

# Flatness-Aware Minimization for Domain Generalization

Xingxuan Zhang, Renzhe Xu, Han Yu, Yancheng Dong, Pengfei Tian, Peng Cui\*  
 Department of Computer Science, Tsinghua University  
 xingxuanzhang@hotmail.com, cuip@tsinghua.edu.cn

## Abstract

Domain generalization (DG) seeks to learn robust models that generalize well under unknown distribution shifts. As a critical aspect of DG, optimizer selection has not been explored in depth. Currently, most DG methods follow the widely used benchmark, DomainBed, and utilize Adam as the default optimizer for all datasets. However, we reveal that Adam is not necessarily the optimal choice for the majority of current DG methods and datasets. Based on the perspective of loss landscape flatness, we propose a novel approach, Flatness-Aware Minimization for Domain Generalization (FAD), which can efficiently optimize both zeroth-order and first-order flatness simultaneously for DG. We provide theoretical analyses of the FAD's out-of-distribution (OOD) generalization error and convergence. Our experimental results demonstrate the superiority of FAD on various DG datasets.

## 1. Introduction

Current neural networks are expected to generalize to unseen distributions in real-world applications, which can break the independent and identically distributional (I.I.D.) assumption of traditional machine learning algorithms [129, 99]. The distribution shift between training and test data may largely deteriorate most current approaches in practice. Hence, instead of generalization within the training distribution, the ability to generalize under distribution shift, namely domain generalization (DG) [73], has attracted increasing attention [99].

Various methods have been proposed to address the DG problem recently, many of which have shown promising performance. As a learning problem mainly in the computer vision field, the DG task relies highly on the chosen optimizer, yet the effect of the optimizer in DG has hardly been studied. Most current algorithms, following the well-known benchmark DomainBed [32], default to optimize




Figure 1. Test accuracy on PACS and DomainNet of different optimizers. Adam, the default optimizer in popular benchmarks shows no advantage compared with its counterparts, while the proposed FAD achieve better OOD generalization performance compared with current optimizers.

models with Adam [48] without considering other optimizers. However, the in-distribution generalization ability of popular optimizers has been investigated from the perspective of loss landscape recently [28, 29] and Adam is found to achieve inferior generalization compared with other optimizers such as SGD [76] due to the sharp minima Adam selected. Thus, investigating the impact of optimizers in DG and clarifying whether Adam should be considered as the default optimizer is of significance.

In this paper, we empirically compare the performance of current optimizers, including Adam, SGD, AdaBelief [131], YOGI [117], AdaHessian [116], SAM [28], and GAM[121]. Through extensive experiments on current DG benchmarks, we show that Adam can hardly surpass other optimizers on most DG datasets. The best choice of the optimizer for different DG methods varies across different datasets, and choosing the right optimizer can help exceed the current leaderboards' limitations.

Recently, the connection between the flatness of the loss landscape and in-distribution generalization ability has been widely studied and verified both theoretically[130] and empirically[25]. Some works [13, 87] show that flatness also leads to superior OOD generalization. However, none of the previous works consider the optimization in DG nor provide theoretical assurance of their approach.

\*Corresponding author

Among flatness-aware optimization methods, sharpness-Aware Minimization (SAM) [28] and its variants [130, 53, 126, 47], which optimize the zeroth-order flatness, achieve SOTA performance on various in-distribution image classification tasks [53, 130]. Most recently, Gradient Norm Aware Minimization (GAM) [121] shows that minimizing the first-order flatness provides a more substantial penalty on the sharpness of minima than zeroth-order flatness yet requires more computation overhead. Due to the first-order Taylor expansions adopted to optimize zeroth-order and first-order flatness, both SAM and GAM lose some of their penalty strength, and thus they can be mutually reinforced by each other. To accelerate the optimization of first-order flatness and combine zero-order flatness, we propose a unified optimization framework, Flatness-Aware minimization for Domain generalization (FAD), which eliminates the considerable computation overhead for Hessian or Hessian-vector products. We theoretically show that the proposed FAD controls the dominant eigenvalue of Hessian of the training loss, which indicates the sharpness of the loss landscape along its most critical direction [42]. We present an OOD generalization bound to show that the proposed FAD guarantees the generalization error on test data and convergence analysis. Through extensive experiments, we evaluate FAD and other optimizers on various DG benchmarks.

We summarize the main contribution as follows.

- We propose a unified optimization framework, Flatness-Aware minimization for Domain generalization (FAD), to optimize zeroth-order and first-order flatness simultaneously. Without calculating Hessian and Hessian-vector products, FAD considerably reduces the computation overhead of first-order flatness.
- We theoretically analyze the OOD generalization error and convergence of FAD. We show that FAD controls the dominant eigenvalue of Hessian and thus the flatness of learned minima.
- We empirically show that Adam, the default optimizer for most current DG benchmarks, can hardly be the optimal choice for most DG datasets. We present the superiority of FAD compared with current optimizers through extensive experiments.
- We empirically validate that FAD finds flatter optima with lower Hessian spectra compared with zeroth- and first-order flatness-aware optimization methods.

## 2. Related Works

### 2.1. Optimizers

Some works [107, 28] have studied the connection between current optimization approaches, such as SGD [76], Adam [49], AdamW [65] and others [26, 63] and in-distribution generalization. Some literature shows that

Adam is more vulnerable to sharp minima than SGD [104], which may result in worse generalization [106, 33, 35]. Some follow-ups [66, 16, 107, 117] propose generalizable optimizers to address this problem. However, the generalization ability and convergence speed is often a trade-off [44, 107, 63, 117, 26]. Different optimizers may favor different tasks and network architectures (e.g., SGD is often chosen for ResNet [34] while AdamW [65] for ViTs [21]). Selecting the right optimizer is critical for performance and convergence while the understanding of its relationship to model generalization remains nascent [28]. [74] finds that fine-tuned SGD outperforms Adam on OOD tasks. Yet all the previous works do not consider the effect of optimizers in current DG benchmarks and evaluation protocols. In this paper, we discuss the selection of the default optimizer in DG benchmarks.

### 2.2. Domain Generalization and Optimization

Domain generalization (DG) aims to improve the generalization ability to novel domains [31, 45, 50]. A common approach is to learn domain-invariant or causal features over multiple source domains [22, 36, 57, 59, 72, 80, 89, 120, 110, 118, 101, 123] or aggregating domain-specific modules [67, 68]. Some works propose to enlarge the input space by augmentation of training data [12, 98, 81, 91, 127, 128, 109]. Exploit regularization with meta-learning [22, 58] and Invariant Risk Minimization (IRM) framework [2] are also proposed for DG. Recently, several works propose to ensemble model weights for better generalization [13, 4, 84, 105, 86, 62, 18] and achieve outstanding performance.

### 2.3. Flatness of loss landscape

Recent works study the relationship between flatter minima lead and better generalization on in-distribution data [44, 130, 41, 79]. [42] reviews the literature related to generalization and the sharpness of minima. It highlights the role of maximum Hessian eigenvalue in deciding the sharpness of minima [44, 103]. Some simple strategies are proposed to optimize maximum Hessian eigenvalue, such as choosing a large learning rate [54, 19, 40] and smaller batch size [94, 54, 39]. Sharpness-Aware Minimization (SAM) [28] and its variants [130, 53, 24, 64, 25, 71, 126, 47, 46], which are representative training algorithm to seek zeroth-order flatness, show outstanding performance on in-distribution generalization. Most recently, several works discuss the relationship between generalization and gradient norm [8, 125]. GAM[121] proposes to optimize first-order flatness for better generalization. [13, 87] show that flatness minima also lead to superior OOD generalization. However, none of the previous works consider the optimization in DG nor provide theoretical assurance of their approach.

### 3. Method

In in-distribution generalization tasks, sharpness-aware optimization methods show outstanding performance [28]. We theoretically show that optimizing both zeroth-order and first-order flatness strengthens out-of-distribution performance via a generalization bound. We propose a unified optimization framework to optimize zeroth-order and first-order flatness effectively and efficiently. We present the connection between Hessian spectra and our regularization term and give the convergence analysis of our method.

**Notations** We use  $\mathcal{X}$  and  $\mathcal{Y}$  to denote the space of input  $X$  and outcome  $Y$ , respectively. We use  $\Delta_{\mathcal{Y}}$  to denote a distribution on  $\mathcal{Y}$ . Let  $S = \{(x_i, y_i)\}_{i=1}^n$  denote the training dataset with  $n$  data-points drawn independently from the training distribution  $\mathcal{D}_{\text{tr}}$ . We consider the covariate shift scenario [10, 92, 111] and assume that the conditional probabilities are invariant across training and test distributions (i.e.,  $P_{\text{tr}}(Y|X) = P_{\text{te}}(Y|X)$ ). The invariant conditional probability is denoted as  $P(Y|X)$ . A prediction model  $f_{\theta} : \mathcal{X} \rightarrow \Delta_{\mathcal{Y}}$  parametrized by  $\theta \in \Theta \subseteq \mathbb{R}^d$  with dimension  $d$  maps an input to a simplex on  $\mathcal{Y}$ , which indicates the predicted probability of each class<sup>1</sup>.  $\Theta$  is the hypothesis set. We further assume that the ground truth model  $\theta^*$  that characterizes  $P(Y|X)$  is in the hypothesis set  $\Theta$ .

Let  $\ell : \Delta_{\mathcal{Y}} \times \Delta_{\mathcal{Y}} \rightarrow \mathbb{R}_+$  define a loss function over  $\Delta_{\mathcal{Y}}$ . For any hypotheses  $\theta_1, \theta_2 \in \Theta$ , the expected loss  $\mathcal{L}_{\mathcal{D}}(\theta_1, \theta_2)$  for distribution  $\mathcal{D}$  is given as  $\mathcal{L}_{\mathcal{D}}(\theta_1, \theta_2) = \mathbb{E}_{x \sim \mathcal{D}} [\ell(f_{\theta_1}(x), f_{\theta_2}(x))]$ . To simplify the notations, we use  $\mathcal{L}_{\text{tr}}$  and  $\mathcal{L}_{\text{te}}$  to denote the expected loss  $\mathcal{L}_{\mathcal{D}_{\text{tr}}}$  and  $\mathcal{L}_{\mathcal{D}_{\text{te}}}$  in training and test distributions, respectively. In addition, we use  $\mathcal{E}_{\text{tr}}(\theta) = \mathcal{L}_{\text{tr}}(\theta, \theta^*)$  and  $\mathcal{E}_{\text{te}}(\theta) = \mathcal{L}_{\text{te}}(\theta, \theta^*)$  to denote the loss of a function  $\theta \in \Theta$  w.r.t. the true labeling function  $\theta^*$  in the two distributions. We use  $\hat{\mathcal{E}}_{\text{tr}}(\theta)$  to denote the empirical loss function w.r.t. the  $n$  samples  $S$  from the training distribution. In addition, we assume that  $\hat{\mathcal{E}}_{\text{tr}}(\theta)$  is twice differentiable throughout the paper and we use  $\nabla \hat{\mathcal{E}}_{\text{tr}}(\theta)$  and  $\nabla^2 \hat{\mathcal{E}}_{\text{tr}}(\theta)$  to denote the derivative and Hessian matrix, respectively.

#### 3.1. A General Flatness-Aware Optimization for DG

In in-distribution generalization tasks, it is shown that the zeroth-order flatness can be not strong enough in cases where  $\rho$  covers multiple minima or the maximum loss in  $\rho$  is misaligned with the uptrend of loss and first-order flatness can help improve it [121]. But both zeroth-order and first-order flatness require first-order approximations in practice. Thus optimizing both zeroth-order and first-order flatness together leads to stronger generalization than both of them

<sup>1</sup>We use  $\Delta_{\mathcal{Y}}$  to denote the non-deterministic function case. This formulation also includes deterministic function cases.

[121]. Yet optimizing first-order flatness requires the calculation of Hessian or Hessian-vector products, which introduces heavy computation overhead. As a result, we aim to accelerate the optimization in a Hessian-free way for OOD generalization.

Formally, the zeroth-order and first-order flatness are defined as follows [28, 130, 121].

**Definition 3.1** ( $\rho$ -zeroth-order flatness). For any  $\rho > 0$ , the  $\rho$ -zeroth-order flatness  $R_{\rho}^{(0)}(\theta)$  of function  $\hat{\mathcal{E}}_{\text{tr}}(\theta)$  at a point  $\theta$  is defined as

$$R_{\rho}^{(0)}(\theta) \triangleq \max_{\theta' \in B(\theta, \rho)} (\hat{\mathcal{E}}_{\text{tr}}(\theta') - \hat{\mathcal{E}}_{\text{tr}}(\theta)), \quad \forall \theta \in \Theta. \quad (1)$$

Here  $\rho$  is the perturbation radius that controls the magnitude of the neighborhood.

**Definition 3.2** ( $\rho$ -first-order flatness). For any  $\rho > 0$ , the  $\rho$ -first-order flatness  $R_{\rho}^{(1)}(\theta)$  of function  $\hat{\mathcal{E}}_{\text{tr}}(\theta)$  at a point  $\theta$  is defined as

$$R_{\rho}^{(1)}(\theta) \triangleq \rho \cdot \max_{\theta' \in B(\theta, \rho)} \|\nabla \hat{\mathcal{E}}_{\text{tr}}(\theta')\|, \quad \forall \theta \in \Theta. \quad (2)$$

Here  $\rho$  is the perturbation radius similar to Definition 3.1.

Based on these definitions, we propose to jointly optimize zeroth-order and first-order flatness through the following linear combination of the two flatness metrics.

$$R_{\rho, \alpha}^{\text{FAD}}(\theta) = \alpha R_{\rho}^{(0)}(\theta) + (1 - \alpha) R_{\rho}^{(1)}(\theta), \quad (3)$$

where  $\alpha$  is a hyper-parameter that controls the trade-off between zeroth-order and first-order flatness. Since both flatness are related to the maximal eigenvalue of the Hessian matrices [130, 121], we can derive that  $R_{\rho, \alpha}^{\text{FAD}}(\theta)$  is also related to the maximal eigenvalue. Specifically, when a point  $\theta^*$  is a local minimum of  $\hat{\mathcal{E}}_{\text{tr}}(\theta)$  and  $\hat{\mathcal{E}}_{\text{tr}}(\theta)$  can be second-order Taylor approximated in the neighbourhood of  $\theta^*$ , then

$$\lambda_{\max} \left( \nabla^2 \hat{\mathcal{E}}_{\text{tr}}(\theta^*) \right) = \frac{R_{\rho, \alpha}^{\text{FAD}}(\theta^*)}{\rho^2 \left( 1 - \frac{\alpha}{2} \right)}. \quad (4)$$

See Appendix A for details. Since the maximal eigenvalue is a proper measure of the curvature of minima [42, 43] and related to generalization abilities [121, 17, 39], optimizing Equation (3) would hopefully improve the DG performances, which we demonstrate in Section 3.2.

Then the overall objective of our method, Flatness-Aware minimization for Domain generalization (FAD), is as follows.

$$\mathcal{L}_{\text{all}}(\theta) = \hat{\mathcal{E}}_{\text{tr}}(\theta) + \beta \cdot R_{\rho, \alpha}^{\text{FAD}}(\theta). \quad (5)$$

The strength of FAD regularization is controlled by a hyperparameter  $\beta$ .

### 3.2. Flatness Penalty and DG Performance

We give a generalization bound to demonstrate that FAD controls the generalization error on test distribution as follows.

**Proposition 3.1.** *Suppose the per-data-point loss function  $\ell$  is differentiable, symmetric, bounded by  $M$ , and obeys the triangle inequality. Suppose  $\theta^* \in \Theta$ . Fix  $\rho > 0$  and  $\theta \in \Theta$ . Then with probability at least  $1 - \delta$  over training set  $S$  generated from the distribution  $\mathcal{D}$ ,*

$$\begin{aligned} & \mathbb{E}_{\epsilon_i \sim N(0, \rho^2 / (\sqrt{d} + \sqrt{\log n})^2)} [\mathcal{E}_{te}(\theta + \epsilon)] \\ & \leq \hat{\mathcal{E}}_r(\theta) + R_{\rho, \alpha}^{\text{FAD}}(\theta) \\ & + \sup_{\theta_1, \theta_2 \in \Theta} |\mathcal{L}_r(\theta_1, \theta_2) - \mathcal{L}_{te}(\theta_1, \theta_2)| + \frac{M}{\sqrt{n}} \\ & + \sqrt{\frac{\frac{1}{4}d \log \left( 1 + \frac{\|\theta\|^2 (\sqrt{d} + \sqrt{\log n})^2}{d\rho^2} \right) + \frac{1}{4} + \log \frac{n}{\delta} + 2 \log(6n + 3d)}{n-1}} \end{aligned} \quad (6)$$

**Remark.** The term  $\sup_{\theta_1, \theta_2 \in \Theta} |\cdot|$  corresponds to the discrepancy distance [69], which measures the covariate shift between the training domain and the unknown test domain. Please note that in DG/OOD generalization tasks, the information on test distribution is unavailable in the optimization phase. Thus one cannot intervene in the third term on RHS. Thus optimizing the regularization of FAD (the second term on RHS) with ERM loss on training data leads to better generalization on test data.

### 3.3. Algorithm Details

Although the regularization of FAD controls the OOD generalization error, directly optimizing it suffers from significant computational costs. Previous works [114, 121] propose to calculate second order gradient with Hessian-vector product, yet the cost increases fast as the model dimension grows. Inspired by [93, 125], we approximate second-order gradient with first-order gradients and optimize zeroth-order and first-order flatness simultaneously as follows.

**Optimization and acceleration** At each round, suppose the perturbation radius is set to  $\rho_t$ , then the gradient of  $\mathcal{L}_{\text{all}}(\theta_t)$  at the point  $\theta_t$  in Equation (5) is approximated as follows.

The gradient of  $R_{\rho_t}^{(0)}(\theta_t)$ . Following [28], let

$$\begin{aligned} \tilde{\theta}_{t,1} &= \theta_t + \rho_t \cdot \frac{\mathbf{g}_{t,0}}{\|\mathbf{g}_{t,0}\|} \quad \text{where} \quad \mathbf{g}_{t,0} = \nabla \hat{\mathcal{E}}_{\text{tr}}(\theta_t), \\ \mathbf{g}_{t,1} &= \nabla \hat{\mathcal{E}}_{\text{tr}}(\tilde{\theta}_{t,1}). \end{aligned} \quad (7)$$

Then the gradient of  $R_{\rho_t}^{(0)}(\theta_t)$  is approximated by

$$\nabla R_{\rho_t}^{(0)}(\theta_t) \approx \mathbf{g}_{t,1} - \mathbf{g}_{t,0}. \quad (8)$$

---

### Algorithm 1 Flatness-Aware Minimization for Domain Generalization (FAD)

---

- 1: **Input:** Batch size  $b$ , Learning rate  $\eta_t$ , Perturbation radius  $\rho_t$ , Trade-off coefficients  $\alpha, \beta$ , Small constant  $\xi$
  - 2:  $t \leftarrow 0, \theta_0 \leftarrow$  initial parameters
  - 3: **while**  $\theta_t$  not converged **do**
  - 4:   Sample  $W_t$  from the training data with  $b$  instances
  - 5:    $\mathbf{g}_{t,0} \leftarrow \nabla \hat{\mathcal{E}}_{W_t}(\theta_t)$
  - 6:    $\tilde{\theta}_{t,1} \leftarrow \theta_t + \rho_t \cdot \mathbf{g}_{t,0} / (\|\mathbf{g}_{t,0}\| + \xi)$
  - 7:    $\mathbf{g}_{t,1} \leftarrow \nabla \hat{\mathcal{E}}_{W_t}(\tilde{\theta}_{t,1})$
  - 8:    $\mathbf{h}_{t,0} \leftarrow \mathbf{g}_{t,1} - \mathbf{g}_{t,0}$                     $\triangleright$  gradient of  $R_{\rho_t}^{(0)}(\theta_t)$
  - 9:    $\tilde{\theta}_{t,2} \leftarrow \theta_t + \rho_t \cdot \mathbf{h}_{t,0} / (\|\mathbf{h}_{t,0}\| + \xi)$
  - 10:    $\mathbf{g}_{t,2} \leftarrow \nabla \hat{\mathcal{E}}_{W_t}(\tilde{\theta}_{t,2})$
  - 11:    $\tilde{\theta}_{t,3} \leftarrow \tilde{\theta}_{t,2} + \rho_t \cdot \mathbf{g}_{t,2} / (\|\mathbf{g}_{t,2}\| + \xi)$
  - 12:    $\mathbf{g}_{t,3} = \nabla \hat{\mathcal{E}}_{W_t}(\tilde{\theta}_{t,3})$
  - 13:    $\mathbf{h}_{t,1} \leftarrow \mathbf{g}_{t,3} - \mathbf{g}_{t,2}$                     $\triangleright$  gradient of  $R_{\rho_t}^{(1)}(\theta_t)$
  - 14:    $\theta_{t+1} \leftarrow \theta_t - \eta_t(\mathbf{g}_{t,0} + \beta(\alpha \mathbf{h}_{t,0} + (1 - \alpha)\mathbf{h}_{t,1}))$
  - 15:    $t \leftarrow t + 1$
  - 16: **end while**
  - 17: **return**  $\theta_t$
- 

The gradient of  $R_{\rho_t}^{(1)}(\theta_t)$ . [121] propose to approximate the gradient through the following procedure

$$\nabla R_{\rho_t}^{(1)}(\theta_t) \approx \rho_t \cdot \nabla \left\| \nabla \hat{\mathcal{E}}_{\text{tr}}(\theta_t^{\text{adv}}) \right\|, \theta_t^{\text{adv}} = \theta_t + \rho_t \cdot \frac{\nabla \left\| \nabla \hat{\mathcal{E}}_{\text{tr}}(\theta_t) \right\|}{\left\| \nabla \left\| \nabla \hat{\mathcal{E}}_{\text{tr}}(\theta_t) \right\| \right\|}.$$

One can approximate  $\nabla \left\| \nabla \hat{\mathcal{E}}_{\text{tr}}(\theta_t) \right\|$  with first-order gradient as follows.

$$\nabla \left\| \nabla \hat{\mathcal{E}}_{\text{tr}}(\theta_t) \right\| \approx \frac{\nabla \hat{\mathcal{E}}_{\text{tr}}(\tilde{\theta}_{t,1}) - \nabla \hat{\mathcal{E}}_{\text{tr}}(\theta_t)}{\rho_t} = \frac{\mathbf{g}_{t,1} - \mathbf{g}_{t,0}}{\rho_t},$$

Thus we can approximate the gradient of  $R_{\rho_t}^{(1)}(\theta_t)$  efficiently and rewrite the term as follows. The details of the derivation can be found in Appendix A.

$$\begin{aligned} \nabla R_{\rho_t}^{(1)}(\theta_t) &\approx \mathbf{g}_{t,3} - \mathbf{g}_{t,2}, \quad \text{where} \\ \tilde{\theta}_{t,2} &= \theta_t + \rho_t \cdot \frac{\mathbf{g}_{t,1} - \mathbf{g}_{t,0}}{\|\mathbf{g}_{t,1} - \mathbf{g}_{t,0}\|}, \quad \mathbf{g}_{t,2} = \nabla \hat{\mathcal{E}}_{\text{tr}}(\tilde{\theta}_{t,2}), \\ \tilde{\theta}_{t,3} &= \tilde{\theta}_{t,2} + \rho_t \frac{\mathbf{g}_{t,2}}{\|\mathbf{g}_{t,2}\|}, \quad \mathbf{g}_{t,3} = \nabla \hat{\mathcal{E}}_{\text{tr}}(\tilde{\theta}_{t,3}). \end{aligned} \quad (9)$$

Please note that the terms  $\mathbf{g}_{t,1}$  and  $\mathbf{g}_{t,1}$  for optimizing zeroth-order flatness  $R_{\rho_t}^{(0)}(\theta_t)$  are preserved in the approximation of  $R_{\rho_t}^{(1)}(\theta_t)$  and thus causes no overhead. Combining the gradient approximation steps as shown in Equations (8) and (9), we can obtain the unified optimization framework as shown in Algorithm 1.

**Convergence analysis** To analyze the convergence property of Algorithm 1, we first introduce the Lipschitz smooth assumption that are adopted commonly in optimization-related literature [1, 112, 130].

**Definition 3.3.** For a function  $J : \Theta \rightarrow \mathbb{R}$ ,  $J$  is  $\gamma_2$ -Lipschitz smooth if

$$\forall \theta_1, \theta_2 \in \Theta, \|\nabla J(\theta_1) - \nabla J(\theta_2)\| \leq \gamma_2 \|\theta_1 - \theta_2\|. \quad (10)$$

Let  $\Delta_t = \mathbf{g}_{t,0} + \beta(\alpha \mathbf{h}_{t,0} + (1-\alpha)\mathbf{h}_{t,1})$  be the estimated gradient at round  $t$  as shown in Line 14 in Algorithm 1. Now we can prove that the algorithm would eventually converge by the following theorem.

**Theorem 3.2.** Suppose  $\hat{\mathcal{E}}_r(\theta)$  is  $\gamma$ -Lipschitz smooth and bounded by  $M$ . For any timestamp  $t \in \{0, 1, \dots, T\}$  and any  $\theta \in \Theta$ , suppose we can obtain noisy and bounded observations  $f_t(\theta)$  of  $\nabla \hat{\mathcal{E}}_r(\theta)$  such that

$$\mathbb{E}[f_t(\theta)] = \nabla \hat{\mathcal{E}}_r(\theta), \quad \|f_t(\theta)\| \leq G. \quad (11)$$

Then with learning rate  $\eta_t = \eta_0/\sqrt{t}$  and perturbation radius  $\rho_t = \rho_0/\sqrt{t}$ ,

$$\sum_{t=1}^T \mathbb{E} \left[ \|\Delta_t\|^2 \right] \leq \frac{C_1 + C_2 \log T}{\sqrt{T}}, \quad (12)$$

for some constants  $C_1$  and  $C_2$  that only depend on  $\gamma_1, \gamma_2, G, M, \eta_0, \rho_0, \alpha, \beta$ .

Thus, the convergence of FAD is guaranteed, and we show FAD’s convergence rate empirically in Section 4.4.

## 4. Experiments

### 4.1. Experimental Settings

We evaluate our optimization method and current optimizers on various DG benchmarks including PACS [55], VLCS [27], OfficeHome [97], TerraIncognita [9], DomainNet [78] and NICO++ [122]. Detailed introductions of these datasets are in Appendix. We consider the generalization accuracy, training time, and hessian spectra of found minima as the evaluation metrics. We show how optimizers contribute when training with ERM and current SOTA DG algorithms. Generally, Adam is hardly the best choice for DG benchmarks.

For a fair comparison, we follow the basic training and evaluation protocol introduced in [32], where the information of test data is unavailable for hyperparameter search. We train all the models on DomainNet for 15,000 iterations as suggested in [13], 10,000 iterations on NICO++, and 5,000 iterations on other datasets unless otherwise noted. For datasets except for NICO++, we follow the leave-one-out protocol in [32] where one domain is chosen as the target domain and the remaining domains as the training domain. For NICO++, we choose two domains as target domains in each run and train models on the remaining four

domains. Following the official combination [122], we select  $\{\text{autumn}, \text{rock}\}$ ,  $\{\text{dim}, \text{grass}\}$ ,  $\{\text{outdoor}, \text{water}\}$  as the target domain pairs. A detailed description of the split of training and test domains is in Appendix B. For all datasets, 20% samples in the training data are used for validation and model selection. ImageNet [20] pretrained ResNet-50 [34] is adopted as the initial model. The search space for the initial learning rate, weight decay, and other hyperparameters is in Appendix B.

### 4.2. Comparison of Optimizers on DG

Here we investigate how the optimizers affect the generalization under domain shifts. We first conduct experiments with various optimizers following the validation with training data protocol in DomainBed [32] to study the generalization of current optimizers and the robustness of the choice of hyperparameters. The results are shown in Table 1. Detailed results of optimizers on all the splits of reported datasets are in Appendix B.

Despite SGD remaining the most popular optimizer for most CNN-based models on in-distribution image recognition benchmarks, such as CIFAR [51], and ImageNet due to its outstanding performance, it fails to show advantages compared with other optimizers on current benchmarks for DG. The previous work [74] shows that SGD with momentum outperforms adaptive-based optimizers such as Adam with an exhaustive hyperparameter search (hundreds of trials for each optimizer on each dataset) to select hyperparameters with good in-distribution performance. Yet for fair comparisons with previous works, we follow the model selection protocol in DomainBed, where only 20 trials of search are conducted for each evaluation, resulting in inadequate searches. In our experiments, SGD fails to outperform adaptive-based optimizers (e.g., AdamW and AdaBelief). This may be because SGD is more sensitive to the choice of hyperparameters [48] and the hyperparameter search space in evaluation protocols for DG is too large to be traversed and fine-tune the hyperparameters within several random runs. The empirical demonstration of the sensitivity of optimizers to hyperparameters is in Appendix B.

Furthermore, it is important to note that different datasets may require different optimization strategies. For example, Adam outperforms SGD considerably on PACS, while the opposite stands for VLCS, OfficeHome, and NICO++. Among none-flatness-aware optimizers, AdaBelief shows the best average performance. AdamW outperforms Adam on 5 out of 6 datasets and achieves higher average accuracy. Thus, Adam is hardly the best choice for DG and choosing the right optimizer can help exceed current leaderboards’ limitations.

Among all optimizers, flatness-aware optimizers (e.g., SAM, GAM, and FAD) show promising results across all the datasets. FAD considerably outperforms all its coun-

Table 1. Comparison of current optimizers on DG datasets. The best results for each dataset are highlighted in bold font.

Algorithm	PACS	VLCS	OfficeHome	TerraInc	DomainNet	NICO++	Avg.
Adam [48]	84.2 $\pm$ 0.6	77.3 $\pm$ 1.3	67.6 $\pm$ 0.4	44.4 $\pm$ 0.8	43.0 $\pm$ 0.1	76.9 $\pm$ 0.0	65.6
AdamW [65]	83.6 $\pm$ 1.5	77.4 $\pm$ 0.8	68.8 $\pm$ 0.6	45.2 $\pm$ 1.4	43.4 $\pm$ 0.1	77.5 $\pm$ 0.1	66.0
SGD [76]	79.9 $\pm$ 1.4	78.1 $\pm$ 0.2	68.5 $\pm$ 0.3	44.9 $\pm$ 1.8	43.2 $\pm$ 0.1	77.2 $\pm$ 0.2	65.3
YOGI [117]	81.2 $\pm$ 0.4	77.6 $\pm$ 0.6	68.3 $\pm$ 0.3	45.4 $\pm$ 0.5	43.5 $\pm$ 0.0	77.9 $\pm$ 0.2	65.7
AdaBelief [131]	84.6 $\pm$ 0.6	78.4 $\pm$ 0.4	68.0 $\pm$ 0.9	45.2 $\pm$ 2.0	43.5 $\pm$ 0.1	77.4 $\pm$ 0.1	66.2
AdaHessian [116]	84.5 $\pm$ 1.0	78.6 $\pm$ 0.8	68.4 $\pm$ 0.9	44.4 $\pm$ 0.5	<b>44.4</b> $\pm$ 0.1	77.7 $\pm$ 0.1	66.3
SAM [28]	85.3 $\pm$ 1.0	78.2 $\pm$ 0.5	68.0 $\pm$ 0.8	<b>45.7</b> $\pm$ 0.9	43.4 $\pm$ 0.1	77.9 $\pm$ 0.1	66.4
GAM [121]	86.1 $\pm$ 0.6	78.5 $\pm$ 0.4	68.2 $\pm$ 1.0	45.2 $\pm$ 0.6	43.8 $\pm$ 0.1	78.0 $\pm$ 0.2	66.6
FAD (Ours)	<b>88.2</b> $\pm$ 0.5	<b>78.9</b> $\pm$ 0.8	<b>69.2</b> $\pm$ 0.5	<b>45.7</b> $\pm$ 1.0	<b>44.4</b> $\pm$ 0.1	<b>79.0</b> $\pm$ 0.1	<b>67.6</b>

terparts on all the benchmarks, indicating that optimizing zeroth-order and first-order flatness simultaneously learns flatter minima and better generalization.

### 4.3. Ensemble with current DG algorithms

As a base optimizer, FAD can be combined with current DG methods with various objective functions. The compatibility with DG methods is critical for the choice of optimizers. In this subsection, we investigate the compatibility of GAM with current DG algorithms and compare it with other optimizers. We also report the results of other DG methods for better comparisons.

Please note that we only consider a single pretrained model (i.e., ImageNet pretrained ResNet-50) as the initial model for fair comparisons in this paper, so we do not combine optimizers with some SOTA methods, such as SIMPLE [62], MIRO [14] with RegNetY-16GF [83], and CARFT [70] with CLIP [82] which involve the knowledge of pretrained models other than ImageNet pretrained ResNet-50. We consider SOTA and representative methods including CORAL [95], SWAD [13], and Fishr [85] as the base approach to compare optimizers and leave more methods, including EoA [4] and RSC [37] in Appendix B.

The results are shown in Table 2. Different optimizers seem to favor different DG methods. AdamW achieves the best performance when combined with SWAD and CORAL, while SGD surpasses Adam and AdamW with Fishr. This indicates that fixing a single optimizer in DG benchmarks may be unfair for some methods that are not favored, and choosing the proper optimizer can further improve DG methods’ performance.

FAD consistently outperforms other optimizers with all the methods across various datasets. Compared with the default optimizer Adam, FAD further improves SWAD by 1.7% on PACS and 1.2% on average and surpasses previous SOTA methods.




Figure 2. Comparison of computation overhead. Red points on folding lines indicate models trained with FAD for various ratios of iterations (From left to right in each figure, the proportions are 100%, 50%, 10%, 5%).

### 4.4. Computation Overhead

Since flatness-aware optimization methods require more computation overhead than traditional optimizers, we report the comparison of FAD with other optimizers with more training iterations. We train current optimizers on DG benchmarks for 5,000, 10,000, and 15,000 iterations. We report the performance and required average training time in Table 3. Most optimizers achieve noticeable improvements in performance with more training iterations on PACS, VLCS, and OfficeHome. We show that FAD still outperforms current optimizers with fewer training iterations. For example, models trained with FAD for 5,000 iterations on PACS and NICO++ outperform those trained with other optimizers for 15,000 iterations. This indicates the effectiveness and efficiency of FAD.

Furthermore, as discussed in Section 3.3, the FAD gradient can be efficiently calculated via first-order gradient approximation. Yet it still introduces extra computation when calculated in each iteration. Here we investigate optimizing the FAD term only for a subset of iterations in each epoch and optimize the remaining iterations simply with SGD on PACS and NICO++. We change the ratio of the subset with FAD from 5% to 100% and report the performance and training speed. We set the batch size to 96 (32 × 3) for PACS and 128 (32 × 4) for NICO++ for all the methods for a fair

Table 2. Ensemble with domain generalization methods. Numbers for methods marked with \* and all the combinations of DG methods with optimizers other than Adam are reproduced results. Other results are from the original literature and DomainBed (donated with †).

Algorithm	PACS	VLCS	OfficeHome	TerraInc	DomainNet	Avg.
MASF [23]	82.7	-	-	-	-	-
DMC [15]	83.4	-	-	-	43.6	-
MetaReg [7]	83.6	-	-	-	43.6	-
ER [124]	85.3	-	-	-	-	-
pAdalN [77]	85.4	-	-	-	-	-
EISNet [100]	85.8	-	-	-	-	-
DSON [90]	86.6	-	-	-	-	-
ERM† [96]	85.5	77.5	66.5	46.1	40.9	63.3
ERM* (with Adam)	84.2	77.3	67.6	44.4	43.0	63.3
IRM† [3]	83.5	78.6	64.3	47.6	33.9	61.6
GroupDRO† [88]	84.4	76.7	66.0	43.2	33.3	60.7
I-Mixup† [108, 113, 102]	84.6	77.4	68.1	47.9	39.2	63.4
MLDG† [56]	84.9	77.2	66.8	47.8	41.2	63.6
MMD† [60]	84.7	77.5	66.4	42.2	23.4	58.8
DANN† [30]	83.7	78.6	65.9	46.7	38.3	62.6
CDANN† [61]	82.6	77.5	65.7	45.8	38.3	62.0
MTL† [11]	84.6	77.2	66.4	45.6	40.6	62.9
SagNet† [75]	86.3	77.8	68.1	48.6	40.3	64.2
ARM† [119]	85.1	77.6	64.8	45.5	35.5	61.7
VREx† [52]	84.9	78.3	66.4	46.4	33.6	61.9
RSC† [38]	85.2	77.1	65.5	46.6	38.9	62.7
Mixstyle [129]	85.2	77.9	60.4	44.0	34.0	60.3
MIRO* [14]	85.4	78.9	69.5	45.4	44.0	64.6
<hr/>						
Adam + SWAD* [13]	86.8	79.1	70.1	46.5	44.1	65.3
AdamW + SWAD	87.0	78.5	70.8	46.9	<b>45.0</b>	65.6
SGD + SWAD	85.2	79.1	71.0	46.7	42.8	65.0
FAD (Ours) + SWAD	<b>88.5</b>	<b>79.8</b>	<b>71.8</b>	<b>47.5</b>	<b>45.0</b>	<b>66.5</b>
<hr/>						
Adam + Fishr* [85]	85.5	78.0	68.2	46.2	44.7	64.5
AdamW + Fishr	85.7	77.5	68.0	46.7	43.6	64.3
SGD + Fishr	84.4	78.5	<b>69.2</b>	46.9	<b>44.4</b>	64.7
FAD (Ours) + Fishr	<b>88.3</b>	<b>79.6</b>	<b>69.2</b>	<b>48.1</b>	43.8	<b>65.8</b>
<hr/>						
Adam + CORAL* [95]	86.0	78.9	68.7	43.7	44.5	64.5
AdamW + CORAL	86.4	<b>79.5</b>	69.8	45.0	<b>44.9</b>	65.1
SGD + CORAL	85.6	78.2	69.5	45.8	44.6	64.7
FAD (Ours) + CORAL	<b>88.5</b>	78.9	<b>70.8</b>	<b>46.1</b>	<b>44.9</b>	<b>65.9</b>

comparison. As shown in Figure 2, FAD consistently outperforms SGD, Adam, and AdamW by considerable margins and surpasses SAM and GAM with less training time when setting the ratio of FAD to 10% and 50%. Detailed results are in Appendix B.

#### 4.5. The Hessian spectra of selected minima

Sharpness-aware optimization methods are shown to decrease the curvature of loss landscape [28, 121]. We investigate how FAD affects the flatness of minima. We consider the maximum eigenvalue of Hessian, which indicates the sharpness across the most critical direction, and the Hessian trace, which measures the expected loss increase under

random perturbations to the weights [42] as the measures of flatness.

We show the Hessian spectra of ResNet-50 trained on art\_painting, cartoon, and photo domains of PACS for 5,000 iterations with SGD, Adam, AdamW, GAM, SAM, and FAD. Following [121], we adopt power iteration [115] to compute the top eigenvalues of Hessian and Hutchinson’s method [6, 5, 114] to compute the Hessian trace. We report the distribution of the top-50 Hessian eigenvalues and Hessian trace for each method at convergence in Figure 3. None-flatness-aware optimizers, including SGD, Adam, and AdamW show high Hessian trace and dominant eigenvalues. In contrast, flatness-aware optimizers find

Table 3. The performance and average training time of current optimizers and FAD when training for 5,000, 10,000, 15,000 iterations. *Training Time* is the average time for finishing each dataset with the given iterations.

Optimizers	Iterations	Training Time /s	PACS	VLCS	OfficeHome	NICO++	Avg.
Adam	5,000	1613.9	84.2	77.3	67.6	76.2	76.3
Adam	10,000	3049.5	82.9	77.7	67.4	76.9	76.2
Adam	15,000	4478.6	84.0	78.4	69.6	77.3	77.3
AdamW	5,000	2145.2	83.6	77.4	68.8	76.9	76.7
AdamW	10,000	4632.8	83.3	77.1	66.9	77.5	76.2
AdamW	15,000	6995.3	83.7	76.3	68.0	77.5	76.4
SGD	5,000	1864.1	79.9	78.1	68.5	76.5	75.8
SGD	10,000	3438.9	81.4	78.1	68.7	77.2	76.4
SGD	15,000	4992.0	82.6	78.7	69.6	77.0	77.0
YOGI	5,000	1388.9	81.2	77.6	68.3	77.1	76.1
YOGI	10,000	2882.9	81.2	77.8	70.2	77.9	76.8
YOGI	15,000	4434.2	83.1	78.1	67.9	77.5	76.7
SAM	5,000	3325.6	85.3	78.2	68.0	77.0	77.1
SAM	10,000	6506.8	85.6	78.3	68.6	77.9	77.6
SAM	15,000	9960.2	86.1	78.9	68.8	78.1	78.0
GAM	5,000	8073.6	86.1	78.5	68.2	77.4	77.6
GAM	10,000	16210.5	86.0	79.2	68.5	78.0	78.0
GAM	15,000	23855.1	86.5	78.6	69.0	77.8	78.0
FAD (Ours)	5,000	5057.8	88.2	78.9	69.2	78.5	78.7
FAD (Ours)	10,000	10450.8	88.5	79.5	69.6	79.0	79.1
FAD (Ours)	15,000	16255.2	88.5	79.0	69.0	79.1	78.9




Figure 3. The distribution of top eigenvalues and the trace of Hessian at convergence on PACS with SGD, Adam, AdamW, GAM, SAM, and FAD. *Trace* indicates Hessian trace and  $\lambda_{\max}$  indicates the maximum eigenvalue of Hessian.

minima with significantly lower Hessian spectra. FAD leads to both the lowest Hessian maximum eigenvalue and trace and thus the lowest curvature compared with other optimizers. More results are shown in Appendix B.

## 5. Discussions

Despite the wide usage in DG benchmarks, Adam is hardly the optimal choice of default optimizer. We show that the best optimizer for different DG methods varies across different datasets. Choosing the right optimizer can help exceed the current benchmarks' limitations. We propose a unified optimization method, flatness-Aware minimization for Domain generalization (FAD), to optimize zeroth-order and first-order flatness efficiently to address the problem. We further show that FAD controls the OOD generalization error and convergences fast. Compared with current optimizers, flatness-aware optimization approaches,

including SAM, GAM, and FAD, leads to flatter minima and better generalization. FAD outperforms its counterparts significantly on most DG datasets.

FAD still has the following limitations which could lead to potential future work. First, the term  $\sup_{\theta_1, \theta_2 \in \Theta} |\mathcal{L}_{\text{tr}}(\theta_1, \theta_2) - \mathcal{L}_{\text{te}}(\theta_1, \theta_2)|$  in Equation (6) indicates the discrepancy between training and test data. Without knowledge of test distribution and extra assumptions, the term cannot be optimized and leads to the optimization of the flatness term. Proper assumptions on the connection between training and test data or the geometry of data may help to learn the connection between flatness and the discrepancy, and thus control the discrepancy for better generalization. Second, We show that Adam is hardly an optimal choice of default optimizer in DG. However, selecting a proper optimizer for each DG method leads to massive overhead. Designing a better DG protocol



or learning an indicator for the fitness of optimizers and methods can significantly contribute.

## Acknowledgement

This work was supported in part by National Key R&D Program of China (No. 2018AAA0102004, No. 2020AAA0106300), National Natural Science Foundation of China (No. U1936219, 62141607), Beijing Academy of Artificial Intelligence (BAAI).

## References

- [1] Zeyuan Allen-Zhu and Yuanzhi Li. Neon2: Finding local minima via first-order oracles. *Advances in Neural Information Processing Systems*, 31, 2018. 5
- [2] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019. 2
- [3] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization, 2019. 7
- [4] Devansh Arpit, Huan Wang, Yingbo Zhou, and Caiming Xiong. Ensemble of averages: Improving model selection and boosting performance in domain generalization. *arXiv preprint arXiv:2110.10832*, 2021. 2, 6
- [5] Haim Avron and Sivan Toledo. Randomized algorithms for estimating the trace of an implicit symmetric positive semi-definite matrix. *Journal of the ACM (JACM)*, 58(2):1–34, 2011. 7
- [6] Zhaojun Bai, Gark Fahey, and Gene Golub. Some large-scale matrix computation problems. *Journal of Computational and Applied Mathematics*, 74(1-2):71–89, 1996. 7
- [7] Yogesh Balaji, Swami Sankaranarayanan, and Rama Chellappa. Metareg: Towards domain generalization using meta-regularization. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. 7
- [8] David GT Barrett and Benoit Dherin. Implicit gradient regularization. *arXiv preprint arXiv:2009.11162*, 2020. 2
- [9] Sara Beery, Grant Van Horn, and Pietro Perona. Recognition in terra incognita. In *Proceedings of the European conference on computer vision (ECCV)*, pages 456–473, 2018. 5
- [10] Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. *Advances in neural information processing systems*, 19, 2006. 3
- [11] Gilles Blanchard, Aniket Anand Deshmukh, Urun Dogan, Gyemin Lee, and Clayton Scott. Domain generalization by marginal transfer learning, 2017. 7
- [12] Fabio M Carlucci, Antonio D’Innocente, Silvia Bucci, Barbara Caputo, and Tatiana Tommasi. Domain generalization by solving jigsaw puzzles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2229–2238, 2019. 2
- [13] Junbum Cha, Sanghyuk Chun, Kyungjae Lee, Han-Cheol Cho, Seunghyun Park, Yunsung Lee, and Sungrae Park. SWAD: Domain generalization by seeking flat minima. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. 1, 2, 5, 6, 7
- [14] Junbum Cha, Kyungjae Lee, Sungrae Park, and Sanghyuk Chun. Domain generalization by mutual-information regularization with pre-trained models. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIII*, pages 440–457. Springer, 2022. 6, 7
- [15] Prithvijit Chattopadhyay, Yogesh Balaji, and Judy Hoffman. Learning to balance specificity and invariance for in and out of domain generalization. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 301–318, Cham, 2020. Springer International Publishing. 7
- [16] Jinghui Chen, Dongruo Zhou, Yiqi Tang, Ziyang Yang, Yuan Cao, and Quanquan Gu. Closing the generalization gap of adaptive gradient methods in training deep neural networks. *arXiv preprint arXiv:1806.06763*, 2018. 2
- [17] Xiangning Chen, Cho-Jui Hsieh, and Boqing Gong. When vision transformers outperform resnets without pre-training or strong data augmentations. *arXiv preprint arXiv:2106.01548*, 2021. 3
- [18] Xu Chu, Yujie Jin, Wenwu Zhu, Yasha Wang, Xin Wang, Shanghang Zhang, and Hong Mei. Dna: Domain generalization with diversified neural averaging. In *International Conference on Machine Learning*, pages 4010–4034. PMLR, 2022. 2
- [19] Jeremy M. Cohen, Simran Kaur, Yuanzhi Li, J. Zico Kolter, and Ameet Talwalkar. Gradient descent on neural networks typically occurs at the edge of stability. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. 2
- [20] 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. 5
- [21] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2020. 2
- [22] Qi Dou, Daniel Coelho de Castro, Konstantinos Kamnitsas, and Ben Glocker. Domain generalization via model-agnostic learning of semantic features. *Advances in Neural Information Processing Systems*, 32, 2019. 2
- [23] Qi Dou, Daniel Coelho de Castro, Konstantinos Kamnitsas, and Ben Glocker. Domain generalization via model-agnostic learning of semantic features. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. 7

- [24] Jiawei Du, Hanshu Yan, Jiashi Feng, Joey Tianyi Zhou, Liangli Zhen, Rick Siow Mong Goh, and Vincent YF Tan. Efficient sharpness-aware minimization for improved training of neural networks. *arXiv preprint arXiv:2110.03141*, 2021. [2](#)
- [25] Jiawei Du, Daquan Zhou, Jiashi Feng, Vincent YF Tan, and Joey Tianyi Zhou. Sharpness-aware training for free. *arXiv preprint arXiv:2205.14083*, 2022. [1](#), [2](#)
- [26] John Duchi, Elad Hazan, and Yoram Singer. Adaptive sub-gradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(7), 2011. [2](#)
- [27] Chen Fang, Ye Xu, and Daniel N Rockmore. Unbiased metric learning: On the utilization of multiple datasets and web images for softening bias. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1657–1664, 2013. [5](#)
- [28] Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. In *International Conference on Learning Representations*, 2021. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#)
- [29] Zhe Gan, Yen-Chun Chen, Linjie Li, Chen Zhu, Yu Cheng, and Jingjing Liu. Large-scale adversarial training for vision-and-language representation learning. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 6616–6628. Curran Associates, Inc., 2020. [1](#)
- [30] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016. [7](#)
- [31] Muhammad Ghifary, W Bastiaan Kleijn, Mengjie Zhang, and David Balduzzi. Domain generalization for object recognition with multi-task autoencoders. In *Proceedings of the IEEE international conference on computer vision*, pages 2551–2559, 2015. [2](#)
- [32] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. In *International Conference on Learning Representations*, 2021. [1](#), [5](#)
- [33] Moritz Hardt, Ben Recht, and Yoram Singer. Train faster, generalize better: Stability of stochastic gradient descent. In *International conference on machine learning*, pages 1225–1234. PMLR, 2016. [2](#)
- [34] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [2](#), [5](#)
- [35] Sepp Hochreiter and Jürgen Schmidhuber. Simplifying neural nets by discovering flat minima. *Advances in neural information processing systems*, 7, 1994. [2](#)
- [36] Shoubo Hu, Kun Zhang, Zhitang Chen, and Laiwan Chan. Domain generalization via multidomain discriminant analysis. In *Uncertainty in Artificial Intelligence*, pages 292–302. PMLR, 2020. [2](#)
- [37] 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. [6](#)
- [38] Zeyi Huang, Haohan Wang, Eric P. Xing, and Dong Huang. Self-challenging improves cross-domain generalization. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 124–140, Cham, 2020. Springer International Publishing. [7](#)
- [39] Stanislaw Jastrzebski, Zachary Kenton, Devansh Arpit, Nicolas Ballas, Asja Fischer, Yoshua Bengio, and Amos Storkey. Three factors influencing minima in SGD. *arXiv preprint arXiv:1711.04623*, 2017. [2](#), [3](#)
- [40] Stanislaw Jastrzebski, Maciej Szymczak, Stanislav Fort, Devansh Arpit, Jacek Tabor, Kyunghyun Cho, and Krzysztof J. Geras. The break-even point on optimization trajectories of deep neural networks. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. [2](#)
- [41] Zhiwei Jia and Hao Su. Information-theoretic local minima characterization and regularization. In *International Conference on Machine Learning*, pages 4773–4783. PMLR, 2020. [2](#)
- [42] Simran Kaur, Jeremy Cohen, and Zachary C Lipton. On the maximum hessian eigenvalue and generalization. *arXiv preprint arXiv:2206.10654*, 2022. [2](#), [3](#), [7](#)
- [43] Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836*, 2016. [3](#)
- [44] Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. In *International Conference on Learning Representations*, 2017. [2](#)
- [45] Aditya Khosla, Tinghui Zhou, Tomasz Malisiewicz, Alexei A Efros, and Antonio Torralba. Undoing the damage of dataset bias. In *European Conference on Computer Vision*, pages 158–171. Springer, 2012. [2](#)
- [46] Minyoung Kim, Da Li, Shell X Hu, and Timothy Hospedales. Fisher sam: Information geometry and sharpness aware minimisation. In *International Conference on Machine Learning*, pages 11148–11161. PMLR, 2022. [2](#)
- [47] Taero Kim, Sungjun Lim, and Kyungwoo Song. Sharpness-aware minimization for worst case optimization. *arXiv preprint arXiv:2210.13533*, 2022. [2](#)
- [48] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [1](#), [5](#), [6](#)
- [49] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *International Conference on Learning Representations*, 2015. [2](#)
- [50] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, et al. Wilds: A benchmark of in-the-wild dis-

- tribution shifts. In *International Conference on Machine Learning*, pages 5637–5664. PMLR, 2021. 2
- [51] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *Citeseer*, 2009. 5
- [52] David Krueger, Ethan Caballero, Joern-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Dinghuai Zhang, Remi Le Priol, and Aaron Courville. Out-of-distribution generalization via risk extrapolation (rex), 2020. 7
- [53] Jungmin Kwon, Jeongseop Kim, Hyunseo Park, and In Kwon Choi. Asam: Adaptive sharpness-aware minimization for scale-invariant learning of deep neural networks. In *International Conference on Machine Learning*, pages 5905–5914. PMLR, 2021. 2
- [54] Aitor Lewkowycz, Yasaman Bahri, Ethan Dyer, Jascha Sohl-Dickstein, and Guy Gur-Ari. The large learning rate phase of deep learning: the catapult mechanism. *arXiv preprint arXiv:2003.02218*, 2020. 2
- [55] 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. 5
- [56] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M. Hospedales. Learning to generalize: Meta-learning for domain generalization, 2017. 7
- [57] 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. 2
- [58] Da Li, Jianshu Zhang, Yongxin Yang, Cong Liu, Yi-Zhe Song, and Timothy M Hospedales. Episodic training for domain generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1446–1455, 2019. 2
- [59] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5400–5409, 2018. 2
- [60] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C. Kot. Domain generalization with adversarial feature learning. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5400–5409, 2018. 7
- [61] Ya Li, Mingming Gong, Xinmei Tian, Tongliang Liu, and Dacheng Tao. Domain generalization via conditional invariant representation, 2018. 7
- [62] Ziyue Li, Kan Ren, XINYANG JIANG, Yifei Shen, Haipeng Zhang, and Dongsheng Li. Simple: Specialized model-sample matching for domain generalization. In *The Eleventh International Conference on Learning Representations*, 2023. 2, 6
- [63] Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. On the variance of the adaptive learning rate and beyond. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. 2
- [64] Yong Liu, Siqi Mai, Xiangning Chen, Cho-Jui Hsieh, and Yang You. Towards efficient and scalable sharpness-aware minimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12360–12370, 2022. 2
- [65] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 2, 6
- [66] Liangchen Luo, Yuanhao Xiong, Yan Liu, and Xu Sun. Adaptive gradient methods with dynamic bound of learning rate. *arXiv preprint arXiv:1902.09843*, 2019. 2
- [67] Massimiliano Mancini, Samuel Rota Buló, Barbara Caputo, and Elisa Ricci. Best sources forward: domain generalization through source-specific nets. In *2018 25th IEEE international conference on image processing (ICIP)*, pages 1353–1357. IEEE, 2018. 2
- [68] Massimiliano Mancini, Samuel Rota Buló, Barbara Caputo, and Elisa Ricci. Robust place categorization with deep domain generalization. *IEEE Robotics and Automation Letters*, 3(3):2093–2100, 2018. 2
- [69] Yishay Mansour, Mehryar Mohri, and Afshin Roshtamzadeh. Domain adaptation: Learning bounds and algorithms. *arXiv preprint arXiv:0902.3430*, 2009. 4
- [70] Xiaofeng Mao, Yuefeng Chen, Xiaojun Jia, Rong Zhang, Hui Xue, and Zhao Li. Context-aware robust fine-tuning. *arXiv preprint arXiv:2211.16175*, 2022. 6
- [71] Peng Mi, Li Shen, Tianhe Ren, Yiyi Zhou, Xiaoshuai Sun, Rongrong Ji, and Dacheng Tao. Make sharpness-aware minimization stronger: A sparsified perturbation approach. *arXiv preprint arXiv:2210.05177*, 2022. 2
- [72] Saeid Motiian, Marco Piccirilli, Donald A Adjeroh, and Gianfranco Doretto. Unified deep supervised domain adaptation and generalization. In *Proceedings of the IEEE international conference on computer vision*, pages 5715–5725, 2017. 2
- [73] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *ICML*, pages 10–18. PMLR, 2013. 1
- [74] Hiroki Naganuma, Kartik Ahuja, Ioannis Mitliagkas, Shiro Takagi, Tetsuya Motokawa, Rio Yokota, Kohta Ishikawa, and Ikuro Sato. Empirical study on optimizer selection for out-of-distribution generalization. In *NeurIPS 2022 Workshop on Distribution Shifts: Connecting Methods and Applications*, 2022. 2, 5
- [75] Hyeonseob Nam, Hyunjae Lee, Jongchan Park, Wonjun Yoon, and Donggeun Yoo. Reducing domain gap by reducing style bias. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8686–8695, 2021. 7
- [76] Yu E Nesterov. A method for solving the convex programming problem with convergence rate. In *Dokl. Akad. Nauk SSSR*, volume 269, pages 543–547, 1983. 1, 2, 6
- [77] O. Nuriel, S. Benaim, and L. Wolf. Permuted adain: Reducing the bias towards global statistics in image classification. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9477–9486, Los Alamitos, CA, USA, jun 2021. IEEE Computer Society. 7
- [78] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source

- domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1406–1415, 2019. 5
- [79] Henning Petzka, Michael Kamp, Linara Adilova, Cristian Sminchisescu, and Mario Boley. Relative flatness and generalization. *Advances in Neural Information Processing Systems*, 34:18420–18432, 2021. 2
- [80] Vihari Piratla, Praneeth Netrapalli, and Sunita Sarawagi. Efficient domain generalization via common-specific low-rank decomposition. In *International Conference on Machine Learning*, pages 7728–7738. PMLR, 2020. 2
- [81] Fengchun Qiao, Long Zhao, and Xi Peng. Learning to learn single domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12556–12565, 2020. 2
- [82] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 6
- [83] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10428–10436, 2020. 6
- [84] Alexandre Ramé, Kartik Ahuja, Jianyu Zhang, Matthieu Cord, Léon Bottou, and David Lopez-Paz. Recycling diverse models for out-of-distribution generalization. *arXiv preprint arXiv:2212.10445*, 2022. 2
- [85] Alexandre Rame, Corentin Dancette, and Matthieu Cord. Fishr: Invariant gradient variances for out-of-distribution generalization. In *International Conference on Machine Learning*, pages 18347–18377. PMLR, 2022. 6, 7
- [86] Alexandre Rame, Matthieu Kirchmeyer, Thibaud Rahier, Alain Rakotomamonjy, Patrick Gallinari, and Matthieu Cord. Diverse weight averaging for out-of-distribution generalization. *arXiv preprint arXiv:2205.09739*, 2022. 2
- [87] Harsh Rangwani, Sumukh K Aithal, Mayank Mishra, Arihant Jain, and Venkatesh Babu Radhakrishnan. A closer look at smoothness in domain adversarial training. In *International Conference on Machine Learning*, pages 18378–18399. PMLR, 2022. 1, 2
- [88] Shiori Sagawa\*, Pang Wei Koh\*, Tatsunori B. Hashimoto, and Percy Liang. Distributionally robust neural networks. In *International Conference on Learning Representations*, 2020. 7
- [89] Seonguk Seo, Yumin Suh, Dongwan Kim, Jongwoo Han, and Bohyung Han. Learning to optimize domain specific normalization for domain generalization. *arXiv preprint arXiv:1907.04275*, 3(6):7, 2019. 2
- [90] Seonguk Seo, Yumin Suh, Dongwan Kim, Geeho Kim, Jongwoo Han, and Bohyung Han. Learning to optimize domain specific normalization for domain generalization. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 68–83, Cham, 2020. Springer International Publishing. 7
- [91] Shiv Shankar, Vihari Piratla, Soumen Chakrabarti, Sidhartha Chaudhuri, Preethi Jyothi, and Sunita Sarawagi. Generalizing across domains via cross-gradient training. *arXiv preprint arXiv:1804.10745*, 2018. 2
- [92] Hidetoshi Shimodaira. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of statistical planning and inference*, 90(2):227–244, 2000. 3
- [93] Sidak Pal Singh and Dan Alistarh. Woodfisher: Efficient second-order approximation for neural network compression. *Advances in Neural Information Processing Systems*, 33:18098–18109, 2020. 4
- [94] Samuel L. Smith and Quoc V. Le. A bayesian perspective on generalization and stochastic gradient descent. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. 2
- [95] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In Gang Hua and Hervé Jégou, editors, *Computer Vision – ECCV 2016 Workshops*, pages 443–450, Cham, 2016. Springer International Publishing. 6, 7
- [96] Vladimir N Vapnik. An overview of statistical learning theory. *IEEE transactions on neural networks*, 10(5):988–999, 1999. 7
- [97] Hemant Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5018–5027, 2017. 5
- [98] Riccardo Volpi, Hongseok Namkoong, Ozan Sener, John C Duchi, Vittorio Murino, and Silvio Savarese. Generalizing to unseen domains via adversarial data augmentation. *Advances in neural information processing systems*, 31, 2018. 2
- [99] Jindong Wang, Cuiling Lan, Chang Liu, Yidong Ouyang, Tao Qin, Wang Lu, Yiqiang Chen, Wenjun Zeng, and Philip Yu. Generalizing to unseen domains: A survey on domain generalization. *IEEE Transactions on Knowledge and Data Engineering*, 2022. 1
- [100] Shujun Wang, Lequan Yu, Caizi Li, Chi-Wing Fu, and Pheng-Ann Heng. Learning from extrinsic and intrinsic supervisions for domain generalization. In *ECCV*, 2020. 7
- [101] Xinyi Wang, Michael Saxon, Jiachen Li, Hongyang Zhang, Kun Zhang, and William Yang Wang. Causal balancing for domain generalization. *arXiv preprint arXiv:2206.05263*, 2022. 2
- [102] Yufei Wang, Haoliang Li, and Alex C. Kot. Heterogeneous domain generalization via domain mixup. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3622–3626, 2020. 7
- [103] Yeming Wen, Kevin Luk, Maxime Gazeau, Guodong Zhang, Harris Chan, and Jimmy Ba. An empirical study of large-batch stochastic gradient descent with structured covariance noise. *arXiv preprint arXiv:1902.08234*, 2019. 2

- [104] Ashia C Wilson, Rebecca Roelofs, Mitchell Stern, Nati Srebro, and Benjamin Recht. The marginal value of adaptive gradient methods in machine learning. *Advances in neural information processing systems*, 30, 2017. [2](#)
- [105] Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, et al. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *International Conference on Machine Learning*, pages 23965–23998. PMLR, 2022. [2](#)
- [106] Zeke Xie, Qian-Yuan Tang, Yunfeng Cai, Mingming Sun, and Ping Li. On the power-law spectrum in deep learning: A bridge to protein science. *arXiv preprint arXiv:2201.13011*, 2022. [2](#)
- [107] Zeke Xie, Xinrui Wang, Huishuai Zhang, Issei Sato, and Masashi Sugiyama. Adaptive inertia: Disentangling the effects of adaptive learning rate and momentum. In *International Conference on Machine Learning*, pages 24430–24459. PMLR, 2022. [2](#)
- [108] Minghao Xu, Jian Zhang, Bingbing Ni, Teng Li, Chengjie Wang, Qi Tian, and Wenjun Zhang. Adversarial domain adaptation with domain mixup, 2019. [7](#)
- [109] Qinwei Xu, Ruipeng Zhang, Ya Zhang, Yanfeng Wang, and Qi Tian. A fourier-based framework for domain generalization. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14378–14387, 2021. [2](#)
- [110] Renzhe Xu, Peng Cui, Zheyang Shen, Xingxuan Zhang, and Tong Zhang. Why stable learning works? a theory of covariate shift generalization. *arXiv preprint arXiv:2111.02355*, 2021. [2](#)
- [111] Renzhe Xu, Xingxuan Zhang, Zheyang Shen, Tong Zhang, and Peng Cui. A theoretical analysis on independence-driven importance weighting for covariate-shift generalization. In *International Conference on Machine Learning*, pages 24803–24829. PMLR, 2022. [3](#)
- [112] Yi Xu, Rong Jin, and Tianbao Yang. First-order stochastic algorithms for escaping from saddle points in almost linear time. *Advances in neural information processing systems*, 31, 2018. [5](#)
- [113] Shen Yan, Huan Song, Nanxiang Li, Lincan Zou, and Liu Ren. Improve unsupervised domain adaptation with mixup training, 2020. [7](#)
- [114] Zhewei Yao, Amir Gholami, Kurt Keutzer, and Michael W Mahoney. Pyhessian: Neural networks through the lens of the hessian. In *2020 IEEE international conference on big data (Big data)*, pages 581–590. IEEE, 2020. [4](#), [7](#)
- [115] Zhewei Yao, Amir Gholami, Qi Lei, Kurt Keutzer, and Michael W Mahoney. Hessian-based analysis of large batch training and robustness to adversaries. *Advances in Neural Information Processing Systems*, 31, 2018. [7](#)
- [116] Zhewei Yao, Amir Gholami, Sheng Shen, Mustafa Mustafa, Kurt Keutzer, and Michael Mahoney. Adahessian: An adaptive second order optimizer for machine learning. In *proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 10665–10673, 2021. [1](#), [6](#)
- [117] Manzil Zaheer, Sashank Reddi, Devendra Sachan, Satyen Kale, and Sanjiv Kumar. Adaptive methods for nonconvex optimization. *Advances in neural information processing systems*, 31, 2018. [1](#), [2](#), [6](#)
- [118] Hanlin Zhang, Yi-Fan Zhang, Weiyang Liu, Adrian Weller, Bernhard Schölkopf, and Eric P Xing. Towards principled disentanglement for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8024–8034, 2022. [2](#)
- [119] Marvin Mengxin Zhang, Henrik Marklund, Nikita Dhawan, Abhishek Gupta, Sergey Levine, and Chelsea Finn. Adaptive risk minimization: A meta-learning approach for tackling group shift, 2021. [7](#)
- [120] Xingxuan Zhang, Peng Cui, Renzhe Xu, Linjun Zhou, Yue He, and Zheyang Shen. Deep stable learning for out-of-distribution generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5372–5382, 2021. [2](#)
- [121] Xingxuan Zhang, Renzhe Xu, Han Yu, Hao Zou, and Peng Cui. Gradient norm aware minimization seeks first-order flatness and improves generalization. *arXiv preprint arXiv:2303.03108*, 2023. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#)
- [122] Xingxuan Zhang, Linjun Zhou, Renzhe Xu, Peng Cui, Zheyang Shen, and Haoxin Liu. Nico++: Towards better benchmarking for domain generalization. *arXiv preprint arXiv:2204.08040*, 2022. [5](#)
- [123] YiFan Zhang, Xue Wang, Jian Liang, Zhang Zhang, Liang Wang, Rong Jin, and Tieniu Tan. Free lunch for domain adversarial training: Environment label smoothing. *arXiv preprint arXiv:2302.00194*, 2023. [2](#)
- [124] Shanshan Zhao, Mingming Gong, Tongliang Liu, Huan Fu, and Dacheng Tao. Domain generalization via entropy regularization. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 16096–16107. Curran Associates, Inc., 2020. [7](#)
- [125] Yang Zhao, Hao Zhang, and Xiuyuan Hu. Penalizing gradient norm for efficiently improving generalization in deep learning. In *International Conference on Machine Learning*, pages 26982–26992. PMLR, 2022. [2](#), [4](#)
- [126] Qihuang Zhong, Liang Ding, Li Shen, Peng Mi, Juhua Liu, Bo Du, and Dacheng Tao. Improving sharpness-aware minimization with fisher mask for better generalization on language models. *arXiv preprint arXiv:2210.05497*, 2022. [2](#)
- [127] Kaiyang Zhou, Yongxin Yang, Timothy Hospedales, and Tao Xiang. Deep domain-adversarial image generation for domain generalisation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13025–13032, 2020. [2](#)
- [128] Kaiyang Zhou, Yongxin Yang, Timothy Hospedales, and Tao Xiang. Learning to generate novel domains for domain generalization. In *European Conference on Computer Vision*, pages 561–578. Springer, 2020. [2](#)
- [129] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain generalization with mixstyle. In *International Conference on Learning Representations*, 2021. [1](#), [7](#)
- [130] Juntang Zhuang, Boqing Gong, Liangzhe Yuan, Yin Cui, Hartwig Adam, Nicha C Dvornek, James s Duncan, Ting Liu, et al. Surrogate gap minimization improves sharpness-

aware training. In *International Conference on Learning Representations*, 2022. [1](#), [2](#), [3](#), [5](#)

- [131] Juntang Zhuang, Tommy Tang, Yifan Ding, Sekhar C Tatikonda, Nicha Dvornek, Xenophon Papademetris, and James Duncan. Adabelief optimizer: Adapting stepsizes by the belief in observed gradients. *Advances in neural information processing systems*, 33:18795–18806, 2020. [1](#), [6](#)