

Sharpness-Aware Gradient Matching for Domain Generalization

Pengfei Wang^{1,2} Zhaoxiang Zhang^{1,2,3*} Zhen Lei^{1,2,3} Lei Zhang^{1*}

¹The Hong Kong Polytechnic University ²Center for Artificial Intelligence and Robotics, HKISI, CAS

³State Key Laboratory of Multimodal Artificial Intelligence Systems, CASIA

pengfei.wang@connect.polyu.hk, zhaoxiang.zhang@ia.ac.cn,

zlei@nlpr.ia.ac.cn, cslzhang@comp.polyu.edu.hk

Abstract

The goal of domain generalization (DG) is to enhance the generalization capability of the model learned from a source domain to other unseen domains. The recently developed Sharpness-Aware Minimization (SAM) method aims to achieve this goal by minimizing the sharpness measure of the loss landscape. Though SAM and its variants have demonstrated impressive DG performance, they may not always converge to the desired flat region with a small loss value. In this paper, we present two conditions to ensure that the model could converge to a flat minimum with a small loss, and present an algorithm, named Sharpness-Aware Gradient Matching (SAGM), to meet the two conditions for improving model generalization capability. Specifically, the optimization objective of SAGM will simultaneously minimize the empirical risk, the perturbed loss (i.e., the maximum loss within a neighborhood in the parameter space), and the gap between them. By implicitly aligning the gradient directions between the empirical risk and the perturbed loss, SAGM improves the generalization capability over SAM and its variants without increasing the computational cost. Extensive experimental results show that our proposed SAGM method consistently outperforms the state-of-the-art methods on five DG benchmarks, including PACS, VLCS, OfficeHome, TerraIncognita, and DomainNet. Codes are available at <https://github.com/Wang-pengfei/SAGM>.

1. Introduction

Deep learning methods have achieved remarkable success in various computer vision tasks when the source data and target data are independently and identically distributed (i.i.d). However, the performance of deep models trained in a source domain can drop significantly when applied to unseen target domains. Domain generalization (DG) aims

to train a model from a set of source data such that it can be well generalized to new domains without retraining. Over the past decade, research on DG has led to a plethora of methods, including those based on domain alignment [35, 16, 32, 2, 52], meta-learning [29, 3, 11, 51], and data augmentation [53, 43, 7]. Though numerous DG approaches have been proposed, a recent study called DomainBed [18] reveals that under a fair evaluation protocol, the naive empirical risk minimization (ERM) method can even outperform most existing DG methods. Unfortunately, simply minimizing the empirical loss on a complex and non-convex loss landscape is typically insufficient to achieve a good generalization capability [23, 17, 22, 15]. As a result, ERM tends to overfit the training data set and converge to sharp local minima [15].

Recent studies such as Sharpness-Aware Minimization (SAM) [15] try to improve the model generalization ability by minimizing the sharpness measure of the loss landscape. Denote by $\mathcal{L}(\theta)$ the loss to be minimized (e.g., the cross-entropy loss for classification), where θ is the parameters of the neural network. SAM first adversarially computes a weight perturbation ϵ that maximizes the empirical risk $\mathcal{L}(\theta)$ and then minimizes the loss of the perturbed network, i.e., $\mathcal{L}(\theta + \epsilon)$. Specifically, the objective of SAM is to minimize the maximum loss around the model parameter θ . Due to the high complexity of this min-max optimization problem, SAM chooses to minimize an approximation of $\mathcal{L}(\theta)$, denoted by $\mathcal{L}_p(\theta)$ (perturbed loss, see Section 3 for details). However, minimizing $\mathcal{L}_p(\theta)$ is not guaranteed to converge to flat minimum regions [55].

While $\mathcal{L}_p(\theta)$ may not characterize well the sharpness of loss surface, the surrogate gap $h(\theta) \triangleq \mathcal{L}_p(\theta) - \mathcal{L}(\theta)$ can better describe it. Intuitively, the loss surface will become flatter when $h(\theta)$ is closer to zero because the perturbation parameters around θ will have a similar loss value. Zhuang *et al.* [55] proved that $h(\theta)$ is an equivalent measure of the dominant eigenvalue of Hessian (which is the measure of sharpness) at a local minimum. Therefore, we can seek flat minima for better generalization ability by mini-

*Corresponding author

mizing the surrogate gap $h(\theta)$. Unlike SAM which only optimizes the perturbation loss $\mathcal{L}_p(\theta)$, GSAM [55] jointly minimizes $\mathcal{L}_p(\theta)$ and the surrogate gap $h(\theta)$. However, GSAM minimizes $h(\theta)$ by increasing the loss $\mathcal{L}(\theta)$, which will reduce the generalization performance. Zhang *et al.* [55] have shown that when the perturbation amplitude ρ is sufficiently small, the surrogate gap is always non-negative, *i.e.*, $\mathcal{L}_p(\theta) \geq \mathcal{L}(\theta), \forall \theta$. Therefore, increasing $\mathcal{L}(\theta)$ will also increase the difficulty of optimizing $\mathcal{L}_p(\theta; \mathcal{D})$.

We propose two conditions that should be met to obtain a model with good generalization performance. (i) First, the loss within a neighborhood of the desired minimum should be sufficiently low. (ii) Second, the minimum is within a flat loss surface. More specifically, condition (i) implies a low training loss and represents good performance on the source training data, while condition (ii) reduces the performance gap of the model on training and testing data.

Based on the above analysis, we propose a new DG method, namely Sharpness-Aware Gradient Matching (SAGM), to simultaneously minimize three objectives, the empirical risk $\mathcal{L}(\theta)$, the perturbed loss $\mathcal{L}_p(\theta)$ and the surrogate gap $h(\theta)$. By minimizing $\mathcal{L}(\theta)$ and $\mathcal{L}_p(\theta)$, we search for a region with low loss, which satisfies condition (i). By minimizing $h(\theta)$, we avoid steep valleys, which meets condition (ii). However, optimizing these three objectives simultaneously is difficult due to the inevitable gradient conflicts during training. Fortunately, we find that when the gradient directions of $\mathcal{L}(\theta)$ and $\mathcal{L}_p(\theta)$ are consistent, the gradient direction of $h(\theta)$ is also consistent with them, and hence the gradient descent can be effectively applied to all the three losses. Therefore, we transform the optimization objective into minimizing $\mathcal{L}(\theta)$, $\mathcal{L}_p(\theta)$, and the angle between their gradients, and achieve this goal by implicitly aligning the gradient directions between the $\mathcal{L}(\theta)$ and $\mathcal{L}_p(\theta)$. The proposed SAGM improves the generalization performance of the model by facilitating the model to converge to a flat region with a small loss value. Compared with SAM, our SAGM does not increase the computational cost. In addition, SAGM can be combined with previous data augmentation methods, such as Mixstyle [54] for further performance improvements.

The contributions of this work are summarized as follows. First, we analyze the limitations of SAM-like methods and propose two conditions to ensure the model convergence to a flat region with a small loss. Second, we propose the SAGM algorithm to improve the DG capability of deep models. Finally, we demonstrate the superior performance of SAGM to state-of-the-arts on five DG benchmarks.

2. Related Work

Domain Generalization (DG). DG methods aim to learn a model from one or several observed source data sets so that it can be generalized to unseen target domains. Ex-

isting DG methods handle domain shift from various perspectives, including domain alignment [35, 16, 32, 2, 52], meta learning [29, 3, 11, 51], data augmentation [53, 43, 7], disentangled representation learning [38, 24, 48] and capturing causal relation [1, 28]. Most of the relevant works to this paper address the DG problem from a gradient perspective [29, 27, 44, 34]. For example, Mansilla *et al.* [34] proposed a gradient surgery strategy to tackle DG problems. They designed an AND gate on the inter-domain gradients. The values of the same sign in the gradient are retained, and the values of different signs are set to zero. While these methods achieve good performance on the training set, they could fail when generalizing to unseen domains because there is no guarantee that the trained models could converge to flat loss regions.

Sharpness-Aware Minimization (SAM). SAM [15] was proposed to solve the sharp minima problem by modifying the objective to minimize a perturbed loss, which is defined as the maximum loss within a neighborhood in the parameter space. Liu *et al.* and Du *et al.* proposed LookSAM [33] and Efficient SAM (ESAM) [12], respectively, to address the computation issue of SAM. These two methods can approach the performance of SAM with less computation, but they inherit the key problem of SAM, *i.e.*, the perturbed loss $\mathcal{L}_p(\theta)$ might disagree with sharpness. GSAM [55] was then proposed to address this issue. However, GSAM minimizes the surrogate gap $h(\theta)$ at the price of increasing $\mathcal{L}(\theta)$. This limits the potential improvement of generalization capability because increasing $\mathcal{L}(\theta)$ is equivalent to increasing the bound of $\mathcal{L}_p(\theta)$. In this paper, we propose a novel SAGM method to facilitate the model converge to a flat region with better generalization ability.

Sharpness and Generalization. The first study to reveal the relationship between sharpness and model generalization ability can be dated back to [20]. Subsequently, extensive theoretical and empirical studies have been carried out on the relationship between sharpness and generalization from the perspective of loss surface geometry under the i.i.d. assumption [23, 13, 17, 22, 10, 15]. For example, Keskar *et al.* [23] proposed a sharpness measure and revealed the negative correlation between the sharpness measure and the generalization abilities. Dinh *et al.* [10] further argued that the sharpness measure can be related to the spectrum of the Hessian, whose eigenvalues encode the curvature information of the loss landscape. As for the out-of-distribution (OOD) conditions, a recent work SWAD [8] has theoretically shown that a flatter minimum could lead to a smaller DG gap, which provides a theoretical basis for our method. Based on this theoretical finding, Cha *et al.* [8] modified the stochastic weight averaging (SWA) method [22], one of the popular sharpness-aware solvers, by introducing a dense and overfit-aware stochastic weight sampling strategy. Unfortunately, this ensemble-like weighting

method could not promote the model to find flatter minima during the training process.

3. Preliminaries

Throughout the paper, we use θ to denote the weight parameters of a neural network f . Suppose there are a set of S training source domains, denoted by $\mathcal{D} = \{\mathcal{D}_i\}_{i=1}^S$. In the DG setting, usually we assume there is one target domain, denoted by \mathcal{T} . From each source domain \mathcal{D}_i , suppose there are n training samples, each consisting of an input x and a target label y , i.e., $(x_j^i, y_j^i)_{j=1}^n \sim \mathcal{D}_i$. The training loss over all training domains \mathcal{D} is defined as follows:

$$\mathcal{L}(\theta; \mathcal{D}) = \frac{1}{Sn} \sum_{i=1}^S \sum_{j=1}^n \ell(f(x_j^i; \theta), y_j^i), \quad (1)$$

where ℓ donates the cross-entropy loss. The training of the network is typically a non-convex optimization problem, aiming to search for an optimal weight vector $\hat{\theta}$ that has the lowest empirical risk $\mathcal{L}(\theta; \mathcal{D})$.

The conventional optimization paradigm in Equation 1 is called *empirical risk minimization* (ERM). In practice, ERM tends to overfit the training set and converge to sharp minima, which can lead to unsatisfactory performance on an unseen domain. The SAM method [15] is proposed to find a flatter region around the minimum with low loss values. To achieve this goal, SAM solves the following min-max optimization problem:

$$\min_{\theta} \max_{\epsilon: \|\epsilon\| \leq \rho} \mathcal{L}(\theta + \epsilon; \mathcal{D}), \quad (2)$$

where ρ is a predefined constant controlling the radius of the neighborhood. Given θ , the maximization in Equation 2 finds the weight perturbation ϵ in the Euclidean ball with radius ρ that maximizes the empirical loss.

Since the maximization in Equation 2 is generally costly, an approximate maximizer was proposed in [15] by invoking the Taylor expansion of the empirical loss. For a small ρ , using Taylor expansion around θ , the inner maximization in Equation 2 turns into a linearly constrained optimization with a solution as follows:

$$\epsilon = \max_{\epsilon: \|\epsilon\| \leq \rho} \mathcal{L}(\theta + \epsilon; \mathcal{D}) \approx \rho \frac{\nabla \mathcal{L}(\theta; \mathcal{D})}{\|\nabla \mathcal{L}(\theta; \mathcal{D})\|}. \quad (3)$$

As a result, the optimization problem of SAM reduces to

$$\min_{\theta} \mathcal{L}(\theta + \hat{\epsilon}; \mathcal{D}), \text{ where } \hat{\epsilon} \triangleq \rho \frac{\nabla \mathcal{L}(\theta; \mathcal{D})}{\|\nabla \mathcal{L}(\theta; \mathcal{D})\|}. \quad (4)$$

For simplicity of notation, the loss function of SAM can be written as the following perturbed loss:

$$\mathcal{L}_p(\theta; \mathcal{D}) = \mathcal{L}(\theta + \hat{\epsilon}; \mathcal{D}). \quad (5)$$

In other words, SAM seeks for a minimum on the surface of loss $\mathcal{L}_p(\theta; \mathcal{D})$ rather than the original loss $\mathcal{L}(\theta; \mathcal{D})$.

4. Sharpness-Aware Gradient Matching

In this section, we first analyze the limitations of SAM-like methods [15, 55] and propose two conditions to ensure the model convergence to a flat loss region. Then we propose a Sharpness-Aware Gradient Matching (SAGM) algorithm that meets these two conditions to achieve good generalization performance. Finally, we provide some analysis to better understand the proposed SAGM method.

4.1. The perturbed loss $\mathcal{L}_p(\theta; \mathcal{D})$ is not always sharpness-aware

Though SAM aims to find a flat region with low loss values, it may not always find such solutions. A recent work [55] has proven that for some fixed ρ , there is no linear relationship between the perturbed loss and σ_{max} (the dominant eigenvalue of the Hessian, a measure of sharpness). This means that a smaller $\mathcal{L}_p(\theta; \mathcal{D})$ does not guarantee a flatter region. A toy example is shown in Fig. 1 to illustrate this. We see that the loss surface of local minimum θ_2 is flatter than that of θ_1 . However, SAM will favor θ_1 over θ_2 because $\mathcal{L}_p(\theta_1; \mathcal{D}) < \mathcal{L}_p(\theta_2; \mathcal{D})$, resulting in the selection of a sharper minimum.

Fortunately, we find that the surrogate gap, which is defined as:

$$h(\theta) \triangleq \mathcal{L}_p(\theta; \mathcal{D}) - \mathcal{L}(\theta; \mathcal{D}), \quad (6)$$

can describe the sharpness much better. Intuitively, the surrogate gap measures the difference between the maximum loss within the neighborhood and the loss at the minimum point. Zhuang *et al.* [55] showed that $h(\theta)$ can be used to measure the sharpness. Specifically, for a local minimum θ^* , consider the neighborhood centered at θ^* with a small radius ρ , the surrogate gap $h(\theta^*)$ approximately has the following linear relationship with σ_{max} :

$$\sigma_{max} \approx 2h(\theta^*)/\rho^2. \quad (7)$$

Based on the above findings, Zhuang *et al.* [55] proposed the Surrogate Gap Guided Sharpness-Aware Minimization (GSAM) to simultaneously minimize the perturbation loss $\mathcal{L}_p(\theta; \mathcal{D})$ and the surrogate gap $h(\theta)$:

$$\min_{\theta} (\mathcal{L}_p(\theta; \mathcal{D}), h(\theta)). \quad (8)$$

In particular, GSAM decomposes $\nabla \mathcal{L}(\theta; \mathcal{D})$ into two components that are respectively parallel and orthogonal to $\nabla \mathcal{L}_p(\theta; \mathcal{D})$. As shown in Fig. 2 (a), the final gradient direction of GSAM is as follow:

$$\nabla \mathcal{L}_{GSAM} = \nabla \mathcal{L}_p(\theta; \mathcal{D}) - \beta \nabla \mathcal{L}(\theta; \mathcal{D})_{\perp}, \quad (9)$$

where β is a hyperparameter to scale the stepsize of the ascent step. Note that the horizontal component $\nabla \mathcal{L}(\theta; \mathcal{D})_{\perp}$ keeps $\mathcal{L}_p(\theta; \mathcal{D})$ unchanged because they are perpendicular

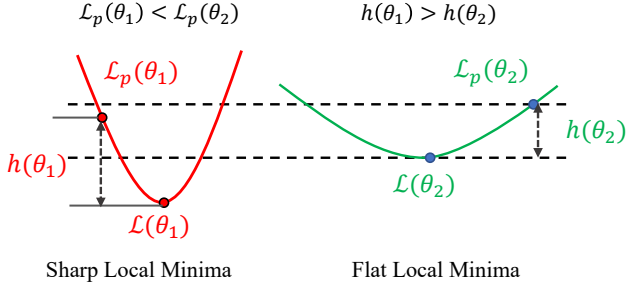


Figure 1: Consider a sharp local minimum θ_1 (red) and a flat local minimum θ_2 (green). Since $\mathcal{L}_p(\theta_1; \mathcal{D}) < \mathcal{L}_p(\theta_2; \mathcal{D})$ (we omit symbol \mathcal{D} in the figure for simplicity), SAM prefers θ_1 to θ_2 . However, θ_1 is a sharper minimum than θ_2 , and hence $\mathcal{L}_p(\theta; \mathcal{D})$ disagrees with sharpness. Instead, the surrogate gap $h(\theta)$ can better describe the sharpness of loss surface. The smaller value of $h(\theta_2)$ correctly indicates that θ_2 is a flatter minimum than θ_1 . Best viewed in color.

to each other. However, $\nabla \mathcal{L}(\theta; \mathcal{D})_{\perp}$ will increase $\mathcal{L}(\theta; \mathcal{D})$ because there is a gradient conflict between them. In other words, GSAM reduces the surrogate gap $h(\theta)$ at the price of increasing $\mathcal{L}(\theta; \mathcal{D})$. We hence argue that GSAM can only lead to sub-optimal generalization ability. Since the surrogate gap is always non-negative, the increase of $\mathcal{L}(\theta; \mathcal{D})$ will make the decrease of $\mathcal{L}_p(\theta; \mathcal{D})$ difficult, which may hurt the generalization performance of the model.

4.2. The objective of SAGM

Inspired by the analysis above, we propose two conditions that should be satisfied to facilitate the model achieving good generalization performance. (i) First, the loss within a neighborhood of the desired minimum should be sufficiently low. (ii) Second, the obtained minimum is within a flat loss surface. In order to meet these two conditions, we propose the SAGM method to simultaneously minimize three objectives, the empirical risk loss $\mathcal{L}(\theta; \mathcal{D})$, the perturbed loss $\mathcal{L}_p(\theta; \mathcal{D})$ and the surrogate gap $h(\theta)$:

$$\min_{\theta} (\mathcal{L}(\theta; \mathcal{D}), \mathcal{L}_p(\theta; \mathcal{D}), h(\theta)). \quad (10)$$

Intuitively, we search for a region with a low loss by minimizing $\mathcal{L}(\theta; \mathcal{D})$ and $\mathcal{L}_p(\theta; \mathcal{D})$, and we search for a local minimum with a flat surface by minimizing $h(\theta)$. A low loss implies small training errors and represents good performance on the source training data, while a flat loss surface reduces the generalization gap between training and testing performances. When both objectives are reached, the trained model is expected to have high prediction accuracy and good generalization capability.

To achieve our goal, we design a new optimization ob-

jective as follows:

$$\min_{\theta} \mathcal{L}(\theta; \mathcal{D}) + \mathcal{L}_p(\theta - \alpha \nabla \mathcal{L}(\theta; \mathcal{D}); \mathcal{D}), \quad (11)$$

where α is a hyperparameter. We can rewrite Equation 11 as follows:

$$\min_{\theta} \mathcal{L}(\theta; \mathcal{D}) + \mathcal{L}(\theta + \hat{\epsilon} - \alpha \nabla \mathcal{L}(\theta; \mathcal{D}); \mathcal{D}), \quad (12)$$

where $\hat{\epsilon} \triangleq \rho \frac{\nabla \mathcal{L}(\theta; \mathcal{D})}{\|\nabla \mathcal{L}(\theta; \mathcal{D})\|}$ (please refer to Equation 3 in Section 3 for more details about $\hat{\epsilon}$). Then we have

$$\min_{\theta} \mathcal{L}(\theta; \mathcal{D}) + \mathcal{L}\left(\theta + \left(\frac{\rho}{\|\nabla \mathcal{L}(\theta; \mathcal{D})\|} - \alpha\right) \nabla \mathcal{L}(\theta; \mathcal{D}); \mathcal{D}\right). \quad (13)$$

As can be seen from Equation 13, our method only needs to calculate losses $\mathcal{L}(\theta; \mathcal{D})$ and $\mathcal{L}\left(\theta + \left(\frac{\rho}{\|\nabla \mathcal{L}(\theta; \mathcal{D})\|} - \alpha\right) \nabla \mathcal{L}(\theta; \mathcal{D}); \mathcal{D}\right)$. Note that the latter loss has the same computational cost as $\mathcal{L}_p(\theta; \mathcal{D})$. Considering the fact that $\mathcal{L}(\theta; \mathcal{D})$ has already been calculated when we calculate $\nabla \mathcal{L}(\theta; \mathcal{D})$, SAGM has basically the same computational cost as SAM.

4.3. Analysis and algorithm of SAGM

In this section, we provide some analyses to better understand the proposed SAGM method and explain how SAGM satisfies both the two conditions in Section 4.2.

For the second term in Equation 11, we can perform the first order Taylor expansion around $\theta + \hat{\epsilon}$ as follows:

$$\begin{aligned} & \min_{\theta} \mathcal{L}_p(\theta - \alpha \nabla \mathcal{L}(\theta; \mathcal{D}); \mathcal{D}) \\ &= \min_{\theta} \mathcal{L}_p(\theta; \mathcal{D}) - \alpha \nabla \mathcal{L}_p(\theta; \mathcal{D}) \cdot \nabla \mathcal{L}(\theta; \mathcal{D}) + O(\alpha \nabla \mathcal{L}(\theta; \mathcal{D})) \\ &\approx \min_{\theta} \mathcal{L}_p(\theta; \mathcal{D}) - \alpha \nabla \mathcal{L}_p(\theta; \mathcal{D}) \cdot \nabla \mathcal{L}(\theta; \mathcal{D}). \end{aligned} \quad (14)$$

Therefore, the objective of SAGM in Equation 11 can be rewritten as:

$$\min_{\theta} \mathcal{L}(\theta; \mathcal{D}) + \mathcal{L}_p(\theta; \mathcal{D}) - \alpha \nabla \mathcal{L}_p(\theta; \mathcal{D}) \cdot \nabla \mathcal{L}(\theta; \mathcal{D}). \quad (15)$$

This equation reveals that we want to *minimize* the loss $\mathcal{L}(\theta; \mathcal{D})$ and $\mathcal{L}_p(\theta; \mathcal{D})$, and *maximize* the inner product of $\nabla \mathcal{L}_p(\theta; \mathcal{D})$ and $\nabla \mathcal{L}(\theta; \mathcal{D})$. Minimizing $\mathcal{L}(\theta; \mathcal{D})$ and $\mathcal{L}_p(\theta; \mathcal{D})$ is intuitive and it satisfies condition (i). Below we will build the connection between minimizing $h(\theta)$ and maximizing $\nabla \mathcal{L}_p(\theta; \mathcal{D}) \cdot \nabla \mathcal{L}(\theta; \mathcal{D})$.

First of all, if $\nabla \mathcal{L}_p(\theta; \mathcal{D})$ and $\nabla \mathcal{L}(\theta; \mathcal{D})$ are along similar directions, their inner product is greater than zero, and achieves the peak when their directions are the same. Therefore, maximizing $\nabla \mathcal{L}_p(\theta; \mathcal{D}) \cdot \nabla \mathcal{L}(\theta; \mathcal{D})$ is actually constraining the directions of these two gradients to be consistent, that is, to perform gradient matching on the gradient of the two losses. We then investigate the relationship between gradient matching and minimizing $h(\theta)$. As shown

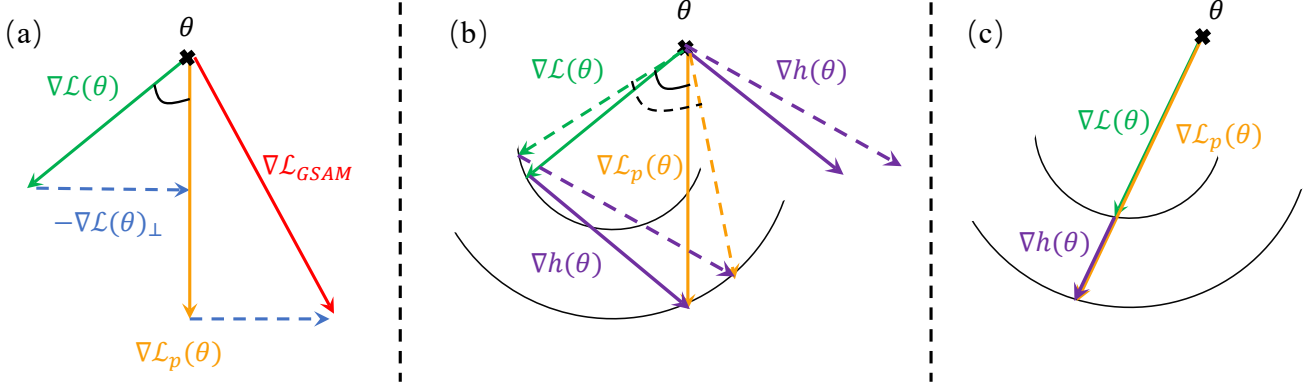


Figure 2: (a) GSAM [55] decomposes $\nabla\mathcal{L}(\theta; \mathcal{D})$ into parallel and vertical components by projecting it onto $\nabla\mathcal{L}_p(\theta; \mathcal{D})$, *i.e.*, $\nabla\mathcal{L}_{GSAM} = \nabla\mathcal{L}_p(\theta; \mathcal{D}) - \beta\nabla\mathcal{L}(\theta; \mathcal{D})_{\perp}$. (b) When the angle between $\nabla\mathcal{L}(\theta; \mathcal{D})$ and $\nabla\mathcal{L}_p(\theta; \mathcal{D})$ is large, there is an inevitable gradient conflict between $\nabla\mathcal{L}(\theta; \mathcal{D})$, $\nabla\mathcal{L}_p(\theta; \mathcal{D})$, and $\nabla h(\theta)$. (c) When the gradient directions of $\nabla\mathcal{L}(\theta; \mathcal{D})$ and $\nabla\mathcal{L}_p(\theta; \mathcal{D})$ are the same, $\mathcal{L}(\theta; \mathcal{D})$, $\mathcal{L}_p(\theta; \mathcal{D})$ and $h(\theta)$ all descend in the most efficient way. Best viewed in color.

in Fig. 2 (b), we draw two circles with $\|\nabla\mathcal{L}_p(\theta; \mathcal{D})\|$ and $\|\nabla\mathcal{L}(\theta; \mathcal{D})\|$ as radii respectively, and $\nabla h(\theta)$ is the line connecting the respective points on the two circles. Then the gradient direction of $h(\theta)$ is determined by $\nabla\mathcal{L}_p(\theta; \mathcal{D})$ and $\nabla\mathcal{L}(\theta; \mathcal{D})$. If the angle between $\nabla\mathcal{L}_p(\theta; \mathcal{D})$ and $\nabla\mathcal{L}(\theta; \mathcal{D})$ is too large, it is hard to optimize $\mathcal{L}(\theta; \mathcal{D})$, $\mathcal{L}_p(\theta; \mathcal{D})$ and $h(\theta)$ simultaneously because of the serious gradient conflict among them. As shown in Fig. 2 (c), when the gradient directions of $\mathcal{L}(\theta; \mathcal{D})$ and $\mathcal{L}_p(\theta; \mathcal{D})$ are consistent, the gradient direction of $h(\theta)$ will also be consistent with them. At this time, $\mathcal{L}(\theta; \mathcal{D})$, $\mathcal{L}_p(\theta; \mathcal{D})$ and $h(\theta)$ will descend consistently and efficiently. Therefore, our optimization objective in Equation 11 can be well achieved by simultaneously minimizing the empirical risk loss $\mathcal{L}(\theta; \mathcal{D})$, the perturbed loss $\mathcal{L}_p(\theta; \mathcal{D})$ and performing gradient matching on the gradient of the two losses.

The algorithm of SAGM is summarized in Algorithm 1.

5. Experiments

5.1. Experiment setups and implementation details

Dataset and protocol. Following [18], we exhaustively evaluate our method and competing methods on five benchmarks: PACS [30] (9,991 images, 7 classes, and 4 domains), VLCS [14] (10,729 images, 5 classes, and 4 domains), OfficeHome [47] (15,588 images, 65 classes, and 4 domains), TerraIncognita [4] (24,788 images, 10 classes, and 4 domains), and DomainNet [37] (586,575 images, 345 classes, and 6 domains). For a fair comparison, we follow the training and evaluation protocol in DomainBed [18], including the dataset splits, hyperparameter (HP) search and model selection on the validation set, except for the HP search space. Since the original DomainBed setting requires very heavy computation resources, we use a reduced HP search

Algorithm 1: The Algorithm of SAGM

Input: Domains \mathcal{D} ; initial weight θ_0 ; learning rate γ , batch size, dropout rate, and weight decay; hyperparameter α and radius ρ in Equation 13; total number of iterations T .

Output: Model trained with SAGM.

```

1 for  $i \leftarrow 1$  to  $T$  do
2   Sample mini-batch  $\mathcal{B} \in \mathcal{D}$ ;
3   Compute the training loss gradient  $\nabla\mathcal{L}(\theta_t; \mathcal{B})$ ;
4   Compute  $\mathcal{L}_{SAGM}$  according to Equation 13:
      $\mathcal{L}_{SAGM} = \mathcal{L}(\theta_t; \mathcal{B}) + \mathcal{L}(\theta_t + (\frac{\rho}{\|\nabla\mathcal{L}(\theta_t; \mathcal{B})\|} - \alpha)\nabla\mathcal{L}(\theta_t; \mathcal{B}); \mathcal{B})$ ;
5   Update weights:
6     Vanilla  $\theta_{t+1} = \theta_t - \gamma\nabla\mathcal{L}(\theta_t; \mathcal{B})$ 
7     SAM  $\theta_{t+1} = \theta_t - \gamma\nabla\mathcal{L}_p(\theta_t; \mathcal{B})$ 
8     GSAM  $\theta_{t+1} = \theta_t - \gamma\nabla\mathcal{L}_{GSAM}$  (Eq. 9)
9     SAGM  $\theta_{t+1} = \theta_t - \gamma\nabla\mathcal{L}_{SAGM}$ ;
10   $i = i + 1$ ;
11 end
```

space to reduce computational costs as suggested in [8].

Evaluation protocol. In all the experiments, the performances of competing methods are evaluated by the leave-one-out cross-validation scheme. In each case, a single domain is used as the target (test) domain and the others are used as the source (training) domains. The performance is calculated by averaging over three different trials with different train-validation splits. The final accuracy is reported as the average of all possible out-of-domain settings.

Implementation details. We use ResNet-50 [19] pre-trained on ImageNet [40] as the default backbone. The model is optimized by using the Adam [26] optimizer.

Table 1: **Comparison with state-of-the-art domain generalization methods.** Out-of-domain accuracies on five domain generalization benchmarks are shown. The **best results** are highlighted in bold. The results marked by †, ‡ are copied from Gulrajani and Lopez-Paz [18] and Cha *et al.* [8], respectively. Average accuracies and standard errors are calculated from three trials. The method ‘Miro (with CLIP)’ uses a pre-trained large-scale CLIP model [39] to achieve stronger generalization ability. Following [8], we do not use strategies such as warm up for all methods for a fair comparison.

Algorithm	PACS	VLCS	OfficeHome	TerraInc	DomainNet	Avg.
MMD† [31]	84.7±0.5	77.5±0.9	66.3±0.1	42.2±1.6	23.4±9.5	58.8
Mixstyle‡ [53]	85.2±0.3	77.9±0.5	60.4±0.3	44.0±0.7	34.0±0.1	60.3
GroupDRO† [41]	84.4±0.8	76.7±0.6	66.0±0.7	43.2±1.1	33.3±0.2	60.7
IRM† [1]	83.5±0.8	78.5±0.5	64.3±2.2	47.6±0.8	33.9±2.8	61.6
ARM† [51]	85.1±0.4	77.6±0.3	64.8±0.3	45.5±0.3	35.5±0.2	61.7
VREx† [28]	84.9±0.6	78.3±0.2	66.4±0.6	46.4±0.6	33.6±2.9	61.9
CDANN† [32]	82.6±0.9	77.5±0.1	65.8±1.3	45.8±1.6	38.3±0.3	62.0
AND-mask [42]	84.4±0.9	78.1±0.9	65.6±0.4	44.6±0.3	37.2±0.6	62.0
DANN† [16]	83.6±0.4	78.6±0.4	65.9±0.6	46.7±0.5	38.3±0.1	62.6
RSC† [21]	85.2±0.9	77.1±0.5	65.5±0.9	46.6±1.0	38.9±0.5	62.7
MTL† [5]	84.6±0.5	77.2±0.4	66.4±0.5	45.6±1.2	40.6±0.1	62.9
Mixup† [49]	84.6±0.6	77.4±0.6	68.1±0.3	47.9±0.8	39.2±0.1	63.4
MLDG† [29]	84.9±1.0	77.2±0.4	66.8±0.6	47.7±0.9	41.2±0.1	63.6
ERM [46]	85.5±0.2	77.3±0.4	66.5±0.3	46.1±1.8	43.8±0.1	63.9
Fish [44]	85.5±0.3	77.8±0.3	68.6±0.4	45.1±1.3	42.7±0.2	63.9
SagNet† [36]	86.3±0.2	77.8±0.5	68.1±0.1	48.6±1.0	40.3±0.1	64.2
SelfReg [25]	85.6±0.4	77.8±0.9	67.9±0.7	47.0±0.3	42.8±0.0	64.2
CORAL† [45]	86.2±0.3	78.8±0.6	68.7±0.3	47.6±1.0	41.5±0.1	64.5
mDSDI [6]	86.2±0.2	79.0±0.3	69.2±0.4	48.1±1.4	42.8±0.1	65.1
Miro [9] (with CLIP [39])	85.4±0.4	79.0±0.0	70.5±0.4	50.4±1.1	44.3±0.2	65.9
SAM ‡[15]	85.8±0.2	79.4±0.1	69.6±0.1	43.3±0.7	44.3±0.0	64.5
GSAM [55]	85.9±0.1	79.1±0.2	69.3±0.0	47.0±0.8	44.6±0.2	65.1
SAGM (ours)	86.6±0.2	80.0±0.3	70.1±0.2	48.8±0.9	45.0±0.2	66.1

The hyperparameter α in Equation 13 is tuned in [0.001, 0.0005]. The hyperparameter ρ is set to 0.05 following SAM [15]. The other hyperparameters, such as batch size, learning rate, dropout rate, and weight decay, are tuned in the same search space as that proposed in Cha *et al.* [8]. The optimal parameter settings on each dataset are provided in the **supplementary file**.

5.2. Main results

We report the average out-of-domain performances of state-of-the-art DG methods on five benchmarks in Table 1. Due to the limit of space, the accuracies for each domain are reported in the **supplementary file**. Since SAGM does not rely on domain labels, it can be applied to other robustness tasks that do not contain domain labels. We also show the generalization capability of SAGM on ImageNet and ImageNet-R in the **supplementary file**.

Comparison with conventional DG methods. We first

compare SAGM with conventional DG methods [31, 53, 41, 1, 51, 28, 32, 42, 16, 21, 5, 49, 29, 44, 46, 36, 25, 45, 6] on the top panel of Table 1. One can see that SAGM consistently outperforms ERM [46] on all benchmarks, resulting in +2.2% average improvement. Meanwhile, SAGM outperforms the leading conventional DG method on each individual dataset: +0.3% on PACS over SagNet [36] (86.6% → 86.3%); +1.0% on VLCS over mDSDI [6] (79.0% → 80.0%); +0.9% on OfficeHome over mDSDI [6] (69.2% → 70.1%); +0.7% on TerraIncognita over mDSDI [6] (48.1% → 48.8%); and +2.2% on DomainNet over mDSDI [6] and SelfReg [25] (42.8% → 45.0%). It is worth mentioning that the performance improvement of SAGM on DomainNet is very significant, as DomainNet is the largest dataset with 586,575 images from 6 domains.

Comparison with Miro (with CLIP) [9]. SAGM even has an overall performance advantage over Miro, which exploits the large-scale pre-trained CLIP [39] model to im-

prove the generalization capability. As shown in Table 1, SAGM outperforms Miro on PACS, VLCS and DomainNet, resulting in +0.2% average improvement.

Comparison with sharpness-aware methods. We then compare SAGM with the leading sharpness-aware methods, including SAM [15] and GSAM [55]. As shown in the bottom panel of Table 1, SAGM consistently outperforms SAM and GSAM on all the five datasets, leading to an average performance improvement of 1.6% and 1.0% over SAM and GSAM, respectively.

Considering the extensive experiments on 5 datasets and 22 target domains, the above results clearly demonstrate the effectiveness and superiority of SAGM in improving the domain generalization capability of deep models.

5.3. In-domain and out-of-domain performance

To show that the performance improvement of SAGM is brought by satisfying the two conditions we proposed in Section 4.2, we compare SAGM with three classes of generalization methods in more details: data augmentation methods, gradient-based methods, and sharpness-aware methods. In these experiments, we split in-domain datasets into training (60%), validation (20%), and test (20%) splits, and the average results over three runs on PACS are shown in Table 2. It can be seen that the conventional methods can help in-domain generalization, performing better than ERM on in-domain test set. However, their out-of-domain performances can be worse than ERM.

Comparison with data augmentation methods. The data augmentation methods, including Mixup [49], CutMix [50] and Mixstyle [53], can improve the in-domain performance over ERM by 0.7%, 1.0% and 0.6%, but their out-of-domain performances are degraded by 0.5%, 1.5% and 1.4%, respectively. This is because these data augmentation methods cannot ensure the model to find a flat minima. On the contrary, our SAGM method seeks a flat minimum to reduce the performance gap between the source and target domains. Actually, we can combine SAGM with data augmentation methods such as Mixstyle to further improve the performance. Specifically, Mixstyle can transform data \mathcal{D} into \mathcal{D}_{new} , which has the same semantics as \mathcal{D} but different styles. We can then get a stronger model by minimizing (i) the loss $\mathcal{L}(\theta; \mathcal{D})$ and $\mathcal{L}_p(\theta; \mathcal{D}_{new})$, as well as (ii) the angle between $\nabla \mathcal{L}_p(\theta; \mathcal{D}_{new})$ and $\nabla \mathcal{L}(\theta; \mathcal{D})$. For SAGM+Mixstyle, we insert Mixstyle block after the 1st, 2nd, and 3rd residual blocks following Mixstyle [53]. As shown in Table 2, SAGM+Mixstyle obtains further performance improvement over SAGM and Mixstyle.

Comparison with gradient-based methods. From Table 2, we can see that those methods solving the DG problem from the gradient perspective, including Fish [44], AND-mask [42] and MLDG [29], achieve good in-domain performance, but they fail to improve out-of-domain performance

Table 2: **In-domain and out-of-domain performance of different generalization methods on PACS.** The scores are averaged over all settings using different target domains. \uparrow and \downarrow indicate statistically significant improvement and degradation from ERM.

Method	Out-of-domain	In-domain
ERM	85.3±0.4	96.6±0.0
ERM+Mixup	84.8±0.3↓	97.3±0.1↑
ERM+CutMix	83.8±0.4↓	97.6±0.1↑
ERM+Mixstyle	83.9±0.3↓	97.2±0.2↑
Fish	85.3±0.2(-)	97.4±0.3↑
AND-mask	83.7±0.5↓	97.1±0.2↑
MLDG	84.0±0.2↓	97.4±0.2↑
SAM	85.6±0.1↑	97.4±0.1↑
GSAM	85.8±0.2↑	97.2±0.3↑
SAGM	86.2±0.2↑	97.6±0.1↑
SAGM+Mixstyle	87.2±0.2↑	97.8±0.1↑

because they cannot guarantee the model converge to a flat region. For example, AND-mask and MLDG improve the in-domain performances by 0.5% and 0.8% but degrade the out-of-domain performances by 1.6% and 1.3%. In contrast, our method achieves better generalization performance under both settings by performing gradient matching of $\mathcal{L}(\theta; \mathcal{D})$ and $\mathcal{L}_p(\theta; \mathcal{D})$ during training.

Comparison with sharpness-aware methods. It can be seen that the sharpness-aware method SAM [15] can increase both in-domain and out-of-domain performances over ERM, but the improvement on out-of-domain performance is not statistically significant. This suggests that SAM may not find a really flat minimum. As a comparison, GSAM [55] has a smaller intra-domain and out-of-domain performance gap than SAM. These experimental results are consistent with our theoretical analysis in Section 4. As a comparison, SAGM improves both the in-domain and out-of-domain performances over SAM and GSAM, achieving 0.8% gain on the out-of-domain performance and 1.0% gain on the in-domain performance over ERM.

5.4. Ablation study

In the following, we conduct an ablation study to investigate the effectiveness of the optimization objective of SAGM. Specifically, we compare the performance of SAGM with ERM, SAM, GSAM, as well as ERM+SAM, whose optimization objective is $\mathcal{L}_p(\theta; \mathcal{D}) + \mathcal{L}(\theta; \mathcal{D})$. Compared with SAGM (refer to Equation 15), ERM+SAM does not have the gradient matching term so that we can use it to investigate the effectiveness of gradient matching in SAGM. Experiments are conducted on the DomainBed benchmark.

The following observations can be made from Table 3. First, both SAM, GSAM, and SAGM outperform ERM by significant margins, which verifies the effectiveness of solv-

Table 3: **Ablation study on SAGM.** The scores are averaged over all target domain cases. The performances of ERM, SAM, GSAM, ERM+SAM are all optimized by HP searches of DomainBed.

	optimization objective	PACS	VLCS	OfficeHome	TerraInc	DomainNet	Avg.
ERM	$\mathcal{L}(\theta; \mathcal{D})$	85.5±0.2	77.3±0.4	66.5±0.3	46.1±1.8	43.8±0.1	63.9
SAM	$\mathcal{L}_p(\theta; \mathcal{D})$	85.8±0.2	79.4±0.1	69.6±0.1	43.3±0.7	44.3±0.0	64.5
GSAM	$\mathcal{L}_p(\theta; \mathcal{D}) \downarrow -\mathcal{L}(\theta; \mathcal{D}) \uparrow$	85.9±0.1	79.1±0.2	69.3±0.0	47.0±0.8	44.6±0.2	65.1
ERM + SAM	$\mathcal{L}_p(\theta; \mathcal{D}) + \mathcal{L}(\theta; \mathcal{D})$	85.6±0.2	78.6±0.2	69.5±0.4	46.4±1.2	44.2±0.0	64.8
SAGM	$\mathcal{L}(\theta; \mathcal{D}) + \mathcal{L}_p(\theta - \alpha \nabla \mathcal{L}(\theta; \mathcal{D}); \mathcal{D})$	86.6±0.2	80.0±0.3	70.1±0.2	48.8±0.4	45.0±0.1	66.1

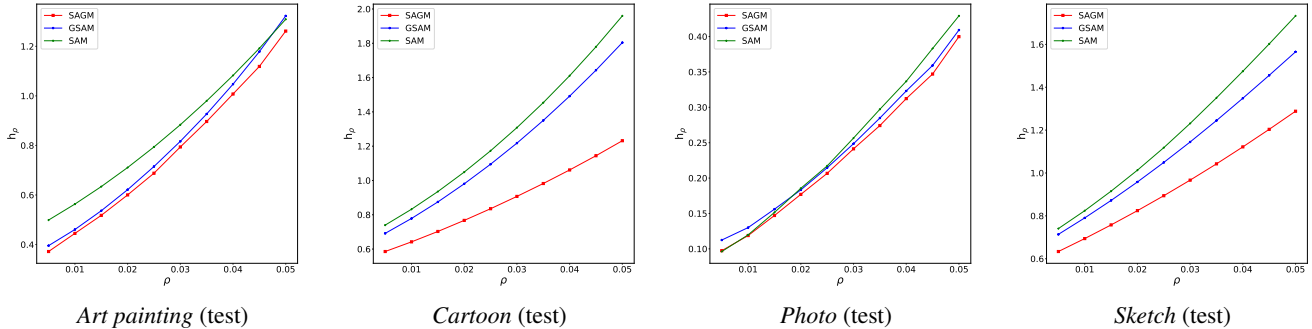


Figure 3: **Local sharpness comparison.** We plot the local sharpness via the loss gap of SAM, GSAM and SAGM by varying radius ρ on different target domains of PACS. For each figure, Y-axis indicates the sharpness $h_\rho(\theta)$ and X-axis indicates the radius ρ . One can see that SAGM finds flatter minima than SAM and GSAM.

ing the DG problem by seeking a flat minima. Second, SAGM achieves an average performance improvement of 1.3% over ERM+SAM, which illustrates the importance of gradient matching in our designed objective function. Third, SAGM consistently outperforms SAM and GSAM methods by a large margin. These experimental results verify the effectiveness of the proposed SAGM method.

5.5. Local sharpness analysis

In this experiment, we show the local sharpness of the trained models to demonstrate that SAGM not only achieves superior DG performance to SAM and GSAM, but also converges to a flatter region in the loss landscape. To begin with, we quantify the local sharpness of a model parameter θ by assuming that flat minima will have smaller loss value changes within their neighborhoods than sharp minima. For the given model parameter θ , we compute the expected loss value changes between θ and parameters $\theta + \rho \frac{\nabla \mathcal{L}(\theta; \mathcal{D})}{\|\nabla \mathcal{L}(\theta; \mathcal{D})\|}$ with radius ρ , i.e., $h_\rho(\theta) = \mathcal{L}(\theta + \rho \frac{\nabla \mathcal{L}(\theta; \mathcal{D})}{\|\nabla \mathcal{L}(\theta; \mathcal{D})\|}; \mathcal{D}) - \mathcal{L}(\theta; \mathcal{D})$. Specifically, we compare the $h_\rho(\theta)$ of SAM, GSAM and our SAGM by varying radius ρ on each target domain of PACS.

From Fig. 3, one can see that GSAM finds a flatter region than SAM in most experiments, while SAGM finds the flat-test minima in all experiments. It is noted that when the

target domain is ‘Photo (test)’, the value of $h_\rho(\theta)$ is much smaller than other domains. This is because SAM, GSAM and our SAGM all achieve over 97.0% accuracy on this domain (please refer to the **supplementary file** for the performance on each domain of PACS), which leads to smaller loss value changes. Nonetheless, the trend of ‘Photo (test)’ is consistent with other domains.

6. Conclusion

In this paper, we analyzed the limitations of SAM-like algorithms for domain generalization (DG), and proposed two conditions that could ensure the model converges to a flat region with improved generalization ability. We then proposed a novel algorithm named Sharpness-Aware Gradient Matching (SAGM), which minimized simultaneously the empirical loss, perturbed loss and the gap between them, to meet these two conditions. The proposed SAGM method demonstrated superior performance to state-of-the-art DG methods, including SAM and GSAM, on the five DG benchmarks. Specifically, SAGM achieved 66.1% average performance on DomainBed without using any additional information. It even outperformed the Miro [9] method that is equipped with the pre-trained CLIP model, demonstrating its strong DG capability.

Acknowledgement. This work was supported by the InnoHK program.

References

- [1] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019. 2, 6
- [2] Hyojin Bahng, Sanghyuk Chun, Sangdoon Yun, Jaegul Choo, and Seong Joon Oh. Learning de-biased representations with biased representations. In *International Conference on Machine Learning (ICML)*, 2020. 1, 2
- [3] Yogesh Balaji, Swami Sankaranarayanan, and Rama Chellappa. Metareg: Towards domain generalization using meta-regularization. *Neural Information Processing Systems*, 31:998–1008, 2018. 1, 2
- [4] Sara Beery, Grant Van Horn, and Pietro Perona. Recognition in terra incognita. In *European Conference on Computer Vision*, 2018. 5
- [5] Gilles Blanchard, Aniket Anand Deshmukh, Urun Dogan, Gyemin Lee, and Clayton Scott. Domain generalization by marginal transfer learning. *Journal of Machine Learning Research*, 22(2):1–55, 2021. 6
- [6] Manh-Ha Bui, Toan Tran, Anh Tran, and Dinh Phung. Exploiting domain-specific features to enhance domain generalization. *Neural Information Processing Systems*, 2021. 6
- [7] Fabio M Carlucci, Antonio D’Innocente, Silvia Bucci, Barbara Caputo, and Tatiana Tommasi. Domain generalization by solving jigsaw puzzles. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2229–2238, 2019. 1, 2
- [8] Junbum Cha, Sanghyuk Chun, Kyungjae Lee, Hanchchol Cho, Seunghyun Park, Yunsung Lee, and Sungrae Park. Swad: Domain generalization by seeking flat minima. *Advances in Neural Information Processing Systems*, 34:22405–22418, 2021. 2, 5, 6
- [9] Junbum Cha, Kyungjae Lee, Sungrae Park, and Sanghyuk Chun. Domain generalization by mutual-information regularization with pre-trained models. *European Conference on Computer Vision (ECCV)*, 2022. 6, 8
- [10] Laurent Dinh, Razvan Pascanu, Samy Bengio, and Yoshua Bengio. Sharp minima can generalize for deep nets. In *International Conference on Machine Learning*, pages 1019–1028. PMLR, 2017. 2
- [11] Qi Dou, Daniel C Castro, Konstantinos Kamnitsas, and Ben Glocker. Domain generalization via model-agnostic learning of semantic features. *Neural Information Processing System*, 2019. 1, 2
- [12] 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
- [13] Gintare Karolina Dziugaite and Daniel M Roy. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. In *Uncertainty in Artificial Intelligence*, 2017. 2
- [14] Chen Fang, Ye Xu, and Daniel N Rockmore. Unbiased metric learning: On the utilization of multiple datasets and web images for softening bias. In *International Conference on Computer Vision*, 2013. 5
- [15] 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, 6, 7
- [16] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *Journal of machine learning research*, 17(1):2096–2030, 2016. 1, 2, 6
- [17] Timur Garipov, Pavel Izmailov, Dmitrii Podoprikin, Dmitry P Vetrov, and Andrew Gordon Wilson. Loss surfaces, mode connectivity, and fast ensembling of dnns. In *Neural Information Processing Systems*, 2018. 1, 2
- [18] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. In *International Conference on Learning Representations*, 2021. 1, 5, 6
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Computer Vision and Pattern Recognition*, 2016. 5
- [20] Sepp Hochreiter and Jürgen Schmidhuber. Simplifying neural nets by discovering flat minima. In *Advances in neural information processing systems*, pages 529–536, 1995. 2
- [21] Zeyi Huang, Haohan Wang, Eric P Xing, and Dong Huang. Self-challenging improves cross-domain generalization. *European Conference on Computer Vision*, 2020. 6
- [22] Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. *Conference on Uncertainty in Artificial Intelligence*, 2018. 1, 2
- [23] 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. 1, 2
- [24] Aditya Khosla, Tinghui Zhou, Tomasz Malisiewicz, Alexei A Efros, and Antonio Torralba. Undoing the damage of dataset bias. In *ECCV*, pages 158–171. Springer, 2012. 2
- [25] Daehee Kim, Youngjun Yoo, Seunghyun Park, Jinkyu Kim, and Jaekoo Lee. Selfreg: Self-supervised contrastive regularization for domain generalization. In *International Conference on Computer Vision*, 2021. 6
- [26] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015. 5
- [27] Masanori Koyama and Shoichiro Yamaguchi. Out-of-distribution generalization with maximal invariant predictor. 2020. 2
- [28] David Krueger, Ethan Caballero, Joern-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Remi Le Priol, and Aaron Courville. Out-of-distribution generalization via risk extrapolation (rex). *arXiv preprint arXiv:2003.00688*, 2020. 2, 6
- [29] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy Hospedales. Learning to generalize: Meta-learning for domain generalization. In *AAAI Conference on Artificial Intelligence*, volume 32, 2018. 1, 2, 6, 7
- [30] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generaliza-

- tion. In *International Conference on Computer Vision*, 2017. 5
- [31] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *Computer Vision and Pattern Recognition*, 2018. 6
- [32] Ya Li, Mingming Gong, Xinmei Tian, Tongliang Liu, and Dacheng Tao. Domain generalization via conditional invariant representations. In *AAAI Conference on Artificial Intelligence*, volume 32, 2018. 1, 2, 6
- [33] 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
- [34] Lucas Mansilla, Rodrigo Echeveste, Diego H Milone, and Enzo Ferrante. Domain generalization via gradient surgery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6630–6638, 2021. 2
- [35] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *International Conference on Machine Learning*, pages 10–18. PMLR, 2013. 1, 2
- [36] Hyeonseob Nam, HyunJae Lee, Jongchan Park, Wonjun Yoon, and Donggeun Yoo. Reducing domain gap by reducing style bias. In *Computer Vision and Pattern Recognition*, 2021. 6
- [37] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *International Conference on Computer Vision*, 2019. 5
- [38] Xingchao Peng, Zijun Huang, Ximeng Sun, and Kate Saenko. Domain agnostic learning with disentangled representations. In *International Conference on Machine Learning*, pages 5102–5112. PMLR, 2019. 2
- [39] 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
- [40] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015. 5
- [41] Shiori Sagawa*, Pang Wei Koh*, Tatsunori B. Hashimoto, and Percy Liang. Distributionally robust neural networks. In *International Conference on Learning Representations*, 2020. 6
- [42] Soroosh Shahtalebi, Jean-Christophe Gagnon-Audet, Touraj Laleh, Mojtaba Faramarzi, Kartik Ahuja, and Irina Rish. Sand-mask: An enhanced gradient masking strategy for the discovery of invariances in domain generalization. *arXiv preprint arXiv:2106.02266*, 2021. 6, 7
- [43] 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. 1, 2
- [44] 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. 2, 6, 7
- [45] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *European Conference on Computer Vision*, 2016. 6
- [46] V Vapnik. *Statistical learning theory*. NY: Wiley, 1998. 6
- [47] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *Computer Vision and Pattern Recognition*, 2017. 5
- [48] Guoqing Wang, Hu Han, Shiguang Shan, and Xilin Chen. Cross-domain face presentation attack detection via multi-domain disentangled representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6678–6687, 2020. 2
- [49] Minghao Xu, Jian Zhang, Bingbing Ni, Teng Li, Chengjie Wang, Qi Tian, and Wenjun Zhang. Adversarial domain adaptation with domain mixup. In *AAAI Conference on Artificial Intelligence*, 2020. 6, 7
- [50] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6023–6032, 2019. 7
- [51] Marvin Zhang, Henrik Marklund, Abhishek Gupta, Sergey Levine, and Chelsea Finn. Adaptive risk minimization: A meta-learning approach for tackling group shift. *arXiv preprint arXiv:2007.02931*, 2020. 1, 2, 6
- [52] Shanshan Zhao, Mingming Gong, Tongliang Liu, Huan Fu, and Dacheng Tao. Domain generalization via entropy regularization. *Neural Information Processing Systems*, 33, 2020. 1, 2
- [53] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain generalization with mixstyle. In *International Conference on Learning Representations*, 2021. 1, 2, 6, 7
- [54] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain generalization with mixstyle. *arXiv preprint arXiv:2104.02008*, 2021. 2
- [55] Juntang Zhuang, Boqing Gong, Liangzhe Yuan, Yin Cui, Hartwig Adam, Nicha Dvornek, Sekhar Tatikonda, James Duncan, and Ting Liu. Surrogate gap minimization improves sharpness-aware training. *arXiv preprint arXiv:2203.08065*, 2022. 1, 2, 3, 5, 6, 7