each source domain $d \in \{d_1, d_2, \dots, d_{tr}\}$, where $x_i^d, u_i^d \in R^{(3 \times H \times W)}$ are labeled and unlabeled input images. The unlabeled part contains more data than the labeled part i.e., $n_U \gg n_L$. Our goal is to use the training (source) domains to learn a domain-generalizable model F that can accurately predict on out-of-distribution data, whose examples are drawn from $\mathcal{P}(\mathcal{X}^{d_{te}}, \mathcal{Y}^{d_{te}})$, where d_{te} denotes the target domain. Our model $F = \mathcal{C}(\mathcal{M}(f))$ comprises of feature extractor f , feature modulator \mathcal{M} , and classifier \mathcal{C} .

Our Motivation: We observe that the feature spaces of certain classes can be closely aligned due to shared domain characteristics or overlapping class features (see Fig. 1, A). Some class prototypes, such as \mathbf{P}_1 and \mathbf{P}_c , are closer to each other than to other classes, making classification more challenging. A higher degree of class separation in the feature space leads to a more accurate classification. Our goal is to enhance generalization by: 1) training the feature extractor to effectively separate classes, regardless of domain differences, and 2) training the classifier to focus on unique class-specific features while disregarding domain-related and shared class information.

To achieve this, we propose feature modulation, a strategy that blends features from similar classes with the orig-

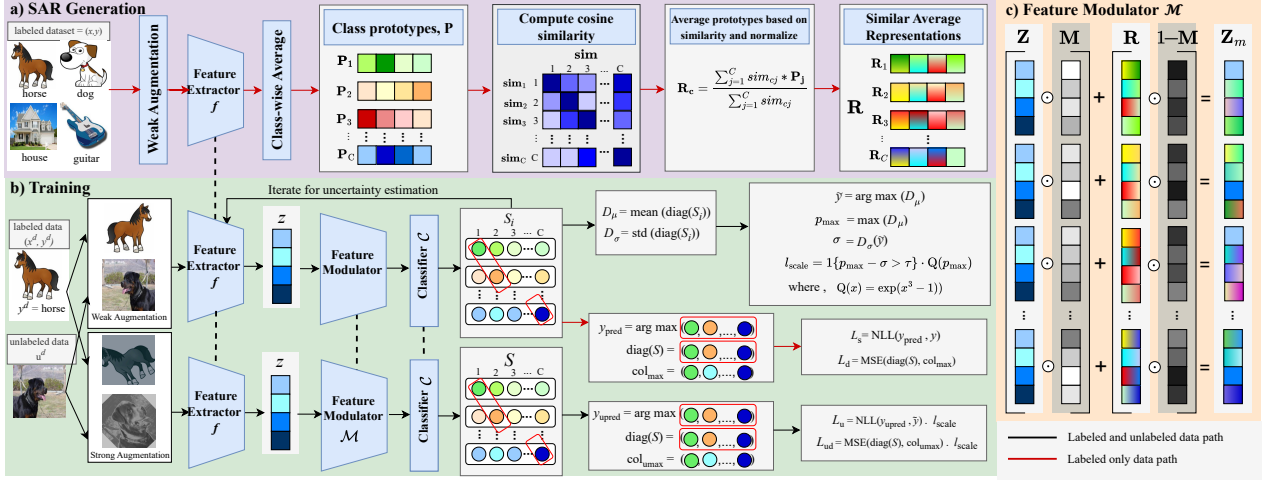


Figure 3. Proposed Architecture: a) Similar Average Representation (SAR (Sec. 3.1.1)) generation. b) Complete model during training. c) Feature Modulator (Sec. 3.1.2) combines instance-specific features (z) and SAR (\mathbf{R}) to soft-eliminate domain noise using our modulation strategy. This results in modulated features \mathbf{Z}_m . We input this \mathbf{Z}_m to the shared classifier \mathcal{C} and get the probabilities for each class. With these probabilities, we compute the MSE and NLL losses for both labeled and unlabeled data. For unsupervised losses, we consider loss scaling (Sec. 3.1.4) to maximize unlabeled data utilization. Note that our method is agnostic of domain labels.

inally extracted features, exposing the classifier to more challenging classification scenarios where class boundaries become less distinct. For instance, cats and dogs share several visual features, causing their feature spaces to be closely aligned, which can lead to misclassification. To address this, we introduce a modulation process where a portion of the features from one class (e.g., a cat) is incorporated into another class (e.g., a dog) while maintaining the original class label. This modulation is guided by a learnable matrix and the cosine similarity between class prototypes. By shifting features toward their SARs (Sec. 3.1.1), the classifier learns to focus on distinct class characteristics while disregarding shared features and domain information. Training on these feature-modulated samples forces the classifier to become more robust in distinguishing between visually similar classes, ultimately improving classification performance and generalization.

3.1. Feature Modulation for SSDG

An overview of our method is shown in Fig. 3. Our approach consists of four key components: Similar Average Representations (SAR) \mathbf{R} (see Fig. 3.a), the feature modulator \mathcal{M} (see Fig. 3.c), uncertainty estimation of pseudo-labels, and the unlabeled loss scaling function l_{scale} (see Fig. 3). First, we compute SAR (shown by the green square in Fig. 1.B) for each class using features extracted from labeled data, as described in Sec. 3.1.1. The feature modulator \mathcal{M} serves as an intermediary between the feature extractor and the classifier (see Fig. 3.b), reducing the propagation of domain-specific information from the feature extractor to the classifier. The modulating matrix \mathbf{M} is a

learnable (i.e., dynamically updating) matrix that updates when minimizing the training loss. The modulation process adjusts the extracted instance-specific features z to enhance class-discriminative characteristics while suppressing domain-specific variations (modulated features z_m are shown by the gray starts in Fig. 1, D). This is achieved using SAR along with a learnable feature modulation matrix \mathbf{M} , which guides the transformation to ensure improved generalization across domains.

We then compute pseudo-labels for unlabeled data using the modulated features of weakly augmented unlabeled samples, as described in Sec. 3.1.3. With improved pseudo-label accuracy resulting from reduced domain interference, we employ the loss scaling function l_{scale} Eq. (7) to leverage less confident pseudo-labels. To further refine pseudo-label selection, we incorporate uncertainty estimation (Sec. 3.1.4) using the Monte Carlo (MC) dropout method [10]. This enhances the impact of pseudo-labels with low uncertainty while mitigating the effect of unreliable predictions.

Finally, we apply the feature modulator \mathcal{M} during inference to minimize domain-specific information in test data, ensuring better generalization across unseen domains.

3.1.1. Similar Average Representations (SAR)

To train the classifier to distinguish between closely located classes regardless of domain differences, we introduce SAR as a representative of such classes in the feature space. Recent works [11, 33] leverage domain-aware class prototypes to mitigate misclassification during pseudo-labeling, which often arises due to domain shifts. In contrast, our approach extends beyond pseudo-labeling by utilizing class prototypes to suppress domain-specific information during both

training and testing.

We model each source image as a combination of class-specific information—essential for classification—and domain-specific information, such as background, colors, and texture. For instance, images from the photo domain contain rich color and texture details, whereas the sketch domain provides a more abstract representation with minimal color and texture. By averaging features across multiple domains, we reduce the influence of domain-specific attributes. Recent work [12] shows that averaging feature vectors leads to a more compact representative class embedding compared to alternative methods.

Unlike existing methods that compute domain-aware class prototypes based on domain labels, we define class prototypes \mathbf{P} solely based on class labels, making them domain-agnostic. Given that $d^L = \{(x^d, y^d)\}$ denotes the set of labeled data, we denote features K_c of all labeled samples belonging to class c across all domains as follows:

$$K_c = \{f(x_i) \mid y_i = c\}_{i=1}^{|d^L|} \quad (1)$$

where f is the feature extractor. The class prototype \mathbf{P}_c , of each class $c \in \{1, 2, \dots, C\}$, is obtained by averaging K_c :

$$\mathbf{P}_c = \frac{1}{|K_c|} \sum_{i=1}^{|K_c|} K_c[i]. \quad (2)$$

Next, we compute the cosine similarity between class prototypes to obtain a similarity matrix \mathbf{Sim} . Using this matrix, we derive the SAR for each class c , denoted as \mathbf{R}_c . This is achieved by weighing each prototype \mathbf{P}_j according to its similarity Sim_{cj} (see Fig.3.a) with class c , ensuring that SAR is more influenced by most similar classes, where \mathbf{P}_j denotes the class prototype of class j :

$$\mathbf{R}_c = \frac{\sum_{j=1}^C \text{Sim}_{cj} \cdot \mathbf{P}_j}{\sum_{j=1}^C \text{Sim}_{cj}}. \quad (3)$$

We use SAR in the feature modulation process (see Fig. 3), where instance-specific features z are adjusted towards their respective similar class prototypes (Fig. 1). During training, as the classifier improves in separating classes, inter-class similarity naturally decreases. Consequently, the contribution of other class prototypes to SAR diminishes, leading to a gradual convergence of \mathbf{R}_c towards \mathbf{P}_c .

3.1.2. Feature Modulation

The motivation behind this modulation is similar to how the human brain reacts to an unseen target: unfamiliar features are interpreted by replacing them with familiar ones. In the same way, feature modulation modifies the original domain-sensitive features by mixing them with a portion of SARs. Since SARs are a modification of \mathbf{P}_c , which is created by averaging features from multiple source domains,

they contain less domain-specific information. We introduce a feature-modulating matrix \mathbf{M} , which is a learnable matrix that trains to highlight class-relevant features while down-weighting domain information containing features.

During training, we shift the features toward SARs and force the classifier to predict the true class. This way, the feature extractor achieves better class separation while enabling the classifier to distinguish between classes more effectively, even when their embeddings are close. This process also suppresses domain information, preventing the classifier from being misled by domain variations. \mathbf{M} learns with the use of our novel diagonal loss (Eq 8,9). Feature modulation is applied to all source domain data, reducing the transfer of domain information to classifier (Fig. 3).

To minimize domain influence, it is important to detect features that vary greatly across domains, as these are more domain-specific. We initialize the feature-modulating matrix \mathbf{M} using the variance of the feature set K_c for each class c in the labeled dataset. Initially, Features with higher variance are given lower initial weights, ensuring they contribute less to the final representation, but this \mathbf{M} learns during training, as we can't exactly define the domain information only using the labeled data. Let

$$\mathbf{V} = [\mathbf{V}_1, \mathbf{V}_2, \dots, \mathbf{V}_C],$$

where $\mathbf{V}_c = \text{Variance}(K_c)$ for class c . We normalize \mathbf{V} to the range $[0, 1]$ and initialize \mathbf{M} as:

$$\mathbf{M} = \mathbf{1} - \frac{\mathbf{V} - \min(\mathbf{V})}{\max(\mathbf{V}) - \min(\mathbf{V})}. \quad (4)$$

Tab. 3 shows that this initialization outperforms random (normal distribution) initialization. During training, \mathbf{M} learns to further suppress features rich in domain information. This improves the reliability of pseudo-labels, enhancing the use of unlabeled data.

The modulated feature representation \mathbf{Z}_m is defined as:

$$\mathbf{Z}_m = \mathcal{M}(\mathbf{Z}) = \mathbf{M} \odot \mathbf{Z} + (\mathbf{1} - \mathbf{M}) \odot \mathbf{R}, \quad (5)$$

where \mathbf{Z} replicates the instance-specific feature z for each class C , and \mathbf{R} is the SAR from Sec. 3.1.1. Here, $\mathbf{M} \odot \mathbf{Z}$ reduces domain information, while $(\mathbf{1} - \mathbf{M}) \odot \mathbf{R}$ shifts features closer to SARs. Each row of \mathbf{M} controls how much instance-specific information contributes to \mathbf{Z}_m for each class (Fig. 3). If a feature carries strong domain bias, its weight in \mathbf{M} is close to zero; if it is class-specific, its weight is higher, preserving its role in the modulated features. In practice, we identify feature positions with high domain sensitivity through \mathbf{M} . Features with small values in \mathbf{M} are adjusted by adding a portion, $(\mathbf{1} - \mathbf{M})$ of SARs, which strengthens class-related signals while reducing domain information. This effect is most pronounced when modulating features toward their correct class, leading to better class separation and more robust classification.

3.1.3. Pseudo-Labeling using Modulated Features

To generate more accurate pseudo-labels for unlabeled samples, we modulate the features of unlabeled data points to each class using Eq. 5 (see Fig. 3). Following standard SSDG practice [11, 12, 45], we construct the unlabeled dataset $U = d^L \cup d^U$, where weakly augmented unlabeled samples are considered alongside labeled samples without using their labels.

For a given unlabeled sample u_i , we compute a modulated feature representation for each class c as $\mathcal{M}(f(u_i), \mathbf{R})$. Since shifting the features towards the SAR of class c mitigates domain-related information for that class, the mitigation of domain information is more effective when applied to an unlabeled data of class c . So, the classifier achieves the highest prediction confidence when the features are modulated toward the true class. To identify this class, we pass the modulated feature for each class c through the classifier and compute the softmax scores:

$$S_i = \text{softmax}(\mathcal{C}(\mathcal{M}(f(\alpha(u_i))), \mathbf{R})), \quad (6)$$

where α denotes weak augmentation. The j^{th} row of S_i (see Fig.3.b) represents the prediction probabilities for each class after modulating the features of u_i toward class j .

To determine the predicted class, we extract the diagonal values of S_i , denoted as $\text{diag}(S_i)$, which forms an array of size C containing the confidence scores for each class. The highest confidence score $p_{\max} = \max(\text{diag}(S_i))$

We assign the class corresponding to p_{\max} as the pseudo-label \tilde{y}_i for u_i : $\tilde{y}_i = \text{argmax}(\text{diag}(S_i))$. This approach leverages the classifier’s predictions on modulated features, which contain reduced domain-specific information, leading to more reliable pseudo-labels. As illustrated in Fig. 2, this process improves the accuracy of pseudo-labeling, enhancing the effectiveness of the SSL setting.

3.1.4. Uncertainty Estimation of PLs and Loss Scaling

Unlike the fixed pseudo-label threshold used in Fix-Match [27], we estimate the uncertainty score (US) of pseudo-labels to introduce an adaptive threshold that considers both prediction probability and uncertainty. To achieve this, we integrate dropout layers into the feature extractor, applying a dropout probability of 0.05 after the last batch normalization layer, following [17].

To quantify uncertainty, we perform Monte Carlo (MC) sampling by passing each input through the network multiple times with dropout enabled. We obtain a distribution of 5 predictions for each sample and compute the mean p_{\max} and standard deviation σ of the prediction probabilities.

Increasing the utilization of unlabeled data enhances domain generalization by mitigating model overfitting. However, existing methods discard a significant portion of unlabeled data during training, as they only consider predictions that exceed a high confidence threshold (typically

0.95) for pseudo-labeling. Additionally, the contribution of these pseudo-labels to training remains uniform, regardless of their prediction probability.

To address this limitation, we introduce a loss scaling function $Q(x) = \exp(x^3 - 1)$ that dynamically adjusts the contribution of each pseudo-label based on its prediction probability p_{\max} (Sec. 3.1.3). This function ensures that higher-confidence pseudo-labels have a greater influence on training while still allowing lower-confidence pseudo-labels to contribute in a controlled manner.

Our feature modulation approach improves the accuracy of pseudo-labels, enabling us to lower the confidence threshold (0.75) and effectively leverage pseudo-labels with lower confidence scores. By incorporating uncertainty estimation, we refine the pseudo-labeling process, ensuring that samples with high uncertainty are assigned labels more conservatively. This adaptive mechanism allows us to extract useful information from previously unutilized data, further enhancing the semi-supervised learning process.

$$l_{\text{scale}} = \mathbb{1}\{p_{\max} - \sigma > \tau\} \cdot Q(p_{\max}). \quad (7)$$

Here, l_{scale} determines whether a pseudo-label is used in training based on its confidence-adjusted probability, incorporating both the mean prediction probability p_{\max} and uncertainty measure σ . This approach ensures that pseudo-labels with lower confidence but high uncertainty value still contributes to the training, leading to more effective utilization of unlabeled data.

3.2. Training

The training process consists of four loss terms, supervised loss L_s , unsupervised loss L_u , diagonal maximizing supervised loss L_d , and diagonal maximizing unsupervised loss L_{ud} . Supervised loss trains the model to give a correct prediction for weakly augmented labeled data x_i , $\alpha(x_i)$. We get prediction as $y_{\text{pred}} = \text{argmax}(\text{diag}(S_i))$ where $S_i = \log_softmax(\mathcal{C}(\mathcal{M}(f(\alpha(x_i))), \mathbf{R}))$. With y_i denoting the correct label for x_i , supervised loss $L_s = \text{NLL}(y_{\text{pred}}, y_i)$.

Unsupervised loss trains the model to predict the same label as the pseudo-label \tilde{y}_i for strongly augmented unlabeled data u_i , $\mathcal{A}(u_i)$. We get prediction as $y_{\text{upred}} = \text{argmax}(\text{diag}(S_{ui}))$, where $S_{ui} = \log_softmax(\mathcal{C}(\mathcal{M}(f(\mathcal{A}(u_i))), \mathbf{R}))$. Then with l_{scale} (Eq.(7)) unsupervised loss, $L_u = \text{NLL}(y_{\text{upred}}, \tilde{y}_i) * l_{\text{scale}}$

According to our concept of feature modulation, we emphasize the class-specific features while de-emphasizing the domain information related to class c . So, the classifier should give the maximum prediction for a respective class when the features are modulated to that class. We introduce diagonal maximizing loss to minimize the difference between the maximum prediction for each class col_{\max} and the prediction for that class from the modulated features to

Method	5 labels per class				10 labels per class			
	PACS	OfficeHome	VLCS	DigitsDG	PACS	OfficeHome	VLCS	DigitsDG
FixMatch	73.4 \pm 1.3	55.1 \pm 0.5	69.9 \pm 0.6	56.0 \pm 2.2	76.6 \pm 1.2	57.8 \pm 0.3	70.0 \pm 2.1	66.4 \pm 1.4
FlexMatch	69.5 \pm 1.1	50.8 \pm 0.4	56.0 \pm 1.1	56.5 \pm 1.8	72.7 \pm 1.2	53.7 \pm 0.7	56.2 \pm 2.1	68.9 \pm 1.2
StyleMatch	78.4 \pm 1.1	<u>56.3</u> \pm 0.3	72.5 \pm 1.3	55.7 \pm 1.6	79.4 \pm 0.9	<u>59.7</u> \pm 0.2	73.3 \pm 0.6	64.8 \pm 1.9
FBCSA	77.3 \pm 1.1	55.8 \pm 0.2	71.3 \pm 0.7	62.0 \pm 1.5	78.2 \pm 1.2	59.0 \pm 0.4	72.2 \pm 1.0	<u>70.4</u> \pm 1.4
DGWM	<u>77.9</u> \pm 0.8	56.2 \pm 0.2	<u>75.2</u> \pm 0.9	57.4 \pm 1.5	78.4 \pm 1.0	59.7 \pm 0.3	<u>75.2</u> \pm 0.7	68.4 \pm 1.5
FM (Ours)	75.4 \pm 1.6	57.2 \pm 0.3	75.3 \pm 0.8	<u>59.7</u> \pm 1.9	<u>78.7</u> \pm 1.0	61.0 \pm 0.3	75.5 \pm 0.1	73.0 \pm 1.3

Table 1. Averaged top-1 accuracies over 5 independent trials in 5 labels and 10 labels setting for the major DG datasets. We achieve 3.3% and 4.3% increase in averaged accuracy between datasets for 10 labels and 5 labels settings over the FixMatch baseline respectively. (Average over 5 independent seeds is reported.) Note that StyleMatch [45], FBCSA [11], and DGWM [12] utilize domain labels of unlabeled data and they are unable to perform well when there are no domain labels.

their true class $\text{diag}(S_i)$. This loss mainly trains the feature-modulating matrix M to emphasize class-specific features and de-emphasize domain information containing features.

For labeled data, we define L_d using S_i computed with labeled data x_i and col_{\max} , where col_{\max} = Maximum of each coloum in S_i (see Fig. 3).

$$L_d = \text{MSE}(\text{diag}(S_i), \text{col}_{\max}). \quad (8)$$

For unlabeled data, we define L_{ud} using S_{ui} computed with unlabeled data u_i and col_{\max} , where col_{\max} = Maximum of each coloum in S_{ui} (see Fig. 3).

$$L_{ud} = \text{MSE}(\text{diag}(S_{ui}), \text{col}_{\max}) * l_{\text{scale}}. \quad (9)$$

This loss effectively trains the modulating matrix M to optimally weigh the contribution of instance-specific features and the SAR to the final modulated feature.

The final loss $L = L_s + L_u + \beta L_d + \gamma L_{ud}$, where β and γ denotes hyperparameters. We attained consistent results for $\beta = 1$ $\gamma = 0.5$. See suppl. for more details.

3.3. Inference

During the inference, we perform feature modulation to mitigate the effect of unseen domain information. We feed the modulated features of test images x^t . When modulating to class c , we mitigate effect of domain information related to class c as the modulating matrix is already trained to blend more weight of domain related information from SARs. So, the domain information of unseen domain get blended by a known generalized features, reducing the effect of domain shift. Finally, we take the maximum prediction out of them as the final prediction $y_{\text{pred}}^t = \text{argmax}(D^t)$, where $D^t = \text{diag}(\text{softmax}(\mathcal{C}(\mathcal{M}(f(x^t))))))$.

4. Experiments

Datasets: We conducted experiments on four widely used DG datasets: PACS [20], OfficeHome [30], DigitsDG [44], and VLCS [29]. The PACS dataset emphasizes model robustness against significant style shifts, particularly from

realistic photos to abstract sketches and cartoons. OfficeHome [30] contains more complexity in terms of object variety and background, offering a more challenging evaluation of generalization capabilities. In the VLCS [29] dataset, the class imbalance is prominent and the performance of the proposed method proves its robustness for class imbalance. DigitsDG [44] is a synthetic DG dataset with digit classification tasks across four domains. See suppl. for a detailed description of datasets and details of the class imbalance.

Training and Implementation Settings: We conducted our training following the same training settings as in StyleMatch [45], FBCSA [11] and DGWM [12]. We used ImageNet [7] pretrained ResNet-18 [16] as the feature extractor for all the experiments. We followed leave-one-domain-out training protocol which is widely used in DG and SSDG methods. Our experiments were conducted in 5-label and 10-label settings. A learning rate of 0.03 was applied across all runs, with training completed over 20 epochs for each dataset. These learning rates are decayed using cosine annealing. For mini-batch construction, we randomly sampled 16 labeled and 16 unlabeled samples per source domain. We report top-1 accuracy averaged over 5 independent trials.

Main Results As shown in Tab 1, our method achieves consistent improvements of **3.3%** and **4.3%** in averaged accuracy across datasets compared to the FixMatch [27] for 10 and 5 labels settings, respectively. Our method has notable gains in OfficeHome, VLCS and DigitsDG, surpassing all the SOTA methods. **PACS:** We obtain 2% and 2.1% average gains over FixMatch for 5 and 10 labels settings, respectively. **OfficeHome:** We attain the best accuracies of 57.2% and 61.0% for 5 and 10 labels setting. Our approach significantly surpasses all the SOTA methods for both settings. **VLCS:** Overcoming the challenge of class imbalance in the VLCS dataset, we attain a gain of 5.4% compared to FixMatch in both settings. Furthermore our method outperforms all the SOTA methods with accuracies of 75.5% and 75.3% for 5 and 10 labels settings respectively. **DigitsDG:** Our approach exceeds FixMatch by a significant gain of 6.6% and 3.7% for 10 and 5 labels settings respectively.

Method	τ	Acc.	Keep Rate	PL Acc.
Baseline FixMatch	0.95	57.8	63.5	87.1
Baseline + Loss Scaling	0.95	60.2	52.2	88.3
Baseline + Feature Modulator (\mathcal{M})	0.75	60.8	78.5	85.0
Baseline + Uncertainty Score (US)	0.75	58.5	63.7	87.6
Baseline + Loss Scaling + \mathcal{M}	0.75	60.6	75.8	87.7
Baseline + Loss Scaling + US	0.75	57.7	76.4	84.7
Baseline + \mathcal{M} + US	0.75	57.4	79.3	82.4
Baseline + \mathcal{M} + Loss Scaling + US	0.75	61.0	67.6	90.4

Table 2. Contribution of key components; Loss scaling alone degrades performance due to noisy pseudo-labels, especially with lower confidence thresholds. The feature modulator improves accuracy by accurate pseudo-labels but has a low keep rate, indicating poor unlabeled data use. Combining both enhances pseudo-label accuracy and unlabeled data utilization.

Feature Modulator (\mathcal{M})	Accuracy	Keep Rate	PL Accuracy
Normal Init. + SAR + $L_d + L_{ud}$	47.8	62.3	87.2
Variance Init. + AP + $L_d + L_{ud}$	60.5	80.5	85.3
Variance Init. + SAR	60.4	66.7	90.3
Variance Init. + SAR + $L_d + L_{ud}$	61.0	67.6	90.4

Table 3. Comparison of feature modulator components at $\tau = 0.75$. Variance initialization outperforms normal initialization, SAR outperforms average prototypes (AP), and diagonal loss is imperative for feature modulator.

Conf. Thresh. (τ)	Loss Scaling	Accuracy	Keep Rate	PL Acc.
0.95	✗	59.7	41.9	95.9
0.95	✓	59.8	41.6	95.9
0.80	✓	60.6	63.2	91.6
0.75	✗	57.4	79.3	82.4
0.75	✓	61.0	67.6	90.6
0.70	✓	60.8	71.2	89.4

Table 4. Comparison of our method with various confidence thresholds with and without loss scaling. Experiments show that incorporating loss scaling with a 0.95 confidence threshold limits unlabeled data utilization, causing sub-optimal performance. The optimal balance occurs at a 0.75 threshold; further lowering thresholds degrades the performance.




Figure 4. Comparison of PL accuracy and keep rate with training iterations (20) for target domain Art of OfficeHome dataset. We achieve consistent improvements over FixMatch.

Contribution of Key Components: We conducted a comprehensive analysis of the key components of the proposed method under 10 label settings for the OfficeHome dataset. (Tab 2). The feature modulator alone shows a considerable improvement in accuracy due to the effective emphasis of the class-related features. Even with a higher average accuracy and keep rate, it shows lower pseudo-label accuracy. This illustrates the dilemma between keep rate and pseudo-label accuracy. We combine the loss scaling and uncertainty score to the feature modulator to make the model robust

Method	D1			D1+D2			D1+D2+D3		
	Acc.	KR	PL Acc.	Acc.	KR	PL Acc.	Acc.	KR	PL Acc.
FixMatch	59.5	37.0	93.7	63.6	49.1	89.7	67.2	62.6	88.2
FM (ours)	59.0	41.1	92.1	64.2	53.9	89.1	71.2	63.4	90.7

Table 5. Performance of our method with FixMatch for the increase of source domains (D1-Art, D2-ClipArt, D3-Product) of OfficeHome where target domain is real world. Our method maintains higher average accuracy, while the baseline’s drops.

with other benchmarks, balancing accurate pseudo-labels with higher unlabeled data utilization.

Feature modulator components: We compared proposed feature variance based initialization method for feature modulating matrix \mathbf{M} with normal initialization in Tab 3. The variance initialization has a significant impact compared to the normal initialization. SARs show a better improvement over the average prototypes (AP) generated by averaging all the class prototypes. In the SAR generation, the process of averaging class prototypes based on cosine similarity provides better feature modulation for clear class separation. The impact of the contribution of the diagonal losses on accurate class separation of the classifier is evident from the accuracy improvement.

Impact of confidence threshold and loss scaling: Tab 4 shows performance across thresholds with and without loss scaling, which is more effective at 0.75 than 0.95. Lowering the threshold increases inclusion but also introduces noisy pseudo-labels that degrade learning and limit Keep Rate. Thus, no sharp increase is observed. A threshold of 0.75 achieves the best balance, while further reduction harms performance.

Improved PL accuracy and keep rate: We plot the PL accuracy and keep rate of our method in Fig. 4 compared to FixMatch [27]. Throughout training iterations, we achieve consistent improvements in both the PL accuracy and the keep rate. With improved average accuracy, we verify our claim of balancing the dilemma between accurate pseudo-labels and higher unlabeled data utilization.

5. Conclusion

We address a unique problem in relatively under-explored research area SSDG, where domain labels are unavailable for model training. We address the *dilemma* between more accurate pseudo-labels and a higher keep rate. Existing models struggle to achieve better data utilization with higher pseudo-label accuracy. The proposed feature modulation improves the generalization by pushing the features towards similar average representations and loss scaling improves unlabeled data utilization by effectively using low-confident pseudo-labels. This encourages the classifier to accurately classify closely clustered features despite the domain by minimizing the propagation of domain information. Extensive experiments with major DG benchmark datasets validate our claims.

References

- [1] Abulikemu Abuduweili, Xingjian Li, Humphrey Shi, Cheng-Zhong Xu, and Dejing Dou. Adaptive consistency regularization for semi-supervised transfer learning. In *CVPR*, pages 6923–6932, 2021. 3
- [2] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019. 1
- [3] David Berthelot, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. Remixmatch: Semi-supervised learning with distribution alignment and augmentation anchoring. *arXiv preprint arXiv:1911.09785*, 2019. 3
- [4] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. *NeurIPS*, 32, 2019. 2, 3
- [5] Yuhao Chen, Xin Tan, Borui Zhao, Zhaowei Chen, Renjie Song, Jiajun Liang, and Xuequan Lu. Boosting semi-supervised learning by exploiting all unlabeled data. In *CVPR*, pages 7548–7557, 2023. 2
- [6] Yuhao Chen, Xin Tan, Borui Zhao, Zhaowei Chen, Renjie Song, Jiajun Liang, and Xuequan Lu. Boosting semi-supervised learning by exploiting all unlabeled data. In *CVPR*, pages 7548–7557, 2023. 3
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 7
- [8] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 1
- [9] Qi Dou, Daniel Coelho de Castro, Konstantinos Kamnitsas, and Ben Glocker. Domain generalization via model-agnostic learning of semantic features. *NIPS*, 32, 2019. 1, 2
- [10] Yarín Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR, 2016. 4
- [11] Chamuditha Jayanga Galappaththige, Sanoojan Baliah, Malitha Gunawardhana, and Muhammad Haris Khan. Towards generalizing to unseen domains with few labels. In *CVPR*, pages 23691–23700, 2024. 2, 3, 4, 6, 7
- [12] Chamuditha Jayanga Galappaththige, Zachary Izzo, Xilin He, Honglu Zhou, and Muhammad Haris Khan. Domain-guided weight modulation for semi-supervised domain generalization. In *Proceedings of the Winter Conference on Applications of Computer Vision (WACV)*, pages 6495–6505, 2025. 3, 5, 6, 7
- [13] 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. 1
- [14] Yves Grandvalet and Yoshua Bengio. Semi-supervised learning by entropy minimization. *NeurIPS*, 17, 2004. 3
- [15] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. In *ICLR*, 2020. 1
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 1, 7
- [17] Bum Jun Kim, Hyeeyeon Choi, Hyeonah Jang, Donggeon Lee, and Sang Woo Kim. How to use dropout correctly on residual networks with batch normalization. In *Uncertainty in Artificial Intelligence*, pages 1058–1067. PMLR, 2023. 6
- [18] Doyup Lee, Sungwoong Kim, Ildoo Kim, Yeongjae Cheon, Minsu Cho, and Wook-Shin Han. Contrastive regularization for semi-supervised learning. In *CVPR*, pages 3911–3920, 2022. 3
- [19] 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*, page 896. Atlanta, 2013. 2
- [20] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generalization. In *ICCV*, pages 5542–5550, 2017. 1, 7
- [21] Da Li, Jianshu Zhang, Yongxin Yang, Cong Liu, Yi-Zhe Song, and Timothy M Hospedales. Episodic training for domain generalization. In *ICCV*, pages 1446–1455, 2019. 1, 2
- [22] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *CVPR*, pages 5400–5409, 2018. 1
- [23] Ya Li, Xinmei Tian, Mingming Gong, Yajing Liu, Tongliang Liu, Kun Zhang, and Dacheng Tao. Deep domain generalization via conditional invariant adversarial networks. In *Proceedings of the European conference on computer vision (ECCV)*, pages 624–639, 2018. 1
- [24] Saeid Motiian, Quinn Jones, Seyed Iranmanesh, and Gianfranco Doretto. Few-shot adversarial domain adaptation. *NeurIPS*, 30, 2017. 1
- [25] Lei Qi, Hongpeng Yang, Yinghuan Shi, and Xin Geng. Multitask: Multi-task learning for semi-supervised domain generalization. *ACM Transactions on Multimedia Computing, Communications and Applications*, 20(6):1–21, 2024. 3
- [26] Yang Shu, Zhangjie Cao, Chenyu Wang, Jianmin Wang, and Mingsheng Long. Open domain generalization with domain-augmented meta-learning. In *CVPR*, pages 9624–9633, 2021. 1, 2
- [27] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *NeurIPS*, pages 596–608, 2020. 2, 3, 6, 7, 8
- [28] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *NeurIPS*, 30, 2017. 2
- [29] Antonio Torralba and Alexei A Efros. Unbiased look at dataset bias. In *CVPR*, pages 1521–1528, 2011. 7
- [30] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for

- unsupervised domain adaptation. In *CVPR*, pages 5018–5027, 2017. [7](#)
- [31] Riccardo Volpi, Hongseok Namkoong, Ozan Sener, John C Duchi, Vittorio Murino, and Silvio Savarese. Generalizing to unseen domains via adversarial data augmentation. *NeurIPS*, 31, 2018. [1](#), [2](#)
- [32] Haotao Wang, Chaowei Xiao, Jean Kossaifi, Zhiding Yu, Anima Anandkumar, and Zhangyang Wang. Augmax: Adversarial composition of random augmentations for robust training. *Advances in neural information processing systems*, 34: 237–250, 2021. [1](#)
- [33] Ruiqi Wang, Lei Qi, Yinghuan Shi, and Yang Gao. Better pseudo-label: Joint domain-aware label and dual-classifier for semi-supervised domain generalization. *Pattern Recognition*, 133:108987, 2023. [2](#), [3](#), [4](#)
- [34] Yidong Wang, Hao Chen, Qiang Heng, Wenxin Hou, Yue Fan, Zhen Wu, Jindong Wang, Marios Savvides, Takahiro Shinozaki, Bhiksha Raj, et al. Freematch: Self-adaptive thresholding for semi-supervised learning. *arXiv preprint arXiv:2205.07246*, 2022. [2](#), [3](#)
- [35] Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le. Unsupervised data augmentation for consistency training. *NeurIPS*, 33:6256–6268, 2020. [3](#)
- [36] Qinwei Xu, Ruyang Zhang, Ya Zhang, Yanfeng Wang, and Qi Tian. A fourier-based framework for domain generalization. In *CVPR*, pages 14383–14392, 2021. [1](#), [2](#)
- [37] Bowen Zhang, Yidong Wang, Wenxin Hou, Hao Wu, Jindong Wang, Manabu Okumura, and Takahiro Shinozaki. Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling. *NeurIPS*, 34:18408–18419, 2021. [3](#)
- [38] H Zhang, M Cisse, Y Dauphin, and D Lopez-Paz. mixup: Beyond empirical risk management. In *ICLR*, pages 1–13, 2018. [3](#)
- [39] Kun Zhang, Bernhard Schölkopf, Krikamol Muandet, and Zhikun Wang. Domain adaptation under target and conditional shift. In *ICML*, pages 819–827. Pmlr, 2013. [1](#)
- [40] Lei Zhang, Ji-Fu Li, and Wei Wang. Semi-supervised domain generalization with known and unknown classes. *NeurIPS*, 36, 2024. [2](#), [3](#)
- [41] Xingxuan Zhang, Peng Cui, Renzhe Xu, Linjun Zhou, Yue He, and Zheyang Shen. Deep stable learning for out-of-distribution generalization. In *CVPR*, pages 5372–5382, 2021. [1](#)
- [42] Zhun Zhong, Yuyang Zhao, Gim Hee Lee, and Nicu Sebe. Adversarial style augmentation for domain generalized urban-scene segmentation. *Advances in neural information processing systems*, 35:338–350, 2022. [1](#)
- [43] Beitong Zhou, Jing Lu, Kerui Liu, Yunlu Xu, Zhanzhan Cheng, and Yi Niu. Hypermatch: Noise-tolerant semi-supervised learning via relaxed contrastive constraint. In *CVPR*, pages 24017–24026, 2023. [2](#), [3](#)
- [44] Kaiyang Zhou, Yongxin Yang, Timothy Hospedales, and Tao Xiang. Deep domain-adversarial image generation for domain generalisation. In *AAAI*, pages 13025–13032, 2020. [7](#)
- [45] Kaiyang Zhou, Chen Change Loy, and Ziwei Liu. Semi-supervised domain generalization with stochastic stylematch. *IJCV*, 131(9):2377–2387, 2023. [2](#), [3](#), [6](#), [7](#)
- [46] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Mixstyle neural networks for domain generalization and adaptation. *IJCV*, 132(3):822–836, 2024. [2](#)