

This CVPR 2021 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

AutoDO: Robust AutoAugment for Biased Data with Label Noise via Scalable Probabilistic Implicit Differentiation

Denis Gudovskiy¹ Luca Rigazio² Shun Ishizaka³ Kazuki Kozuka³ Sotaro Tsukizawa³ ¹ Panasonic AI Lab, USA ² AIoli Labs, USA ³ Panasonic Technology Division, Japan denis.gudovskiy@us.panasonic.com luca@aiolilabs.com {ishizaka.shun, kozuka.kazuki, tsukizawa.sotaro}@jp.panasonic.com

Abstract

AutoAugment [4] has sparked an interest in automated augmentation methods for deep learning models. These methods estimate image transformation policies for train data that improve generalization to test data. While recent papers evolved in the direction of decreasing policy search complexity, we show that those methods are not robust when applied to biased and noisy data. To overcome these limitations, we reformulate AutoAugment as a generalized automated dataset optimization (AutoDO) task that minimizes the distribution shift between test data and distorted train dataset. In our AutoDO model, we explicitly estimate a set of per-point hyperparameters to flexibly change distribution of train data. In particular, we include hyperparameters for augmentation, loss weights, and softlabels that are jointly estimated using implicit differentiation. We develop a theoretical probabilistic interpretation of this framework using Fisher information and show that its complexity scales linearly with the dataset size. Our experiments on SVHN, CIFAR-10/100, and ImageNet classification show up to 9.3% improvement for biased datasets with label noise compared to prior methods and, importantly, up to 36.6% gain for underrepresented SVHN classes¹.

1. Introduction

Data augmentation (DA) plays one of the key roles in improving accuracy of deep neural networks (DNNs) [14]. DA increases size and diversity of train dataset and consequently improves generalization to test data distribution. Unfortunately, DA design requires expert domain knowledge, dataset analysis, and numerous costly experiments. In real applications with biased [31] and noisy-label data [39], the handpicking of DA becomes a challenging task.



Figure 1. Shared-policy DA dilemma: the distribution of augmented train data (dashed blue) may not match the test data (solid red) in the latent space: "2" is under-augmented, while "5" is overaugmented. As a result, prior methods cannot match the test distribution and the decision of the learned classifier $f(\theta)$ is inaccurate.

The automation of DA aims to estimate data transformation models without incurring the aforementioned difficulties. The seminal AutoAugment [4] (AA) proposes a DA policy model estimated by reinforcement learning (RL). Though AA outperforms prior classification baselines, its policy search takes thousands of GPU hours. Follow-up works [9, 16, 8] address the search complexity problem while keeping comparable accuracy results.

In this paper, we show that existing methods are not robust to train dataset distortions such as distribution bias and noisy labels. We attribute this to the policy model that *shares parameters* across all data points. We illustrate this with the toy classification task in Figure 1. The distribution of digits "2" and "5" in the train dataset (solid blue) is uneven and less diverse than test distribution (solid red) in the latent space before a linear classifier. Then, the estimated augmentation policies virtually increase train data size (dashed blue line). However, such DA model *evenly* increases diversity of all digits: digit "2" is under-augmented, while digit "5" is over-augmented. This pushes decision boundary of the linear classifier in the wrong direction. A multi-class classifier is even less robust [31], specifically, when it overfits to examples with noisy labels [39].

To overcome these limitations, we propose to estimate

¹Our code is available at github.com/gudovskiy/autodo



Figure 2. General setting: the validation dataset is selected to approximate the test data distribution $Q_{x,y}^{\text{test}}$. The train dataset is subject to the specified data distortions and unknown DA hyperparameters. Unlike prior DA models with shared policies for all train data points, our AutoDO model $g(\lambda)$ estimates hyperparameters $\lambda_{i=1...N} \in \mathbb{R}^{K \times 1}$ for each of N train data points via implicit differentiation framework.

DA hyperparameters for *each train data point*. Moreover, our model includes and jointly optimizes loss weights to capture data biases and soft-labels to address noisy labels. This reformulates the original AA shared-policy search task into a generalized automated dataset optimization (AutoDO). The objective of our AutoDO model is to match the distribution of a small clean unbiased validation dataset with a large distorted train dataset using a set of per-point hyperparameters as illustrated in Figure 2. We optimize our model with large-scale hyperparameters using implicit differentiation [21] and analytically show that it is equivalent to maximizing the Fisher information between empirical datasets. Experiments on class-imbalanced data with noisy labels show the advantages of our approach.

2. Related work

DA has been widely used in computer vision from shallow convolutional neural networks (CNNs) for MNIST [28] to deep CNNs for large-scale ImageNet [14] recognition. Other applications such as segmentation [3] and object detection [20] benefit from DA as well. Recent Mixup [40] and CutMix [37] heuristically mix images to produce better regularization at the semantic level. However, all these methods require to handpick a set of hyperparameters.

Naturally, the idea to estimate DA hyperparameters emerged in AutoAugment [4] (AA) by sequentially sampling policies and using validation accuracy as a reward to update the RL controller. However, AA requires thousands of GPU hours to achieve results superior to the baseline. To speed up policy search, Ho *et al.* [9] propose populationbased training [11] that can be efficiently parallelized across distributed CPUs or GPUs. Unlike AA, this populationbased augmentation (PBA) searches for DA policy schedules rather than individual augmentations. Recent Fast AutoAugment [16] (FAA) further extends this idea by searching for policy schedules that maximize the match between distribution of augmented train data and unaugmented validation data. Inspired by Tran *et al.* [32], Bayesian optimization is adopted by FAA to achieve density matching.

Though the methods in [9, 16] substantially decrease the complexity of AA [4], gradient-based optimization can be scaled up to even larger models and datasets. Furthermore, it is directly supported by differentiable DNNs and ML frameworks. Zhang *et al.* [41] and Lin *et al.* [17] adopt the REINFORCE gradient estimator [35], which further decreases learning time. Authors use the same policy model of [4, 16], but estimate DA hyperparameters online within the bilevel gradient optimization framework. In this framework, an inner objective minimizes conventional train loss, while an outer objective minimizes adversarial [41] or validation [17] loss. Similarly, Hataya *et al.* [8] and Li *et al.* [15] employ the DARTS gradient estimator [19] inspired by its success in a neural architecture search.

Lastly, RandAugment [5] and UniformAugment [18] propose to dramatically decrease policy search space by a reparameterization scheme and to manually select only few DA hyperparameters. These near search-free methods are almost as effective as the computationally expensive searchbased methods [4, 16]. This observation raises questions on the applicability and optimality of the search-based methods. In this paper, we show that these DA methods with the same shared-policy model perform very similarly only on undistorted train datasets. At the same time, the searchbased methods can be significantly more effective and robust in the case of biased datasets with label noise. However, prior search-based methods with the shared-policy model are ill-equipped to handle distorted datasets that is sketched in Figure 1. Motivated by this, we develop our AutoDO model with the following contributions:

- To control train dataset distribution more accurately, we introduce the AutoDO model that jointly optimizes per-point hyperparameters for DA, loss weights and soft-labels using implicit differentiation.
- We analytically show that implicit differentiation minimizes the distribution shift between validation and train data by maximizing Fisher information. We also show a connection to DARTS [19] used in [8, 15].
- Our model outperforms prior methods in image classification by up to 9.3% when the task model is trained on class-imbalanced data with label noise. More importantly, we show that it improves accuracy of underrepresented classes: by up to 36.6% for SVHN.

3. Theoretical background

3.1. Problem statement

Let $(\boldsymbol{x}, \boldsymbol{y})$ be an input-label pair where a given label $\boldsymbol{y} = \mathbf{1}_c \in \mathbb{B}^{C \times 1}$ is one-hot vector (*hard-label*) for classification task with C classes. A validation dataset $\mathcal{D}_{val} = \{(\boldsymbol{x}_i^v, \boldsymbol{y}_i^v)\}_{i \in \mathbb{M}}$ of size M approximates test data \mathcal{D}_{test} . A train dataset $\mathcal{D}_{train} = \{(\boldsymbol{x}_i, \boldsymbol{y}_i)\}_{i \in \mathbb{N}}$ of size N may have a different from \mathcal{D}_{val} distribution and noisy labels \boldsymbol{y} .

There are a vector of parameters θ for a *task model* $f(\theta)$ and a vector λ of *hyperparameters*. The latter only *implic-itly* influences the train loss $\mathcal{L}(\lambda, \theta)$ such that the optimal parameters θ^* can be written as an *implicit function* of hyperparameters $\theta^*(\lambda)$. Then, the *inner objective* optimization can be expressed as

$$\boldsymbol{\theta}^*(\boldsymbol{\lambda}) := \operatorname*{arg\,min}_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\lambda}, \boldsymbol{\theta}), \tag{1}$$

where the train loss is the empirical risk: $\mathcal{L}(\boldsymbol{\lambda}, \boldsymbol{\theta}) = \sum_{i \in \mathbb{N}} \mathcal{L}(\boldsymbol{y}_i, \hat{\boldsymbol{y}}_i)/N = \sum_{i \in \mathbb{N}} \mathcal{L}(\boldsymbol{y}_i, f(\boldsymbol{x}_i, \boldsymbol{\theta}(\boldsymbol{\lambda})))/N.$ Since test distribution $Q_{\boldsymbol{x}, \boldsymbol{y}}^{\text{test}}$ is not known, it is usually

Since test distribution $\hat{Q}_{x,y}^{\text{val}}$ is not known, it is usually replaced by an empirical distribution $\hat{Q}_{x,y}^{\text{val}}$ of the validation dataset \mathcal{D}_{val} [16]. Then, the optimal hyperparameters λ^* minimize an *outer objective* as

$$\lambda^* := \underset{\lambda}{\arg\min} \mathcal{L}_{v}^*(\lambda) = \underset{\lambda}{\arg\min} \mathcal{L}_{v}(\lambda, \theta^*(\lambda)), \quad (2)$$

where the validation dataset loss $\mathcal{L}_{v}(\boldsymbol{\lambda}, \boldsymbol{\theta}^{*}(\boldsymbol{\lambda})) = \sum_{i \in \mathbb{M}} \mathcal{L}(\boldsymbol{y}_{i}^{v}, \hat{\boldsymbol{y}}_{i}^{v})/M = \sum_{i \in \mathbb{M}} \mathcal{L}(\boldsymbol{y}_{i}^{v}, f(\boldsymbol{x}_{i}^{v}, \boldsymbol{\theta}^{*}(\boldsymbol{\lambda})))/M$, and it does not *explicitly* depend on $\boldsymbol{\lambda}$.

A solution for (1-2) problem statement can be found by *bilevel* (nested) optimization procedure [17, 15, 8], where the inner objective is the conventional training to find θ^* and the outer objective is a hyperparameter optimization (HO) to estimate λ^* . Unlike search-free methods [5, 18], HO can find nearly optimal λ^* in automated fashion. However, the solution to bilevel optimization can be computationally costly for large-scale hyperparameters. We find a scalable solution using gradient-based HO for fully-differentiable DNN models. Specifically, our AutoDO model relies on an implicit differentiation framework.

3.2. Solution using implicit differentiation

Since differentiation of (1) is known [23], we are only interested in obtaining $\partial \mathcal{L}_v / \partial \lambda$ in (2). This derivative at a point $(\lambda, \theta^*(\lambda))$ can be found using a chain rule as

$$\frac{\partial \mathcal{L}_{v}}{\partial \boldsymbol{\lambda}} = \frac{\partial \mathcal{L}_{v}}{\partial \boldsymbol{\lambda}} + \frac{\partial \mathcal{L}_{v}}{\partial \boldsymbol{\theta}^{*}(\boldsymbol{\lambda})} \frac{\partial \boldsymbol{\theta}^{*}(\boldsymbol{\lambda})}{\partial \boldsymbol{\lambda}} = \frac{\partial \mathcal{L}_{v}}{\partial \boldsymbol{\theta}^{*}(\boldsymbol{\lambda})} \frac{\partial \boldsymbol{\theta}^{*}(\boldsymbol{\lambda})}{\partial \boldsymbol{\lambda}}, \quad (3)$$

where the direct derivative $\partial \mathcal{L}_{v} / \partial \lambda = 0$, because \mathcal{L}_{v} by definition does not explicitly depend on λ . The derivative $\partial \mathcal{L}_{v} / \partial \theta^{*}(\lambda)$ can be computed using conventional differentiation. However, the derivative $\partial \theta^{*}(\lambda) / \partial \lambda$ is not known because implicit function $\theta^{*}(\lambda)$ is not defined.

Recently, Lorraine *et al.* [21] proposed to use the implicit function theorem (IFT) [2] to find $\partial \theta^*(\lambda)/\partial \lambda$. Let $S(\lambda, \theta) = \partial \mathcal{L}(\lambda, \theta)/\partial \theta : \Lambda \times \Theta \to \Theta$ be a continuously differentiable function (C^1) . For a fixed point $(\dot{\lambda}, \dot{\theta})$, if $S(\dot{\lambda}, \dot{\theta}) = 0$ and det $J_{\theta}S(\dot{\lambda}, \dot{\theta}) \neq 0$: a) there exists an implicit function $\theta = s(\lambda)$ for $||\lambda - \dot{\lambda}|| \leq r_1$ and $||\theta - \dot{\theta}|| \leq r_2$, b) the function $\theta = s(\lambda)$ is of class C^1 , and its Jacobian is defined as

$$\partial \boldsymbol{\theta}(\boldsymbol{\lambda}) / \partial \boldsymbol{\lambda} = - \left[J_{\boldsymbol{\theta}} S(\boldsymbol{\lambda}, \boldsymbol{\theta}) \right]^{-1} J_{\boldsymbol{\lambda}} S(\boldsymbol{\lambda}, \boldsymbol{\theta}).$$
 (4)

Using the definition of $S(\lambda, \theta)$ and (4), we rewrite (3) as

$$\frac{\partial \mathcal{L}_{v}}{\partial \boldsymbol{\lambda}} = -\frac{\partial \mathcal{L}_{v}}{\partial \boldsymbol{\theta}} \left[\frac{\partial^{2} \mathcal{L}}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^{T}} \right]^{-1} \frac{\partial^{2} \mathcal{L}}{\partial \boldsymbol{\theta} \partial \boldsymbol{\lambda}^{T}}.$$
 (5)

The difficulty in implicit differentiation (5) is to find inverse of Hessian $H_{\theta}^{-1} = \left[\partial^2 \mathcal{L}/(\partial \theta \partial \theta^T)\right]^{-1}$, which scales quadratically $(H_{\theta} \in \mathbb{R}^{L \times L})$ with the size of task model parameters $\theta \in \mathbb{R}^{L \times 1}$. Fortunately, ML frameworks [23] are able to calculate Jacobian-vector products to reduce memory. Computationally, H_{θ}^{-1} can be approximated by conjugate gradients [24, 25], quasi-Newton methods [1], or Neumann series [13, 21]. We rely on the latter approach.

3.3. Implicit differentiation meets density matching

One may ask what is the meaning of implicit differentiation (5) in the context of the bilevel framework (1-2)? The classification loss in (1) typically minimizes Kullback-Leibler (KL) divergence between joint train data distribution $Q_{x,y}$ with density q(x, y) and the learned model distribution $P_{x,y}(\theta)$ with $p(x, y|\theta)$ density. However, it can be shown that KL objective learns only a *conditional distribution* $p(y|x, \theta)$ of the train data distribution $Q_{y|x}$ [7]. In other words, $f(\theta)$ is a strictly *discriminative model*.

The KL loss (1) for the empirical train distribution Q_x with hard-labels y can be simplified to

$$\mathcal{L}(\boldsymbol{\lambda}, \boldsymbol{\theta}) = -\sum_{i \in \mathbb{N}} \log p(\boldsymbol{y}_i | \boldsymbol{x}_i, \boldsymbol{\theta}(\boldsymbol{\lambda})) / N.$$
 (6)

By substituting (6) into each term in (5) and rewriting them as gradients for *i*th data point, it can be shown that $u_j^{v}(\theta) = -\partial \mathcal{L}_v(j)/\partial \theta = \nabla_{\theta} \log p(y_j^{v} | x_j^{v}, \theta(\lambda))$ and $u_i(\theta) = -\partial \mathcal{L}(i)/\partial \theta = \nabla_{\theta} \log p(y_i | x_i, \theta(\lambda))$ are the Fisher scores w.r.t. validation and train data points [33], respectively. Moreover, the expectation of Hessian $-\mathbb{E}_{\hat{Q}_x}[H_{\theta}]$ is equal to Fisher information metric $\mathcal{I}_{\theta} = \sum_{i \in \mathbb{N}} u_i(\theta) u_i(\theta)^T / N$ [22].

Then, the gradient of (5) for empirical datasets simply maximizes Fisher kernel between \hat{Q}_x^{val} and \hat{Q}_x as

$$\nabla_{\boldsymbol{\lambda}} \mathbb{E}_{\hat{Q}_{\boldsymbol{x}}^{\mathrm{val}}, \hat{Q}_{\boldsymbol{x}}} \left[\mathcal{L}_{\mathrm{v}} \right] = \mathbb{E}_{\hat{Q}_{\boldsymbol{x}}^{\mathrm{val}}} \left[\boldsymbol{u}^{\mathrm{v}}(\boldsymbol{\theta}) \right] \mathcal{I}_{\boldsymbol{\theta}}^{-1} \mathbb{E}_{\hat{Q}_{\boldsymbol{x}}} \left[\boldsymbol{u}(\boldsymbol{\theta}) \boldsymbol{u}(\boldsymbol{\lambda})^{T} \right],$$
(7)

where the detailed derivations are presented in Appendix C.



Figure 3. Data pipeline of AutoDO model $g(\lambda)$ for joint dataset optimization: augmentation, loss reweighting and soft-labeling are accomplished by, correspondingly, $g_A(\lambda^A)$, $g_W(\lambda^W)$ and $g_S(\lambda^S)$ using per-point hyperparameter vectors $\lambda = [\lambda^A; \lambda^W; \lambda^S] \in \mathbb{R}^{K \times 1}$.



Figure 4. Diagram of the proposed AutoDO for loss reweighting and soft-labeling: conventional cross-entropy loss for classification is replaced by a weighted Kullback-Leibler (KL) divergence.

Therefore, the implicit differentiation in (5) finds the gradient (7) for an outer objective (2) that changes hyperparameters λ in the direction to match probability density of \mathcal{D}_{val} by \mathcal{D}_{train} density on a Riemannian manifold with a local metric given by Fisher information [6]. If replace \mathcal{I}_{θ} with the identity matrix and recall the IFT condition $\partial \mathcal{L}(\hat{\lambda}, \hat{\theta}) / \partial \hat{\theta} = 0$, (7) is equivalent to the DARTS gradient estimator [19] used in prior gradient-based methods [8, 15].

Typically, we can assume a setup in Figure 2 with a *hybrid* model where the discriminative model $f(\theta)$ is separated from an additional model $g(\lambda)$. The latter changes the data distribution $p(\boldsymbol{x}, \boldsymbol{y} | \boldsymbol{\lambda})$. Hence, $g(\boldsymbol{\lambda})$ can be viewed as a *generative model* with the outer objective (2) to minimize KL divergence between $\hat{Q}_{\boldsymbol{x},\boldsymbol{y}}^{\text{val}}$ and $\hat{Q}_{\boldsymbol{x},\boldsymbol{y}}$ distributions [10].

4. The proposed AutoDO model

We use the theoretical framework introduced in Section 3 to *jointly optimize* \mathcal{D}_{train} hyperparameters. Figure 2 shows a general data setup: \mathcal{D}_{val} is selected to approximate \mathcal{D}_{test} distribution, and \mathcal{D}_{train} is subject to data distortions and unknown hyperparameters. In addition to a task model $f(\theta)$, there is the proposed AutoDO model $g(\lambda)$. The latter is parameterized by a vector λ and can be applied to input or output of $f(\theta)$. In contrast to prior models, we propose to use per-point dataset hyperparameters $\lambda_i \in \mathbb{R}^{K \times 1}$ that form a matrix $\Lambda \in \mathbb{R}^{K \times N}$ for dataset \mathcal{D}_{train} of size N.

Figure 3 shows a diagram of our pipeline, where the AutoDO model $g(\lambda)$ is divided into application-specific

sub-models: $g_A(\lambda^A)$ for augmentation, $g_W(\lambda^W)$ for loss reweighting, and $g_S(\lambda^S)$ for soft-labeling. Hence, a generic vector λ_i is a concatenation of application-specific vectors: $\lambda_i \in \mathbb{R}^{K \times 1} = [\lambda_i^A; \lambda_i^W; \lambda_i^S] \in \mathbb{R}^{(A+W+S) \times 1}$.

A conventional data sampler uniformly generates train data points x_i that are passed to the augmentation submodel to perform $g_A(x_i, \lambda_i^A)$ in Figure 3. This augmentation block outputs a modified train data point x_i^A to the task model for training procedure. Simultaneously, soft-labels and loss weights are propagated to calculate a weighted Kullback-Leibler (KL) divergence loss $w_i \mathcal{L}(y_i^S, \hat{y}_i)$. This part is detailed in Figure 4.

4.1. Augmentation sub-model

Our augmentation sub-model contains binary probabilities $b_i \in \mathbb{B}^{A \times 1}$ and continuous magnitudes $m_i \in \mathbb{R}^{A \times 1}$ that parameterize A operations for each *i*th data point. We argue that it is less restrictive to model each data point distribution compared to finite-length policies in [4, 16].

We use a multivariate Gaussian prior for magnitudes: $m_i \sim \operatorname{rng} \frac{M}{10} \mathcal{N}(\mu, \Sigma_i)$, where M normalizes magnitudes as in [5]. The mean vector $\mu \in \mathbb{R}^{A \times 1}$ is usually consists of zeros or ones depending on operation. Only the covariance matrix $\Sigma_i \in \mathbb{R}^{A \times A}$ is parameterized by AutoDO. Although it is possible to model multivariate covariance at the expense of quadratically-growing hyperparameters, we experiment with an univariate option: $\Sigma_i = \sigma(\operatorname{diag}(\lambda_i^{A_m}))$, where $\lambda_i^{A_m} \in \mathbb{R}^{A \times 1}$ is the magnitude hyperparameter vector and σ is a conventional sigmoid nonlinearity. We initialize the magnitude hyperparameters to zeros, which sets the variances of Σ_i to be in the middle of available range (rng).

While the widely-used reparameterization trick allows to obtain gradients for magnitudes m_i with continuous Gaussian distribution, it is not trivial for the binomial probability distribution of each operation: $b_i \sim \text{Bern}(\sigma(\lambda_i^{A_b})) \in \{0, 1\}$. We use Gumbel-softmax [12] to relax backpropagation through the discrete Bernoulli distribution. Initially, we set probability of enabling the *a*th operation b_i^a to 25%.

Finally, we model a sequence $(a = 1 \dots A)$ of augmentation operations $\mathcal{O} : \mathcal{X} \to \mathcal{X}$ on the image space \mathcal{X} as

$$\boldsymbol{x}_{i}^{A}(a) = \begin{cases} \mathcal{O}(\boldsymbol{x}_{i}^{A}(a-1), m_{i}^{a}), & \text{if } b_{i}^{a} = 1\\ \boldsymbol{x}_{i}^{A}(a-1), & \text{otherwise,} \end{cases}$$
(8)

where $\boldsymbol{x}_{i}^{A}(0) = \boldsymbol{x}_{i}$ is the input to augmentation sub-model.

The number of hyperparameters in this particular submodel $\lambda_i^A = [\lambda_i^{A_m}; \lambda_i^{A_b}] \in \mathbb{R}^{2A \times 1}$ scales linearly with the number of augmentations A and train data $i \in \mathbb{N}$. In general, it is possible to jointly model augmentations not only along A-dimension using Σ_i , but also along sample dimension using e.g. Gaussian mixture model.

4.2. Reweighting and soft-labeling sub-models

Our loss reweighting sub-model multiplies *i*th train loss by a scalar: $g_W(\mathcal{L}_i, \lambda_i^W) = w_i \mathcal{L}_i$, where the weights are parameterized as: $w_i = 1.44 \times \text{softplus}(\lambda_i^W)$. Hyperparameters λ_i^W are initialized to zeros such that the weights w_i are ones at the start and their output is in range $[0:\infty)$.

Although soft-label estimation in the presence of noise has been recently studied in [30, 36], we approach it as part of joint optimization procedure with other types of applications. Moreover, we apply implicit differentiation rather than the alternating optimization as in [30, 36].

Our soft-label sub-model consists of hyperparameters $\lambda_i^S \in \mathbb{R}^{C \times 1}$ that are estimated for each label in $\mathcal{D}_{\text{train}}$ from the noise-free \mathcal{D}_{val} . Similarly to [36], we set soft-labels as follows: $y_i^S = g_S(y_i, \lambda_i^S) = \operatorname{softmax}(\lambda_i^S)$. But, unlike [36], we initialize λ_i^S at epoch e = 0 in a way to output smooth-labels [29]: $y_i^S(0) = (1 - \alpha)y_i + \alpha/C$. Our hyperparameters at the initialization step are: $\lambda_i^S(0) = (y_i - 0.5) \log(1 - C - C/\alpha)$, where $y_i = \mathbf{1}_c \in \mathbb{B}^{C \times 1}$ are potentially noisy hard-labels and α is a small constant. The α constant is usually in the range $\alpha = [0.05 - 0.2]$ [29].

Lastly, we replace the conventional asymmetric KL loss in (1) with its symmetric version. This step is necessary because we simultaneously optimize conditional distribution $p(y|x, \theta)$ and joint data distribution $p(x, y|\lambda)$.

4.3. AutoDO optimization and complexity

Our optimization procedure is described in Alg. 1 with two additional tweaks compared to the theoretical description in Section 3. First, we adopt widely-used stochastic optimization using mini-batches \mathcal{B} , which speeds up learning process at the expense of less accurate gradients. Second, the IFT in (4) requires to be at a minima point. Hence, we start HO only after reaching epoch E (line 10).

We estimate complexity of Alg. 1 in terms of forward and backward passes through a DNN. The conventional inner-objective training (lines 3-9) requires 2N passes per epoch. The outer-objective implicit differentiation in lines 12-19 require (5 + T)N passes per epoch, where T is the number of Neumann iterations. We set T = 5 using ablation study in [21] and $E = 0.5 \times \text{epochs}$. This results in a factor of $(1 + 0.5(5 + T)/2) = 3.5 \times$ increase in computation (lines 3-19) compared to only conventional training (lines 3-9). In comparison, DARTS [19] in [8, 15], without a finite-difference approximation, has T = 0 and a factor of

Algorithm 1 AutoDO bilevel optimization of $f(\theta)$, $g(\lambda)$

1:	Initialize parameters $ heta$ and hyperparameters λ
2:	for epoch = $1 \dots$ epochs do
3:	for $batch = 1 \dots batches do$
4:	sample batch $\{(\boldsymbol{x}, \boldsymbol{y})\}_{b_i \in \mathcal{B}}$ from $\mathcal{D}_{ ext{train}}$
5:	augment data $oldsymbol{x}^A = g_A(oldsymbol{x},oldsymbol{\lambda}^A)$
6:	predict $\hat{\boldsymbol{y}} = f(\boldsymbol{x}^A, \boldsymbol{\theta})$
7:	generate soft-labels $oldsymbol{y}^S = g_S(oldsymbol{y},oldsymbol{\lambda}^S)$
8:	calculate $ abla_{oldsymbol{ heta}}\left(w\mathcal{L}(oldsymbol{y}^S, \hat{oldsymbol{y}}) ight)$
9:	update θ using task optimizer
10:	if $epoch > E$ then
11:	for $batch = 1 \dots batches$ do
12:	sample $\{(\boldsymbol{x}, \boldsymbol{y})\}_{b_i \in \mathcal{B}}$ from $\mathcal{D}_{ ext{train}}$
13:	sample $\{(\boldsymbol{x}^{\mathrm{v}}, \boldsymbol{y}^{\mathrm{v}})\}_{b_{j} \in \mathcal{B}}$ from $\mathcal{D}_{\mathrm{val}}$
14:	predict $\hat{y}^{v} = f(x^{v}, \theta)$
15:	predict $\hat{\boldsymbol{y}} = f(\boldsymbol{x}^A, \boldsymbol{\theta}) = f(g_A(\boldsymbol{x}, \boldsymbol{\lambda}^A), \boldsymbol{\theta})$
16:	generate soft-labels $oldsymbol{y}^S = g_S(oldsymbol{y},oldsymbol{\lambda}^S)$
17:	calculate $\nabla_{\lambda} \mathcal{L}_{v}$ using (5), where
18:	$\mathcal{L} = w\mathcal{L}(m{y}^S, \hat{m{y}})$ and $\mathcal{L}_{ ext{v}} = \mathcal{L}_{ ext{v}}(m{y}^{ ext{v}}, \hat{m{y}}^{ ext{v}})$
19:	update λ using HO optimizer

 $(1 + 0.5 \times 5/2) = 2.25 \times$ increase in corresponding computation. Runtime estimates are compared in Appendix A.

5. Experiments

5.1. Experimental setup

We conduct SVHN, CIFAR-10/100 and ImageNet classification experiments. The code is in PyTorch [23] with the Kornia library [26] for differentiable image operations. We evaluate DA methods using both the near-perfect original train datasets and their distorted versions. To do that, we introduce a *class imbalance* ratio (IR) that scales down the number of available train images for a subset of classes. Class imbalance is defined as the ratio of images with $\{1 \dots C/2\}$ labels to $\{C/2 + 1 \dots C\}$, where *C* is the total number of classes. Next, we introduce a *label noise* ratio (NR) that defines a proportion of randomly flipped labels in the train dataset. For example, the notation IR-NR=100-0.1 for SVHN in Table 1 means that 10% (NR=0.1) of random images have randomly flipped labels and the number of images with digits $\{0 \dots 4\}$ is $100 \times$ larger than for $\{5 \dots 9\}$.

For comparison, there is a variety of gradient-free [4, 9, 16], gradient-based [17, 8, 15] and search-free [5, 18] methods described in Section 2. We narrow it down to only a single method from each type by excluding [4, 9] due to high complexity and [17, 8, 18] due to publicly unavailable code. Hence, we define the following experimental configurations: baseline with standard augmentations, optimizationfree RandAugment [5] (RAA), Bayesian Fast AutoAugment [16] (FAA), gradient-based DADA [15], and our AutoDO. As part of ablation studies, we enable AutoDO

Table 1. WRNet28-10 SVHN top-1 test error rate, $\mu_{\pm\sigma}$ %.

Alg./IR-NR	1-0.0	100-0.0	1-0.1	100-0.1
Baseline	$3.6_{\pm 0.10}$	$13.6_{\pm 0.69}$	$5.3{\scriptstyle \pm 0.27}$	$20.0{\scriptstyle \pm 1.92}$
RAA [5]	$2.7{\scriptstyle \pm 0.04}$	$10.9{\scriptstyle \pm 0.66}$	$3.4{\scriptstyle \pm 0.11}$	$13.6{\scriptstyle \pm 0.96}$
FAA [16]	$2.8{\scriptstyle \pm 0.02}$	$11.5{\scriptstyle \pm 0.32}$	$3.7{\scriptstyle\pm0.08}$	$15.3{\scriptstyle \pm 1.07}$
DADA [15]	$2.9{\scriptstyle \pm 0.03}$	$12.2{\scriptstyle \pm 0.54}$	$4.1{\scriptstyle \pm 0.13}$	$16.5{\scriptstyle\pm1.51}$
$oldsymbol{\lambda}^{A_{ ext{SHA}}}$ (ours)	$2.8_{\pm 0.10}$	$12.6{\scriptstyle\pm1.53}$	$3.0_{\pm 0.17}$	$13.7_{\pm 0.77}$
$\boldsymbol{\lambda}^{A}$ (ours)	$2.7{\scriptstyle \pm 0.09}$	$10.2{\scriptstyle \pm 0.50}$	$3.0{\scriptstyle \pm 0.07}$	$12.3{\scriptstyle \pm 0.80}$
$\boldsymbol{\lambda}^{A,W}$ (ours)	$2.8{\scriptstyle \pm 0.04}$	$6.1{\scriptstyle \pm 0.22}$	$2.8{\scriptstyle \pm 0.07}$	$8.1{\scriptstyle \pm 0.14}$
$\boldsymbol{\lambda}^{A,W,S}$ (ours)	$2.5{\scriptstyle \pm 0.04}$	$\textbf{5.3}_{\pm 0.21}$	$2.6{\scriptstyle \pm 0.05}$	$6.3{\scriptstyle \pm 0.57}$

Table 2. WRNet28-10 CIFAR-10 top-1 test error rate, $\mu_{\pm\sigma}$ %.

Alg./IR-NR	1-0.0	10-0.0	1-0.1	10-0.1
Baseline	$6.0_{\pm 0.10}$	$16.3{\scriptstyle \pm 0.54}$	$12.8{\scriptstyle \pm 0.14}$	22.8±1.09
RAA [5]	5.2 ± 0.08	$13.4{\scriptstyle \pm 0.30}$	$7.8{\scriptstyle \pm 0.16}$	$17.4{\scriptstyle \pm 0.21}$
FAA [16]	$5.0{\scriptstyle \pm 0.14}$	$13.3{\scriptstyle \pm 0.38}$	8.2 ± 0.28	$17.9{\scriptstyle \pm 0.32}$
DADA [15]	$5.5{\scriptstyle \pm 0.12}$	$14.2{\scriptstyle \pm 0.15}$	$10.6{\scriptstyle \pm 0.08}$	$20.2{\scriptstyle \pm 0.83}$
$\boldsymbol{\lambda}^{A_{ ext{SHA}}}$ (ours)	$5.9_{\pm 0.22}$	$17.1_{\pm 0.87}$	$7.6_{\pm 0.28}$	$19.4{\scriptstyle\pm1.23}$
$\boldsymbol{\lambda}^A$ (ours)	5.2 ± 0.20	$13.0_{\pm 0.47}$	$7.1{\scriptstyle \pm 0.16}$	$16.6{\scriptstyle \pm 0.20}$
$oldsymbol{\lambda}^{A,W}$ (ours)	$5.4{\scriptstyle \pm 0.17}$	$12.9{\scriptstyle \pm 0.40}$	6.8 ± 0.07	$15.1{\scriptstyle \pm 0.19}$
$\boldsymbol{\lambda}^{A,W,S}$ (ours)	$\textbf{4.9}_{\pm 0.09}$	$11.8{\scriptstyle \pm 0.35}$	$\textbf{5.8}_{\pm 0.29}$	$13.1{\scriptstyle \pm 0.16}$

features one-by-one with the following notation: $\lambda^{A_{SHA}}$ - shared (across all data points) augmentation hyperparameters, λ^{A} - per-point augmentation hyperparameters, λ^{W} - loss reweighting, and λ^{S} - soft-labeling.

The baseline configuration uses only standard augmentations such as random cropping, horizontal flipping and erasing [42]. We adopt the publicly available policy models for RAA, FAA and DADA. Using Section 4.1 format [μ , rng], the list of augmentation operations optimized by AutoDO includes: rotation - [0°, 30°], scale - [1.0, 0.5], translateX/Y - [0.0, 0.45], and shearX/Y - [0.0, 0.3]. This list has been selected from the ablation study of the most helpful operations in Cubuk *et al.* [5]. Other operations in AutoDO are applied as in the RAA model.

We use the same learning hyperparameters for all methods as specified in the Appendix A. AutoDO optimization starts after the 100th epoch for ImageNet and 50th epoch for others (parameter *E* in Alg. 1). The evaluated architectures are ResNet-18 for ImageNet, Wide-ResNet-28-10 for SVHN and CIFAR-10/100. We run each experiment four times for SVHN and CIFAR-10/100 and report the top-1 mean error rate μ and standard deviation σ on test datasets. We run only one experiment on large-scale ImageNet and report top-1 error rate. Due to a lack of test labels for ImageNet, we use a validation dataset as a test dataset.

We follow [16] and split original train dataset $\mathcal{D}_{\text{train}}$ into validation $\mathcal{D}_{\text{val}}^i$ and train $\mathcal{D}_{\text{train}}^i$ datasets using K-fold stratified shuffling [27] for *i*th experiment. The size of $\mathcal{D}_{\text{val}}^i$ is

Table 3. WRNet28-10 CIFAR-100 top-1 test error rate, $\mu_{\pm\sigma}$ %.

Alg./IR-NR	1-0.0	10-0.0	1-0.1	10-0.1
Baseline	$24.9{\scriptstyle\pm0.19}$	$45.9_{\pm 0.86}$	$36.6{\scriptstyle \pm 0.65}$	$54.5{\scriptstyle \pm 0.74}$
RAA [5]	$23.0{\scriptstyle \pm 0.15}$	$41.7{\scriptstyle\pm0.29}$	$33.7{\scriptstyle\pm0.66}$	$50.1{\scriptstyle \pm 0.52}$
FAA [16]	$22.7{\scriptstyle\pm0.23}$	$40.0{\scriptstyle \pm 0.27}$	$31.4{\scriptstyle \pm 0.25}$	$48.9{\scriptstyle \pm 0.51}$
DADA [15]	$23.1{\scriptstyle \pm 0.05}$	$41.1{\scriptstyle \pm 0.22}$	$32.4{\scriptstyle\pm0.48}$	$49.6{\scriptstyle \pm 0.44}$
$\boldsymbol{\lambda}^{A_{ ext{SHA}}}$ (ours)	$24.7{\scriptstyle\pm0.94}$	41.0±1.29	31.4±1.75	$48.9{\scriptstyle\pm1.80}$
$\boldsymbol{\lambda}^A$ (ours)	$23.1{\scriptstyle \pm 0.25}$	$40.8{\scriptstyle \pm 0.43}$	$30.6{\scriptstyle \pm 0.48}$	$47.4{\scriptstyle \pm 0.33}$
$\boldsymbol{\lambda}^{A,W}$ (ours)	$\textbf{22.2}_{\pm 0.22}$	$39.0{\scriptstyle \pm 0.30}$	$30.6{\scriptstyle \pm 0.50}$	$48.0{\scriptstyle \pm 0.38}$
$\boldsymbol{\lambda}^{A,W,S}$ (ours)	$22.4{\scriptstyle \pm 0.25}$	$36.7{\scriptstyle\pm0.32}$	$27.3{\scriptstyle \pm 0.29}$	$39.6{\scriptstyle \pm 0.06}$

32% for SVHN, 20% for CIFAR-10/100 and ImageNet of the original \mathcal{D}_{train} . After the split, we distort \mathcal{D}^i_{train} by introducing class imbalance (IR) and label noise (NR). We estimate AutoDO hyperparameters of \mathcal{D}^i_{train} using \mathcal{D}^i_{val} , train on \mathcal{D}^i_{train} and test on \mathcal{D}_{test} . Unlike [16], we never train on the original $\mathcal{D}_{train} = \mathcal{D}^i_{train} \cup \mathcal{D}^i_{val}$ (except for ablations), which leads to lower accuracy compared to ones reported in the corresponding papers due to smaller train dataset size. We choose this approach $(\mathcal{D}^i_{val} \xrightarrow{\lambda} \mathcal{D}^i_{train})$ to obtain undistorted \mathcal{D}^i_{val} and distorted \mathcal{D}^i_{train} for our experiments.

5.2. Quantitative results

Tables 1-4 present test error rates (lower is better) for the selected datasets and model configurations. We clearly see that all DA algorithms outperform the baselines with only standard augmentations. This confirms the common statement that DA improves generalization to test data.

In the case of undistorted train datasets (IR-NR=1-0.0), prior DA algorithms and our AutoDO with augmentations only (λ^A) tend to perform close to each other, if take into account σ error bars. This empirically confirms our proposition in Section 2 that search-free methods [5, 18] are as effective as search-based methods for nearly perfect datasets. The enabled additional AutoDO features along with perpoint augmentations ($\lambda^{A,W,S}$) result in moderate improvements that reach 0.2% for SVHN, 0.1% for CIFAR-10 and 0.5% for CIFAR-100 compared to the best prior methods.

In ablation study, we observe degradation in results by 0.1-1.6% for AutoDO with the shared hyperparameters ($\lambda^{A_{\text{SHA}}}$) compared to per-point model (λ^{A}). The former configuration is equivalent to a single-policy model, while, for example, a typical RAA [5], FAA [16] and DADA [15] consist of 25-125 policies. This experimentally verifies the advantage of per-point AutoDO model.

Next, we compare robustness of all methods to train dataset distortions. In the case of class imbalance only (IR=10/100, NR=0.0), AutoDO achieves at least 5.6%, 1.5%, 3.3% and 1.2% gain compared to prior methods for SVHN, CIFAR-10/100, and ImageNet, respectively. For instance, the AutoDO loss reweighting sub-model decreases

	Alg./IR-NR	10-0.0	10-0.1	
	Baseline	43.1	46.2	
	RAA [5]	42.0	45.4	
	FAA [16]	41.6	44.7	
	DADA [15]	41.9	44.9	
	$\lambda^{A,W,S}$ (ours)	40.4	44.1	
90	hM	man	when when the	
\$ 80 \$	AMM	mm	Whan	
Accura	Auto		Maria	AutoDO
60	Start	(c) baselii	FAA (b) ne (a)	infiisti (C)
T	20 40 60 80	100 12 Epoch	0 140 160	180

Table 4. ResNet18 ImageNet top-1 test error rate, %.

Figure 5. SVHN top-1 test accuracy learning curve for the models trained on a dataset with $100 \times$ class imbalance and 10% label noise: (a) baseline, (b) FAA [16], and (c) our AutoDO ($\lambda^{A,W,S}$). Our optimization starts at epoch E = 50. It prevents overfitting to the distorted train data and improves generalization to test data.

SVHN error rate from 10.2% for our model with λ^A to 6.1% for the model with $\lambda^{A,W}$. When we introduce 10% labels noise only (IR-NR=1-0.1), our improvements are, correspondingly, 0.8%, 2.0% and 4.1% for SVHN and CIFAR-10/100 compared to the best prior methods. It is mostly achieved by enabling the soft-label estimation sub-model. For example, CIFAR-100 test error rate decreased from 30.6% for our model with $\lambda^{A,W}$ to 29.4% for $\lambda^{A,W,S}$.

When we distort train data by both class imbalance and label noise (IR=10/100, NR=0.1), AutoDO achieves at least 7.3%, 4.3%, 9.3% and 0.6% gain for SVHN, CIFAR-10/100, and ImageNet, respectively. To sum up, all components of our model gradually improve classification accuracy for all datasets. The only exception is CIFAR-100, where the soft-label sub-model degrades mean accuracy by 0.2% in the case of undistorted train dataset. This can be related to suboptimal initialization constant $\alpha = 0.1$ for 100 classes. Also, we notice that the gain on ImageNet is lower (0.6-1.2%) compared to small-scale datasets. The latter might be related to a relatively shallow ResNet18 model, which has been chosen to decrease experiment time.

Figure 5 illustrates the typical learning process on distorted train data, specifically, for SVHN with $100 \times$ classimbalance and 10% label noise. It is evident that FAA [16] (b) improves accuracy compared to the baseline (a), but

Table 5. AutoDO on \mathcal{D}_{val}^i vs. \mathcal{D}_{test} , test error rate, $\mu_{\pm\sigma}$ %.

Dataset/Split	$\mathcal{D}^i_{\mathrm{val}} \xrightarrow{\boldsymbol{\lambda}} \mathcal{D}^i_{\mathrm{train}}$	$\mathcal{D}_{\text{test}} \xrightarrow{\boldsymbol{\lambda}} \mathcal{D}^{i}_{\text{train}}$
SVHN	$2.48{\scriptstyle \pm 0.04}$	$2.48{\scriptstyle \pm 0.05}$
CIFAR-10	$4.92{\scriptstyle \pm 0.09}$	5.11 ± 0.12
CIFAR-100	$22.20{\scriptstyle\pm0.22}$	$22.12{\scriptstyle \pm 0.30}$

Table 6. Split as in [16]: $\mathcal{D}_{val}^i \xrightarrow{\lambda} \mathcal{D}_{train}$, test error rate, $\mu_{\pm \sigma} \%$.

Dataset/Alg.	RAA [5]	FAA [16]	AutoDO
SVHN	$2.46{\scriptstyle \pm 0.05}$	$2.49{\scriptstyle \pm 0.05}$	$2.39{\scriptstyle \pm 0.02}$
CIFAR-10	$4.70{\scriptstyle \pm 0.13}$	$4.37 \scriptstyle \pm 0.11$	$4.56{\scriptstyle \pm 0.16}$
CIFAR-100	$21.85{\scriptstyle\pm0.21}$	$20.42{\scriptstyle\pm0.21}$	$20.49{\scriptstyle\pm0.27}$

it quickly starts to overfit to noisy labels [39] and overrepresented classes. The AutoDO optimization starts at epoch E = 50 and avoids overfitting to the distorted train data by selecting proper per-point hyperparameter vectors $\lambda^{A,W,S}$. Hence, we conclude that AutoDO is more robust to distorted data. Learning curves for train and test losses are given in the Appendix B.

Ablation study: overfitting to test data. In our data setup in Figure 2, we assume that small validation dataset approximates test data distribution $Q_{x,y}^{\text{test}}$. To verify this and to check how well the learned hyperparameters generalize to unseen test data, we replace $\mathcal{D}_{\text{val}}^{i}$ with $\mathcal{D}_{\text{test}}$ in AutoDO estimation step (line 13 in Alg. 1). Table 5 compares AutoDO results in both setups. Due to minor differences in error rates, we conclude that, indeed, $Q_{x,y}^{\text{test}} \approx Q_{x,y}^{\text{val}}$ and no significant overfitting to $\mathcal{D}_{\text{val}}^{i}$ happens.

Ablation study: training and estimating λ on all train data: $\mathcal{D}_{val}^i \xrightarrow{\lambda} \mathcal{D}_{train}$, where $\mathcal{D}_{train} = \mathcal{D}_{train}^i \cup \mathcal{D}_{val}^i$. We conduct experiments for the Lim *et al.* [16] data split: AutoDO dataset hyperparameters are estimated from \mathcal{D}_{val}^i and applied to the original \mathcal{D}_{train} . This increases train dataset size compared to experiments in Tables 1-3. Unlike [16] with explicit steps for policy merge after each data split step, we randomly sample mini-batches for \mathcal{D}_{train}^i and \mathcal{D}_{val}^i from \mathcal{D}_{train} in lines 12-13 of Alg. 1. Table 6 contains results for the original train datasets. Although error rates decrease by 0.1% for SVHN, 0.3% for CIFAR-10 and 1.9% for CIFAR-100 due to larger train datasets, there is no significant difference in results between the best prior methods and AutoDO for nearly perfect train data.

5.3. Qualitative results

We present additional evidence that our AutoDO model solves the shared-policy DA limitations sketched in Figure 1. To do that, we select three models from Table 1 for SVHN: (a) baseline with standard augmentations only, (b) shared-policy FAA [16], and (c) our per-point AutoDO model with all features ($\lambda^{A,W,S}$). These models have been



Figure 6. Confusion matrices (top) of SVHN test data and t-SNE clusters of penultimate layer embeddings (bottom) for the models trained on a dataset with 100× class imbalance and 10% label noise: (a) baseline, (b) FAA [16], and (c) our AutoDO ($\lambda^{A,W,S}$). Dots and balls represent correctly and incorrectly classified images for t-SNE visualizations, respectively. The underrepresented classes {5...9} have on average 53.5% accuracy for baseline (a), 59.6% for FAA [16] (b), and 90.1% for our method (c).

trained on a distorted dataset with 100× class imbalance and 10% label noise, and tested on $\mathcal{D}_{test}.$

First, we calculate confusion matrices of test data for each trained model. Figure 6 (top) shows that the sharedpolicy FAA [16] model improves the average top-1 accuracy compared to the baseline (from 78.4% to 83.1%), but the improvement for the class-imbalanced digits $\{5...9\}$ (from 53.5% to 59.6%) is approximately the same as the average gain. In contrast, our AutoDO model aligns the distribution of the underrepresented digits with the unbiased test data and their accuracy improves from 53.5% to 90.1% with average accuracy of 93.0%. The standard deviation of the inter-class accuracy drops from 20% for (a) and 19.2% for (b) to only 4.4% for our AutoDO model (c).

Second, we cluster test data embeddings of the trained models using the popular t-SNE [34] method. These embeddings are the output of WRNet28-10 [38] penultimate layer that precedes a linear classifier. This setup exactly matches Figure 1 setting. T-SNE clusters are shown in Figure 6 (bottom), where dots and balls represent correctly and incorrectly classified images, respectively. While FAA [16] (b) has significantly less misclassified digits (number of balls) compared to the baseline (a), it is unable to clearly separate clusters of underrepresented digits *e.g.* "6", "8" and "9". These clusters are heavily mixed due to the limitations of shared-policy DA model. Our AutoDO model solves this

issue by per-point augmentation and loss reweighting: "6", "8" and "9" clusters are much better separated, which explains 36.6% increase in accuracy. Furthermore, our softlabel estimation allows to enlarge margins between not only clusters of underrepresented digits $\{5...9\}$, but also wellrepresented digits $\{0...4\}$.

6. Conclusions

We introduced the AutoDO model that addresses the limitations of existing policy-based AutoAugment models by incorporating the set of hyperparameters independently for each train data point. We jointly optimized these large-scale hyperparameters using an elegant low-complexity implicit differentiation framework. Our probabilistic interpretation using Fisher information matched the objective of minimizing the distribution shift between small noise-free unbiased validation dataset and large distorted train dataset.

AutoDO components such as augmentation hyperparameters, loss weights and soft-labels showed their effectiveness on SVHN, CIFAR-10/100 and ImageNet classification in ablation studies. For class-imbalanced train datasets with label noise, our AutoDO improved the average accuracy and, more importantly, aligned the precision on underrepresented classes by a significant margin compared to previous methods. Hence, our AutoDO is more robust in real applications with imperfect train datasets.

References

- Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. Deep equilibrium models. In Advances in Neural Information Processing Systems, 2019. 3
- [2] Augustin-Louis Cauchy. Résumé d'un mémoir sur la mécanique céleste et sur un nouveau calcul appelé calcul des limites. *Oeuvres Complétes d'Augustun Cauchy*, 12:48–112, 1831. 3
- [3] Liang-Chieh Chen, G. Papandreou, I. Kokkinos, Kevin Murphy, and A. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018. 2
- [4] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. AutoAugment: Learning augmentation policies from data. arXiv:1805.09501, 2018. 1, 2, 4, 5
- [5] Ekin D. Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V. Le. RandAugment: Practical automated data augmentation with a reduced search space. *arXiv*:1909.13719, 2019. 2, 3, 4, 5, 6, 7
- [6] Pedro Domingos. Every model learned by gradient descent is approximately a kernel machine. *arXiv:2012.00152*, 2020.
 4
- [7] Denis Gudovskiy, Alec Hodgkinson, Takuya Yamaguchi, and Sotaro Tsukizawa. Deep active learning for biased datasets via fisher kernel self-supervision. In *Proceedings* of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2020. 3
- [8] Ryuichiro Hataya, Jan Zdenek, Kazuki Yoshizoe, and Hideki Nakayama. Faster autoAugment: Learning augmentation strategies using backpropagation. arXiv:1911.06987, 2019. 1, 2, 3, 4, 5
- [9] Daniel Ho, Eric Liang, Ion Stoica, Pieter Abbeel, and Xi Chen. Population based augmentation: Efficient learning of augmentation policy schedules. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2019. 1, 2, 5
- [10] Tommi Jaakkola and David Haussler. Exploiting generative models in discriminative classifiers. In *Proceedings of the* 11th Conference on Neural Information Processing Systems, 1999. 4
- [11] Max Jaderberg, Valentin Dalibard, Simon Osindero, Wojciech M Czarnecki, Jeff Donahue, Ali Razavi, Oriol Vinyals, Tim Green, Iain Dunning, Karen Simonyan, et al. Population based training of neural networks. *arXiv:1711.09846*, 2017.
- [12] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with Gumbel-softmax. arXiv:1611.01144, 2016. 4
- [13] Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2017.
 3
- [14] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet classification with deep convolutional neural networks. In Advances in Neural Information Processing Systems, 2012. 1, 2

- [15] Yonggang Li, Guosheng Hu, Yongtao Wang, Timothy Hospedales, Neil M. Robertson, and Yongxin Yang. DADA: Differentiable automatic data augmentation. In *Proceedings* of the European Conference on Computer Vision (ECCV), 2020. 2, 3, 4, 5, 6, 7
- [16] Sungbin Lim, Ildoo Kim, Taesup Kim, Chiheon Kim, and Sungwoong Kim. Fast AutoAugment. In Advances in Neural Information Processing Systems, 2019. 1, 2, 3, 4, 5, 6, 7, 8
- [17] Chen Lin, Minghao Guo, Chuming Li, Xin Yuan, Wei Wu, Junjie Yan, Dahua Lin, and Wanli Ouyang. Online hyperparameter learning for auto-augmentation strategy. In *The IEEE International Conference on Computer Vision (ICCV)*, 2019. 2, 3, 5
- [18] Tom Ching LingChen, Ava Khonsari, Amirreza Lashkari, Mina Rafi Nazari, Jaspreet Singh Sambee, and Mario A. Nascimento. UniformAugment: A search-free probabilistic data augmentation approach. *arXiv*:2003.14348, 2020. 2, 3, 5, 6
- [19] Hanxiao Liu, Karen Simonyan, and Yiming Yang. DARTS: Differentiable architecture search. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019. 2, 4, 5
- [20] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C. Berg. Ssd: Single shot multibox detector. In *Proceedings* of the European Conference on Computer Vision (ECCV), 2016. 2
- [21] Jonathan Lorraine, Paul Vicol, and David Duvenaud. Optimizing millions of hyperparameters by implicit differentiation. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020. 2, 3, 5
- [22] Alexander Ly, Maarten Marsman, A J Verhagen, Raoul Grasman, and Eric-Jan Wagenmakers. A tutorial on Fisher information. *Journal of Mathematical Psychology*, 2017. 3
- [23] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. In Autodiff workshop at Advances in Neural Information Processing Systems, 2017. 3, 5
- [24] Fabian Pedregosa. Hyperparameter optimization with approximate gradient. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2016. 3
- [25] Aravind Rajeswaran, Chelsea Finn, Sham Kakade, and Sergey Levine. Meta-learning with implicit gradients. Advances in Neural Information Processing Systems, 2019. 3
- [26] Edgar Riba, Dmytro Mishkin, Daniel Ponsa, Ethan Rublee, and Gary Bradski. Kornia: an open source differentiable computer vision library for PyTorch. In *Proceedings of* the Winter Conference on Applications of Computer Vision (WACV), 2020. 5
- [27] Mohammad Shahrokh Esfahani and Edward R. Dougherty. Effect of separate sampling on classification accuracy. *Bioin-formatics*, 30(2):242–250, 2013.
- [28] Patrice Y. Simard, Dave Steinkraus, and John Platt. Best practices for convolutional neural networks applied to visual document analysis. In *Proceedings of the International Conference on Document Analysis and Recognition (ICDAR)*, 2003. 2

- [29] Christian Szegedy, V. Vanhoucke, S. Ioffe, Jon Shlens, and Z. Wojna. Rethinking the inception architecture for computer vision. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016. 5
- [30] Daiki Tanaka, Daiki Ikami, Toshihiko Yamasaki, and Kiyoharu Aizawa. Joint optimization framework for learning with noisy labels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 5
- [31] Philipp Terhörst, Jan Niklas Kolf, Marco Huber, Florian Kirchbuchner, Naser Damer, Aythami Morales, Julian Fierrez, and Arjan Kuijper. A comprehensive study on face recognition biases beyond demographics. arXiv:2103.01592, 2021. 1
- [32] Toan Tran, Trung Pham, Gustavo Carneiro, Lyle Palmer, and Ian Reid. A bayesian data augmentation approach for learning deep models. In Advances in Neural Information Processing Systems. 2017. 2
- [33] Laurens van der Maaten. Learning discriminative Fisher kernels. In Proceedings of the International Conference on Machine Learning (ICML), 2011. 3
- [34] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. Journal of Machine Learning Research, 2008. 8
- [35] Ronald J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. In *Machine Learning*, 1992. 2
- [36] Kun Yi and Jianxin Wu. Probabilistic end-to-end noise correction for learning with noisy labels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 5
- [37] Sangdoo Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. CutMix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2019. 2
- [38] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In Proceedings of the British Machine Vision Conference (BMVC), 2016. 8
- [39] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017. 1, 7
- [40] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In Proceedings of the International Conference on Learning Representations (ICLR), 2018. 2
- [41] Xinyu Zhang, Qiang Wang, Jian Zhang, and Zhao Zhong. Adversarial AutoAugment. In Proceedings of the International Conference on Learning Representations (ICLR), 2020. 2
- [42] Z. Zhong, L. Zheng, Guoliang Kang, Shaozi Li, and Y. Yang. Random erasing data augmentation. In *Proceedings of the* AAAI Conference on Artificial Intelligence, 2020. 6