

# Robust Generalization against Photon-Limited Corruptions via Worst-Case Sharpness Minimization

Zhuo Huang<sup>1,†</sup>, Miaoxi Zhu<sup>2,†</sup>, Xiaobo Xia<sup>1</sup>, Li Shen<sup>3</sup>, Jun Yu<sup>4,✉</sup>,  
Chen Gong<sup>5</sup>, Bo Han<sup>6</sup>, Bo Du<sup>2</sup>, Tongliang Liu<sup>1</sup>

## Abstract

Robust generalization aims to tackle the most challenging data distributions which are rare in the training set and contain severe noises, i.e., photon-limited corruptions. Common solutions such as distributionally robust optimization (DRO) focus on the worst-case empirical risk to ensure low training error on the uncommon noisy distributions. However, due to the over-parameterized model being optimized on scarce worst-case data, DRO fails to produce a smooth loss landscape, thus struggling on generalizing well to the test set. Therefore, instead of focusing on the worst-case risk minimization, we propose SharpDRO by penalizing the sharpness of the worst-case distribution, which measures the loss changes around the neighbor of learning parameters. Through worst-case sharpness minimization, the proposed method successfully produces a flat loss curve on the corrupted distributions, thus achieving robust generalization. Moreover, by considering whether the distribution annotation is available, we apply SharpDRO to two problem settings and design a worst-case selection process for robust generalization. Theoretically, we show that SharpDRO has a great convergence guarantee. Experimentally, we simulate photon-limited corruptions using CIFAR10/100 and ImageNet30 datasets and show that SharpDRO exhibits a strong generalization ability against severe corruptions and exceeds well-known baseline methods with large performance gains.

## 1. Introduction

Learning against corruptions has been a vital challenge in the practical deployment of computer vision models, as learning models are much more fragile to subtle noises than human perception systems [12, 15, 31]. During the training, the encountered corruptions are essentially perceived as




Figure 1. Illustration of photon-limited corruptions.

distribution shift, which would significantly hinder the prediction results [29, 34, 51, 58–60]. Therefore, to mitigate the performance degradation, enhancing generalization to corrupted data distributions has drawn lots of attention [1, 45].

In the real world, noise corruptions are commonly known as photon-limited imaging problems [21, 26, 35, 50] which arises due to numerous small corruption photons arriving to a image sensor. Consequently, different numbers of captured photons would form different levels of corruptions severity, further producing multiple data distributions and imposing varied impacts on learning models [14]. Specifically, the encountered photon-limited corruption  $\mathcal{E}$  is a composition of multiple noise photon  $u$ , which is triggered by some discrete factors with a certain probability during a time interval. For example, a photon  $u$  can be triggered by each platform changing, re-distribution, transmission, etc. More photons are captured, and severer corruption would be applied to the image. Therefore, the severity  $s$  of the photon-limited corruptions  $\mathcal{E}$  can be modeled by Poisson distribution, i.e.,  $s \sim P(s; \lambda) = \frac{e^{-\lambda} \lambda^s}{s!}$ , which is illustrated in Figure 1. As a result, the real-world training set is not completely composed of clean data, but contains corrupted data with a smaller proportion as the severity goes stronger.

Dealing with such a realistic problem by vanilla empirical risk minimization can achieve satisfactory averaged accuracy on the whole training set. However, due to the extremely limited number of severely-corrupted data, the learning model would produce large training errors on the corrupted distributions, further hindering the robust performance under challenging real-world situations. A popular approach to

<sup>†</sup>Equal contribution. <sup>1</sup>Sydney AI Centre, The University of Sydney; <sup>2</sup>National Engineering Research Center for Multimedia Software, Institute of Artificial Intelligence, School of Computer Science and Hubei Key Laboratory of Multimedia and Network Communication Engineering, Wuhan University; <sup>3</sup>JD Explore Academy; <sup>4</sup>Department of Automation, University of Science and Technology of China; <sup>5</sup>Key Laboratory of Intelligent Perception and Systems for High-Dimensional Information of Ministry of Education, School of Computer Science and Engineering, Nanjing University of Science and Technology; <sup>6</sup>Department of Computer Science, Hong Kong Baptist University. Correspondence to Jun Yu (harryjun@ustc.edu.cn).

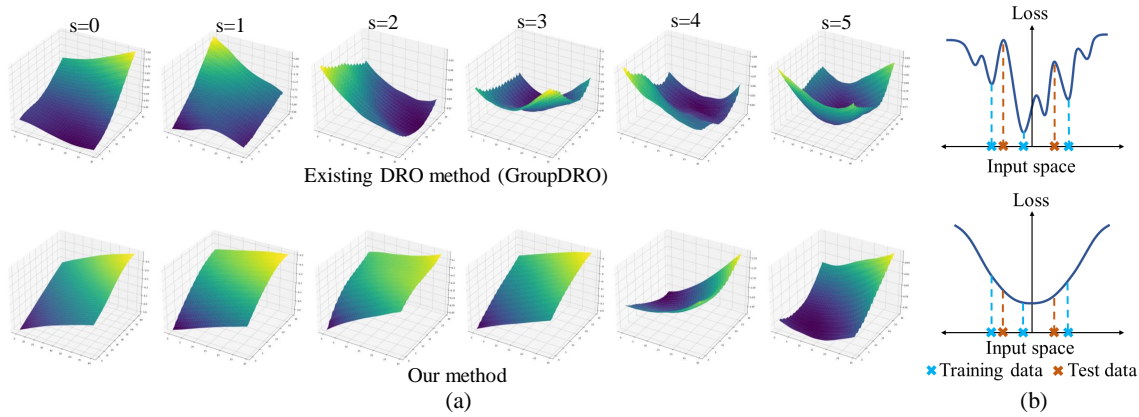


Figure 2. Illustration of our motivation. (a) Loss surface visualization of GroupDRO and the proposed SharpDRO. The columns from left to right stand for corrupted distributions with severity  $s = 0$  to 5. (b) Illustration of why a sharp loss surface hinders generalization to test data.

achieve low error on the scarce corrupted data is distributionally robust optimization (DRO) [39, 42, 45, 54, 55, 65], which commonly optimizes the model parameter  $\theta$  by optimizing:

$$\min_{\theta \in \Theta} \sup_{Q \in \mathcal{Q}} \mathbb{E}_{(x,y) \sim Q} [\mathcal{L}(\theta; (x, y))], \quad (1)$$

where  $\mathcal{Q}$  denotes the uncertainty set that is utilized to estimate the possible test distribution. Intuitively, DRO assumes that  $\mathcal{Q}$  consists of multiple sub-distributions, among which exists a worst-case distribution  $Q$ . By concentrating on the risk minimization of the worst-case distribution, DRO hopes to train a robust model which can deal with the potential distribution shift during the test phase. However, existing DRO methods usually leverage over-parameterized models to focus on a small portion of worst-case training data. Therefore, the worst-case data contaminated with severe corruption is highly possible to get stuck into sharp minima. As shown in the upper of Figure 2 (a), a stronger corruption would cause the existing method to learn a sharper loss surface. Consequently, optimization via DRO fails to produce a flat loss landscape over the corrupted distributions, which leads to a large generalization gap between training and test set [6, 22].

To remedy this defect, in this paper, we propose SharpDRO method to focus on learning a flat loss landscape of the worst-case data, which can largely mitigate the training-test generalization gap problem of DRO. Specifically, we adopt the largest loss difference formed by applying weight perturbation [11, 57] to measure the sharpness of the loss function. Intuitively, a sharp loss landscape is sensitive to noise and cannot generalize well on the test set. On the contrary, a flat loss landscape produces consistent loss values and is robust against perturbations (Figure 2 (b)). By minimizing the sharpness, we can effectively enhance the generalization performance [6, 22]. However, directly applying sharpness minimization on multiple distributions would yield poor results [4], as the computed sharpness could be influenced by the largest data distribution, and thus cannot generalize well to small corrupted data. Therefore, we only focus on worst-case sharpness minimization. In this way, as the lower of Figure 2 (a) shows, SharpDRO successfully produces a flat loss surface, thus achieving robust generalization on the

severely corrupted distributions.

In addition, identification of the worst-case distribution requires expensive annotations, which are not always practically feasible [30]. In this paper, we apply SharpDRO to solve two problem settings: 1) *Distribution-aware robust generalization* which assumes that distribution indexes are accessible, and 2) *Distribution-agnostic robust generalization* where the distributions are no longer identifiable, making the worst-case data hard to find. Existing approaches such as Just Train Twice (JTT) require two-stage training which is rather inconvenient. To tackle this challenge, we propose a simple (Out-of-distribution) OOD detection [15, 19, 20, 29, 32, 52, 53] process to detect the worst-case data, which can be further leveraged to enable worst-case sharpness minimization. Through constructing training sets according to the Poisson distributed noisy distribution using CIFAR10/100 and ImageNet30, we show that SharpDRO can achieve robust generalization results on both two problem settings, surpassing well-known baseline methods by a large margin.

To sum up, our main contributions are three-fold:

- We proposed a sharpness-based DRO method that overcomes the poor worst-case generalization performance of distributionally robust optimization.
- We apply the proposed SharpDRO to both distribution-aware and distribution-agnostic settings, which brings a practical capability to our method.
- Theoretically, we show that SharpDRO has a convergence rate of  $\mathcal{O}(\frac{\kappa^2}{\sqrt{MT}})$ . Empirically, we conduct extensive analysis to validate its generalization capability on photon-limited corruptions.

In the following, we first introduce the background in section 2. Then, we specify our SharpDRO over two problem settings in Section 3. Moreover, we give a detailed optimization process and analysis in Section 3.3. Further, experiments are demonstrated to validate SharpDRO in Section 4. At last, we conclude this paper in Section 5.

## 2. Robust Generalization Methods

Due to the practical significance of robust generalization, various approaches have been proposed to deal with distribution shift. Here we briefly introduce three typical baseline methods, namely Invariant Risk Minimization, Risk Extrapolation, and GroupDRO.

**Invariant Risk Minimization (IRM)** [1, 5, 7] aims to extract the invariant feature across different distributions (also denoted as environments). Specifically, the learning model is separated into a feature extractor  $G$  and a classifier  $C$ . IRM assumes an invariant model  $C \circ G$  over various environments can be achieved if the classifier  $C$  constantly stays optimal. Then, the learning objective is formulated as:

$$\begin{aligned} \min_{C^* \circ G} \{ \mathcal{L}_{\text{IRM}} := \sum_{e \in \mathcal{E}} \mathcal{L}^e(C^* \circ G) \} \\ \text{s. t. } C^* \in \arg \min_G \mathcal{L}^e(C \circ G), \text{ for all } e \in \mathcal{E}, \end{aligned} \quad (2)$$

where  $C^*$  stands for the optimal classifier, and  $e$  denotes a specific environment from a given environmental set  $\mathcal{E}$ . By solving Eq. 2, the feature extractor  $G$  can successfully learn invariant information without being influenced by the distribution shift between different environments.

**Risk Extrapolation (REx)** [25] targets at generalization to out-of-distribution (OOD) environments. Inspired by the discovery that penalizing the loss variance across distributions helps achieve improved performance on OOD generalization, REx proposes to optimizing via:

$$\min_{\theta \in \Theta} \{ \mathcal{L}_{\text{REx}} := \sum_{e \in \mathcal{E}} \mathcal{L}^e(\theta) + \beta \text{Var}(\mathcal{L}^e, \dots, \mathcal{L}^m) \}, \quad (3)$$

where  $\beta$  controls the penalization level. Intuitively, REx seeks to achieve risk fairness among all  $m$  training environments, so as to increase the similarity of different learning tasks. As a result, the training model can capture the invariant information that helps generalize to unseen distributions.

**GroupDRO** [13, 42, 45] deal with the situation when the correlation between class label  $y$  and unknown attribute  $a$  differs in the training and test set. Such a difference is called spurious correlation which could seriously misguide the model prediction. As a solution, GroupDRO considers each combination of class and attribute as a group  $g$ . By conducting risk minimization through:

$$\min_{\theta \in \Theta} \{ \mathcal{L}_{\text{GroupDRO}} := \max_g \mathbb{E}_{(x,y) \sim P_g} [\mathcal{L}(\theta; (x,y))] \}, \quad (4)$$

the worst-case group from distribution  $P_g$  is emphasized, thus breaking the spurious correlation.

**Discussion:** IRM and REx both focus on learning invariant knowledge across various environments. However, when the training set contains extremely imbalanced noisy distributions, as shown in Figure 1, the invariant learning results would be greatly misled by the most dominating distribution. Thus, the extracted invariant feature would be questionable for generalization against distribution shift. Although emphasizing the risk minimization of worst-case data via GroupDRO can alleviate the imbalance problem, its generalization

performance is still sub-optimal when facing novel test data. However, SharpDRO can not only focus on the uncommon corrupted data but also effectively improve the generalization performance on the test set by leveraging worst-case sharpness minimization.

Our investigated problem is closely related to OOD generalization which is a broad field that contains many popular research topics, such as **Domain Generalization** [3, 18, 36, 38, 41, 43, 47, 66], **Causal Invariant Learning** [1, 25, 27, 46, 62, 63]. Generally, existing works mainly studies two types of research problem: 1) mitigating domain shift between training and test dataset; and 2) breaking the spurious correlation between causal factors. However, as generalization against corruptions **does not introduce any domain shift or spurious correlation**, such a problem cannot be naively solved by domain generalization methods or causal representation learning techniques. Therefore, in this paper, we focus on complementing this rarely-explored field and propose SharpDRO to enforce robust generalization against corruption. In the next section, we elaborate on the methodology of SharpDRO.

## 3. Methodology

In robust generalization problems, we are given a training set  $\mathcal{D}^{\text{train}}$  containing  $n$  image examples, each example  $x \in \mathcal{X}$  is given a class label  $y \in \mathcal{Y} = \{1, 2, \dots, c\}$ . Moreover, the training set is corrupted by a certain type of noise whose severity  $s$  follows a Poisson distribution  $P(s; \lambda)$ . Here we assume  $\lambda = 1$  which indicates that the mean number of the noise photon  $u$  that occurred during a time interval is 1. Therefore, the distribution  $P$  of the whole training set is composed of  $S$  sub-distributions  $P_s, s \in \{1, 2, \dots, S\}$  with varied levels of corruption. Our goal is to learn a robust model  $\theta \in \Theta$  that can achieve good generalization performance on challenging data distributions  $P_s$  with large severity.

The general objective of SharpDRO is formulated as:

$$\begin{aligned} \min_{\theta} \{ \mathcal{L}_{\text{SharpDRO}} := \mathbb{E}_{(x,y) \sim Q} [\mathcal{L}(\theta; (x,y))] + \\ \mathbb{E}_{(x,y) \sim Q} [\mathcal{R}(\theta; (x,y))] \}, \end{aligned} \quad (5)$$

where the first term denotes the risk minimization using loss function  $\mathcal{L}$ , meanwhile a worst-case distribution  $Q$  is selected based on the model prediction. The second term  $\mathcal{R}$  indicates the sharpness minimization which aims to maximally improve the generalization performance on the worst-case distribution  $Q$ . Specifically, as shown in Figure 3, the sharpness gradually increases as the corruption severity enlarges. Therefore, to accomplish robust generalization, we are motivated to emphasize the worst-case distribution. As a result, we can produce much smaller sharpness compared to other methods, especially on severely corrupted distributions.

In the following, we first introduce worst-case sharpness for robust generalization. Then, we demonstrate worst-case

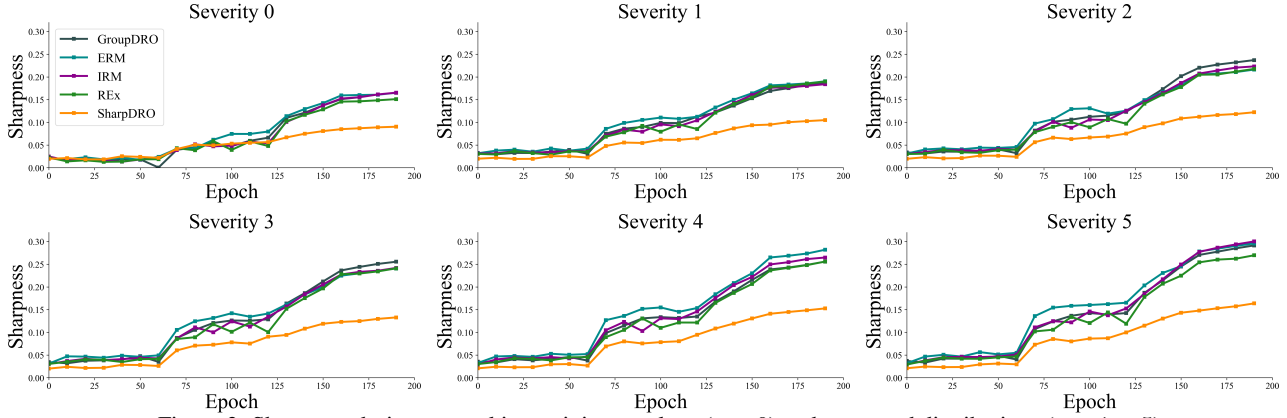


Figure 3. Sharpness during networking training on clean ( $s = 0$ ) and corrupted distributions ( $s = 1$  to 5).

data selection on two scenarios. Finally, we provide a detailed optimization process and convergence analysis.

### 3.1. Sharpness for Robust Generalization

The main challenge of robust generalization is that the training distribution is extremely imbalanced, as shown in Figure 1. The performance on the abundant clean data is quite satisfactory, but robustness regarding the corrupted distribution is highly limited, owing to the severe disturbance of corruption as well as the insufficiency of noisy data. To enhance the generalization performance, we leverage sharpness to fully exploit the worst-case data. Specifically, sharpness [11, 23, 33, 57, 70] is measured by **the largest loss change** when model weight  $\theta$  is perturbed with  $\epsilon$ , formally:

$$\mathcal{R} := \max_{\|\epsilon\|_2 \leq \rho} \{\mathcal{L}(\theta + \epsilon; (x, y)) - \mathcal{L}(\theta; (x, y))\}, \quad (6)$$

where  $\rho$  is a scale parameter to control the perturbation magnitude. By supposing  $\epsilon$  is small enough, we can have:

$$\mathcal{L}(\theta + \epsilon) - \mathcal{L}(\theta) \approx \nabla \mathcal{L}(\theta) \epsilon. \quad (7)$$

Further, we hope to obtain the largest loss change to find the optimal weight perturbation  $\epsilon^*$ , which is be computed as:

$$\epsilon^* := \arg \max_{\|\epsilon\|_2 \leq \rho} \nabla \mathcal{L}(\theta) \epsilon. \quad (8)$$

By following dual norm problem, the optimal  $\epsilon^*$  can be solved as  $\rho \text{sign}(\nabla \mathcal{L}(\theta))$  [11], which is essentially the  $\infty$ -norm of the gradient  $\nabla \mathcal{L}$  multiplied with a scale parameter  $\rho$ . Hence, common sharpness minimization aims to minimize:

$$\mathbb{E}_{(x, y) \sim Q} \mathcal{R} := \mathcal{L}(\theta + \rho \text{sign}(\nabla \mathcal{L}(\theta; (x, y)))) - \mathcal{L}(\theta; (x, y)). \quad (9)$$

The intuition is that the perturbation along the gradient norm direction increases the loss value significantly. When training on corrupted distributions, the scarce noisy data scatter sparsely in the high-dimensional space. As a consequence, the neighbor of each datum could not be sufficiently explored, thus producing a sharp loss curve. During test, the unseen noisy data is likely to fall on an unexplored point with a large loss, further causing inaccurate model predictions.

Therefore, instead of directly applying sharpness minimization on the whole dataset, which leads to poor generalization performance [4] (as demonstrated in Section 4.3),

we focus on sharpness minimization over the worst-case distribution  $Q$ . By conducting the worst-case sharpness minimization, we can enhance the flatness of our classifier. Consequently, when predicting unknown data during the test phase, a flat loss landscape is more likely to produce a low loss than a sharp one, hence our SharpDRO can generalize better than other DRO methods. However, the robust performance largely depends on the worst-case distribution  $Q$ , so next, we explain our worst-case data selection.

### 3.2. Worst-Case Data Selection

Generally, the worst-case data selection focuses on finding the most uncertain data distribution  $Q$  from the uncertainty set  $\mathcal{Q}$ , which is a  $f$ -divergence ball from the training distribution  $P$  [2, 10, 17]. Most works assume each distribution is distinguishable from the other. However, when the distribution index is not available, it would be very hard to select worst-case data. In this section, we investigate two situations: distribution-aware robust generalization and distribution-agnostic robust generalization.

#### 3.2.1 Distribution-Aware Robust Generalization

When given annotations to denote different severity of corruptions, the image data  $x$  is paired with class label  $y$  and distribution index  $s$ . Then, the worst-case distribution  $Q$  can be found by identifying the sub-distribution  $P_s \in P$  that yields the largest loss. Hence, we can optimize through:

$$\min_{\theta} \left\{ \max_{\omega_s; Q = \{\omega_s P_s\}_{s=1}^S} \left\{ \sum_{(x_i, y_i) \in Q} [\mathcal{L}(\theta, \omega_s; (x_i, y_i))] \right\} \right. \quad (10) \\ \left. + \sum_{(x_i, y_i) \in Q} [\mathcal{R}(\theta, \omega_s; (x_i, y_i))] \right\},$$

where  $\omega_s$  belongs to a  $(S - 1)$ -dimensional probability simplex. The first term simply recovers the learning target of GroupDRO [17, 45] and helps find the worst-case distribution  $Q$ . Then, by emphasizing the selected  $P_s$ , the second sharpness minimization term can act as a sharpness regularizer. As a result, SharpDRO can learn a flatter loss surface on the worst-case data, thus generalize better compared to GroupDRO, as discussed in Section 4.



### 3.2.2 Distribution-Agnostic Robust Generalization

Due to the annotations being extremely expensive in the real world, a practical challenge is how to learn a robust model without a distribution index. Unlike JTT [30] which trains the model through two stages, we aim to solve this problem more efficiently by detecting the worst-case data during network training. As the corrupted data essentially lie out-of-distribution, so we are motivated to conduct OOD detection [16, 28, 29, 32, 56] to find the worst-case data.

Particularly, we re-utilize the previously computed weight perturbation  $\epsilon^*$  to compute an OOD score:

$$\omega_i = \max f(\theta; (x_i)) - \max f(\theta + \epsilon^*; (x_i)), \quad (11)$$

where  $f(\cdot)$  stands for the  $c$ -dimensional label prediction in the label space, whose maximum value is considered as prediction confidence. Intuitively, as the model is much more robust to the clean distribution than the corrupted distribution, the prediction of clean data usually exhibits more stability than scarce noisy data when facing perturbations. Hence, if an example comes from a rarely explored distribution, its prediction certainty would deviate significantly from the original value, thus producing a large OOD score, as shown in Section 4.3. Note that the major difference is that we target generalization on worst-case data, but OOD detection aims to exclude OOD data.

To this end, we can construct our worst-case dataset as  $Q := \left\{ \sum_{i=1}^M \bar{\omega}_i \cdot (x_i, y_i) : \bar{\omega}_i = \frac{\omega_i}{\sum_{i=1}^M \omega_i} \right\}$ , where normalization on  $\omega_i$  is performed simultaneously. Then, the learning target of the distribution-agnostic setting becomes:

$$\min_{\theta} \left\{ \max_{\bar{\omega}_i} \left\{ \sum_{(x_i, y_i) \in Q} [\mathcal{L}(\theta, \bar{\omega}_i; (x, y))] \right\} + \sum_{(x_i, y_i) \in Q} [\mathcal{R}(\theta, \bar{\omega}_i; (x, y))] \right\}, \quad (12)$$

Therefore, the worst-case data can be selected by focusing on the examples with large OOD scores. In this way, our

---

#### Algorithm 1 Optimization process of SharpDRO

---

- 1: Training set  $\mathcal{D}^{\text{train}} = \{x_i, y_i\}_{i=1}^M$  containing Poisson distributed noisy corruptions; Model parameter  $\theta \leftarrow \theta_0$ ; Weighting parameter  $\omega \leftarrow \omega_0$ ; Learning rate:  $\eta_{\theta}, \eta_{\omega}$ .
  - 2: **for**  $t \in 0, 1, \dots, T-1$  **do**
  - 3:   **if** *Distribution-aware* **then**
  - 4:      $\triangleright$  Loss maximization via optimizing  $\omega_{t+1}$
  - 5:      $\omega_{t+1} := \arg \max_{\omega} \left\{ \mathbb{E}_{(x, y) \sim \omega P_s} [\mathcal{L}(\theta, \omega; (x, y))] \right\}$ ;
  - 6:   **else if** *Distribution-agnostic* **then**
  - 7:      $\triangleright$  OOD detection for computing  $\omega_{t+1}$
  - 8:     Update  $\omega_{t+1}$  via Eq. 11;
  - 9:   **end if**
  - 10:    $\triangleright$  Optimize variable  $\theta$
  - 11:    $\theta_{t+1} = \arg \min_{\theta} \left\{ \mathbb{E}_{(x, y) \sim \omega_t P} [\mathcal{L}(\theta, \omega_t) + \mathcal{R}(\theta, \omega_t)] \right\}$
  - 12: **end for**
- 

sharpDRO can be successfully deployed into the distribution-agnostic setting to ensure robust generalization, whose effectiveness is demonstrated in Sections 4.2 and 4.3.

### 3.3. Optimization for SharpDRO

In both distribution-aware and distribution-agnostic scenarios, the worst-case data distribution is identified using distribution weighting parameter  $\omega_s$  and OOD score  $\omega$ , respectively. Intuitively, their effect is similar: finding the worst data distribution which yields the maximum loss. Therefore, without loss of generality, we consider the maximization in Eqs. 10 and 12 as the optimization on the same weighting parameter  $\omega$  and the samples  $(x, y)$  can be considered i.i.d. from  $\mathcal{Q}$  weighted by  $\omega$  to compute the loss value  $\mathcal{L}$ . Moreover, the sharpness regularization can be reformulated in the same way as Eq. 6 by including  $\omega$ :  $\mathcal{R}(\theta, \omega; (x, y)) = \max_{\|\epsilon\|_2 \leq \rho} \{ \mathcal{L}(\theta + \epsilon, \omega; (x, y)) - \mathcal{L}(\theta, \omega; (x, y)) \}$ . Therefore, our general learning objective can be formulated as a bi-level optimization problem:

$$\min_{\theta} \mathbb{E}_{(x, y) \sim \mathcal{Q}} [\mathcal{L}(\theta, \omega^*; (x, y))] + \mathcal{R}(\theta, \omega^*; (x, y)) \quad (13)$$

$$\text{s. t. } \omega^* = \arg \max_{\omega} \mathbb{E}_{(x, y) \sim \mathcal{Q}} [\mathcal{L}(\theta, \omega; (x, y))]. \quad (14)$$

The optimization process is shown in Algorithm 1. Specifically, we first update the weighting parameter  $\omega$  based on the empirical risk term  $\mathcal{L}$  using stochastic gradient ascent. Then, by leveraging the updated  $\omega$ , we optimize the general objective which contains both risk minimization of  $\mathcal{L}$  and worst-case sharpness minimization of  $\mathcal{R}$ . We iterate these processes until convergence, hoping to minimize the risk on target loss function  $\mathcal{L}$  with the worst-case data distribution.

**Convergence Analysis:** First we give some brief notations:  $\mathbb{L}(\theta, \omega) := \mathbb{E}_{(x, y) \sim \mathcal{Q}} \mathcal{L}(\theta, \omega; (x, y))$ . The worst-case data distribution which has the maximum loss is denoted by  $\omega^*(\theta) := \arg \max_{\omega} \mathbb{L}(\theta, \omega)$ . We can obtain the convergence to a stationary point of  $\mathbb{L}^*(\theta) := \max_{\omega} \mathbb{L}(\theta, \omega) = \mathbb{L}(\theta, \omega^*(\theta))$  by averaged gradient  $\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla \mathbb{L}^*(\theta_t)\|^2$ .

**Theorem 1 (Informal).** *Assuming the loss function  $\mathbb{L}$  is  $l$ -Lipschitz smooth, satisfies  $\mu$ -Polyak-Lojasiewicz (PL) condition on the second variable  $\omega$ , and has unbiased estimation about the gradient as well as  $\sigma^2$  bounded variance, we can get the convergence rate during  $T$  iterations:*

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla \mathbb{L}^*(\theta_t)\|^2 &\leq 320 \sqrt{\frac{3\kappa^4 l (\mathbb{E}[\mathbb{L}^*(\theta_0)] - \min_{\theta} \mathbb{E}[\mathbb{L}^*(\theta)]) \sigma^2}{11MT}} \\ &= \mathcal{O} \left( \frac{\kappa^2}{\sqrt{MT}} \right), \end{aligned}$$

where the conditional number  $\kappa = l/\mu$  and  $M$  means the sample batch (here we can choose  $M = 1$ )<sup>1</sup>.

---

<sup>1</sup>The resulting bound here means our SharpDRO can converge to the  $\epsilon$ -stationary point in  $\frac{1}{\epsilon^2}$  iterations. Moreover, we will present our proof details in the **Appendix**, including the definitions, assumptions, and lemmas.

Table 1. Quantitative comparisons on distribution-aware robust generalization setting. Averaged accuracy (%) with standard deviations are computed over three independent trials.

data	Type	Method	Corruption Severity					
			0	1	2	3	4	5
CIFAR10	Gaussian	ERM	90.9 ± 0.02	89.2 ± 0.02	86.4 ± 0.03	85.9 ± 0.01	83.5 ± 0.01	78.8 ± 0.01
		IRM	91.8 ± 0.01	90.3 ± 0.01	89.5 ± 0.01	86.7 ± 0.02	81.8 ± 0.02	80.0 ± 0.02
		REx	91.3 ± 0.03	89.5 ± 0.02	88.1 ± 0.02	86.7 ± 0.02	83.3 ± 0.01	80.5 ± 0.02
		GroupDRO	90.2 ± 0.03	89.1 ± 0.02	88.4 ± 0.04	84.3 ± 0.01	83.0 ± 0.02	78.2 ± 0.02
		SharpDRO	<b>92.9 ± 0.02</b>	<b>91.3 ± 0.02</b>	<b>90.5 ± 0.01</b>	<b>88.4 ± 0.02</b>	<b>86.9 ± 0.01</b>	<b>84.7 ± 0.01</b>
	JPEG	ERM	91.0 ± 0.02	89.0 ± 0.02	86.2 ± 0.02	83.1 ± 0.02	82.5 ± 0.03	81.4 ± 0.03
		IRM	90.2 ± 0.02	88.2 ± 0.02	86.7 ± 0.03	84.0 ± 0.02	82.9 ± 0.03	81.6 ± 0.02
		REx	89.6 ± 0.03	89.2 ± 0.02	86.0 ± 0.03	85.8 ± 0.03	82.7 ± 0.03	81.9 ± 0.02
		GroupDRO	90.3 ± 0.02	88.6 ± 0.03	86.5 ± 0.03	84.2 ± 0.02	83.2 ± 0.02	82.1 ± 0.02
		SharpDRO	<b>91.2 ± 0.01</b>	<b>89.3 ± 0.02</b>	<b>87.6 ± 0.02</b>	<b>86.6 ± 0.02</b>	<b>85.8 ± 0.03</b>	<b>84.5 ± 0.01</b>
CIFAR100	Gaussian	ERM	68.2 ± 0.01	64.8 ± 0.01	60.6 ± 0.01	56.9 ± 0.01	53.9 ± 0.01	50.2 ± 0.03
		IRM	64.7 ± 0.02	64.7 ± 0.01	62.2 ± 0.01	54.5 ± 0.02	53.4 ± 0.03	50.4 ± 0.01
		REx	68.0 ± 0.03	65.1 ± 0.03	61.8 ± 0.01	56.8 ± 0.01	53.2 ± 0.01	51.5 ± 0.01
		GroupDRO	66.1 ± 0.01	61.7 ± 0.02	59.3 ± 0.03	53.6 ± 0.01	54.0 ± 0.02	50.6 ± 0.02
		SharpDRO	<b>71.2 ± 0.02</b>	<b>70.1 ± 0.01</b>	<b>68.6 ± 0.01</b>	<b>58.8 ± 0.01</b>	<b>57.5 ± 0.02</b>	<b>53.8 ± 0.03</b>
	JPEG	ERM	64.7 ± 0.01	62.8 ± 0.02	57.2 ± 0.02	54.7 ± 0.03	54.0 ± 0.02	50.6 ± 0.03
		IRM	64.2 ± 0.02	62.8 ± 0.04	58.0 ± 0.01	56.3 ± 0.02	55.0 ± 0.02	51.8 ± 0.03
		REx	63.5 ± 0.03	62.2 ± 0.02	58.4 ± 0.03	56.5 ± 0.03	55.1 ± 0.02	52.2 ± 0.01
		GroupDRO	63.4 ± 0.02	61.6 ± 0.04	58.6 ± 0.02	57.2 ± 0.02	55.8 ± 0.03	53.1 ± 0.02
		SharpDRO	<b>65.3 ± 0.02</b>	<b>63.0 ± 0.02</b>	<b>59.8 ± 0.02</b>	<b>58.8 ± 0.02</b>	<b>57.8 ± 0.03</b>	<b>55.3 ± 0.03</b>
ImageNet30	Gaussian	ERM	87.5 ± 0.01	84.6 ± 0.01	81.9 ± 0.01	76.5 ± 0.01	71.2 ± 0.01	65.3 ± 0.01
		IRM	86.6 ± 0.01	84.4 ± 0.03	80.6 ± 0.01	75.2 ± 0.01	70.7 ± 0.03	64.8 ± 0.01
		REx	86.3 ± 0.01	83.8 ± 0.03	81.1 ± 0.02	75.6 ± 0.02	71.5 ± 0.01	66.1 ± 0.03
		GroupDRO	85.1 ± 0.02	84.2 ± 0.01	81.2 ± 0.03	76.3 ± 0.03	72.0 ± 0.02	66.3 ± 0.01
		SharpDRO	<b>88.4 ± 0.02</b>	<b>87.6 ± 0.01</b>	<b>83.3 ± 0.01</b>	<b>79.1 ± 0.02</b>	<b>73.5 ± 0.03</b>	<b>68.7 ± 0.01</b>
	JPEG	ERM	86.8 ± 0.03	85.3 ± 0.03	83.2 ± 0.02	82.6 ± 0.01	80.4 ± 0.04	78.2 ± 0.02
		IRM	86.2 ± 0.03	85.1 ± 0.02	83.8 ± 0.03	83.2 ± 0.03	81.6 ± 0.02	79.1 ± 0.01
		REx	85.8 ± 0.02	85.3 ± 0.02	83.5 ± 0.02	82.9 ± 0.02	81.4 ± 0.02	78.2 ± 0.02
		GroupDRO	86.7 ± 0.02	84.9 ± 0.02	84.1 ± 0.02	84.5 ± 0.02	82.3 ± 0.02	79.0 ± 0.02
		SharpDRO	<b>87.4 ± 0.02</b>	<b>86.4 ± 0.03</b>	<b>86.2 ± 0.03</b>	<b>85.6 ± 0.02</b>	<b>83.9 ± 0.02</b>	<b>82.1 ± 0.03</b>

## 4. Experiments

In our experiments, we first give details about our experimental setup. Then, we conduct quantitative experiments to compare to proposed SharpDRO with the most popular baseline methods by considering both distribution-aware and distribution agnostic settings, which shows the capability of SharpDRO to tackle the most challenging distributions. Finally, we conduct qualitative analyses to investigate the effectiveness of SharpDRO in achieving robust generalization.

### 4.1. Experimental Setup

For distribution-aware situation, we choose GroupDRO [45], IRM [1], REx [25], and ERM for comparisons. As for distribution-agnostic situation, we pick JTT [30] and Environment Inference for Invariant Learning (EIIL) [7] for baseline methods<sup>2</sup>. For each problem setting, we construct corrupted dataset using CIFAR10/100 [24] and ImageNet30 [44]. Specifically, we following [14] to perturb the image data with severity level varies from 1 to 5 by using two types of corruption: “Gaussian Noise” and “JPEG Compression”. Moreover, the clean data are considered as having a corruption severity of 0. For each corrupted distribution, we sample them with different probabilities by following Poisson distribution  $P(s; \lambda = 1)$ , *i.e.*, for  $s$  varies from 0 to 5, the sample probabilities

<sup>2</sup>We do not include sharpness minimization SAM [11] in the comparisons because its OOD generalization performance is limited [4], but we carefully analyze SAM in Section 4.3

are  $\{0.367, 0.367, 0.184, 0.061, 0.015, 0.003\}$ , respectively. Then, we test the robust performance on each data distribution. For hyper-parameter  $\rho$ , we follow [11] by setting it to 0.05 to control the magnitude of  $\epsilon^*$ . For each experiment, we conduct three independent trials and report the average test accuracy with standard deviations. Please see more details on experimental setting, results on other corruptions, and practical implementation in the **Appendix**.

### 4.2. Quantitative Comparisons

In this part, we focus on three questions: 1) Can SharpDRO perform well on two situations of robust generalization? 2) Does SharpDRO generalize well on the most severely corrupted distributions? and 3) Is SharpDRO able to tackle different types of corruption? To answer these questions, we conduct experiments on both two settings by testing on different corruption types and severity levels.

**Distribution-Aware Robust Generalization** As shown in Table 1, we can see that SharpDRO surpasses other methods with larger performance gains as the corruption severity goes stronger. Especially in CIFAR10 dataset on “Gaussian Noise” corruption, improvement margin between SharpDRO and second-best method is 1.1% with severity of 0, which is further increased to about 4.2% with severity of 5, which indicates the capability of SharpDRO on generalization against severe corruptions. Moreover, SharpDRO frequently outperforms other methods on all scenarios, which manifests the robustness of SharpDRO against various corruption types.

Table 2. Quantitative comparisons on distribution-agnostic robust generalization setting. Averaged accuracy (%) with standard deviations are computed over three independent trails.

Dataset	Type	Method	Corruption Severity					
			0	1	2	3	4	5
CIFAR10	Gaussian	JTT	89.9 ± 0.02	88.8 ± 0.02	86.5 ± 0.02	86.1 ± 0.02	83.4 ± 0.03	79.8 ± 0.02
		EIIL	88.6 ± 0.02	87.5 ± 0.03	86.3 ± 0.03	85.4 ± 0.02	83.2 ± 0.03	78.8 ± 0.01
		SharpDRO	<b>91.3 ± 0.01</b>	<b>90.2 ± 0.02</b>	<b>88.7 ± 0.01</b>	<b>87.3 ± 0.02</b>	<b>84.2 ± 0.02</b>	<b>84.3 ± 0.01</b>
	JPEG	JTT	88.9 ± 0.03	87.2 ± 0.04	85.3 ± 0.03	83.5 ± 0.03	81.0 ± 0.04	78.8 ± 0.03
		EIIL	89.4 ± 0.03	87.6 ± 0.03	85.3 ± 0.04	83.2 ± 0.03	81.6 ± 0.02	79.1 ± 0.01
		SharpDRO	<b>90.3 ± 0.02</b>	<b>88.2 ± 0.02</b>	<b>87.2 ± 0.02</b>	<b>85.3 ± 0.02</b>	<b>84.2 ± 0.02</b>	<b>82.3 ± 0.03</b>
CIFAR100	Gaussian	JTT	68.0 ± 0.02	65.3 ± 0.02	61.3 ± 0.01	56.3 ± 0.01	54.2 ± 0.03	51.2 ± 0.02
		EIIL	67.2 ± 0.01	66.2 ± 0.02	61.0 ± 0.02	55.8 ± 0.02	54.6 ± 0.03	52.1 ± 0.02
		SharpDRO	<b>69.6 ± 0.03</b>	<b>68.0 ± 0.02</b>	<b>63.6 ± 0.03</b>	<b>58.2 ± 0.02</b>	<b>56.5 ± 0.03</b>	<b>54.1 ± 0.03</b>
	JPEG	JTT	63.1 ± 0.03	61.0 ± 0.04	58.4 ± 0.03	56.2 ± 0.02	54.5 ± 0.04	51.9 ± 0.03
		EIIL	63.6 ± 0.02	61.5 ± 0.03	58.8 ± 0.03	<b>57.7 ± 0.02</b>	54.2 ± 0.03	52.2 ± 0.04
		SharpDRO	<b>64.2 ± 0.02</b>	<b>62.5 ± 0.03</b>	<b>60.1 ± 0.03</b>	57.6 ± 0.03	<b>57.2 ± 0.02</b>	<b>54.8 ± 0.02</b>
ImageNet30	Gaussian	JTT	87.3 ± 0.02	84.5 ± 0.02	82.3 ± 0.04	75.6 ± 0.01	72.1 ± 0.04	66.5 ± 0.02
		EIIL	<b>88.2 ± 0.02</b>	85.2 ± 0.03	81.3 ± 0.02	74.5 ± 0.02	71.5 ± 0.02	65.0 ± 0.04
		SharpDRO	87.1 ± 0.02	<b>86.9 ± 0.02</b>	<b>83.5 ± 0.03</b>	<b>78.0 ± 0.02</b>	<b>74.9 ± 0.02</b>	<b>68.4 ± 0.03</b>
	JPEG	JTT	85.4 ± 0.02	84.3 ± 0.02	82.5 ± 0.02	80.0 ± 0.02	78.8 ± 0.03	77.4 ± 0.02
		EIIL	85.3 ± 0.02	84.2 ± 0.03	83.2 ± 0.03	80.6 ± 0.02	78.4 ± 0.02	77.2 ± 0.02
		SharpDRO	<b>86.6 ± 0.03</b>	<b>85.7 ± 0.02</b>	<b>85.0 ± 0.02</b>	<b>84.3 ± 0.02</b>	<b>83.2 ± 0.03</b>	<b>82.5 ± 0.02</b>

**Distribution-Agnostic Robust Generalization** In Table 2, we can see a similar phenomenon as in Table 1 that the more severe corruptions are applied, the larger performance gains SharpDRO achieves. Especially, in the ImageNet30 dataset corrupted by “JPEG Compression”, SharpDRO shows about 1.2% performance gains upon the second-best method with severity 0, which is further increased to almost 5.1% with severity 5. Moreover, SharpDRO is general to all three corruption types, as it surpasses other methods in most cases. Therefore, the proposed method can perfectly generalize to worst-case data even without the distribution annotations.

### 4.3. Qualitative Analysis

To investigate the effectiveness of SharpDRO, we first conduct an ablation study to show that the worst-case sharpness minimization is essential for achieving generalization with robustness. Then, we utilize gradient norm, an important criterion to present training stability, to validate that our method is stable for severely corrupted distributions. Then, we analyze the hyper-parameter  $\rho$  and OOD score  $\bar{w}$  to disclose the effectiveness of sharpness minimization and worst-case data selection. Finally, another second-order method SAM [11, 37, 48, 49, 71] is investigated to discover the efficiency property of SharpDRO. All analyses are conducted using CIFAR10 with “Gaussian Noise” corruption.

Table 3. Ablation study. “w/o data selection” denotes training without worst-case data selection, which recovers SAM [11], and “w/o sharp min” indicates training without sharpness minimization, which is the same as GroupDRO [45].

Method	Corruption Severity					
	0	1	2	3	4	5
w/o data selection (SAM)	93.2	90.5	87.6	82.1	80.5	75.4
w/o sharp min (GroupDRO)	90.2	89.1	88.4	84.3	83.0	78.2
SharpDRO	92.9	91.3	90.5	88.4	86.9	84.7

**Ablation Study** By eliminating the worst-case data selection, we recover the original sharpness minimization method

SAM [11]. Then, we remove the sharpness minimization module, which is basically training via GroupDRO. The ablation results are shown in Table 3. We can see that deploying SAM on the whole training dataset can achieve improved results on the clean dataset. However, the robust performance on corrupted distributions are even worse than GroupDRO. This could be because that sharpness is easy to be dominated by principle distributions, which is misleading for generalization to small distributions. Thus, the sharpness of corrupted data would be sub-optimal. As for GroupDRO, it fails to produce a flat loss surface for worst-case data, hence cannot generalize as well as the proposed SharpDRO.

**Distributional Stability** To show our method can be stable even in the most challenging distributions, we show the gradient norm on a validation set including corruption severity from 0 to 5. As shown in Figure 4, SharpDRO not only produces the smallest norm value but also can ensure almost equal gradient norm across all corrupted distributions, which indicates that SharpDRO is the most distributionally stable method among all compared methods.

**Parameter Analysis** To understand how the scale parameter  $\rho$  affects our generalization performance, we conduct sensitivity analysis by changing this value and show the test results of different distributions. In figure 5 (a), we find an interesting discovery that as  $\rho$  increases, which indicates the perturbation magnitude  $\epsilon^*$  enlarges, would enhance the generalization of severely corrupted data but degrades the performance of slightly corrupted data. This might be because the exploration of hard distributions needs to cover wide range of neighborhood to ensure generalization. On the contrary, exploration too far on easy distributions can reach out-of-distribution, thus causing performance degradation. Therefore, for practitioners who aim to generalize on small and difficult datasets, we might be able to enhance perfor-




Figure 4. Gradient norm comparisons between different methods over all corrupted distributions.




Figure 5. (a) Sensitivity of  $\rho$  whose value is set to  $\{0.01, 0.05, 0.1, 0.5, 1, 2\}$ . (b) Distribution of the normalized OOD score  $\bar{\omega}$  on distribution  $s = 0$  to 5.

mance by aggressively setting a large perturbation scale.

**OOD Score Analysis** The OOD score is leveraged to select worst-case data for the distribution-agnostic setting. To show its effectiveness in selecting the noisy data, we plot the value distribution of OOD scores from all corrupted distributions in epoch 30 in Figure 5 (b). We can see the tendency that a severer corruption has larger OOD scores. Therefore, our OOD score is a valid criterion to select worst-case data. Note that during the training process, the worst-case data would be **gradually learned**, thus the OOD score can become **smaller**, which explains why the value distribution of our score is not as separable as OOD detection does.

**Training Efficiency Analysis** It is clear that the proposed SharpDRO method is a second-order optimization method. Hence, when compared to first-order methods such as GroupDRO and REx, computational cost is the price to pay for achieving improved generalization performance<sup>3</sup>. However, to further explore the advantage of SharpDRO compared to other second-order method, here we use SAM [11] as a competitor, and show their computational time as well as worst-case accuracy ( $s = 5$ ) in Figure 6. We can see that on all three datasets, our SharpDRO requires nearly the same time to train, and significantly outperforms the worst-case performance of SAM, owing to our efficient worst-case data selection which is vital for robust generalization against severe corruptions.

## 5. Conclusion

In this paper, we proposed a SharpDRO approach to enhance the generalization performance of DRO methods.

<sup>3</sup>Note that our method can be deployed with existing efficient sharpness-based methods [8, 9, 67, 69].

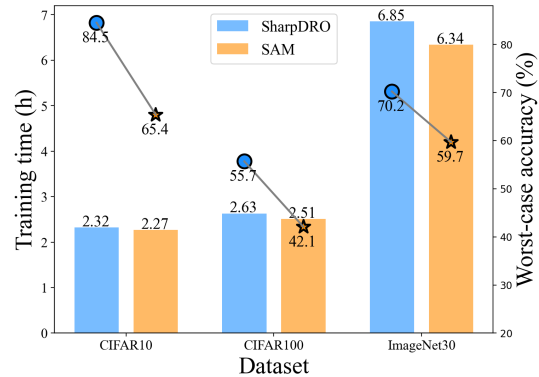


Figure 6. Efficiency comparison between SharpDRO and SAM.

Specifically, we focus on minimizing the sharpness of worst-case data to learn flat loss surfaces. As a result, SharpDRO is more robust to severe corruptions compared to other methods. Moreover, we apply SharpDRO to distribution-aware and distribution-agnostic settings and proposed an OOD detection process to select the worst-case data when the distribution index is not known. Extensive quantitative and qualitative experiments have been conducted to show that SharpDRO can deal with the most challenging corrupted distributions and achieve improved generalization results compared to well-known baseline methods.

## 6. Acknowledgements

Li Shen is supported by STI 2030-Major Projects No. 2021ZD0201405. Zhuo Huang was supported by JD Technology Scholarship for Postgraduate Research in Artificial Intelligence No. SC4103. Xiaobo Xia was supported by Australian Research Council Projects DE-190101473 and Google PhD Fellowship. Jun Yu was supported by Natural Science Foundation of China (62276242), CAAI-Huawei MindSpore Open Fund (CAAIXSJLJJ-2021-016B, CAAIXSJLJJ-2022-001A), Anhui Province Key Research and Development Program (202104a05020007), USTC-IAT Application Sci. & Tech. Achievement Cultivation Program (JL06521001Y), Sci. & Tech. Innovation Special Zone (20-163-14-LZ-001-004-01). Chen Gong was supported by NSF of China No. 61973162, NSF of Jiangsu Province No. BZ2021013, NSF for Distinguished Young Scholar of Jiangsu Province No. BK20220080, and CAAI-Huawei MindSpore Open Fund. Bo Han was supported by NSFC Young Scientists Fund No. 62006202, Guangdong Basic and Applied Basic Research Foundation No. 2022A1515011652 and RGC Early Career Scheme No. 22200720. Tongliang Liu was partially supported by Australian Research Council Projects IC-190100031, LP-220100527, DP-220102121, and FT-220100318.



## References

- [1] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019. 1, 3, 6
- [2] Aharon Ben-Tal, Dick Den Hertog, Anja De Waegenare, Bertrand Melenberg, and Gijs Rennen. Robust solutions of optimization problems affected by uncertain probabilities. *Management Science*, 59(2):341–357, 2013. 4
- [3] Fabio M Carlucci, Antonio D’Innocente, Silvia Bucci, Barbara Caputo, and Tatiana Tommasi. Domain generalization by solving jigsaw puzzles. In *CVPR*, pages 2229–2238, 2019. 3
- [4] Junbum Cha, Sanghyuk Chun, Kyungjae Lee, Han-Cheol Cho, Seunghyun Park, Yunsung Lee, and Sungrae Park. Swad: Domain generalization by seeking flat minima. In *Advances in Neural Information Processing Systems*, volume 34, pages 22405–22418, 2021. 2, 4, 6
- [5] Shiyu Chang, Yang Zhang, Mo Yu, and Tommi Jaakkola. Invariant rationalization. In *International Conference on Machine Learning*, pages 1448–1458. PMLR, 2020. 3
- [6] Pratik Chaudhari, Anna Choromanska, Stefano Soatto, Yann LeCun, Carlo Baldassi, Christian Borgs, Jennifer Chayes, Levent Sagun, and Riccardo Zecchina. Entropy-sgd: Biasing gradient descent into wide valleys. In *ICLR*, 2017. 2
- [7] Elliot Creager, Jörn-Henrik Jacobsen, and Richard Zemel. Environment inference for invariant learning. In *ICML*, pages 2189–2200. PMLR, 2021. 3, 6
- [8] 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. In *ICLR*, 2022. 8
- [9] Jiawei Du, Daquan Zhou, Jiashi Feng, Vincent YF Tan, and Joey Tianyi Zhou. Sharpness-aware training for free. *arXiv preprint arXiv:2205.14083*, 2022. 8
- [10] John Duchi and Hongseok Namkoong. Learning models with uniform performance via distributionally robust optimization. *arXiv preprint arXiv:1810.08750*, 2018. 4
- [11] Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. In *ICLR*, 2020. 2, 4, 6, 7, 8
- [12] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *ICLR*, 2015. 1
- [13] Tatsunori Hashimoto, Megha Srivastava, Hongseok Namkoong, and Percy Liang. Fairness without demographics in repeated loss minimization. In *ICML*, pages 1929–1938. PMLR, 2018. 3
- [14] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *ICLR*, 2019. 1, 6
- [15] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *ICLR*, 2017. 1, 2
- [16] Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep anomaly detection with outlier exposure. In *ICLR*, 2019. 5
- [17] Weihua Hu, Gang Niu, Issei Sato, and Masashi Sugiyama. Does distributionally robust supervised learning give robust classifiers? In *International Conference on Machine Learning*, pages 2029–2037. PMLR, 2018. 4
- [18] Zhuo Huang, Xiaobo Xia, Li Shen, Bo Han, Mingming Gong, Chen Gong, and Tongliang Liu. Harnessing out-of-distribution examples via augmenting content and style. In *ICLR*, 2023. 3
- [19] Zhuo Huang, Chao Xue, Bo Han, Jian Yang, and Chen Gong. Universal semi-supervised learning. In *NeurIPS*, volume 34, 2021. 2
- [20] Zhuo Huang, Jian Yang, and Chen Gong. They are not completely useless: Towards recycling transferable unlabeled data for class-mismatched semi-supervised learning. *IEEE Transactions on Multimedia*, 2022. 2
- [21] Atul Ingle, Trevor Seets, Mauro Buttafava, Shantanu Gupta, Alberto Tosi, Mohit Gupta, and Andreas Velten. Passive inter-photon imaging. In *CVPR*, pages 8585–8595, 2021. 1
- [22] 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 *ICLR*, 2017. 2
- [23] Minyoung Kim, Da Li, Shell X Hu, and Timothy Hospedales. Fisher sam: Information geometry and sharpness aware minimisation. In *ICML*, pages 11148–11161. PMLR, 2022. 4
- [24] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 6
- [25] 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*, pages 5815–5826. PMLR, 2021. 3, 6
- [26] Chengxi Li, Xiangyu Qu, Abhiram Gnanasambandam, Omar A Elgandy, Jiaju Ma, and Stanley H Chan. Photon-limited object detection using non-local feature matching and knowledge distillation. In *CVPR*, pages 3976–3987, 2021. 1
- [27] Ya Li, Xinmei Tian, Mingming Gong, Yajing Liu, Tongliang Liu, Kun Zhang, and Dacheng Tao. Deep domain generalization via conditional invariant adversarial networks. In *ECCV*, pages 624–639, 2018. 3
- [28] Yewen Li, Chaojie Wang, Xiaobo Xia, Tongliang Liu, Xin Miao, and Bo An. Out-of-distribution detection with an adaptive likelihood ratio on informative hierarchical vae. In *NeurIPS*, 2022. 5
- [29] Shiyu Liang, Yixuan Li, and Rayadurgam Srikant. Enhancing the reliability of out-of-distribution image detection in neural networks. In *ICLR*, 2018. 1, 2, 5
- [30] Evan Z Liu, Behzad Haghgoo, Annie S Chen, Aditi Raghunathan, Pang Wei Koh, Shiori Sagawa, Percy Liang, and Chelsea Finn. Just train twice: Improving group robustness without training group information. In *ICML*, pages 6781–6792. PMLR, 2021. 2, 5, 6
- [31] Feng Liu, Qizhou Wang, Bo Han, Tongliang Liu, Chen Gong, Gang Niu, Mingyuan Zhou, Masashi Sugiyama, et al. Probabilistic margins for instance reweighting in adversarial training. In *NeurIPS*, volume 34, pages 23258–23269, 2021. 1

- [32] Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li. Energy-based out-of-distribution detection. In *NeurIPS*, volume 33, pages 21464–21475, 2020. 2, 5
- [33] Yong Liu, Siqi Mai, Xiangning Chen, Cho-Jui Hsieh, and Yang You. Towards efficient and scalable sharpness-aware minimization. In *CVPR*, pages 12360–12370, 2022. 4
- [34] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *ICML*, pages 97–105. PMLR, 2015. 1
- [35] Florian Luisier, Thierry Blu, and Michael Unser. Image denoising in mixed poisson–gaussian noise. *IEEE Transactions on image processing*, 20(3):696–708, 2010. 1
- [36] Divyat Mahajan, Shruti Tople, and Amit Sharma. Domain generalization using causal matching. In *ICML*, pages 7313–7324. PMLR, 2021. 3
- [37] Peng Mi, Li Shen, Tianhe Ren, Yiyi Zhou, Xiaoshuai Sun, Rongrong Ji, and Dacheng Tao. Make sharpness-aware minimization stronger: A sparsified perturbation approach. In *NeurIPS*, 2022. 7
- [38] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *ICML*, pages 10–18. PMLR, 2013. 3
- [39] Hongseok Namkoong and John C Duchi. Stochastic gradient methods for distributionally robust optimization with f-divergences. In *NeurIPS*, volume 29, 2016. 2
- [40] Maher Nouiehed, Maziar Sanjabi, Tianjian Huang, Jason D Lee, and Meisam Razaviyayn. Solving a class of non-convex min-max games using iterative first order methods. In *NeurIPS*, volume 32, 2019.
- [41] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *ICCV*, pages 1406–1415, 2019. 3
- [42] Vihari Piratla, Praneeth Netrapalli, and Sunita Sarawagi. Focus on the common good: Group distributional robustness follows. In *ICLR*, 2021. 2, 3
- [43] Fengchun Qiao, Long Zhao, and Xi Peng. Learning to learn single domain generalization. In *CVPR*, pages 12556–12565, 2020. 3
- [44] 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. 6
- [45] Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. In *International Conference on Learning Representations*, 2020. 1, 2, 3, 4, 6, 7
- [46] Bernhard Schölkopf, Francesco Locatello, Stefan Bauer, Nan Rosemary Ke, Nal Kalchbrenner, Anirudh Goyal, and Yoshua Bengio. Toward causal representation learning. *Proceedings of the IEEE*, 109(5):612–634, 2021. 3
- [47] Yang Shu, Zhangjie Cao, Chenyu Wang, Jianmin Wang, and Mingsheng Long. Open domain generalization with domain-augmented meta-learning. In *CVPR*, pages 9624–9633, 2021. 3
- [48] Hao Sun, Li Shen, Qihuang Zhong, Liang Ding, Shixiang Chen, Jingwei Sun, Jing Li, Guangzhong Sun, and Dacheng Tao. Adasam: Boosting sharpness-aware minimization with adaptive learning rate and momentum for training deep neural networks. *arXiv preprint arXiv:2303.00565*, 2023. 7
- [49] Yan Sun, Li Shen, Tiansheng Huang, Liang Ding, and Dacheng Tao. FedSpeed: Larger local interval, less communication round, and higher generalization accuracy. In *ICLR*, 2023. 7
- [50] Klaus E Timmermann and Robert D Nowak. Multiscale modeling and estimation of poisson processes with application to photon-limited imaging. *IEEE Transactions on Information Theory*, 45(3):846–862, 1999. 1
- [51] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *CVPR*, pages 7167–7176, 2017. 1
- [52] Qizhou Wang, Feng Liu, Yonggang Zhang, Jing Zhang, Chen Gong, Tongliang Liu, and Bo Han. Watermarking for out-of-distribution detection. In *NeurIPS*, 2022. 2
- [53] Qizhou Wang, Junjie Ye, Feng Liu, Quanyu Dai, Marcus Kalander, Tongliang Liu, Jianye Hao, and Bo Han. Out-of-distribution detection with implicit outlier transformation. In *ICLR*, 2023. 2
- [54] Zhenyi Wang, Li Shen, Le Fang, Qiuling Suo, Tiehang Duan, and Mingchen Gao. Improving task-free continual learning by distributionally robust memory evolution. In *ICML*, pages 22985–22998. PMLR, 2022. 2
- [55] Zhenyi Wang, Xiaoyang Wang, Li Shen, Qiuling Suo, Kaiqiang Song, Dong Yu, Yan Shen, and Mingchen Gao. Meta-learning without data via wasserstein distributionally-robust model fusion. In *UAI*, pages 2045–2055. PMLR, 2022. 2
- [56] Hongxin Wei, Renchunzi Xie, Hao Cheng, Lei Feng, Bo An, and Yixuan Li. Mitigating neural network overconfidence with logit normalization. In *ICML*, 2022. 5
- [57] Dongxian Wu, Shu-Tao Xia, and Yisen Wang. Adversarial weight perturbation helps robust generalization. In *NeurIPS*, volume 33, pages 2958–2969, 2020. 2, 4
- [58] Xiaobo Xia, Jiale Liu, Jun Yu, Xu Shen, Bo Han, and Tongliang Liu. Moderate coreset: A universal method of data selection for real-world data-efficient deep learning. In *ICLR*, 2023. 1
- [59] Xiaobo Xia, Tongliang Liu, Bo Han, Nannan Wang, Mingming Gong, Haifeng Liu, Gang Niu, Dacheng Tao, and Masashi Sugiyama. Part-dependent label noise: Towards instance-dependent label noise. In *NeurIPS*, volume 33, pages 7597–7610, 2020. 1
- [60] Xiaobo Xia, Tongliang Liu, Nannan Wang, Bo Han, Chen Gong, Gang Niu, and Masashi Sugiyama. Are anchor points really indispensable in label-noise learning? In *NeurIPS*, volume 32, 2019. 1
- [61] Junchi Yang, Antonio Orvieto, Aurelien Lucchi, and Niao He. Faster single-loop algorithms for minimax optimization without strong concavity. In *AISTATS*, pages 5485–5517. PMLR, 2022.
- [62] Mengyue Yang, Furui Liu, Zhitang Chen, Xinwei Shen, Jianye Hao, and Jun Wang. Causalvae: Disentangled rep-

- resentation learning via neural structural causal models. In *CVPR*, pages 9593–9602, 2021. 3
- [63] Zhongqi Yue, Qianru Sun, Xian-Sheng Hua, and Hanwang Zhang. Transporting causal mechanisms for unsupervised domain adaptation. In *ICCV*, pages 8599–8608, 2021. 3
- [64] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *BMVC*, 2016.
- [65] Runtian Zhai, Chen Dan, Zico Kolter, and Pradeep Ravikumar. Doro: Distributional and outlier robust optimization. In *ICML*, pages 12345–12355. PMLR, 2021. 2
- [66] Xingxuan Zhang, Linjun Zhou, Renzhe Xu, Peng Cui, Zheyuan Shen, and Haoxin Liu. Towards unsupervised domain generalization. In *CVPR*, pages 4910–4920, June 2022. 3
- [67] Zhiyuan Zhang, Ruixuan Luo, Qi Su, and Xu Sun. Gasam: Gradient-strength based adaptive sharpness-aware minimization for improved generalization. *arXiv preprint arXiv:2210.06895*, 2022. 8
- [68] Yang Zhao, Hao Zhang, and Xiuyuan Hu. Penalizing gradient norm for efficiently improving generalization in deep learning. In *ICML*, 2022.
- [69] Yang Zhao, Hao Zhang, and Xiuyuan Hu. Ss-sam: Stochastic scheduled sharpness-aware minimization for efficiently training deep neural networks. *arXiv preprint arXiv:2203.09962*, 2022. 8
- [70] Yaowei Zheng, Richong Zhang, and Yongyi Mao. Regularizing neural networks via adversarial model perturbation. In *CVPR*, pages 8156–8165, 2021. 4
- [71] Qihuang Zhong, Liang Ding, Li Shen, Peng Mi, Juhua Liu, Bo Du, and Dacheng Tao. Improving sharpness-aware minimization with fisher mask for better generalization on language models. In *EMNLP*, 2022. 7