Improving Out-of-Distribution Generalization by Adversarial Training with Structured Priors

Qixun Wang[1*] **Yifei Wang** [2*] **Hong Zhu**[3] **Yisen Wang**[1,4†]
[1] Key Lab. of Machine Perception (MoE),
School of Intelligence Science and Technology, Peking University
[2] School of Mathematical Sciences, Peking University
[3] Huawei Noah's Ark Lab
[4] Institute for Artificial Intelligence, Peking University

Abstract

Deep models often fail to generalize well in test domains when the data distribution differs from that in the training domain. Among numerous approaches to address this Out-of-Distribution (OOD) generalization problem, there has been a growing surge of interest in exploiting Adversarial Training (AT) to improve OOD performance. Recent works have revealed that the robust model obtained by conducting sample-wise AT also retains transferability to biased test domains. In this paper, we empirically show that sample-wise AT has limited improvement on OOD performance. Specifically, we find that AT can only maintain performance at smaller scales of perturbation while Universal AT (UAT) is more robust to larger-scale perturbations. This provides us with clues that adversarial perturbations with universal (low dimensional) structures can enhance the robustness against large data distribution shifts that are common in OOD scenarios. Inspired by this, we propose two AT variants with low-rank structures to train OOD-robust models. Extensive experiments on DomainBed benchmark show that our proposed approaches outperform Empirical Risk Minimization (ERM) and sample-wise AT. Our code is available at `https://github.com/NOVAglow646/NIPS22-MAT-and-LDAT-for-OOD`.

1 Introduction

Existing deep learning methods have achieved good performance on visual classification tasks under the same distribution of training sets and test sets. However, when the data distribution of the test set is different from that of the training set, the classification performance of the deep neural networks (DNNs) may decrease sharply [1]. This is mainly because DNNs may capture spurious features such as the background and style information to assist the fast fitting during the training process [2]. However, in real-world scenarios, test data may differ from training data in the background and style information, thus DNNs that rely on unstable spurious features to make predictions will fail. Solving the above problem is known as the out-of-distribution (OOD) generalization.

Another scenario where DNNs may fail is that they are often vulnerable to adversarial examples [3]. Adversarial training (AT) is originally proposed as an effective way to defend against adversarial attacks [4]. Moreover, there is work showing that adversarial training helps to solve the OOD generalization problem because OOD data can be seen as stronger perturbations to some extent [5]. The reason why AT can defend against adversarial attacks meanwhile benefit OOD generalization is that it can make DNNs robust to the interference of spurious features, such as randomly injected noise

*Equal Contribution.
†Corresponding author: Yisen Wang (yisen.wang@pku.edu.cn).

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

(in adversarial examples) or the spurious correlation between labels and background information (in OOD generalization). In other words, AT enables DNNs to make predictions using intrinsic features rather than spurious features.

A potential problem, however, is that existing AT methods ignore the specific design of perturbations when used for solving OOD generalization problems. They usually simply conduct sample-wise AT [6], which only brings limited performance improvement to OOD generalization. The essential reason for the failure of this type of approach is that the perturbations it uses cannot distinguish invariant and spurious features. As a result, it improves the robustness at the expense of the decreasing standard accuracy [7]. Moreover, we empirically find that when adapting Universal AT (UAT [8]) to OOD problems, i.e., conducting AT with domain-wise perturbations, it shows stronger input-robustness when facing larger-scale perturbations compared to the sample-wise AT (see Section 3.2). Since the sample injected with large-scale perturbations can be regarded as OOD samples [5], we draw inspiration from this phenomenon that AT with universal (low-dimensional) structures can be the key to solving OOD generalization. Therefore, we propose to use structured low-rank perturbations related to domain information in AT, which can help the model to filter out background and style information, thus benefiting OOD generalization. We make the following contributions in our work:

- We identify the limitations of sample-wise AT on OOD generalization through a series of experiments. To alleviate this problem, we further propose two simple but effective AT variants with structured priors to improve OOD performances.
- We theoretically prove that our proposed structured AT approach can accelerate the convergence of reliance on spurious features to 0 when using finite-time-stopped gradient descent, thus enhancing the robustness of the model against spurious correlations.
- By conducting experiments on the DomainBed benchmark [9], we demonstrate that our methods outperform ERM and sample-wise AT on various OOD datasets.

2 Related Work

Solving OOD Generalization with AT. According to [3], the performance of deep models is susceptible to small-scale perturbations injected in the input images, even if these perturbations are imperceptible to humans. Adversarial training (AT) is an effective approach to improve the robustness to input perturbations [4, 10, 11]. However, many recent works have begun to focus on the connection between AT and OOD due to the fact that OOD data can be regarded as one kind of large-scale perturbation. These works seek to exploit the robustness provided by AT to improve OOD generalization. For instance, [6] applied sample-wise AT to OOD generalization. They theoretically found that if a model is robust to input perturbation on training samples, it also generalizes well on OOD data. [5] theoretically established a link between the objective of AT and the OOD robustness. They revealed that the AT procedure can be regarded as a heuristic solution to the worst-case problem around the training domain distribution. Nevertheless, the discussion of [6] and [5] is restricted to the framework of using Wasserstein distance to measure the distribution shift, which is less practical for the real-world OOD setting where domain shifts are diverse. Additionally, they only studied the case of sample-wise AT and did not further investigate the effect of different forms of AT (not sample-wise) on OOD performance. Other works such as [12] focus on the structure design of the perturbations. They used multi-scale perturbations within one sample, but they did not exploit the universal information within one training domain. In our work, we focus on real-world OOD scenarios where there are additional clues lying in the distribution shifts, i.e, the low-rank structures in the spurious features (such as background and style information) across one domain. We further design a low-rank structure in the perturbations to specifically eliminate such low-rank spurious correlations.

OOD Evaluation Benchmark. The DomainBed benchmark [9] provides a fair way of evaluating different state-of-the-art OOD methods, which has been widely accepted by the community. By conducting rigorous experiments in a consistent setting, they revealed that many algorithms that claim to outperform previous methods cannot even outperform ERM. Unlike previous works using AT to address OOD generalization, such as [6] and [5], we adopt the Domainbed benchmark for a fair comparison of our approach with existing state-of-the-art methods in this paper.

3 Weakness of Sample-wise AT for OOD Generalization

3.1 Preliminaries

Out-of-distribution (OOD) Generalization. Assuming $x \in \mathcal{X}$ as the random data in the input space $\mathcal{X}$ and $y \in \mathcal{Y}$ as the target random data in the label space $\mathcal{Y}$, we have the predictor $f = w \circ \phi(x)$ where $\phi : \mathcal{X} \to \mathcal{Z}$ denotes the feature extractor and $w : \mathcal{Z} \to \mathcal{Y}$ denotes the classifier.

Now we give the formal definition of the OOD generalization problem. We have a set of m training domains $\mathcal{E} = \{E_1, E_2, ..., E_m\}$, where each domain E_e is characterized by a input dataset $E_e := \{(x_i^e, y_i^e)\}_{i=1}^{n_e}$ containing n_e i.i.d input samples drawn from the distribution of $\mathcal{P}_e$, and a test domain E_{m+1} with data following the distribution of $\mathcal{P}_{te}$, where $\mathcal{P}_{te} \neq \mathcal{P}_i,\ i = 1, 2, ..., m$. $\mathcal{L} : \mathcal{X} \to \mathbb{R}^+$ denotes the loss function. The ultimate goal of OOD generalization is to find an optimal predictor f that minimizes the risk on the unseen test domain:

$$\min_f \mathbb{E}_{(x,y)\sim\mathcal{P}_{te}(x,y)}[\mathcal{L}(f(x), y)]. \tag{1}$$

Adversarial Training (AT)[3]. According to [4], AT can be expressed as the following optimization problem:

$$\min_f \mathbb{E}_{(x,y)\sim\mathcal{P}(x,y)}[\max_{\delta\in\mathcal{S}} \mathcal{L}(f(x+\delta), y)] \ \text{ s.t. } \|\delta\|_p \leq \epsilon, \tag{2}$$

where $\delta \in \mathcal{S}$ is the random injected perturbation with l_p norm bounded by ϵ. The inner maximization problem can be optimized by fast gradient sign method (FGSM [13]), a simple one-step scheme:

$$x = x + \epsilon \text{sgn}(\nabla_x \mathcal{L}(f(x), y)), \tag{3}$$

where $\text{sgn}(\cdot)$ is the sign function, or by projected gradient descent (PGD [4]), a more powerful multi-step variant:

$$x^{t+1} = \prod_{\mathcal{S}}(x^t + \gamma \text{sgn}(\nabla_x \mathcal{L}(f(x), y))), \tag{4}$$

where $\prod_{\mathcal{S}}$ is the projection operator onto the set $\mathcal{S}$, γ is the step size and t denotes the iteration.

3.2 Weakness of AT for OOD Generalization

We now highlight some weaknesses of sample-wise AT for OOD generalization based on a series of empirical evidence. We first conduct a toy experiment on the DomainBed benchmark [9] to evaluate the OOD performance of AT. We run ERM and AT on four OOD datasets: PACS [14], OfficeHome [15], VLCS [16], and NICO [17] with a fixed set of hyperparameters (detailed experimental settings can be found in Appendix C.1). The results are shown in Table 1. We can see that the improvement of OOD performance by AT is limited with an average improvement of only 0.1%.

Table 1: Test accuracy (%) on four OOD datasets on DomainBed benchmark with a fixed set of hyperparameters. The improvement of AT is marginal.

	Datasets				
Algorithm	PACS	OfficeHome	VLCS	NICO	avg
ERM	79.7 ± 0.0	59.6 ± 0.0	74.4 ± 1.0	**70.7 ± 1.0**	71.1
AT	**81.5 ± 0.4**	**59.9 ± 0.4**	**75.3 ± 0.7**	68.2 ± 2.2	**71.2**

We further investigate the reason behind the limitations of performance improvements on OOD datasets of AT. Although previous works have revealed that the robust features obtained by AT can improve OOD generalization ([6] [5] [18]), we find that sample-wise AT only tolerates small-scale perturbations. Thus, we design an experiment on NICO dataset with multiple scales of perturbations. The scale is calculated with the l_2 norm of the perturbation matrix (experiment details are shown in Appendix C.1). As shown in Figure 1, AT suffers severe performance degradation when using large perturbations. This provides clues to understanding the failure of AT in OOD scenarios. The

[3]For simplicity, we denote 'AT' for sample-wise AT by default in the rest of the paper.

distribution shifts in OOD data usually have much larger scales than the invisible perturbations commonly used in AT. Hence, AT methods designed for small perturbations cannot handle these large-scale domain shifts that often appear in OOD data. However, our experiment shows that this problem can be alleviated by adapting universal AT (UAT [8]) to the OOD setting, i.e., using a perturbation for each domain.

Figure 1 shows that UAT remains its generalization performance when the perturbation scale is large. There are two empirical explanations for this: First, the background and style information usually have a low-rank structure, such as the grassland and snowfield that have recurring parts. Second, similar spurious features often appear within one specific domain, such as PACS [14] and VLCS [16] datasets. As stated in [8], the universal perturbation lies in a low dimensional space. Hence using universal (domain-wise) perturbations will help to resist such low-rank shifts and improve the robustness of the model.

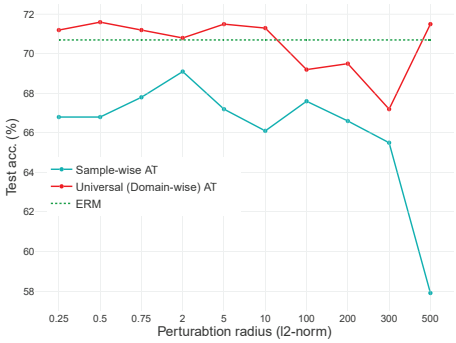

Figure 1: Test accuracy (%) of AT, UAT (l_2 norm), and ERM on NICO dataset.

Inspired by this, we proposed two new AT variants with more sophisticated low-rank structures on different dimensions to improve OOD generalization in the next section.

4 The Proposed Structured AT Method

In order to construct low-rank structured perturbations, we start by analyzing the structure of sample-wise perturbations. Assume that each input data x has a shape of $N \times N \times C$. N is the size of the input image and C is the number of channels. For simplicity, we assume $C = 3$. We reparameterize the sample-wise perturbations as a series of 2-D matrices $\{D_1^1, D_2^1, ..., D_m^1\}$, $\{D_1^2, D_2^2, ..., D_m^2\}$, $\{D_1^3, D_2^3, ..., D_m^3\}$ where $D_e^c \in R^{n_e \times N^2}$ denotes the perturbations in the e-th domain for the input channel c, n_e is the number of the samples in the domain E_e, and m is the number of domains. The i-th row of D_e^c represents the c-th channel of the i-th sample in the domain E_e (see the first column in Figure 2 for illustration). By such reparameterization, it is natural to find that there are two orientations to reduce the rank of the perturbations:

1. **Along the dimension of the number of samples** (along the red arrow in the upper left corner of Figure 2). This corresponds to reducing the number of the perturbations used within one domain.
2. **Along the dimension of the input scale** (along the blue arrow in the upper left corner of Figure 2). This corresponds to reducing the rank of the perturbation used for a specific input sample.

In the following parts, we propose two AT variants with structured priors that reduce the rank in these two directions.

4.1 MAT: Adversarial Training with Combinations of Multiple Perturbations

In this part, we propose domain-wise Multiple-perturbation Adversarial Training (MAT). It aims to conduct rank minimization along the dimension of the number of samples. Instead of using sample-wise perturbations, MAT constructs a combination of multiple perturbations and shares this mixed perturbation within a domain. Specifically, we choose to train the linear combination of k perturbations for each domain E_e to conduct AT. Here k is a hyperparameter and k is far less than the number of samples in domain E_e. The optimization problem can be reformulated as:

$$\min_f \sum_e \mathbb{E}_{(x,y)\sim\mathcal{P}_e(x,y)}[\mathcal{L}(f(x+\delta^e), y)], \tag{5}$$

$$\text{s.t. } \delta^e = \sum_{i=1}^{k} \alpha_i^{e*}\delta_i^{e*},\ \|\delta_i^{e*}\|_p \le \epsilon,\ \sum_{i=1}^{k}\alpha_i^{e*} = 1,\ \alpha_i^{e*} \ge 0 \text{ for } i = 1, 2, ..., k, \tag{6}$$

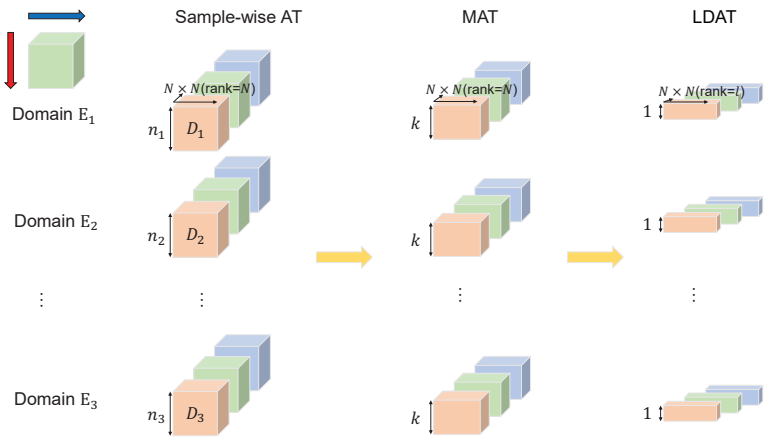

Figure 2: Illustration of how our proposed structured AT reduces the rank of the perturbations comparing to sample-wise AT. The left column shows how we reparameterize sample-wise AT. The perturbations are segmented by domains. A block represents the n_e perturbations injected in a channel of the samples in domain E_e. This figure shows the case where the input image has three channels (RGB). The red and the blue arrow in the upper left shows the two orientations to reduce the rank of the perturbations, i.e, along the dimension of the total number of the perturbations and along the dimension of the rank of a single perturbation, respectively. The mid column illustrates that MAT reduces the number of the perturbations used for domain E_e from n_e to k ($k \ll n_e$). The right column shows that LDAT further reduces the number of the perturbations from k to 1. Moreover, it reduces the rank of a specific perturbation from N to l ($l \ll N$).

where

$$\alpha_i^{e*},\ \delta_i^{e*} = \underset{\alpha_i^e,\ \delta_i^e}{\operatorname{argmax}} \mathbb{E}_{(x,y)\sim\mathcal{P}_e(x,y)}[\mathcal{L}(f(x + \sum_{i=1}^{k} \alpha_i^e \delta_i^e), y)]. \tag{7}$$

Here e denotes the subscript of a training domain and α_i^e is the weight that can be learned for each perturbation δ_i^e. The detailed training procedure of MAT is in Algorithm 1. We first initialize k perturbations δ_i^e and their correspondent coefficients α_i^e for each training domain e with Gaussian noise. Then we transform δ_i^e and α_i^e to make sure $\sum_{i=1}^{k} \alpha_i^e = 1$, $\alpha_i^e \geq 0$, and $||\delta_i^e||_2 \leq \epsilon$. For the inner maximization, we conduct a one-step gradient ascent to optimize δ_i and α_i.

MAT works as a low-rank version of sample-wise AT. In sample-wise AT, we maintain n_e perturbations for each domain E_e, where n_e is the number of training samples in domain E_e. As for MAT, it reduces the number of perturbations available to samples from n_e to k and obtains low-rank structures (see the third column in Figure 2 for illustration). Therefore, it fulfills rank reduction along the sample-number dimension.

4.2 LDAT: Adversarial Training with Low-rank Decomposed Perturbations

Based on MAT, we further propose Adversarial Training with Low-rank Decomposed perturbations (LDAT). Analogous to MAT, LDAT still shares one perturbation in a specific domain. Moreover, LDAT imposes a low-rank constraint on the perturbation itself, which corresponds to the dimension of the input scale. Technically, we obtain the domain-wise low-rank perturbation matrix $\delta \in \mathcal{R}^{N\times N\times C}$ by multiplying two matrices: $\delta = AB$. Here $A \in \mathcal{R}^{N\times l\times C}$ and $B \in \mathcal{R}^{l\times N\times C}$ where l is a hyperparameter and $l \ll N$. Since $\text{rank}(AB) \leq \text{rank}(A)$ and $\text{rank}(AB) \leq \text{rank}(B)$ hold for arbitrary matrices A, B, we have $\text{rank}(\delta) \leq l$. Therefore LDAT reduces the rank of the perturbation from a large value N to a relatively small value l (see the last column in Figure 2 for illustration).

Algorithm 1 Detailed Training Procedure of MAT

Input:

Labeled training data of m domains $E_1, ..., E_m$, where $E_e := \{(x_i^e, y_i^e)\}_{i=1}^{n_e}$,
number of the perturbations to be combined k, perturbation weight α learning rate η,
FGSM step size γ, perturbation radius ϵ,
number of training epochs T, learning rate for model parameters r, batch size b.

Output:

Updated model f_θ with parameter θ.

Randomly initiate θ, perturbation δ_i^e, weight α_i^e such that $\sum_{i=1}^k \alpha_i^e = 1$, $\alpha_i^e \geq 0$, $||\delta_i^e||_2 \leq \epsilon$, $\forall i \in \{1, ..., k\}$ and $\forall e \in \{1, ..., m\}$.
for iterations in $1, 2, ..., T$ **do**
  for e in $1, 2, ..., m$ **do**
    Randomly select batch $\mathcal{B}^e = \{(x_u^e, y_u^e)\}_{u=1}^b$ from domain E_e.
    Compute the adversarial sample: $x_u^{e'} = x_u^e + \sum_{j=1}^k \alpha_j^e \delta_j^e, \forall u \in \{1, ..., b\}$
    Update δ^e by $\delta_i^e \leftarrow \delta_i^e + \gamma \frac{1}{b} \sum_{u=1}^b \nabla_{\delta_i^e} \mathcal{L}(f_\theta(x_u^{e'}), y_u^e), \forall i \in \{1, ..., k\}, \forall u \in \{1, ..., b\}$.
    Update α_i^e by $\alpha_i^e \leftarrow \alpha_i^e + \eta \frac{1}{b} \sum_{u=1}^b \nabla_{\alpha_i^e} \mathcal{L}(f_\theta(x_u^{e'}), y_u^e), \forall i \in \{1, ..., k\}, \forall u \in \{1, ..., b\}$.
    Project δ_i^e to the l_2 ball of radius ϵ.
    Compute the adversarial sample: $x_u^{e'} = x_u^e + \sum_{j=1}^k \alpha_j^e \delta_j^e, \forall u \in \{1, ..., b\}$
    Update model parameter: $\theta \leftarrow \theta - r \frac{1}{b} \sum_{u=1}^b \nabla_\theta \mathcal{L}(f_\theta(x_u^{e'}), y_u^e), \forall u \in \{1, ..., b\}$.
  end for
end for

The formal definition of the LDAT objective is:

$$\min_f \sum_e \mathbb{E}_{(x,y)\sim\mathcal{P}_e(x,y)}[\mathcal{L}(f(x+\delta^e), y)], \text{ s.t. } \delta^e = A^{e*}B^{e*}, \|\delta^e\|_p \leq \epsilon, \tag{8}$$

where

$$A^{e*},\ B^{e*} = \underset{A^e, B^e}{\text{argmax}} \mathbb{E}_{(x,y)\sim\mathcal{P}_e(x,y)}[\mathcal{L}(f(x+A^eB^e), y)],\ A^e \in \mathcal{R}^{N\times l\times C},\ B^e \in \mathcal{R}^{l\times N\times C}. \tag{9}$$

We provide the detailed training procedure of LDAT in Appendix D due to the space limitation of the main text. In comparison to MAT, LDAT reduces the number of perturbations available to the samples in a domain from k to 1. In addition, it reduces the rank of the perturbation for a single channel of a sample from N to l.

4.3 Theoretical Analysis

In this part, we theoretically explain why the domain-wise perturbation proposed in MAT and LDAT can help to improve the robustness of the model against spurious correlations following [2] and [19]. In general, we prove that MAT and LDAT can prevent the model from relying more on spurious features to make predictions as the spurious correlations in the training data increase. Consequently, the model trained with MAT or LDAT will generalize better on OOD data.

Notations. Let $x \in \mathcal{X}$ denote the random data in the input space $\mathcal{X}$ and let $y \in \mathcal{Y}$ denote the target random data in label space $\mathcal{Y}$. For simplicity, let $\mathcal{Y} \in \{1, -1\}$ in this section. Let $\mathbb{D}$ denote an underlying class of distributions over $\mathcal{X} \times \mathcal{Y}$. Let x_{inv} and x_{sp} denote the invariant features and the spurious features respectively. Also for simplicity, assume that there exists an identity mapping $\Phi : \mathcal{X}_{inv} \times \mathcal{X}_{sp} \to \mathcal{X}$ such that each $\mathcal{D} \in \mathbb{D}$ is induced by a distribution over $\mathcal{X}_{inv} \times \mathcal{X}_{sp}$ (so x can be denoted as $x = (x_{inv}, x_{sp})$). Let x_{sp} take values in $\{+\beta, -\beta\}$ for some $\beta > 0$.

A Simple OOD Task. Consider a simple OOD task where we have two training domains representing the grass and desert backgrounds respectively. Both domains have two classes: the cow class and the camel class. In the grass/desert domain, the cow/camel class predominates. During test time, the correlation between the labels and the background flips. We can abstract this cow-camel dataset into the following model: a training dataset $\mathcal{S}$ with four groups of data points drawn from the four quadrants of the feature space $\{-1, +1\} \times \{-\beta, +\beta\}$ respectively (shown in Figure 3). We set the invariant features $x_{inv} = y$ and the spurious features x_{sp} to be $y\beta$ with probability $p \in [0.5, 1)$ and

$-y\beta$ with probability $1-p$. Note that p measures the intensity of spurious correlations in a certain environment. When $p = 0.5$, there are no correlations between the labels and the spurious features.

Consider a linear classifier $h(x) = w_{inv}x_{inv} + w_{sp}x_{sp}$. Following [19], let us consider MAT/LDAT trained with gradient descent algorithm stopped in finite time t. In order to characterize the dependence of the model on spurious features during the training process, we investigate the convergence rate of $\frac{w_{sp}(t)\beta}{|w_{inv}(t)x_{inv}|}$ to 0 on the above dataset, which denotes the ratio between the output of the spurious component to that of the invariant component. We prove that after adding the domain-wise perturbations in finite-time-stopped gradient descent, the lower bound of the convergence rate of this ratio does not increase monotonically with p. Hence, the model will not learn a large prediction weight based on spurious features even if the spurious correlation is strong (p is large).

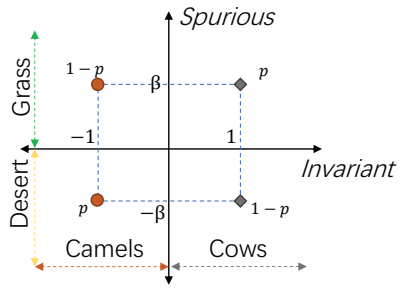

Figure 3: Illustration of the simple OOD task.

In the following theorem, we denote the domain-wise perturbation in MAT/LDAT as δ. Theorem 4.1 applies to both MAT and LDAT since they both use domain-wise perturbations. See Appendix A for a formal statement and full proof of Theorem 4.1.

Theorem 4.1. ***(informal)*** *Let $\mathcal{H}$ be the set of linear classifiers $h(x) = w_{inv}(t)x_{inv} + w_{sp}(t)x_{sp}$. Consider the above 2-D OOD dataset $\mathcal{S}$. Assume that the empirical distribution of x_{inv} given $x_{sp} \cdot y > 0$ is identical to the empirical distribution of x_{inv} given $x_{sp} \cdot y < 0$. δ is the optimal perturbation obtained by optimizing the object in Eq. (7) or Eq. (9). Let $w_{inv}(t)x_{inv} + w_{sp}(t)x_{sp}$ be initialized to the origin, and trained with MAT/LDAT to minimize the exponential loss on $\mathcal{S}$. Then, for any $(x, y) \in \mathcal{S}$, we have:*

$$\Omega(\mathbb{E}_{(x_{inv},y)\sim\mathcal{D}_{inv}}[\frac{\frac{1}{\beta+\delta y}\ln[\frac{c_1+p}{c_2+p^{\frac{1}{2}-\epsilon}(1-p)^{\frac{1}{2}+\epsilon}}]}{M\ln(t+1)}]) \leq \frac{w_{sp}(t)\beta}{|w_{inv}(t)x_{inv}|}, \tag{10}$$

where $\epsilon := \frac{\delta y}{2\beta}$ is a real number close to 0, $c_1 := \frac{2(2M(1+\delta)-1)}{(\beta+\delta y)^2}$, $c_2 := \frac{2(2M(1+\delta)-1)}{(\delta y+\beta)^{\frac{3}{2}-\epsilon}(\beta-\delta y)^{\frac{1}{2}+\epsilon}}$. $M = \max_{x\in S} \hat{w} \cdot x$ denotes the maximum value of the margin of the max-margin classifier $\hat{w}$ on $\mathcal{S}$. $\Omega(\cdot)$ is the lower bound of a given function within a constant factor. Therefore, the lower bound of the convergence rate does not increase monotonically with p under the condition that $2\epsilon c_1 + c_2 + \frac{3}{4} + \frac{3}{2}\epsilon < 0$.

To sum up, since we can prevent this lower bound from growing monotonically with p, we accelerate the convergence rate of $\frac{w_{sp}(t)\beta}{|w_{inv}(t)x_{inv}|}$ to 0 when there is stronger spurious correlation (larger p). Recall that the ratio $\frac{w_{sp}(t)\beta}{|w_{inv}(t)x_{inv}|}$ reflects the degree of reliance on spurious features. Therefore, faster convergence of this ratio to 0 (smaller lower bound) means that the model will end up relying less on the spurious correlations within a finite training time. In other words, the OOD robustness can be enhanced by using domain-wise perturbations.

Remark. Here, we demonstrate that MAT and LDAT show stronger OOD robustness compared to ERM. We compare the result in Theorem 4.1 to that in Theorem 2 of [2]. The full statement of the Theorem 2 in [2] is in Appendix B. According to the Theorem 2 in [2], even if the max-margin classifier does not rely on x_{sp} for any level of spurious correlation $p \in [0.5, 1)$, ERM trained by gradient descent stopped in finite time still fails to avoid using spurious features. Moreover, when conducting ERM with finite-time-stopped gradient descent, the lower bound of the convergence rate of $\frac{w_{sp}(t)\beta}{|w_{inv}(t)x_{inv}|}$ to 0 is

$$\Omega(\frac{\ln\frac{c+p}{c+\sqrt{p(1-p)}}}{M\ln t}) \leq \frac{w_{sp}(t)\beta}{|w_{inv}(t)x_{inv}|}, \tag{11}$$

where $c := \frac{2(2M-1)}{\beta^2}$, and M follows the definition in Theorem 4.1. This lower bound grows monotonically with p, thus ERM will have slower convergence for larger spurious correlations.

However, with domain-wise perturbations, we can modify the lower bound so that it does not increase monotonically with the spurious correlation p. Thus, we can draw the conclusion that using a perturbation for each domain is helpful to reduce dependence on spurious features compared to ERM.

5 Experiments

5.1 Experimental Setup

We conduct experiments on the DomainBed benchmark [9], a testbed for OOD generalization that implements consistent experimental protocols across various approaches to ensure fair comparisons. We evaluate on PACS [14], OfficeHome [15], VLCS [16], NICO [17], and Colored MNIST [1]. There are several changes in our experimentation setting comparing to DomainBed:

1. **Backbone Network.** We use ResNet-18 [20] as our backbone for datasets excluding Colored MNIST instead of ResNet-50 used in [9] for efficiency.
2. **Hyperparameter Search Space.** We use a smaller hyperparameter search space than [9]. We conduct a random search of 8 trials for PACS, OfficeHome, and VLCS while 6 trials for NICO and Colored MNIST in the hyperparameter search space, instead of 20 trials adopted in [9] for feasibility. See Appendix C.2 for more details.

Model Selection Strategy. Since hyperparameter choice has a significant impact on the OOD performance, it is critical to use appropriate model selection method. For PACS, OfficeHome, and VLCS datasets, we use training-domain validation proposed in [9] since it is more in line with the OOD scenario. For NICO, we adopt OOD validation following [21]. For Colored MNIST, we use test-domain validation [9] since it can enlarge the gaps in OOD performance among the algorithms while the gap induced by training-domain validation on Colored MNIST is marginal.

Hyperparameters for MAT and LDAT. To retain low-rank structures in perturbations, we set the upper bound of the search space of the perturbation number k in MAT to be 20. Similarly, the upper bound of the rank of the perturbation used in LDAT l is 20. Specifically, the search space of k and l is $\{5, 10, 15, 20\}$ (except on CMNIST, where $k \in [5, 20]$ and $l \in [10, 20]$). The complete setup of the hyperparameters for MAT and LDAT is provided in Appendix C.2.

5.2 OOD Performance on Benchmark datasets.

Table 2 summarizes the results on the five OOD datasets. The results of other approaches for PACS, OfficeHome, NICO, and Colored MNIST datasets are adopted from [21]. The results on VLCS of other algorithms are missing (denoted as "-") because [21] does not experiment on this dataset.

Comparison with ERM and Sample-wise AT. From Table 2, we observe that both MAT and LDAT outperform ERM (on both our runs and the results in [21]) and AT on average. In particular, MAT achieves consistently better results than ERM on all five datasets. Additionally, the average performance of AT is worse than ERM, which is consistent with our observations in Section 3.2.

Comparison with Existing State-of-the-Art Approaches. Although the results from [21] use a different training protocol from ours: they use a larger search space and 20 random search for the hyperparameter combinations, the comparison between ERM ([21]) and ERM (our runs) indicates that their corresponding performances are close. A similar comparison has been made in [22]. We find that MAT outperforms all previous algorithms and LDAT ranked fourth among all methods, merely after VREx [23] and IRM [1] (see avg^1 in Table 2). And even when excluding Colored MNIST (toy example), our methods still outperform ERM by $0.4 \sim 1.4\%$, whereas other methods show no improvement over ERM (see avg^3). From these results, we can see that the promotion of our proposed methods is higher than the other works, and our methods clearly outperform ERM. We also extend our evaluation to compare with adversarial augmentation based method [5] in Appendix C.3. Single-training domain generalization experiments are shown in Appendix C.4, which shows that our methods can maintain OOD performance without the reliance on multi-source training data.

Comparison between MAT and LDAT. From Table 2 we can see that MAT outperforms LDAT on average. Since LDAT reduces the number of the perturbations used in a domain from k to 1 (shown in Figure 2), LDAT can be regarded as a low-rank version of MAT. This indicates that the oversimplified perturbations may be less effective than the ones maintaining some flexibility.

Table 2: Test accuracy (%) on OOD datasets within DomainBed benchmark using ResNet-18. Here "avg^1" denotes the average accuracy on PACS, OfficeHome, NICO, CMNIST datasets and "avg^2" denotes the average accuracy on all five datasets. "avg^3" denotes the average accuracy on the other four datasets except for CMNIST. "avg^4" denotes the average accuracy on the other three datasets except for CMNIST and VLCS. The best results are in **bold**.

	Datasets								
Algorithm	PACS	OfficeHome	VLCS	NICO	CMNIST	avg^1	avg^2	avg^3	avg^4
ERM (Our runs)	81.7 $\pm$ 0.3	62.1 $\pm$ 0.1	74.4 $\pm$ 1.0	73.2 $\pm$ 1.9	28.1 $\pm$ 1.5	61.3	63.9	72.3	72.9
AT (Our runs)	82.6 $\pm$ 0.4	62.1 $\pm$ 0.3	**76.2** $\pm$ **0.3**	69.7 $\pm$ 1.6	29.1 $\pm$ 1.5	60.9	64.3	71.5	72.7
ERM[21]	81.5 $\pm$ 0.0	63.3 $\pm$ 0.2	-	71.4 $\pm$ 1.3	29.9 $\pm$ 0.1	61.5	-	72.1	-
RSC[24]	**82.8** $\pm$ 0.4	62.9 $\pm$ 0.4	-	69.7 $\pm$ 0.3	28.6 $\pm$ 1.5	61.0	-	71.8	-
MMD[25]	81.7 $\pm$ 0.2	63.8 $\pm$ 0.1	-	68.3 $\pm$ 1.8	50.7 $\pm$ 0.1	66.1	-	71.3	-
SagNet[26]	81.6 $\pm$ 0.4	62.7 $\pm$ 0.4	-	69.3 $\pm$ 1.0	30.5 $\pm$ 0.7	61.0	-	71.2	-
CORAL[27]	81.6 $\pm$ 0.6	63.8 $\pm$ 0.3	-	68.3 $\pm$ 1.4	30.0 $\pm$ 0.5	61.0	-	71.2	-
IRM[1]	81.1 $\pm$ 0.3	63.0 $\pm$ 0.2	-	67.6 $\pm$ 1.4	60.2 $\pm$ 2.4	68.0	-	70.6	-
VREx[23]	81.8 $\pm$ 0.1	63.5 $\pm$ 0.1	-	71.0 $\pm$ 1.3	56.3 $\pm$ 1.9	68.2	-	72.1	-
GroupDRO[28]	80.4 $\pm$ 0.3	63.2 $\pm$ 0.2	-	71.8 $\pm$ 0.8	32.5 $\pm$ 0.2	62.0	-	71.8	-
DANN[29]	81.1 $\pm$ 0.4	62.9 $\pm$ 0.6	-	68.6 $\pm$ 1.1	24.5 $\pm$ 0.8	59.3	-	70.9	-
MTL[30]	81.2 $\pm$ 0.4	62.9 $\pm$ 0.2	-	70.2 $\pm$ 0.6	29.3 $\pm$ 0.1	60.9	-	71.4	-
Mixup[31]	79.8 $\pm$ 0.6	63.3 $\pm$ 0.5	-	66.6 $\pm$ 0.9	27.6 $\pm$ 1.8	59.3	-	69.9	-
ANDMask[32]	79.5 $\pm$ 0.0	62.0 $\pm$ 0.3	-	72.2 $\pm$ 1.2	27.2 $\pm$ 1.4	60.2	-	71.2	-
MLDG[33]	73.0 $\pm$ 0.4	52.4 $\pm$ 0.2	-	51.6 $\pm$ 6.1	32.7 $\pm$ 1.1	52.4	-	59.0	-
MAT (Our work)	82.3 $\pm$ 0.5	**64.5** $\pm$ **2.1**	74.6 $\pm$ 0.8	74.2 $\pm$ 1.5	**65.4** $\pm$ **8.1**	**71.6**	**72.2**	**73.7**	**73.9**
LDAT (Our work)	82.6 $\pm$ 0.5	61.0 $\pm$ 0.9	75.3 $\pm$ 0.3	**74.4** $\pm$ **1.6**	52.5 $\pm$ 5.4	67.6	69.1	72.7	73.3

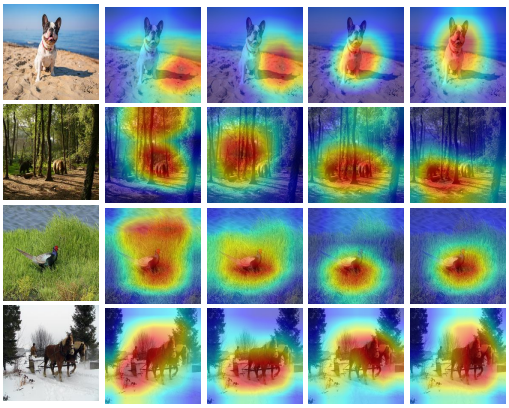

Figure 4: The pixel attention heatmap of ERM, AT, MAT and LDAT on NICO dataset. The redder part indicates that the model relies more on this part to make predictions.

Since both MAT and LDAT outperform most existing state-of-the-art methods and they both exploit low-rank structures, these two methods mutually corroborate the effectiveness of low-rank structure for OOD generalization. The respective advantages of the two methods are as follows:

- MAT is a complex and high-rank version of LDAT, which has a stronger ability to describe more complex spurious background information. As shown by the attention heatmap in Figure 4, MAT can better capture the object than LDAT when faced with a more complex background (the example in the second row of Figure 4).

- LDAT costs less memory than MAT during the training process, although there is no significant difference in training time between the two methods. When the memory is limited, LDAT is preferred.

5.3 Empirical Understanding

Visualization. To empirically show that MAT and LDAT can reduce the reliance on spurious features, we visualize the pixel attention heatmap of ERM, AT, MAT, and LDAT on NICO dataset using GradCam [34]. It reflects the contribution of different components of the feature map to the prediction results. We pick the model with the best performance for each method. The results in Figure 4 indicate that the model trained by MAT and LDAT focuses more on the object itself, while ERM and AT adopt the background information that spuriously correlates to the class to make predictions.

Parameter Analysis. The number of the perturbations k used in a domain in MAT and the rank of the perturbation l in LDAT are two key hyperparameters. We conduct further experiments to analyze the impact on the performances of k and l. We adopt a fixed set of parameters except for k and l and evaluate on PACS dataset. The results in Figure 5 show that MAT and LDAT are able to keep their performances over ERM as long as k and l are far less than the number of the samples N. Additionally, we can observe from the trend that when k and l are too small ($= 5$), the performances degenerate. This implies the oversimplified structures of the perturbations can be less effective for generalization. When k and l take larger value (about 1000), the performance will drop (see Table 3). For the selection of optimal parameters (range), we observe that the parameters that are good on one dataset also work well on others (see Table 6 in the appendix), so in practice we adopt the strategy of searching for the optimal parameters roughly on one dataset and then applying them to other datasets. Additional analysis on the impact of the learning rate for the perturbations is in Appendix C.2.

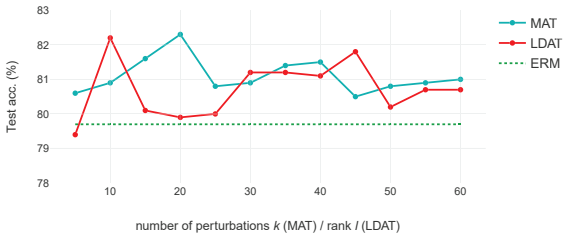

Figure 5: The performance at different values of the number of perturbations to combined within one domain k (in MAT) and the rank of the perturbation in a domain l (in LDAT).

Table 3: The test accuracy (%) on CMNIST with different k and l.

Algorithm	$k \in [5, 20], l \in [10, 20]$	k or $l = 200$	k or $l = 500$	k or $l = 1000$
MAT	65.4 ± 8.1	34.9 ± 20.2	25.6 ± 8.5	23.4 ± 10.8
LDAT	52.5 ± 5.4	24.9 ± 8.9	19.0 ± 6.6	10.3 ± 0.1

6 Conclusion

In this work, we empirically revealed the limitations of sample-wise AT on OOD tasks. Due to the lack of constraints on the perturbation and the utilization of domain features, sample-wise AT fails to generalize well when facing large-scale perturbations which is close to the real-world OOD scenarios. We further proposed two AT variants with structured priors, named MAT and LDAT, which add low-rank perturbations to improve model's robustness against the distribution shift of spurious correlations. We theoretically proved the domain-wise perturbations used in MAT and LDAT can benefit OOD generalization, and validated the effectiveness of the proposed methods on OOD tasks through a series of experiments on Domainbed benchmark.

Acknowledgment

Yisen Wang is partially supported by the NSF China (No. 62006153), Project 2020BD006 supported by PKU-Baidu Fund, Open Research Projects of Zhejiang Lab (No. 2022RC0AB05), and Huawei Technologies Inc. Qixun Wang is partially supported by the State Key Development Program Grand No. 2020YFB1708002.

References

[1] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. In *ICML*, 2020.

[2] Vaishnavh Nagarajan, Anders Andreassen, and Behnam Neyshabur. Understanding the failure modes of out-of-distribution generalization. *arXiv preprint arXiv:2010.15775*, 2020.

[3] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.

[4] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.

[5] Riccardo Volpi, Hongseok Namkoong, Ozan Sener, John Duchi, Vittorio Murino, and Silvio Savarese. Generalizing to unseen domains via adversarial data augmentation. *arXiv preprint arXiv:1805.12018*, 2018.

[6] Mingyang Yi, Lu Hou, Jiacheng Sun, Lifeng Shang, Xin Jiang, Qun Liu, and Zhiming Ma. Improved ood generalization via adversarial training and pretraing. In *ICML*, 2021.

[7] Ludwig Schmidt, Shibani Santurkar, Dimitris Tsipras, Kunal Talwar, and Aleksander Madry. Adversarially robust generalization requires more data. In *NeurIPS*, 2018.

[8] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Omar Fawzi, and Pascal Frossard. Universal adversarial perturbations. In *CVPR*, 2017.

[9] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. *arXiv preprint arXiv:2007.01434*, 2020.

[10] Yisen Wang, Xingjun Ma, James Bailey, Jinfeng Yi, Bowen Zhou, and Quanquan Gu. On the convergence and robustness of adversarial training. In *ICML*, 2019.

[11] Dongxian Wu, Shu-Tao Xia, and Yisen Wang. Adversarial weight perturbation helps robust generalization. In *NeurIPS*, 2020.

[12] Charles Herrmann, Kyle Sargent, Lu Jiang, Ramin Zabih, Huiwen Chang, Ce Liu, Dilip Krishnan, and Deqing Sun. Pyramid adversarial training improves vit performance. *arXiv preprint arXiv:2111.15121*, 2021.

[13] I. J. Goodfellow, J. Shlens, and C. Szegedy. Explaining and harnessing adversarial examples. *Computer Science*, 2014.

[14] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generalization. In *ICCV*, 2017.

[15] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *CVPR*, 2017.

[16] Chen Fang, Ye Xu, and Daniel N Rockmore. Unbiased metric learning: On the utilization of multiple datasets and web images for softening bias. In *ICCV*, 2013.

[17] Yue He, Zheyan Shen, and Peng Cui. Towards non-iid image classification: A dataset and baselines. *Pattern Recognition*, 110:107383, 2021.

[18] Klim Kireev, Maksym Andriushchenko, and Nicolas Flammarion. On the effectiveness of adversarial training against common corruptions. *arXiv preprint arXiv:2103.02325*, 2021.

[19] Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The implicit bias of gradient descent on separable data. *The Journal of Machine Learning Research*, 2018.

[20] Shaoqing Ren, Jian Sun, K He, and X Zhang. Deep residual learning for image recognition. In *CVPR*, 2016.

[21] Nanyang Ye, Kaican Li, Lanqing Hong, Haoyue Bai, Yiting Chen, Fengwei Zhou, and Zhenguo Li. Ood-bench: Benchmarking and understanding out-of-distribution generalization datasets and algorithms. *arXiv preprint arXiv:2106.03721*, 2021.

[22] 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.

[23] 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). In *ICML*, 2021.

[24] Zeyi Huang, Haohan Wang, Eric P Xing, and Dong Huang. Self-challenging improves cross-domain generalization. In *ECCV*, 2020.

[25] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *CVPR*, 2018.

[26] Hyeonseob Nam, HyunJae Lee, Jongchan Park, Wonjun Yoon, and Donggeun Yoo. Reducing domain gap via style-agnostic networks. *arXiv preprint arXiv:1910.11645*, 2(7):8, 2019.

[27] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *ECCV*, 2016.

[28] Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. Distributionally robust neural networks. In *ICLR*, 2019.

[29] 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.

[30] Gilles Blanchard, Aniket Anand Deshmukh, Urun Dogan, Gyemin Lee, and Clayton Scott. Domain generalization by marginal transfer learning. *arXiv preprint arXiv:1711.07910*, 2017.

[31] Shen Yan, Huan Song, Nanxiang Li, Lincan Zou, and Liu Ren. Improve unsupervised domain adaptation with mixup training. *arXiv preprint arXiv:2001.00677*, 2020.

[32] Giambattista Parascandolo, Alexander Neitz, Antonio Orvieto, Luigi Gresele, and Bernhard Schölkopf. Learning explanations that are hard to vary. *arXiv preprint arXiv:2009.00329*, 2020.

[33] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Learning to generalize: Meta-learning for domain generalization. In *AAAI*, 2018.

[34] Jacob Gildenblat and contributors. Pytorch library for cam methods. `https://github.com/jacobgil/pytorch-grad-cam`, 2021.

[35] 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.

[36] Chris Xing Tian, Haoliang Li, Xiaofei Xie, Yang Liu, and Shiqi Wang. Neuron coverage-guided domain generalization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to [Yes] , [No] , or [N/A] . You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? [Yes] See Section .
- Did you include the license to the code and datasets? [No] The code and the data are proprietary.
- Did you include the license to the code and datasets? [N/A]

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] See Section 1.
   (b) Did you describe the limitations of your work? [Yes] See Section 5.2. We mention that the oversimplified perturbation (LDAT) will be less effective than the ones with more flexibility (MAT).
   (c) Did you discuss any potential negative societal impacts of your work? [No]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] Full assumptions are in Appendix A.
   (b) Did you include complete proofs of all theoretical results? [Yes] See Appendix A.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No] It will release upon acceptance.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section C.1 and C.2.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Table 2 for an example.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] We use DomainBed benchmark and GradCam and cite their creators.
   (b) Did you mention the license of the assets? [No] All assets we use are open source.
   (c) Did you include any new assets either in the supplemental material or as a URL? [No]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [No] All datasets we use are open source.
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [No]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [No]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [No]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [No]

A Proof of Theorem 4.1

Theorem 4.1 **(formal)** Let $\mathcal{H}$ be the set of linear classifiers $h(x) = w_{inv}(t)x_{inv} + w_{sp}(t)x_{sp}$. Consider any task that satisfies all the constraints in Section 3.1. in [2]. Consider a dataset $\mathcal{S}$ drawn from $\mathcal{D}$ such that the empirical distribution of x_{inv} given $x_{sp} \cdot y > 0$ (denoted as $(x_{inv}, y) \sim \mathcal{D}_{inv}$) is identical to the empirical distribution of x_{inv} given $x_{sp} \cdot y < 0$. δ is the optimal perturbation obtained by optimizing object (7) or (9). δ can be seen as a random variable.

Let $w_{inv}(t)x_{inv} + w_{sp}(t)x_{sp}$ be initialized to the origin, and trained with MAT/LDAT with an infinitesimal learning rate to minimize the exponential loss on $\mathcal{S}$. Then, for any $(x, y) \in \mathcal{S}$, we have:

$$\Omega(\mathbb{E}_{(x_{inv},y)\sim\mathcal{D}_{inv}}[\frac{\frac{1}{\beta+\delta y}\ln[\frac{c_1+p}{c_2+p^{\frac{1}{2}-\epsilon}(1-p)^{\frac{1}{2}+\epsilon}}]}{M\ln(t+1)}]) \leq \frac{w_{sp}(t)\beta}{|w_{inv}(t)x_{inv}|}, \tag{12}$$

where $\epsilon := \frac{\delta y}{2\beta}$ is a real number close to 0, $c_1 := \frac{2(2M(1+\delta)-1)}{(\beta+\delta y)^2}$, $c_2 := \frac{2(2M(1+\delta)-1)}{(\delta y+\beta)^{\frac{3}{2}-\epsilon}(\beta-\delta y)^{\frac{1}{2}+\epsilon}}$. $M = \underset{x\in S}{max}\ \hat{w} \cdot x$ denotes the maximum value of the margin of the max-margin classifier $\hat{w}$ on $\mathcal{S}$. $\Omega(\cdot)$ is the lower bound of a given function within a constant factor. Therefore, the lower bound of the convergence rate does not increase monotonically with p under the condition that $2\epsilon c_1 + c_2 + \frac{3}{4} + \frac{3}{2}\epsilon < 0$.

Proof. For brevity, we use w_c and w_s to represent w_{inv} and w_{sp} respectively. Also, we use x_c to represent x_{inv}. We use x_e to represent x_{sp} since the spurious feature is correlated to the environment E_e. Let $\mathcal{S}_{min}$ and $\mathcal{S}_{maj}$ denote the subset of datapoints in $\mathcal{S}$ where $x_e \cdot y < 0$ and $x_e \cdot y > 0$ respectively. Let $\mathcal{D}_{inv}$ denote the distribution over (x_c, y) induced by drawing (x, y) uniformly from $\mathcal{S}_{min}$. The corresponding marginal distribution of y is denoted as $\mathcal{D}_y$. By the assumption of the theorem, this distribution would be the same if x was drawn uniformly from $\mathcal{S}_{maj}$. Then, the loss function that is being minimized in this setting corresponds to:

$$\begin{aligned}
\mathcal{L}[f(x), y] &= \mathbb{E}_{(x_c,y)\sim\mathcal{D}_{inv}}[e^{-f(x)y}] \\
&= \mathbb{E}_{(x_c,y)\sim\mathcal{D}_{inv}}[e^{-[w_c x_c + w_s x_e + (w_c+w_s)\delta]y}] \\
&= \mathbb{E}_{(x_c,y)\sim\mathcal{D}_{inv}}[pe^{-[w_c x_c + w_s\beta y + (w_c+w_s)\delta]y} + (1-p)e^{-[w_c x_c - w_s\beta y + (w_c+w_s)\delta]y}] \\
&= \mathbb{E}_{(x_c,y)\sim\mathcal{D}_{inv}}[e^{-(x_c+\delta)w_c y}[pe^{-(\delta y+\beta)w_s} + (1-p)e^{(\beta-\delta y)w_s}]]
\end{aligned}$$

The update on w_s can be written as:

$$\begin{aligned}
\Delta w_s &= -\frac{\partial\mathcal{L}[f(x), y]}{\partial w_s} \\
&= -\mathbb{E}_{(x_c,y)\sim\mathcal{D}_{inv}} e^{-(x_c+\delta)w_c y}[-pe^{-(\delta y+\beta)w_s}(\delta y+\beta) + (1-p)e^{(\beta-\delta y)w_s}(\beta-\delta y)]
\end{aligned}$$

Proof of bounds on $w_c(t)x_c$. Using the result of [19] and [2], we get

$$|w_c(t)x_c| \in [0.5\ln(1+t), 2M\ln(1+t)]$$

for a sufficiently large t and for all $x \in \mathcal{S}$.

Proof of the upper bound on w_s. To calculate the lower bound of w_s, we prove the upper bound as auxiliary first. Note that Δw_s decreases monotonically with w_s. Assume that $\beta > |\delta|$ (this is reasonable since the perturbation radius is usually smaller than the scale of the spurious feature). Let $\Delta w_s = 0$, we get

$$w_s = \frac{1}{2\beta}\ln[\frac{p}{1-p}(\frac{\beta+\delta y}{\beta-\delta y})] =: w_0.$$

Since Δw_s decrease monotonically with w_s, (which can be inferred from $\frac{\partial\Delta w_s}{\partial w_s} = -e^{-(x_c+\delta)w_c y}[-pe^{-(\delta y+\beta)w_s}(\delta y+\beta)^2 + (1-p)e^{(\beta-\delta y)w_s}(\beta-\delta y)^2] \leq 0$), when $w_s < w_0$, $\Delta w_s > 0$ and when $w_s > w_0$, $\Delta w_s < 0$. As a result, for any system that is initialized at 0, w_s can never cross the point w_0. Thus, we get the upper bound of $w_s(t)$:

$$w_s(t) < w_0 = \frac{1}{2\beta}\ln[\frac{p}{1-p}(\frac{\beta+\delta y}{\beta-\delta y})].$$

Proof of the lower bound on w_s. We lower bound w_s via the upper bound on w_s as:

$$\Delta w_s > \mathbb{E}_{(x_c,y)\sim\mathcal{D}_{inv}} e^{-(x_c+\delta)w_c y}[p(\delta y+\beta)e^{-(\delta y+\beta)w_s} - (1-p)(\beta-\delta y)[\frac{p(\beta+\delta y)}{(1-p)(\beta-\delta y)}]^{\frac{\beta-\delta y}{2\beta}}]$$
$$= \mathbb{E}_{(x_c,y)\sim\mathcal{D}_{inv}} e^{-(x_c+\delta)w_c y}[p(\delta y+\beta)e^{-(\delta y+\beta)w_s} - [p(\beta+\delta y)]^{\frac{1}{2}-\frac{\delta y}{2\beta}}[(1-p)(\beta-\delta y)]^{\frac{1}{2}+\frac{\delta y}{2\beta}}].$$

Next, using the upper bound on $|w_c(t)x_c|$, we get:

$$\begin{aligned}\Delta w_s > {} & \mathbb{E}_{(x_c,y)\sim\mathcal{D}_{inv}} e^{-2M\ln(1+t)} e^{-2M\delta\ln(1+t)}[p(\delta y+\beta)e^{-(\delta y+\beta)w_s} \\ & - [p(\beta+\delta y)]^{\frac{1}{2}-\frac{\delta y}{2\beta}}[(1-p)(\beta-\delta y)]^{\frac{1}{2}+\frac{\delta y}{2\beta}}] \\ = {} & \mathbb{E}_{(x_c,y)\sim\mathcal{D}_{inv}} \frac{1}{(t+1)^{2M(1+\delta)}}[p(\delta y+\beta)e^{-(\delta y+\beta)w_s} \\ & - [p(\beta+\delta y)]^{\frac{1}{2}-\frac{\delta y}{2\beta}}[(1-p)(\beta-\delta y)]^{\frac{1}{2}+\frac{\delta y}{2\beta}}].\end{aligned}$$

For brevity, we denote the term $p(\delta y+\beta)e^{-(\delta y+\beta)w_s}$ as T and denote the term $[p(\beta+\delta y)]^{\frac{1}{2}-\frac{\delta y}{2\beta}}[(1-p)(\beta-\delta y)]^{\frac{1}{2}+\frac{\delta y}{2\beta}}$ as L. And in the following proof, we omit the expectation marker $\mathbb{E}$ also for simplicity. It is clear that both T and $L > 0$. Here, $T - L > 0$ since $w_s < w_0$. Note that $\Delta w_s = -\frac{\partial\mathcal{L}[f(x),y]}{\partial w_s} = \frac{\partial w_s}{\partial t}$, rearranging this and integrating, we get:

$$\int_0^{w_s} \frac{1}{p(\beta+\delta y)e^{-(\beta+\delta y)w_s} - L} dw_s > \int_0^t \frac{1}{(t+1)^{2M(1+\delta)}} dt,$$
$$\frac{\ln[p(\beta+\delta y)-L] - \ln[p(\beta+\delta y) - e^{(\beta+\delta y)w_s}L]}{(\beta+\delta y)L} > \frac{1}{2M(1+\delta)-1}[1 - \frac{1}{(1+t)^{2M(1+\delta)-1}}].$$

Since for a sufficiently large t, $1 - \frac{1}{(1+t)^{2M(1+\delta)-1}} > \frac{1}{2}$, we have:

$$\ln[\frac{p(\beta+\delta y)-L}{p(\beta+\delta y) - e^{(\beta+\delta y)w_s}L}] > \frac{(\beta+\delta y)L}{2(2M(1+\delta)-1)},$$

we can further lower bound the right hand side by applying the inequality $x \geq \ln(x+1)$ for positive x:

$$\ln[\frac{p(\beta+\delta y)-L}{p(\beta+\delta y) - e^{(\beta+\delta y)w_s}L}] > \ln[\frac{(\beta+\delta y)L}{2(2M(1+\delta)-1)} + 1].$$

Thus,

$$\frac{p(\beta+\delta y)-L}{p(\beta+\delta y) - e^{(\beta+\delta y)w_s}L} > 1 + \frac{(\beta+\delta y)L}{2(2M(1+\delta)-1)}.$$

Note that the denominator on the left side of the inequality is greater than 0 since $T - L > 0$. Rearrange this inequality:

$$e^{w_s(\beta+\delta y)} > \frac{1 + \frac{(\beta+\delta y)^2 p}{2(2M(1+\delta)-1)}}{1 + \frac{(\beta+\delta y)L}{2(2M(1+\delta)-1)}}.$$

Putting $L = [p(\beta+\delta y)]^{\frac{1}{2}-\frac{\delta y}{2\beta}}[(1-p)(\beta-\delta y)]^{\frac{1}{2}+\frac{\delta y}{2\beta}}$ back into the inequality,

$$\begin{aligned} e^{w_s(\beta+\delta y)} &> \frac{1 + \frac{(\beta+\delta y)^2 p}{2(2M(1+\delta)-1)}}{1 + \frac{(\beta+\delta y)[p(\beta+\delta y)]^{\frac{1}{2}-\frac{\delta y}{2\beta}}[(1-p)(\beta-\delta y)]^{\frac{1}{2}+\frac{\delta y}{2\beta}}}{2(2M(1+\delta)-1)}} \\ &= \frac{p + \frac{2(2M(1+\delta)-1)}{(\beta+\delta y)^2}}{p^{\frac{1}{2}-\frac{\delta y}{2\beta}}(1-p)^{\frac{1}{2}+\frac{\delta y}{2\beta}} + \frac{2(2M(1+\delta)-1)}{(\delta y+\beta)^{\frac{3}{2}-\frac{\delta y}{2\beta}}(\beta-\delta y)^{\frac{1}{2}+\frac{\delta y}{2\beta}}}}. \end{aligned}$$

Let $c_1 := \frac{2(2M(1+\delta)-1)}{(\beta+\delta y)^2}$, $c_2 := \frac{2(2M(1+\delta)-1)}{(\delta y+\beta)^{\frac{3}{2}-\epsilon}(\beta-\delta y)^{\frac{1}{2}+\epsilon}}$, $\epsilon := \frac{\delta y}{2\beta}$.

Put the expectation mark back into this inequality, finally, we get the lower bound on w_s:

$$w_s \geq \mathbb{E}_{(x_c,y)\sim\mathcal{D}_{inv}}[\frac{1}{\beta+\delta y}\ln\left[\frac{c_1+p}{c_2 + p^{\frac{1}{2}-\epsilon}(1-p)^{\frac{1}{2}+\epsilon}}\right]]$$

To show that the lower bound on the dependency on spurious correlations induced by MAT and LDAT does not increase monotonically with p under some conditions, we take the derivative of the obtained lower bound $g(p)$ in Theorem 4.1 with respect to p:

$$\frac{\partial g(p)}{\partial p} := \frac{\partial \left(\frac{c_1+p}{c_2+p^{\frac{1}{2}-\epsilon}(1-p)^{\frac{1}{2}+\epsilon}} \right)}{\partial p}$$

Since the denominator of $\frac{\partial g(p)}{\partial p}$ is positive, we pick out the numerator: $c_2 + (\frac{1}{2}+\epsilon)p^{\frac{1}{2}-\epsilon}(1-p)^{\frac{1}{2}+\epsilon} - c_1(\frac{1}{2}-\epsilon)p^{-\frac{1}{2}-\epsilon}(1-p)^{\frac{1}{2}+\epsilon} + c_1(\frac{1}{2}+\epsilon)p^{\frac{1}{2}-\epsilon}(1-p)^{-\frac{1}{2}+\epsilon} + (\frac{1}{2}+\epsilon)p^{\frac{3}{2}-\epsilon}(1-p)^{-\frac{1}{2}+\epsilon}$. In order to study the positive and negative change of the numerator, we continue to derive it with respect to p and obtain

$$p^{-\frac{3}{2}-\epsilon}(1-p)^{-\frac{3}{2}+\epsilon}[(\frac{1}{4}-\epsilon^2)(p+c_1)].$$

Since we assume $\beta > |\delta|$ and $c_1 > 0$, $\frac{\partial^2 g(p)}{\partial p^2} > 0$ and $\frac{\partial g(p)}{\partial p}$ increase with p monotonically.The minimum of $\frac{\partial g(p)}{\partial p}$ is reached when $p = \frac{1}{2}$. This minimum equals to $2\epsilon c_1 + c_2 + \frac{3}{4} + \frac{3}{2}\epsilon$. When $2\epsilon c_1 + c_2 + \frac{3}{4} + \frac{3}{2}\epsilon < 0$, the lower bound does not increase with p monotonically when p is within a certain range ($\in (0.5, 1)$).

□

B Detailed Statement of Theorem 2 in [2]

We now introduce the Theorem 2 in [2]. Before introducing it, we first introduce the concept of the *easy-to-learn tasks* in [2], i.e. tasks with a set of constraints. The motivation of restricting ourselves to the constrained set of tasks is that it prevents us from designing complex examples where ERM is forced to rely on spurious features due to a not-so-fundamental factor. Each constraint forbids a specific failure mode of ERM in OOD scenarios. The Theorem 2 in [2] shows that even under such favorable conditions for ERM, this classical method can also be perturbed by the spurious features.

Notations. For convenience, we will give some notations here again. Consider an input space $\mathcal{X}$ and a label space $\mathcal{Y} \in \{-1, 1\}$. Let $\mathcal{D} \in \mathbb{D}$ denote a distribution over $\mathcal{X} \times \mathcal{Y}$. $p_{\mathcal{D}}$ denotes the probability density function (PDF) of $\mathcal{D}$. Let $\mathbb{H}$ denote a class of classifiers $h : \mathcal{X} \to \mathbb{R}$. Consider a dataset S drawn from $\mathcal{D}$. Let $L_{\mathcal{D}}(h) := \mathbb{E}_{(x,y)\sim\mathcal{D}}[h(x) \cdot y < 0]$ the loss of h on $\mathcal{D}$. Let $h^\star = \arg\min_{h\in\mathbb{H}} \max_{\mathcal{D}\in\mathbb{D}} L_{\mathcal{D}}(h)$ denote the optimal classifier in the worst case. With an abuse of notation, we also denote the PDF of the distribution over $\mathcal{X}_{\text{inv}} \times \mathcal{X}_{\text{sp}}$ as $p_{\mathcal{D}}(\cdot)$. Let $\mathcal{D}_{\text{train}}$ denote the distribution of the pooled training data. Assume that there exists a mapping $\Phi : \mathcal{X}_{\text{inv}} \times \mathcal{X}_{\text{sp}} \to \mathcal{X}$ such that each $\mathcal{D} \in \mathbb{D}$ is induced by a distribution over $\mathcal{X}_{\text{inv}} \times \mathcal{X}_{\text{sp}}$.

Definition B.1. Easy-to-learn tasks. Tasks that satisfy the following constraints are easy-to-learn.

1. **Fully predictive invariant features.** For all $\mathcal{D} \in \mathbb{D}$, $L_{\mathcal{D}}(h^\star) = 0$.
2. **Identical invariant distribution.** Across all $\mathcal{D} \in \mathbb{D}$, $p_{\mathcal{D}}(x_{\text{inv}})$ is identical.
3. **Conditional independence.** For all $\mathcal{D} \in \mathbb{D}$, $x_{\text{sp}} \perp x_{\text{inv}}$.
4. **Two-valued spurious features.** We set $x_{\text{sp}} = \mathbb{R}$ and the support of x_{sp} in D_{train} is $\{-\beta, +\beta\}$.
5. **Identity mapping.** Φ is the identity mapping i.e., $x = (x_{\text{inv}}, x_{\text{sp}})$.

Theorem B.2. ***(The Theorem 2 in [2])*** *Let H be the set of linear classifiers $h(x) = w_{inv} \cdot x_{inv} + w_{sp} x_{sp}$. Then, for any easy-to-learn task, continuous-time gradient descent training of $w_{inv}(t)x_{inv} + w_{sp}(t)x_{sp}$ to minimize the exponential loss, satisfies:*

$$\Omega\left(\frac{\ln \frac{c+p}{c+\sqrt{p(1-p)}}}{M \ln(t+1)}\right) \leq \frac{w_{sp}(t)\beta}{|w_{inv} \cdot x_{inv}|} \leq \mathcal{O}\left(\frac{\ln \frac{p}{1-p}}{\ln(t+1)}\right) \tag{13}$$

where $M = max_{x\in S} \hat{w}x$ where $\hat{w}$ is the max-margin classifier on S. $c := \frac{2(2M-1)}{\beta^2}$.

Table 4: Hyperparameter setting of the experiment of Table 1, Figure 1, Figure 5, Table 7 and Table 8.

Parameter	Value
learning rate r	0.00005
batch size b	64
weight decay	0.001
drop out	0.1
AT perturbation radius ϵ (excluding Figure 1)	0.1
FGSM step size γ	0.1
perturbation weight α learning rate η (MAT)	0.001
factor matrix A (B) learning rate ρ (LDAT)	0.01

C Experiment Details and Supplementary Experimental Results

C.1 Settings of the Toy Experiments

For the experiment in Table 1, Figure 1, Figure 5, Table 7, and Table 8, we use a fixed set of hyperparameters (see Table 4) instead of conducting a random search of 20 trials over the hyperparameter distribution (the setting in [9]) for efficiency. We report the average across three independent runs. For model selection method, training-domain validation [9] is used for PACS, OfficeHome, and VLCS. For NICO, an OOD validation set is adopted following [21].

C.2 Experiment Setting and Additional Results of Table 2

Overall setup. We conduct a random search of 8 trials for PACS, OfficeHome, VLCS and 6 random trials for NICO and Colored MNIST in the hyperparameter search space, instead of 20 trials adopted in [9] for feasibility. We then average the best results for each hyperparameter combination and dataset (according to each model selection criterion) across test domains (except for Colored MNIST where we test on one biased domain only). Finally, we report the average of this number across three independent runs, and its corresponding standard error. We run all datasets for 8000 epochs during the training process.

Hyperparameter Search Space. We use a smaller hyperparameter search space than that in [9]. The search space for PACS, OfficeHome, VLCS, NICO and Colored MNIST is shown in Table 5. To determine the search space of a hyperparameter for benchmark running, we first fix other parameters and conduct a grid search to determine the approximate range of the better performances. Take the learning rate η of the MAT matrices as an example, we fix $k = 20$ and try different values of η on PACS. We find that the results of $\eta = 0.01$ (82.2 $\pm$ 0.4%) and $\eta = 0.001$ (82.3 $\pm$ 0.5%) are better than that of $\eta = 0.1$ (81.6 $\pm$ 0.2%), so we adopt the random search space of $\{0.01, 0.001\}$. The same is true for the other parameters.

In practical applications, as for the choice of the optimal parameter, we find that through experiments that for MAT and LDAT, the value of the rank k (MAT) and l (LDAT) with good test accuracy (outperforms ERM) on one data set also has good one on other datasets, as shown in the Table 6. In Table 6, $k = 10$ for MAT and $l = 15$ for LDAT outperform ERM on all three datasets, as marked in **bold**. Thus, we could find an optimal set of parameters with the model selection methods and then apply them to other datasets.

Model Selection Stategy. For PACS, OfficeHome and VLCS datasets, we use training-domain validation proposed in [9]. This model selection method first randomly collect 20% of each training domain to form a validation set. Then, it chooses the hyperparameter maximizing the accuracy on the validation set. For NICO, we adopt the OOD validation proposed in [21]. This method chooses the model maximizing the accuracy on a validation set that follows neither the distribution of the training domain or the distribution of the test domain. For Colored MNIST, we use test-domain validation, i.e., using a validation set that follows the distribution of the test domain. This is because it can enlarge the gaps in OOD performance among the algorithms while the gap induced by training-domain validation on Colored MNIST is marginal.

Table 5: Hyperparameter setting of the experiment on PACS, OfficeHome, VLCS, NICO and Colored MNIST of Table 2

Dataset	Parameter	Value
PACS, OfficeHome, VLCS	learning rate r	0.00005
	batch size b	64
	weight decay (ERM, AT)	$10^{\text{Uniform}(-4,-3)}$
	weight decay (MAT, LDAT)	0.001
	drop out (ERM, AT)	RandomChoice([0,0.1,0.5])
	drop out (MAT, LDAT)	0.1
	perturbation number k (MAT)	RandomChoice([5,10,15,20])
	perturbation weight α learning rate η (MAT)	RandomChoice([0.01,0.001])
	perturbation rank l (LDAT)	RandomChoice([5,10,15,20])
	factor matrix A (B) learning rate ρ (LDAT)	RandomChoice([0.1,0.01])
NICO	learning rate r	0.00005
	batch size b	64
	weight decay	$10^{\text{Uniform}(-4,-3)}$
	drop out	RandomChoice([0,0.1,0.5])
	perturbation number k (MAT)	$\text{Uniform}(10, 20)$
	perturbation weight α learning rate η (MAT)	0.001
	perturbation rank l (LDAT)	$\text{Uniform}(10, 20)$
	factor matrix A (B) learning rate ρ (LDAT)	0.01
Colored MNIST	learning rate r	$10^{\text{Uniform}(-4.5,-3.5)}$
	batch size b	$2^{\text{Uniform}(3,9)}$
	weight decay	0
	drop out	RandomChoice([0,0.1,0.5])
	perturbation number k (MAT)	$\text{Uniform}(5, 20)$
	perturbation weight α learning rate η (MAT)	$10^{\text{Uniform}(-3,-2)}$
	perturbation rank l (LDAT)	$\text{Uniform}(10, 20)$
	factor matrix A (B) learning rate ρ (LDAT)	0.01
	AT perturbation radius ϵ (MAT, LDAT)	$10^{\text{Uniform}(-1,2)}$
	FGSM step size γ (MAT)	$10^{\text{Uniform}(-2,1)}$
	FGSM step size γ (AT)	0.1
All except Colored MNIST	AT perturbation radius ϵ	0.1
	FGSM step size γ (AT, MAT)	0.1

Table 6: The test accuracy (%) on PACS, VLCS and NICO with different k and l. The ERM baseline on these datasets is 79.7 $\pm$ 0.4, 74.2 $\pm$ 1.0, 69.7 $\pm$ 1.0 respectively.

Dataset	Algorithm	k or $l = 5$	k or $l = 10$	k or $l = 15$	k or $l = 20$	k or $l = 25$	k or $l = 30$
PACS	MAT	80.6 $\pm$ 0.8	**80.9** $\pm$ 0.2	81.6 $\pm$ 0.3	82.3 $\pm$ 0.5	80.8 $\pm$ 0.1	80.9 $\pm$ 0.4
	LDAT	79.4 $\pm$ 0.5	82.2 $\pm$ 0.6	**80.1** $\pm$ **0.4**	79.9 $\pm$ 0.4	80.0 $\pm$ 0.5	81.2 $\pm$ 0.4
VLCS	MAT	74.2 $\pm$ 0.8	**74.6** $\pm$ **0.5**	74.4 $\pm$ 0.6	74.4 $\pm$ 0.2	72.9 $\pm$ 0.3	74.4 $\pm$ 0.6
	LDAT	74.0 $\pm$ 0.3	74.4 $\pm$ 0.1	**75.3** $\pm$ **0.5**	75.0 $\pm$ 0.5	74.1 $\pm$ 0.4	74.2 $\pm$ 0.7
NICO	MAT	69.8 $\pm$ 1.3	**70.5** $\pm$ **1.2**	71.1 $\pm$ 1.3	69.5 $\pm$ 2.7	71.8 $\pm$ 1.5	69.3 $\pm$ 0.7
	LDAT	66.2 $\pm$ 1.7	67.7 $\pm$ 0.3	**70.0** $\pm$ **1.1**	67.8 $\pm$ 2.0	68.0 $\pm$ 1.3	67.2 $\pm$ 1.5

Backbone Network. We use ResNet-18 [20] pretrained on ImageNet [35] for PACS, OfficeHome and VLCS. We use unpretrained ResNet-18 for NICO since it contains images largely overlapped with ImageNet classes. As for Colored MNIST, We use a small CNN-architecture following [9].

Impact of Learning Rate. We further investigate the impact on OOD performances of the learning rate for the perturbation weights in MAT and the learning rate for the decomposed factors in LDAT. We use the experimental setting introduced in Appendix C.1. The results of MAT and LDAT are shown in Table 7 and 8 respectively. In Table 7 and 8 we observe that the learning rate for the

Table 7: The test accuracy (%) on PACS of MAT when the learning rate (η) for the perturbation weights takes different values. We set the number of perturbations $k = 20$. The other hyperparameters take value in Table 4.

		MAT		
ERM	AT	$\eta = 0.1$	$\eta = 0.01$	$\eta = 0.001$
79.7 ± 0.0	81.5 ± 0.4	81.6 ± 0.2	82.2 ± 0.4	82.3 ± 0.5

Table 8: The test accuracy (%) on PACS of LDAT when the learning rate (ρ) for the decomposed factors takes different values. We set the rank of perturbations $l = 10$. The other hyperparameters take value in Table 4.

		LDAT	
ERM	AT	ρ=0.1	$\rho = 0.01$
79.7 ± 0.0	81.5 ± 0.4	82.2 ± 0.6	82.6 ± 0.2

perturbations has a marginal effect on the OOD accuracy. Both MAT and LDAT outperform ERM and AT on PACS when using different values of learning rate.

C.3 Comparing to Existing Data Augmentation Baseline

To better verify the improvement of our proposed method on the existing adversarial augmentation methods for OOD, we reproduce the algorithm in [5]. [5] proposed a minimax iterative training procedure to generate adversarial data that follows fictitious target distributions (GUT). As discussed in Section 2, their work is restricted in the framework of using Wasserstein distance to measure the distribution shift, which is less practical for the real-world OOD setting where domain shifts are diverse. Additionally, They focus only on sample-wise operations and ignore the use of common features within a domain. The experimental results on NICO dataset is in Table 9. The unique hyperparameters of GUT follow the Settings in [5] except we set the T_{max} to be 5 instead of 15 for efficiency. We can see that both our proposed method outperform GUT.

C.4 Comparing to Existing Data Augmentation Baseline under Single-training Domain Generalization Setting

We conduct experiments to further verify the effectiveness of MAT and LDAT under single-training domain generalization setting, i.e., using only one training domain and generalize on the others. We compare our work with Neuron Coverage-Guided Domain Generalization (NCDG) [36]. The results are in Table 10. Both MAT and LDAT outperform NCDG under the scenario of single-source domain generalization.

D Detailed Description of LDAT

In this section, we describe the detailed training procedure of LDAT (see Algorithm 2). We conduct a single-step gradient ascent for the inner maximization for the perturbations LDAT. We adopt l_2 norm for the perturbations.

Table 9: The test accuracy (%) on NICO.

ERM	MAT	LDAT	GUT
73.2 ± 1.9	**74.2 ± 1.5**	74.4 ± 1.6	66.6 ± 1.7

Table 10: The test accuracy (%) of single-training domain generalization on PACS. Each line represents a case when we train on one domain and test on the other domains. '-' means we train on this domain, and test on the other three domains.

Algorithm	A	C	P	S	avg
MAT	-	73.8	94.1	74	80.6
	78.5	-	94.2	75.9	82.9
	80.9	73.7	-	76.6	77.1
	80.4	76.3	93.3	-	83.3
LDAT	-	74.8	94.2	75.8	81.6
	77.2	-	93.9	75.6	82.3
	78.5	77.9	-	80.4	79
	74.3	76.4	94.7	-	81.8
NCDG	-	68.6	95.0	66.4	76.6
	71.6	-	85.8	71.9	76.4
	68.8	29.8	-	48.6	49.0
	45.6	65.8	47.9	-	53.1

Algorithm 2 Detailed Training Procedure of LDAT

Require:
Labeled training data of m domains $E_1, ..., E_m$, where$E_e := \{(x_i^e, y_i^e)\}_{i=1}^{n_e}$,
rank of the perturbations l, factor A, B learning rate ρ,
FGSM step size γ, perturbation radius ϵ,
number of training epochs T, learning rate for model parameters r, batch size b.

Ensure:
Updated model f_θ with parameter θ.

Randomly initiate θ, perturbation δ^e, factor A^e, B^e such that $||\delta^e||_2 \leq \epsilon, \forall e \in \{1, ..., m\}$.
for iterations in $1, 2, ..., T$ **do**
  for e in $1, 2, ..., m$ **do**
    Randomly select batch $\mathcal{B}^e = \{(x_u^e, y_u^e)\}_{u=1}^b$ from domain E_e.
    Compute the adversarial sample: $x_u^{e'} = x_u^e + A^e B^e, \forall u \in \{1, ..., b\}$
    Update A^e by $A^e \leftarrow A^e + \rho \frac{1}{b} \sum_{u=1}^b \nabla_{A^e} \mathcal{L}(f_\theta(x_u^{e'}), y_u^e), \forall u \in \{1, ..., b\}$.
    Update B^e by $B^e \leftarrow B^e + \rho \frac{1}{b} \sum_{u=1}^b \nabla_{B^e} \mathcal{L}(f_\theta(x_u^{e'}), y_u^e), \forall u \in \{1, ..., b\}$.
    Project δ_i^e to the l_2 ball of radius ϵ.
    Compute the adversarial sample: $x_u^{e'} = x_u^e + A^e B^e, \forall u \in \{1, ..., b\}$
    Update model parameter: $\theta \leftarrow \theta - r \frac{1}{b} \sum_{u=1}^b \nabla_\theta \mathcal{L}(f_\theta(x_u^{e'}), y_u^e), \forall u \in \{1, ..., b\}$.
  end for
end for