

Towards Optimization and Model Selection for Domain Generalization: A Mixup-guided Solution

Wang Lu

*Tsinghua University
Beijing, China*

LUW12@TSINGHUA.ORG.CN

Jindong Wang

*Microsoft Research Asia
Beijing, China*

JINDONG.WANG@MICROSOFT.COM

Yidong Wang

*National Engineering Research Center for Software Engineering
Peking University
Beijing, China*

YIDONGWANG37@GMAIL.COM

Xing Xie

*Microsoft Research Asia
Beijing, China*

XINGX@MICROSOFT.COM

Editor: Thuc Le, Jiuong Li, Robert Ness, Sofia Triantafyllou, Shohei Shimizu, Peng Cui, Kun Kuan, Jian Pei, Fei Wang, Mattia Proserpi

Abstract

The distribution shifts between training and test data typically undermine the performance of deep learning models. In recent years, lots of work pays attention to domain generalization (DG) where distribution shift exists and target data are unseen. Despite the progress in algorithm design, two foundational factors have long been ignored: 1) the optimization for regularization-based objectives (e.g., distribution alignment), and 2) the model selection for DG since no knowledge about the target domain can be utilized. In this paper, we propose Mixup guided optimization and selection techniques for domain generalization. For optimization, we utilize an adapted Mixup to generate an out-of-distribution dataset that can guide the preference direction and optimize with Pareto optimization. For model selection, we generate a validation dataset with a closer distance to the target distribution, and thereby it can better represent the target data. We also present some theoretical insights behind our proposals. Comprehensive experiments on one visual classification benchmark and three time-series benchmarks demonstrate that our model optimization and selection techniques can largely improve the performance of existing domain generalization algorithms and even achieve new state-of-the-art results.

Keywords: Domain generalization, Data augmentation, Pareto optimization, Model selection, Transfer learning

1 Introduction

Deep learning has been widely adopted in daily-life applications (Vaswani et al., 2017), such as face recognition (Wang et al., 2017), speech processing (Han et al., 2020), health-care (Tang et al., 2020), human activity recognition (Robards and Sunehag, 2009), etc. The

success of deep learning methods heavily relies on expensive and laborious labeled data. To reduce such reliance, a common strategy is to leverage data from another domain to improve the model’s generalization performance. However, distribution shift naturally exists among different domains, not to mention that the training and test distributions are also different.

To tackle the distribution shift, transfer learning (Pan and Yang, 2009) is a popular paradigm that reuses the pre-trained models by fine-tuning on the limited target data. Domain adaptation (DA) (Wilson and Cook, 2020) is a branch of transfer learning and receives great attention recently. DA bridges the distribution gap between two domains by instance reweighting or distribution alignment techniques. Although DA has shown effectiveness in multiple fields, it needs access to both source data and target data. In recent years, domain generalization (DG), i.e., out-of-distribution generalization, has attracted increasing interests (Wang et al., 2021). DG aims to learn a model that can generalize to an *unseen* target (i.e., test) domain when given source data only from several different but related domains (Wang et al., 2021; Gulrajani and Lopez-Paz, 2021).

Much prior work has focused on data manipulation (Shankar et al., 2018), representation learning (Jia et al., 2020), and learning strategy (Rame et al., 2021) based DG methods. However, two foundational limitations have long been ignored. First, the regularization items introduced by most DG approaches (Xu et al., 2021; Wang et al., 2021; Planamente et al., 2022) can conflict with both the original goal (i.e. the classification item) and each other during training. For example, CORAL (Sun and Saenko, 2016) introduced the correlation alignment loss for second-order statistics alignment while DANN (Ganin and Lempitsky, 2015) introduced the domain discriminator loss for domain-adversarial training. When optimizing multiple objectives, decreasing the overall objective value can be at the expense of damaging one of the training objectives, e.g., the classification goal (Lv et al., 2021). Therefore, these items for different purposes can have conflicts with the original classification goal, impeding the performance. Second, the traditional model selection methods, e.g. selecting via validation data split from training data (Refaeilzadeh et al., 2009), are not suitable in the DG scenario since target data is unseen and has different distributions from the training data. The popular DomainBed (Gulrajani and Lopez-Paz, 2021) utilized three model selection techniques without any guarantee, but experiments illustrated that all three techniques could not achieve acceptable performance for DG. Due to different distributions between validation and target data, these techniques can only obtain a biased estimation of the target accuracy. Without proper model selection strategies, it is difficult to comprehensively evaluate different algorithms and deploy them in real applications.

To deal with the above two issues, we propose a Mixup-guided optimization and selection solution for domain generalization to tackle these issues. Specifically, we generate two datasets based on Mixup (Zhang et al., 2018), a simple but effective data augmentation method. We slightly adapt Mixup for our purposes and generate one dataset for model optimization (OPTD) and one dataset for model selection (VALD). For model optimization, we first compute the gradients on OPTD with the same model and then utilize the gradients to guide the balance between the classification item and the generalization item. For model selection, since the different training and test distributions make it infeasible to select models based on training data, we replace the traditional validation data split from training data with VALD to choose the best model.

Our contributions can be summarized as follows:

1. We aim to tackle the two fundamental challenges in domain generalization: the optimization of regularization-based DG approaches and model selection. To the best of our knowledge, it is the *first* work towards solving these challenges simultaneously.
2. We propose a simple yet effective Mixup-guided universal ¹ solution to resolve these issues. We also provide some theoretical insights to our solution.
3. Comprehensive experiments across image and time-series classification illustrate that these two techniques can make traditional methods rework and even have better performance than state-of-the-art methods.

2 Related Work

2.1 Domain Generalization

Given one or several different but related domains, domain generalization (DG) aims to learn a model that can generalize well on unseen target domains. Existing domain generalization work can be grouped into three categories (Wang et al., 2021): data manipulation, representation learning, and learning strategy.

Data manipulation mainly focuses on manipulating the inputs to assist learning general representations and it contains two kinds of popular techniques, data augmentation and data generation. Data augmentation is mainly based on augmentation, randomization, and transformation of input data while data generation generates diverse samples to help generalization. CrossGrad (Shankar et al., 2018) parallelly trained a label and a domain classifier on examples perturbed by loss gradients of each other’s objective for better generalization. Peng et al. (Peng et al., 2021) utilized uncertainty assessment to augment the source capacity in both feature and label spaces, and a Bayesian meta-learning framework was used to implement it. Fact (Xu et al., 2021) tried to linearly interpolate between the amplitude spectrums of two images which were thought related to classification.

Learning strategy focuses on exploiting the general learning strategy to promote generalization capability. MLDG (Li et al., 2018) utilized the meta-learning strategy for DG where data from the source domains were split into meta-train and meta-test to simulate the domain shift situation. RSC (Huang et al., 2020) forced the network to activate remaining features that correlated with labels via masking the dominant features for better classification. Fish (Shi et al., 2021) tried to maximize the inner product between gradients from different domains for generalization. These methods are all designed specifically from different points.

Representation learning is the most popular category in DG, including domain-invariant representation learning and feature disentanglement. CORAL (Sun and Saenko, 2016) aligned the second-order statistics of features to alleviate domain shifts. IRM (Arjovsky et al., 2019) enforced the optimal classifier on top of the representation space to be the same across all domains. Mahajan et al. (Mahajan et al., 2021) proposed MatchDG and

1. There exist a lot of regularization-based DG approaches (Xu et al., 2021; Wang et al., 2021; Planamente et al., 2022) and some of them still remain competitive in several benchmarks. For better analysis of the problem, our very specific focuses are the very fundamental approaches such as DANN (Ganin and Lempitsky, 2015) and CORAL (Sun and Saenko, 2016), which prove to be simple and effective. Additionally, our approach can also work with other approaches.

tried to deal with domain generalization from the view of a structural causal model. All these methods attempted to find the invariance behind the classification.

Although these three types of methods all prove their effectiveness in domain generalization, it is often inevitable to introduce some regularized items for better generalization. However, when directly fixing the balance between the regularized item and the classification item, and only considering reducing the overall objective value, some objectives, e.g. classification ability, might be damaged (Lv et al., 2021). Little work pays attention to this field, and our paper tries to utilize Mixup guided optimization to solve it.

2.2 Model Selection

An introduction to model selection was given in (Zucchini, 2000). The most common model selection method for machine learning is cross-validation. It split training data into two parts, one for training and the other for validation. For domain generalization, little work pays attention to model selection. Gulrajani and Lopez-Paz introduced three common model selection methods, namely, training-domain validation set, leave-one-domain-out cross-validation, and test-domain validation set. Moreover, (Li et al., 2022) demonstrated that leave-one-domain-out cross-validation was unbiased and was better than the other two methods. However, training-domain validation set required the assumption that the training and test examples followed similar distributions. Leave-one-domain-out cross-validation sacrificed the quantity of training data that might influence the performance. The last one, test-domain validation set, was impossible for domain generalization where no target data could be seen. Recently, another method (Ye et al., 2021) combined validation accuracy with feature variation and selected the model with high validation accuracy as well as low variation. However, they were only verified in the visual field and required more computational costs. Moreover, it was difficult to obtain a balance between the accuracy and the variation and it was hard to compute exactly. Model selection for domain generalization is still in the infant.

2.3 Model Optimization

Weight-based optimization is an intuitive strategy to optimize multiple objectives at the same time. A more popular choice is to manually tune the hyperparameters to obtain the best weights and form an overall objective. For domain adaptation, there have been lots of methods to dynamically weight the objectives for better performance, such as leveraging gradient magnitude (Chen et al., 2018), task uncertainty (Kendall et al., 2018), or learning non-linear loss combinations by implicit differentiation (Navon et al., 2020). However, for domain generalization, it is just getting started.

When there exist multiple objectives and we cannot further decrease all objectives simultaneously, we obtain a set of so-called Pareto optimal solutions. Multi-objective gradient-based optimization leverages the gradients of objectives to reach the final Pareto optimal solution with specific goals. Mahapatra et al. (Mahajan et al., 2021) proposed Exact Pareto Optimal (EPO) Search to find a preference-specific Pareto optimal solution. Lv et al. (Lv et al., 2021) introduced EPO to domain adaptation and proposed ParetoDA to control the overall optimization direction. However, ParetoDA required access to the target which was impossible for domain generalization.

3 Preliminaries

3.1 Problem Formulation

We follow the definition given in (Wang et al., 2021). In domain generalization, we are given M training (source) domains, $\mathcal{D}^S = \{\mathcal{D}^i | i = 1, 2, \dots, M\}$. Each domain has n_i data, $\mathcal{D}^i = \{(\mathbf{x}_j^i, y_j^i)\}_{j=1}^{n_i}$ where $\mathbf{x}_j^i \in \mathcal{X}^i$ and $y_j^i \in \mathcal{Y}^i$. There also exists an unlabeled target domain which is unseen during training, $\mathcal{D}^T = \{\mathbf{x}_j^T\}_{j=1}^{n_T}$ where $\mathbf{x}_j^T \in \mathcal{X}^T$, and n_T is the number of data in the target. For simplicity, in this paper, we assume that the target only contains one domain and all domains share the same input space and the same label space, i.e. $\mathcal{X}^1 = \mathcal{X}^2 = \dots = \mathcal{X}^M = \mathcal{X}^T = \mathcal{X}$, $\mathcal{Y}^1 = \mathcal{Y}^2 = \dots = \mathcal{Y}^M = \mathcal{Y}^T = \mathcal{Y}$. \mathcal{Y}^T is the target label space. Note that data shifts exist ubiquitously across domains, which means $\mathbb{P}_{XY}^i \neq \mathbb{P}_{XY}^j, i, j \in \{1, 2, \dots, M, T\}$, where \mathbb{P} denotes distribution. The goal of domain generalization is to learn a robust and generalized predictive function: $h : \mathcal{X} \rightarrow \mathcal{Y}$ from M training sources to achieve minimum prediction error on the unseen target domain \mathcal{D}^T , i.e. $\min_h \mathbb{E}_{(\mathbf{x}, y) \in \mathcal{D}^T} [\ell(h(\mathbf{x}), y)]$ where \mathbb{E} is the expectation notation and ℓ is a loss function, e.g. Cross-Entropy loss.

Most of the existing DG approaches have regularization-based objectives (Ganin and Lempitsky, 2015; Sun and Saenko, 2016; Xu et al., 2021; Wang et al., 2021; Planamente et al., 2022), formulated as:

$$\min_h \mathbb{E}_{(\mathbf{x}, y) \sim \mathbb{P}^S} \ell_0(h(\mathbf{x}), y) + \lambda_1 \ell_1 + \dots + \lambda_k \ell_k. \quad (1)$$

ℓ_0 is the classification loss while ℓ_1, \dots, ℓ_k are regularization losses. $\lambda_1, \dots, \lambda_k$ are hyperparameters which are fixed, and k is the number of regularization items. Since \mathcal{D}^T is unseen during training, existing methods (Gulrajani and Lopez-Paz, 2021) typically select the best model for testing according to a validation dataset, \mathcal{D}_{val}^S , split from sources:

$$\arg \max_{h \in \mathcal{H}} \mathbb{E}_{(\mathbf{x}, y) \in \mathcal{D}_{val}^S} Accuracy((h(\mathbf{x}), y)). \quad (2)$$

3.2 Background

Data split We denote the whole distribution of the training domains as $\mathbb{P}_{XY}^S = \sum_{i=1}^m \pi_i \mathbb{P}_{XY}^i$, where $\pi_i > 0$ is the proportion of each domain and $\sum_i \pi_i = 1$. In practice, we typically split randomly \mathcal{D}^S into two parts, one for training (\mathcal{D}_{tra}^S) and the other for validation (\mathcal{D}_{val}^S). \mathcal{D}_{tra}^S and \mathcal{D}_{val}^S share the same distribution, \mathbb{P}_{XY}^S . For simplicity, we use \mathcal{D}^S to denote the training data if no subscript is added.

Mixup Mixup (Zhang et al., 2018; Xu et al., 2021) is a simple but effective data augmentation technique. Mixup incorporates the prior knowledge that linear interpolations of feature vectors should lead to linear interpolations of the corresponding target labels. It generates virtual training examples based on two random data points:

$$\begin{aligned} \lambda &\sim Beta(\alpha, \alpha), \\ \tilde{\mathbf{x}} &= \lambda \mathbf{x}_i + (1 - \lambda) \mathbf{x}_j, \\ \tilde{y} &= \lambda y_i + (1 - \lambda) y_j, \end{aligned} \quad (3)$$

where $Beta(\alpha, \alpha)$ is a Beta distribution and $\alpha \in (0, \infty)$ is a hyperparameter.

4 Our Approach

In this section, we first introduce the Mixup-guided model optimization technique. Then, we introduce our Mixup-guided model selection technique and explain the insights. With these two techniques, we can make traditional methods rework, e.g. DANN, and even achieve better results compared to state-of-the-art methods.

4.1 Gradient-based Model Optimization

We introduce how to reduce conflicts and learn a better generalization model in this part. We first recall some related definitions on Pareto optimal solutions following (Zitzler and Thiele, 1999; Lv et al., 2021).

Consider m objectives, each with a non-negative loss function $\ell_i(\boldsymbol{\theta})$ where $\boldsymbol{\theta}$ is the parameters. There can be no solution that reaches the optima of each objective simultaneously since they can conflict with each other. However, we can still obtain a set of so-called Pareto optimal solutions.

Definition 1 (Pareto dominance). *Suppose two solutions $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \mathbb{R}^d$, define $\boldsymbol{\theta}_1 \prec \boldsymbol{\theta}_2$ if $\ell_i(\boldsymbol{\theta}_1) \leq \ell_i(\boldsymbol{\theta}_2), \forall i \in \{1, 2, \dots, m\}$ and $\ell_i(\boldsymbol{\theta}_1) < \ell_i(\boldsymbol{\theta}_2), \exists i \in \{1, 2, \dots, m\}$. Then we say $\boldsymbol{\theta}_1$ dominates $\boldsymbol{\theta}_2$ in this situation.*

Definition 2 (Pareto optimality). *If a solution $\boldsymbol{\theta}_1$ dominates $\boldsymbol{\theta}_2$, then $\boldsymbol{\theta}_1$ is clearly preferable as it performs better or equally on each objective. A solution $\boldsymbol{\theta}^*$ is Pareto optimality if it is not dominated by any other solutions.*

Definition 3 (Pareto front). *The set of all Pareto optimal solutions in loss space is Pareto front, where each point represents a unique solution.*

Definition 4 (Preference vector). *A Pareto optimal solution can be viewed as an intersection of the Pareto front with a specific direction in loss space. We refer to this direction as the preference vector.*

Now, back to our problem. Assume that there are classification loss, i.e. $\ell_0(h_c(h_f(\mathbf{x})), y)$ and k regularization items, i.e. ℓ_1, \dots, ℓ_k . When optimizing the whole objective, a trade-off is required among the different losses for better generalization since these objectives can have conflicts. Optimizing according to preference directions can be a possible solution (Mahapatra and Rajan, 2020; Lv et al., 2021). However, the previous method on Pareto domain adaptation (Lv et al., 2021) requires access to the unlabeled target distribution, which is unrealistic in the DG scenario. Note that we expect generalization capability in DG. Therefore, we expect that there exists a preference vector reflecting generalization. And we utilize adapted Mixup to generate a dataset (**OPTD**) with a different distribution from sources to compute the preference vector for a better generalization capability.

We add several restrictions to alleviate the noisy generation via Mixup. In particular, we split the data augmentation into two parts, one within the same class but different domains and the other within the same domain but arbitrary classes. The adapted Mixup can enlarge diversity and reduce the influence of redundant domain information (Xu et al., 2021; Zhou et al., 2021; Yao et al., 2022). Each part contains half of OPTD. The first part

can be formulated as:

$$\begin{aligned}\tilde{\mathbf{x}} &= \lambda \mathbf{x}_i + (1 - \lambda) \mathbf{x}_j, \\ \tilde{y} &= y_i = y_j, \\ &\text{where } d_i \neq d_j.\end{aligned}\tag{4}$$

And the second part can be formulated as:

$$\begin{aligned}\tilde{\mathbf{x}} &= \lambda \mathbf{x}_i + (1 - \lambda) \mathbf{x}_j, \\ \tilde{y} &= \lambda y_i + (1 - \lambda) y_j, \\ &\text{where } d_i = d_j.\end{aligned}\tag{5}$$

Given OPTD, we can compute the optimization direction towards the desired Pareto optimal solution now.

Considering $(k + 1)$ losses, classification loss $\ell_s = \ell_0(h_c(h_f(\mathbf{x})), y)$ and regularization loss ℓ_1, \dots, ℓ_k , the update direction \mathbf{d} can be modeled as a convex combination of gradients of these $k + 1$ losses, i.e. $\mathbf{d} = \mathbf{G}\boldsymbol{\omega}$, where $\boldsymbol{\omega} = (w_0, w_1, w_2, \dots, w_k)$, $\sum_{i=0}^k w_i = 1$, and $\mathbf{G} = [\nabla_{\boldsymbol{\theta}_f} \ell_s, \nabla_{\boldsymbol{\theta}_f} \ell_1, \dots, \nabla_{\boldsymbol{\theta}_f} \ell_k]$. Since ℓ_s and ℓ_1, \dots, ℓ_k may optimize different networks, we can only consider the shared network, h_f , among them². The main purpose of gradient-based optimization is to find \mathbf{d} to minimize all the losses and make the direction along with the preferred one. Here, we obtain the first constraint, $\mathbf{d}^T \mathbf{g}_j \geq 0$, where \mathbf{g}_j is the j -th column of \mathbf{G} . This constraint ensures all losses can be minimized simultaneously.

We utilize OPTD as our preference guidance, and follow EPO (Mahapatra and Rajan, 2020) to obtain $\boldsymbol{\omega}$. We denote ℓ_{OPTD} as the classification loss on OPTD, and we can directly obtain the gradient descent direction, $\mathbf{g}_{optd} = \nabla_{\boldsymbol{\theta}_f} \ell_{optd}$. We replace the guidance direction \mathbf{d}_{bal} in EPO by \mathbf{g}_{optd} as dynamical guidance of the optimization direction similar to (Lv et al., 2021). The optimization can be formulated as a linear programming (LP) problem:

$$\begin{aligned}\boldsymbol{\omega}^* &= \arg \max_{\boldsymbol{\omega} \in \Delta^{m-1}} (\mathbf{G}\boldsymbol{\omega})^T (I(\ell_{optd} > 0) \mathbf{g}_{optd} \\ &\quad + I(\ell_{optd} = 0) \mathbf{G}\mathbf{1}/m), \\ \text{s.t. } &(\mathbf{G}\boldsymbol{\omega})^T \mathbf{g}_j \geq I(J \neq \emptyset) (\mathbf{g}_{optd}^T \mathbf{g}_j), \forall j \in \bar{J} - J^*, \\ &(\mathbf{G}\boldsymbol{\omega})^T \mathbf{g}_j \geq 0, \forall j \in J^*.\end{aligned}\tag{6}$$

Δ^{m-1} is m dimensional simple, which means $\boldsymbol{\omega} \in \Delta^{m-1}$ represents $\boldsymbol{\omega} \in \mathbb{R}^m, w_i \geq 0, \sum w_i = 1$. In our situation, $m = k + 1$. $I(\cdot)$ is an indicator function, $\mathbf{1} \in \mathbb{R}^m$ is a vector whose elements are all 1. $J = \{j | \mathbf{g}_{optd}^T \mathbf{g}_j > 0\}$, $\bar{J} = \{j | \mathbf{g}_{optd}^T \mathbf{g}_j < 0\}$, and $J^* = \{j | \mathbf{g}_{optd}^T \mathbf{g}_j = \max_{j'} \mathbf{g}_{optd}^T \mathbf{g}_{j'}\}$. The following theorem ensures that the optimization will not over-fit on OPTD.

Theorem 5 (Theorem 1 in (Lv et al., 2021)). *Let $\boldsymbol{\omega}^*$ be the solution of the problem in Eq. (6), and $\mathbf{d}^* = \mathbf{G}\boldsymbol{\omega}^*$ be the resulted update direction. If $\ell_{optd} = 0$, then the dominating direction \mathbf{d}^* becomes a descent direction, i.e.,*

$$(\mathbf{d}^*)^T \mathbf{g}_j \geq 0, \forall j \in \{1, 2, \dots, k + 1\}.\tag{7}$$

2. When multiple objectives optimize the same network, we can also consider the whole network.

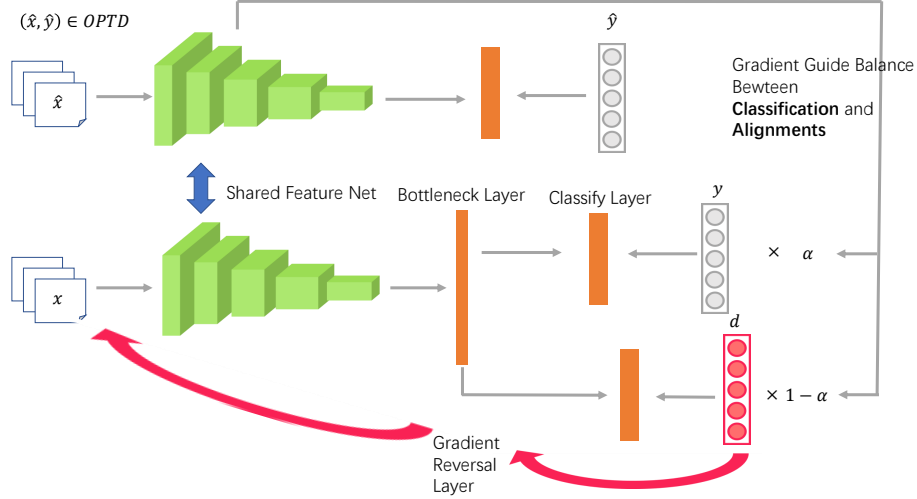


Figure 1: The framework of gradients base model optimization.

On the other hand, if $\ell_{optd} > 0$, let $\gamma^* = (\mathbf{d}^*)^T \mathbf{g}_{optd}$ be the objective value of the problem in Eq. (6). Then,

$$\begin{cases} (\mathbf{d}^*)^T \mathbf{g}_{optd} > 0, & \gamma^* > 0 \\ (\mathbf{d}^*)^T \mathbf{g}_j \geq 0, \forall j \in \{1, 2, \dots, k+1\}, & \gamma^* \leq 0. \end{cases} \quad (8)$$

According to the above theorem and (Lv et al., 2021), we split the learning mechanism into two modes, pure descent mode where \mathbf{d}^* approximates the mean gradient $\mathbf{G}\mathbf{1}/m$, and guidance descent mode where \mathbf{d}^* approximates \mathbf{g}_{optd} . Practically, we utilize a small $\epsilon > 0$ to relax the condition $\ell_{optd} = 0$ or > 0 . $\gamma^* > 0$ forces \mathbf{d}^* to decrease the loss whose gradient is the most consistent with \mathbf{g}_{optd} while $\gamma^* \leq 0$ only requires to decrease the training losses. Therefore, \mathbf{g}_{optd} can dynamically guide the optimization direction towards the desired Pareto solution. We require no prior knowledge of the preference vector nor access to unseen targets. With the best Lp solver (Cohen et al., 2021), we only need $O(m^{2.38})$ to solve Eq. (6), which can be ignored since m is a small integer in our cases.

4.2 Optimization Implementation Based on DANN

Now, we give an implementation of our optimization technique based on DANN. Domain-adversarial neural network (DANN) (Ganin and Lempitsky, 2015) is a classic and effective approach to learn domain-invariant representations for both DA (Ganin and Lempitsky, 2015) and DG (Wang et al., 2021). DANN utilizes adversarial training which contains a feature extractor, a domain discriminator, and a classification network. The domain discriminator tries to discriminate domain labels while the feature network tries to generate features that can be able to confuse the domain discriminator, which thereby learns domain-

invariant representation. It is an adversarial process and can be expressed as:

$$\begin{aligned} \min_{h_f, h_c} \mathbb{E}_{(\mathbf{x}, y) \sim \mathbb{P}^S} \ell_0(h_c(h_f(\mathbf{x})), y) - \ell_1(h_{adv}(h_f(\mathbf{x})), d), \\ \min_{h_{adv}} \mathbb{E}_{(\mathbf{x}, y) \sim \mathbb{P}^S} \ell_1(h_{adv}(h_f(\mathbf{x})), d), \end{aligned} \quad (9)$$

where h_f , h_c , and h_{adv} are the feature extractor, the classification layer, and the domain discriminator, respectively, and d denotes the domain label. To optimize Eq. (9), we need to iteratively optimize h_f , h_c , and h_{adv} . An alternative for the iterative optimization is gradient reversal layer (GRL) (Ganin and Lempitsky, 2015). The key is to solve the problems caused by the negative sign in Eq. (9). GRL acts as an identity transformation that can be ignored during the forward propagation while it takes the gradient from the subsequent level and changes its sign before passing it to the preceding layer that reverses the gradient sign on h_f during the backpropagation. GRL solves the problems caused by the negative sign in Eq. (9).

Obviously, there can exist conflicts between $\ell_0(h_c(h_f(\mathbf{x})), y)$ and $\ell_1(h_{adv}(h_f(\mathbf{x})), d)$. We need a trade-off between these two items for better generalization. To implement it, we just set $k = 1, m = 2$ and let the regularization item be $\ell_1(h_{adv}(h_f(\mathbf{x})), d)$. The framework is shown in Figure 1. We give some theoretical insights as follows.

Proposition 6. *Let \mathcal{X} be a space and \mathcal{H} be a class of hypotheses corresponding to this space. Let \mathbb{Q} and the collection $\{\mathbb{P}^i\}_{i=1}^M$ be distributions over \mathcal{X} and let $\{\varphi_i\}_{i=1}^M$ be a collection of non-negative coefficient with $\sum_i \varphi_i = 1$. Let the object \mathcal{O} be a set of distributions such that for every $\mathbb{S} \in \mathcal{O}$ the following holds*

$$d_{\mathcal{H}\Delta\mathcal{H}}\left(\sum_i \varphi_i \mathbb{P}^i, \mathbb{S}\right) \leq \max_{i,j} d_{\mathcal{H}\Delta\mathcal{H}}(\mathbb{P}^i, \mathbb{P}^j). \quad (10)$$

Then, for any $h \in \mathcal{H}$,

$$\begin{aligned} \varepsilon_{\mathbb{Q}}(h) \leq \lambda' + \sum_i \varphi_i \varepsilon_{\mathbb{P}^i}(h) + \frac{1}{2} \min_{\mathbb{S} \in \mathcal{O}} d_{\mathcal{H}\Delta\mathcal{H}}(\mathbb{S}, \mathbb{Q}) \\ + \frac{1}{2} \max_{i,j} d_{\mathcal{H}\Delta\mathcal{H}}(\mathbb{P}^i, \mathbb{P}^j) \end{aligned} \quad (11)$$

where λ' is the error of an ideal joint hypothesis.

Proof On one hand, with Thm. 7, we have

$$\varepsilon_{\mathbb{Q}}(h) \leq \lambda_1 + \varepsilon_{\mathbb{S}}(h) + \frac{1}{2} d_{\mathcal{H}\Delta\mathcal{H}}(\mathbb{S}, \mathbb{Q}), \forall h \in \mathcal{H}, \forall \mathbb{S} \in \mathcal{O}. \quad (12)$$

On the other hand, with Thm. 7, we have

$$\varepsilon_{\mathbb{S}}(h) \leq \lambda_2 + \varepsilon_{\sum_i \varphi_i \mathbb{P}^i}(h) + \frac{1}{2} d_{\mathcal{H}\Delta\mathcal{H}}\left(\sum_i \varphi_i \mathbb{P}^i, \mathbb{S}\right), \forall h \in \mathcal{H}. \quad (13)$$

Since $\varepsilon_{\sum_i \varphi_i \mathbb{P}^i}(h) = \sum_i \varphi_i \varepsilon_{\mathbb{P}^i}(h)$, and $d_{\mathcal{H}\Delta\mathcal{H}}(\sum_i \varphi_i \mathbb{P}^i, \mathbb{S}) \leq \max_{i,j} d_{\mathcal{H}\Delta\mathcal{H}}(\mathbb{P}^i, \mathbb{P}^j)$, we have

$$\begin{aligned} \varepsilon_{\mathbb{Q}}(h) &\leq \lambda' + \sum_i \varphi_i \varepsilon_{\mathbb{P}^i}(h) + \frac{1}{2} d_{\mathcal{H}\Delta\mathcal{H}}(\mathbb{S}, \mathbb{Q}) \\ &\quad + \frac{1}{2} \max_{i,j} d_{\mathcal{H}\Delta\mathcal{H}}(\sum_i \varphi_i \mathbb{P}^i, \mathbb{S}), \forall h \in \mathcal{H}, \forall \mathbb{S} \in \mathcal{O}. \end{aligned} \tag{14}$$

Eq. (14) holds for all $\mathbb{S} \in \mathcal{O}$. Proof ends. ■

According to Prop. 6 and Eq. (11), domain generalization aims to reduce loss generated by both classification and alignments. Existing simple fixed alignment methods only focus on part of the overall objects which can bring conflicts when optimization. And thereby they may impede the performance of generalization or classification. In our optimization, introducing gradient-based optimization with adapted Mixup, we can simultaneously optimize the classification loss, $\varepsilon_{\mathbb{P}}$, and alignment loss, $d_{\mathcal{H}\Delta\mathcal{H}}$. Consequently, the upper bound of $\varepsilon_T(h)$ can be tighter.

4.3 Mixup-guided Model Selection

In this section, we will introduce how to adapt Mixup to generate a new validation dataset and why it can be better than the original validation data split from the training part³.

Most validation methods typically perform model selection based on the performance on the validation dataset. However, vanilla Mixup generates mixed labels which can be unsuitable to compute the evaluation metric such as accuracy. To deal with this issue, we slightly adapt Mixup and control the mixed data generation process. Vanilla Mixup randomly chooses two samples from all data while ours randomly chooses two samples between the same classes, which makes the label unique and deterministic on one class:

$$\begin{aligned} \tilde{\mathbf{x}} &= \lambda \mathbf{x}_i + (1 - \lambda) \mathbf{x}_j, \\ \tilde{y} &= y_i = y_j. \end{aligned} \tag{15}$$

We generate the same number of samples as the original validation data for fairness, denoted as **VALID** (Mixup based validation dataset). We denote its distribution as $\mathbb{P}_{XY}^{\text{VALID}}$. In the following, we show why it can be better than original validation data split from training data.

Since we select the model according to the accuracy on validation data, it can be better when validation data has a similar distribution to the unseen target. We mainly discuss two cases: (1) Convex case: the unseen target is in the convex combination of the sources; (2) General case: the general situation where the unseen target can be either inside or outside the convex combination of the sources.⁴ To better demonstrate the advantage of our model selection technique, we describe two cases in Figure 2.

3. Note that we utilize the same amount of data (even the completely same data) for training for fairness although we generate some new samples.

4. Please note that we only deal with covariate shift in this paper which means that no conditional shifts exist.




Figure 2: Toy examples of different validation datasets. (a) Case I, the target is the convex combination of the sources. As we can see from Figure 2(a), the distance between origin validation and the target is fixed while the target can be seen as part of VALD (the yellow part). And VALD can even serve as an unbiased estimation of the target, which means it can get a better estimation of the target. (b) Case II, the target is out of the convex combination of the sources. The distance between the target and VALD is still smaller than the fixed distance between the target and origin validation data.

4.3.1 CONVEX CASE

The \mathcal{H} -divergence between two distributions \mathbb{P}, \mathbb{Q} over a space \mathcal{X} w.r.t. a hypothesis class \mathcal{H} can be computed as (Ben-David et al., 2010):

$$d_{\mathcal{H}}(\mathbb{P}, \mathbb{Q}) = 2 \sup_{h \in \mathcal{H}} |Pr_{\mathbb{P}}(\mathbb{I}_h) - Pr_{\mathbb{Q}}(\mathbb{I}_h)|, \quad (16)$$

where $\mathbb{I}_h = \{\mathbf{x} \in \mathcal{X} | h(\mathbf{x}) = 1\}$. We typically consider the $\mathcal{H}\Delta\mathcal{H}$ -divergence in (Ben-David et al., 2010) where the symmetric difference hypothesis class $\mathcal{H}\Delta\mathcal{H}$ is the set of functions characteristic to disagreements between hypotheses.

Theorem 7. (Theorem 2.1 in (Sicilia et al., 2021), modified from Theorem 2 in (Ben-David et al., 2010)). Let \mathcal{X} be a space and \mathcal{H} be a class of hypotheses corresponding to this space. Suppose \mathbb{P} and \mathbb{Q} are distributions over \mathcal{X} . Then for any $h \in \mathcal{H}$, the following holds

$$\varepsilon_{\mathbb{Q}}(h) \leq \lambda'' + \varepsilon_{\mathbb{P}}(h) + \frac{1}{2}d_{\mathcal{H}\Delta\mathcal{H}}(\mathbb{Q}, \mathbb{P}) \quad (17)$$

with λ'' the error of an ideal joint hypothesis for \mathbb{Q}, \mathbb{P} .

Since λ'' is small in the covariate shift, two items, $\varepsilon_{\mathbb{P}}(h)$ and $\frac{1}{2}d_{\mathcal{H}\Delta\mathcal{H}}(\mathbb{Q}, \mathbb{P})$, dominate the target error $\varepsilon_{\mathbb{Q}}(h)$. Here (\mathbb{Q} corresponds to the target \mathbb{P}_{XY}^T while \mathbb{P} corresponds to the

validation distribution. We expect $\varepsilon_{\mathbb{Q}}(h)$ small, and we select the model according to $\varepsilon_{\mathbb{P}}(h)$. Therefore, to obtain more stable and accurate $\varepsilon_{\mathbb{Q}}(h)$, we expect $d_{\mathcal{H}\Delta\mathcal{H}}(\mathbb{Q}, \mathbb{P})$ small. For origin validation data, $\mathbb{P}_{val}^S = \sum_{i=1}^m \pi_i \mathbb{P}_{XY}^i$ where all π are fixed. For \mathbb{Q} , since in the first cases, $\mathbb{P}^T = \sum_{i=1}^m \phi_i \mathbb{P}_{XY}^i$ where ϕ is fixed and $\sum_i \phi_i = 1$. For \mathbb{P}^{VALD} , $\mathbb{P}^{VALD} = \sum_{i=1}^m \varphi_i \mathbb{P}_{XY}^i$ where φ can be dynamic (our adapted Mixup covers all convex combinations) and $\sum_i \varphi_i = 1$. Therefore, $d_{\mathcal{H}\Delta\mathcal{H}}(\mathbb{P}^T, \mathbb{P}_{val}^S) = C$, C is constant while $d_{\mathcal{H}\Delta\mathcal{H}}(\mathbb{P}^T, \mathbb{P}^{VALD})$ can be small enough which means evaluation on VALD achieves a more accurate estimation on the unseen target.

4.3.2 GENERAL CASE

The general case is more difficult.

For VALD, φ can be viewed as dynamic, which means VALD covers the origin validation data distribution and has a tighter upper bound in Eq. (11) (shown in Figure 2(b)).

4.4 Method Summary

Algorithm 1 The process of our methods.

Input: A model h , data of M sources $\{\mathcal{D}_i\}_{i=1}^M$, α

Output: Well trained model h^*

- 1: Initial h_f, h_c .
 - 2: Initial $h^* = h, bestv = 0$.
 - 3: Generate VALD according to Eq. (15).
 - 4: Generate OPTD according to Eq. (4) and Eq. (5).
 - 5: **while** not convergence and not reach the max iteration **do**
 - 6: **if** Update ω **then**
 - 7: Compute $\mathbf{G}, \mathbf{g}_{optd}$ with current h_f .
 - 8: Compute ω according to Eq. (6).
 - 9: Compute ℓ according to ω .
 - 10: Update h according to ℓ .
 - 11: Compute acc_{vald} , accuracy on VALD.
 - 12: **if** $acc_{vald} > bestv$ **then**
 - 13: $bestv = acc_{vald}$.
 - 14: $h^* = h$.
-

Algorithm 1 gives the overall process of our techniques. We update ω every B iterations, where B can be set arbitrarily. As shown in Algorithm 1, we first generate VALD and OPTD. When optimization, we first obtain ω and then utilize ω to weight different objects. After updating the model h , we evaluate it on VALD and record the best one according to the accuracy on VALD.

4.5 Discussion

In the above implementation, we mainly rely on DANN. Actually, our techniques not only work based on DANN but also can be able to enhance other traditional methods, e.g.

CORAL. For CORAL, we only need to replace the adversarial part of DANN with the covariance alignment. We will show the performance for CORAL in experiments.

Surprisingly, our techniques can even play an important role in improving ERM. Each source can be viewed as an independent goal, and thereby we have M objects when we have M sources. OPTD guides the weighting of different sources when training for better generalization and it can be considered as a sample or domain weighting technique.

In our implementation, we mainly utilize the gradients of all data in OPTD for guidance, which can save time. But utilizing the gradients of batch data in OPTD and lasting dynamic changes for ω can be another possible way for the implementation.

5 Experiment

We evaluate the proposed techniques on one visual classification benchmark and three time-series benchmarks.

5.1 Visual Classification

5.1.1 DATASET

PACS (Li et al., 2017) is an object classification benchmark. It contains four domains, including photos, art-paintings, cartoons, and sketches. Among different domains, image styles have large discrepancies. There exist 9,991 images totally and each domain has the same seven classes.

5.1.2 EXPERIMENTAL SETUP

For visual classification, ResNet-18 is applied as the feature net. We compared our technique with seven popular state-of-the-art methods, including ERM, DANN (Ganin and Lempitsky, 2015), CORAL (Sun and Saenko, 2016), Mixup (Zhang et al., 2018), GroupDRO (Sagawa et al., 2020), RSC (Huang et al., 2020), and ANDMask (Parascandolo et al., 2021). For all these methods, we re-implement with Pytorch (Paszke et al., 2019) in the same environment for fairness. We split each source domain with a ratio of 8:2 for training and validation. The best model can be selected via results on validation datasets. In each step, each domain selects 32 samples. The maximum training epoch is set to 120. For all methods, the SGD optimizer with an initial learning rate 10^{-3} and weight decay 5×10^{-4} is used. The learning rate drops by 0.1 at the 70% and 90% of training epochs respectively. We tune hyperparameters for each method and select the best results to report.

5.1.3 EXPERIMENTAL RESULTS

The results on PACS are shown in Table 1. On average, our proposed techniques improve DANN and outperform the second-best method, 0.86%. We observe some more insightful conclusions. (1) Will special designed methods always work? The answer is obviously no. For the first task on PACS, ERM without any artificial designs even performs best. Two latest methods, Mixup and ANDMask, have similar performances to ERM. (2) Is the improvement, 0.8%, significant? This answer is uncertain. PACS is a difficult task in DG. Compared to the baseline, ERM, other state-of-the-art methods only have slight

Table 1: Results on PACS. The **bold** items are the best results.

	A	C	P	S	AVG
ERM	81.84	74.45	<u>96.35</u>	70.40	80.76
CORAL	79.98	74.70	93.77	<u>79.51</u>	81.99
Mixup	79.83	72.06	95.09	76.92	80.97
GroupDRO	78.03	73.25	93.47	80.48	81.31
RSC	<u>81.59</u>	<u>75.64</u>	96.71	72.92	81.71
ANDMask	79.98	74.15	95.87	73.71	80.93
DANN	81.30	75.30	95.51	76.30	<u>82.10</u>
DANN+Ours	81.84	77.01	94.85	78.14	82.96

improvements. The largest improvement cannot reach 1.5%. Therefore, 0.8% can be an acceptable improvement. However, from the view of absolute value, 0.8% is a small value. (3) When do our techniques fail or behave normally? Since our techniques are based on Mixup, they can be inevitably affected by the performance of Mixup. When Mixup performs terribly, ours cannot have significant improvements. However, from Table 1, we can see that ours with DANN still achieve the best. The results demonstrate the effectiveness and superiority of our techniques. For visual classification, it seems that direct Mixup cannot ensure good results. And in the future, we will design better and more suitable data generation methods for visual classification.

5.2 Time-series Classification

5.2.1 DATASETS

UCI daily and sports dataset (**DSADS**) (Barshan and Yksek, 2014) contains data with 19 activities collected from 8 subjects wearing body-worn sensors on 5 body parts. There exist three sensors, accelerometer, gyroscope, and magnetometer. 19 activities include sitting, standing, lying on back and on right side, ascending and descending stairs, standing in an elevator still, moving around in an elevator, walking in a parking lot, walking on a treadmill with a speed of 4 km/h, running on a treadmill with a speed of 8 km/h, exercising on a stepper, exercising on a cross trainer, cycling on an exercise bike in horizontal and vertical positions, rowing, jumping, and playing basketball. We divide DSADS into four domains according to subjects and each domain contains two subjects, $[(0, 1), (2, 3), (4, 5), (6, 7)]$ where the digit is the subject number. Therefore, we construct four domains and different domains have different distributions (In some papers (Wang et al., 2018), it is also called Cross-Person.). We use 0, 1, 2, 3 to denote the four divided domains.

USC-SIPI human activity dataset (**USC-HAD**) (Zhang and Sawchuk, 2012) contains data of 14 subjects (7 male, 7 female, aged from 21 to 49) executing 12 activities with a sensor tied on the front right hip. 12 activities include Walking Forward, Walking Left, Walking Right, Walking Upstairs, Walking Downstairs, Running Forward, Jumping Up, Sitting, Standing, Sleeping, Elevator Up, and Elevator Down. The data dimension is 6 and the sample rate is 100Hz. Similar to DSADS, we divide data into four domains,

$[(1, 11, 2, 0), (6, 3, 9, 5), (7, 13, 8, 10), (4, 12)]$. We try our best to ensure that each domain has a similar number of data.

Table 2: Detailed information on three time-series benchmarks.

Dataset	#Domain	#Sensor	#Class	#Domain Sample	#Total
DSADS	4	3	19	$(285,000) \times 4$	1,140,000
PAMAP2	4	3	12	(592,600; 622,200; 620,000; 623,400)	2,458,200
USC-HAD	4	2	12	(1,401,400;1,478,000;1,522,800;1,038,800)	5,441,000

PAMAP2 physical activity monitoring dataset (**PAMAP2**) (Reiss and Stricker, 2012) contains data of 18 different physical activities, performed by 9 subjects wearing 3 sensors. 18 activities include lying, sitting, standing, walking, running, cycling, Nordic walking, watching TV, computer work, car driving, ascending stairs, descending stairs, vacuum cleaning, ironing, folding laundry, house cleaning, playing soccer, rope jumping, and other (transient activities). The sampling frequency is 100Hz and the data dimension is 27. Similar to DSADS, we divide data into four domains, $[(3, 2, 8), (1, 5), (0, 7), (4, 6)]$.

The detailed information of the datasets is in Table 2.

5.2.2 EXPERIMENTAL SETUP

We adopt sliding window (Bulling et al., 2014) with 50% overlap to construct samples. We compared our technique with four popular state-of-the-art methods⁵:

- ERM, a method that combines all source data together and directly trains the model.
- DANN (Ganin and Lempitsky, 2015), a method that learns domain-invariant features in an adversarial way.
- ANDMask (Parascandolo et al., 2021), a method that learns domain-invariant features based on gradients.
- GILE (Qian et al., 2021), a method that utilizes VAE to decouple domain and classification features.

For these comparison methods, we split data of the source domains into the training splits and validation splits. The training splits are used to train the model while the validation splits are utilized to select the best model. In practice, 80% of all source data serve as training while the rest are for validation. Although our techniques do not require the validation splits (we generate VALD), we still only utilize the same training splits as comparison methods for fairness. For testing, all methods, including ours, report the performance on all data of the target domain.

We implement all methods with PyTorch (Paszke et al., 2019). For GILE, we directly utilize their public code. For the other methods, we utilize the same architecture that contains two blocks, and each has one convolution layer, one pool layer, and one batch normalization layer. Another single fully-connect layer serves as the classification layer. We utilize a batch with 32 samples for each domain in each iteration, and the maximum

5. Our proposed techniques can be embedded in many methods, and they can be viewed as plugins. Therefore, we do not compare ours to lots of methods but focus on the methods enhanced by our techniques.

training epoch is set to 150 (15000 iterations). An Adam optimizer with a learning rate 10^{-2} and weight decay 5×4^{-4} is used for optimization. We tune hyperparameters for each method and report the average results of three trials.

5.2.3 EXPERIMENTAL RESULTS

Table 3: Results on DSADS, USC-HAD, and PAMAP2. The **bold** items are the best results.

Methods	DSADS					USC-HAD					PAMAP2				
	0	1	2	3	AVG	0	1	2	3	AVG	0	1	2	3	AVG
ERM	89.69	81.45	81.05	78.20	82.60	80.33	59.88	74.15	73.93	72.07	87.28	73.10	49.03	78.76	72.04
ANDMask	85.35	73.07	85.04	82.06	81.38	79.51	61.53	76.32	65.52	70.72	88.22	79.11	53.35	83.22	75.97
GILE	79.67	75.00	77.00	67.00	74.65	78.67	63.00	77.00	61.67	70.08	83.33	68.67	44.00	76.67	68.25
DANN	87.54	81.27	78.42	83.03	82.57	81.33	64.02	72.91	66.37	71.16	88.93	75.60	47.35	86.78	74.66
DANN+Ours	93.33	88.77	91.75	84.78	89.66	81.98	64.32	74.84	78.40	74.89	89.23	81.36	61.71	89.28	80.40

The results on DSADS, USC-HAD, PAMAP2 are shown in Table 3. On average, our proposed techniques substantially improve DANN, and outperform the second-best methods: 7.36% for DSADS, 2.82% for USC-HAD, and 4.43% for PAMAP2. Compared to vanilla DANN, ours has a larger improvement, showing the advantage of our model optimization and selection techniques.

We observe some more insightful conclusions. (1) When do our techniques work? Our methods work in almost all situations if a correct way can be adopted. From all three tables, we can see that our techniques improve vanilla DANN on every task, which demonstrates superiority. DANN+Ours performs slightly worse than GILE in the third task for USC-HAD, it can be due to that vanilla DANN performs worst in this task. Our method can only improve DANN, but cannot completely get rid of influence from DANN. To pursue better performance, we will introduce our techniques to some latest methods, which is our future work. (2) When do our techniques perform mediocly? Our techniques have dramatic improvements for some tasks, e.g. the first task for DSADS, while these two techniques perform mediocly for some other tasks, e.g. the second task for USC-HAD. There can be many factors being able to affect performance, e.g. randomness, task difficulty, sample volume, and so on. For example, when the target domain is far away from sources, OPTD and VALD cannot represent them. In some extreme situations, the performance on target data even cannot equal the generalization capability, since it is a really difficult problem. Moreover, Mixup can be too simple to generate good enough data. (3) Can alignments or some other generalization methods always have improvements? The answer is obviously no. Many factors influence the final generalization, e.g. data quantity, diversity, and distribution discrepancy. On some tasks, such as the first task for DSADS and the third task for USC-HAD, ERM even performs better than DANN. These results demonstrate that generalization methods, e.g. alignments, may have a negative influence on classification capability, which proves our motivation again.

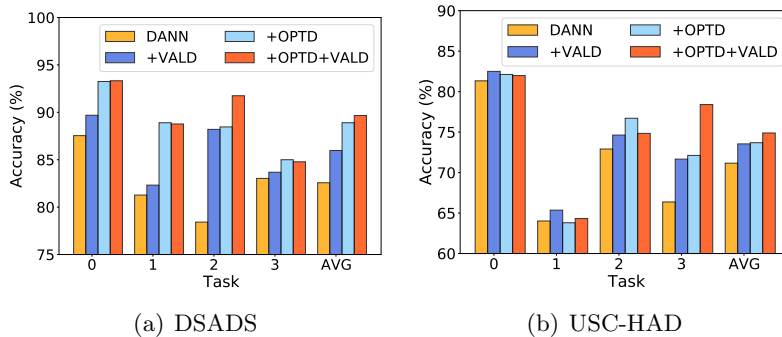


Figure 3: Ablation study on DSADS and USC-HAD.

Table 4: The results for our techniques with CORAL.

DSADS					
ERM	89.69	81.45	81.05	78.20	82.60
CORAL	91.10	85.79	87.28	82.24	86.60
CORAL+Ours	93.11	93.07	89.43	84.43	90.01
USC-HAD					
ERM	80.33	59.88	74.15	73.93	72.07
CORAL	81.88	60.23	74.94	68.81	71.47
CORAL+Ours	82.05	64.87	78.17	76.61	75.43
PAMAP2					
ERM	87.28	73.10	49.03	78.76	72.04
CORAL	88.26	79.91	59.39	85.85	78.35
CORAL+Ours	89.84	79.27	57.81	88.03	78.74

5.3 Analysis

5.3.1 ABLATION STUDY

We perform ablation study in this section and the results are shown in Figure 3. In Figure 3(a), with OPTD or VALD, we can see that our optimization technique achieves a remarkable improvement compared to vanilla DANN. Moreover, with both two techniques, there exists another improvement on average. However, when taking a closer look at each task, we can find that OPTD or VALD can lead slight performance drops in some tasks. The above phenomenon can be normal since we have analyzed above that there can be many factors, e.g. distribution discrepancy, influence performance, and Mixup cannot always work well due to its simplicity and uncertainty. Moreover, just as Figure 3(a), drops are so little that they can be ignored and there exists an improvement overall, which demonstrates the effects of both OPTD and VALD. These results prove that both OPTD and VALD have positive effects on performance for DG. Similar arguments can be concluded in Figure 3(b).

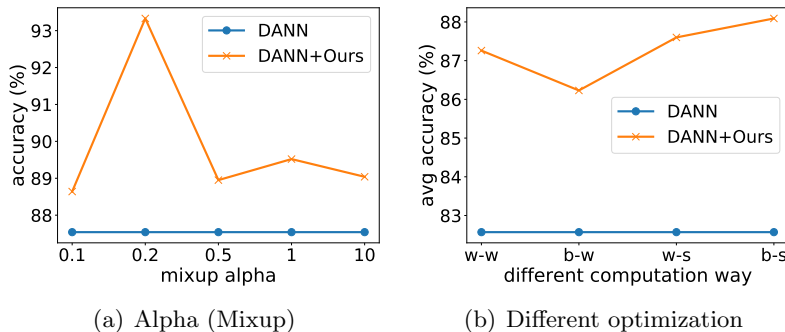


Figure 4: Parameter sensitivity on DSADS.

5.3.2 DISCUSSION

To demonstrate that our techniques are not only designed for DANN but also can be utilized in some other methods, we also embed them into CORAL, another popular domain generalization method. The results are shown in Table 4. In Table 4, we can see that our techniques have significant improvements on DSADS and USC-HAD while they also have a slight improvement on PAMAP2, which demonstrate that our techniques can be universal and useful. We have some more observations. 1) For some tasks, e.g. the first task and the second task for PAMAP2, our techniques even reduce performances, which confirms our analysis mentioned above again. Domain generalization is a difficult problem and there can be many factors that can have an influence on the final performance. 2) CORAL with our techniques even have better performance compared to DANN with our techniques, which demonstrates the effectiveness and implicit room of improvement for our techniques. We will plug our techniques into more latest methods in the future.

5.3.3 PARAMETER SENSITIVITY ANALYSIS

We evaluate the parameter sensitivity of our technique in Figure 4. There are mainly two kinds of hyperparameters in our techniques, α for Mixup and computation ways for optimization.⁶ In Figure 4(a), we can see that results with different Mixup hyperparameters all have improvements compared to vanilla DANN, which demonstrates the superiority of our techniques. In Figure 4(b), different computation ways for optimization mean different ways to compute gradients and different ways to weigh the objects. w-w represents computing the mean gradients of the whole OPTD and viewing the classification of sources as a whole objective (in this case, there are two objectives in total). b-w represents computing the mean gradients of a batch in OPTD and viewing the classification of sources as a whole objective. w-s represents computing the mean gradients of the whole OPTD and viewing the classification of each source as an independent objective (in this case, there are four objectives in total). From Figure 4(b), we can see that our techniques achieve remarkable improvements whatever computation way adopted, which demonstrates the robustness and

6. For simplicity, we utilize the same α for OPTD and VALD. We believe there can be further improvements with finer tuning.

superiority of our techniques. In a nutshell, the results demonstrate that our techniques can be effective and robust that can be easily applied to methods in domain generalization.

5.4 Discussion

Please note that this paper is not for finding state-of-the-art methods in domain generalization, but pays all attention to exploring optimization and model selection for DG. Optimization and model selection is a significant topic in DG but it is still in the infant. To the best of our knowledge, this paper is the *first* work towards solving these challenges simultaneously. We choose DANN and CORAL to implement our techniques for the following three reasons.

- Firstly, DANN and CORAL are two traditional methods and their performances are satisfactory according to DomainBed (Gulrajani and Lopez-Paz, 2021).
- Secondly, they are simple and easy to embed our two techniques.
- Thirdly, with our two techniques, these two methods have shown remarkable performance.

We believe that more techniques will emerge in this extremely important field.

6 Conclusion

In this paper, we proposed two Mixup based techniques for model optimization and selection in domain generalization, which are two emerging fields lacking enough studies. Specifically, on one hand, we utilize adapted Mixup to generate OPTD and then utilize gradients of OPTD to guide the balance among different objects for domain generalization. On the other hand, we generate VALD via another adapted Mixup, and select the best model with VALD. Extensive experiments on one visual classification benchmark and three time-series benchmarks demonstrated the effectiveness of these two techniques. In the future, we plan to generate more out of distribution data to guide the optimization and selection of generalized learning.

References

- Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.
- Billur Barshan and Murat Cihan Yükses. Recognizing daily and sports activities in two open source machine learning environments using body-worn sensor units. *The Computer Journal*, 57(11):1649–1667, 2014.
- Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79(1):151–175, 2010.
- Andreas Bulling, Ulf Blanke, and Bernt Schiele. A tutorial on human activity recognition using body-worn inertial sensors. *ACM Computing Surveys (CSUR)*, 46(3):1–33, 2014.

- Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich. Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In *International Conference on Machine Learning*, pages 794–803. PMLR, 2018.
- Michael B Cohen, Yin Tat Lee, and Zhao Song. Solving linear programs in the current matrix multiplication time. *Journal of the ACM (JACM)*, 68(1):1–39, 2021.
- Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR, 2015.
- Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- J Han, K Qian, M Song, Z Yang, Z Ren, S Liu, J Liu, H Zheng, W Ji, T Koike, et al. An early study on intelligent analysis of speech under covid-19: Severity, sleep quality, fatigue, and anxiety. In *INTERSPEECH*, pages 4946–4950, 2020.
- Zeyi Huang, Haohan Wang, Eric P Xing, and Dong Huang. Self-challenging improves cross-domain generalization. In *European Conference on Computer Vision*, pages 124–140. Springer, 2020.
- Yunpei Jia, Jie Zhang, Shiguang Shan, and Xilin Chen. Single-side domain generalization for face anti-spoofing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8484–8493, 2020.
- Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7482–7491, 2018.
- Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generalization. In *Proceedings of the IEEE international conference on computer vision*, pages 5542–5550, 2017.
- Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Learning to generalize: Meta-learning for domain generalization. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Da Li, Henry Gouk, and Timothy Hospedales. Finding lost dg: Explaining domain generalization via model complexity. *arXiv preprint arXiv:2202.00563*, 2022.
- Fangrui Lv, Jian Liang, Kaixiong Gong, Shuang Li, Chi Harold Liu, Han Li, Di Liu, and Guoren Wang. Pareto domain adaptation. In *Advances in Neural Information Processing Systems*, 2021.
- Divyat Mahajan, Shruti Tople, and Amit Sharma. Domain generalization using causal matching. In *International Conference on Machine Learning*, pages 7313–7324. PMLR, 2021.

- Debabrata Mahapatra and Vaibhav Rajan. Multi-task learning with user preferences: Gradient descent with controlled ascent in pareto optimization. In *International Conference on Machine Learning*, pages 6597–6607. PMLR, 2020.
- Aviv Navon, Idan Achituve, Haggai Maron, Gal Chechik, and Ethan Fetaya. Auxiliary learning by implicit differentiation. In *International Conference on Learning Representations*, 2020.
- Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2009.
- Giambattista Parascandolo, Alexander Neitz, Antonio Orvieto, Luigi Gresele, and Bernhard Schölkopf. Learning explanations that are hard to vary. In *9th International Conference on Learning Representations, ICLR*, 2021.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*, volume 32, pages 8026–8037, 2019.
- Xi Peng, Fengchun Qiao, and Long Zhao. Out-of-domain generalization from a single source: A uncertainty quantification approach. *arXiv preprint arXiv:2108.02888*, 2021.
- Mirco Planamente, Chiara Plizzari, Emanuele Alberti, and Barbara Caputo. Domain generalization through audio-visual relative norm alignment in first person action recognition. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1807–1818, 2022.
- Hangwei Qian, Sinno Jialin Pan, and Chunyan Miao. Latent independent excitation for generalizable sensor-based cross-person activity recognition. In *AAAI*, volume 35, pages 11921–11929, 2021.
- Alexandre Rame, Corentin Dancette, and Matthieu Cord. Fishr: Invariant gradient variances for out-of-distribution generalization. *arXiv:2109.02934*, 2021.
- Payam Refaeilzadeh, Lei Tang, and Huan Liu. Cross-validation. *Encyclopedia of database systems*, 5:532–538, 2009.
- Attila Reiss and Didier Stricker. Introducing a new benchmarked dataset for activity monitoring. In *2012 16th international symposium on wearable computers*, pages 108–109. IEEE, 2012.
- Matthew W Robards and Peter Sunehag. Semi-markov kmeans clustering and activity recognition from body-worn sensors. In *2009 Ninth IEEE International Conference on Data Mining*, pages 438–446. IEEE, 2009.
- Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. In *International Conference on Learning Representations (ICLR)*, 2020.

- Shiv Shankar, Vihari Piratla, Soumen Chakrabarti, Siddhartha Chaudhuri, Preethi Jyothi, and Sunita Sarawagi. Generalizing across domains via cross-gradient training. In *International Conference on Learning Representations*, 2018.
- Yuge Shi, Jeffrey Seely, Philip HS Torr, N Siddharth, Awni Hannun, Nicolas Usunier, and Gabriel Synnaeve. Gradient matching for domain generalization. *arXiv preprint arXiv:2104.09937*, 2021.
- Anthony Sicilia, Xingchen Zhao, and Seong Jae Hwang. Domain adversarial neural networks for domain generalization: When it works and how to improve. *arXiv preprint arXiv:2102.03924*, 2021.
- Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *ECCV*, pages 443–450. Springer, 2016.
- Fengyi Tang, Lifan Zeng, Fei Wang, and Jiayu Zhou. Adversarial precision sensing with healthcare applications. In *2020 IEEE International Conference on Data Mining (ICDM)*, pages 521–530. IEEE, 2020.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Jindong Wang, Yiqiang Chen, Lisha Hu, Xiaohui Peng, and S Yu Philip. Stratified transfer learning for cross-domain activity recognition. In *2018 IEEE international conference on pervasive computing and communications (PerCom)*, pages 1–10. IEEE, 2018.
- Jindong Wang, Cuiling Lan, Chang Liu, Yidong Ouyang, Wenjun Zeng, and Tao Qin. Generalizing to unseen domains: A survey on domain generalization. *arXiv preprint arXiv:2103.03097*, 2021.
- Wen Wang, Ruiping Wang, Shiguang Shan, and Xilin Chen. Discriminative covariance oriented representation learning for face recognition with image sets. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5599–5608, 2017.
- Garrett Wilson and Diane J Cook. A survey of unsupervised deep domain adaptation. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 11(5):1–46, 2020.
- Qinwei Xu, Ruipeng Zhang, Ya Zhang, Yanfeng Wang, and Qi Tian. A fourier-based framework for domain generalization. In *CVPR*, pages 14383–14392, 2021.
- Huaxiu Yao, Yu Wang, Sai Li, Linjun Zhang, Weixin Liang, James Zou, and Chelsea Finn. Improving out-of-distribution robustness via selective augmentation. *arXiv preprint arXiv:2201.00299*, 2022.
- Haotian Ye, Chuanlong Xie, Tianle Cai, Ruichen Li, Zhenguo Li, and Liwei Wang. Towards a theoretical framework of out-of-distribution generalization. *Advances in Neural Information Processing Systems*, 34, 2021.

- Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations*, 2018.
- Mi Zhang and Alexander A Sawchuk. Usc-had: a daily activity dataset for ubiquitous activity recognition using wearable sensors. In *Proceedings of the 2012 ACM conference on ubiquitous computing*, pages 1036–1043, 2012.
- Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain generalization with mixstyle. In *International Conference on Learning Representations*, 2021.
- Eckart Zitzler and Lothar Thiele. Multiobjective evolutionary algorithms: a comparative case study and the strength pareto approach. *IEEE transactions on Evolutionary Computation*, 3(4):257–271, 1999.
- Walter Zucchini. An introduction to model selection. *Journal of mathematical psychology*, 44(1):41–61, 2000.