

---

# AutoBalance: Optimized Loss Functions for Imbalanced Data

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 Imbalanced datasets are commonplace in modern machine learning problems. The  
2 presence of under-represented classes or groups with sensitive attributes results in  
3 concerns about generalization and fairness. Such concerns are further exacerbated  
4 by the fact that large capacity deep nets can perfectly fit the training data and appear  
5 to achieve perfect accuracy and fairness during training, but perform poorly during  
6 test. To address these challenges, we propose AutoBalance, a bi-level optimization  
7 framework that automatically designs a training loss function to optimize a blend  
8 of accuracy and fairness-seeking objectives. Specifically, a lower-level problem  
9 trains the model weights, and an upper-level problem tunes the loss function by  
10 monitoring and optimizing the desired objective over the validation data. Our  
11 loss design enables personalized treatment for classes/groups by employing a  
12 parametric cross-entropy loss and individualized data augmentation schemes. We  
13 evaluate the benefits and transferability of our approach to the application scenarios  
14 of imbalanced and group-sensitive classification. Extensive empirical evaluations  
15 demonstrate the benefits of AutoBalance over state-of-the-art approaches. Our  
16 experimental findings are complemented with theoretical insights on loss function  
17 design and the benefits of the train-validation split.

## 18 Organization of Supplementary

- 19 1. Section **A** is the extended related works. Here, we provided a discussion for the related  
20 works of data augmentation and algorithm for Implicit Differentiation [49].
- 21 2. Section **B** provides the standard error of Table 1,2,3 and 4 over 5 runs as well as several  
22 extra experiments. We also extend the discussion of Fig. 3 in its caption.
- 23 3. Section **C** extends the discussion of Section 4 on the importance of validation in multi-  
24 objective learning and generalization.
- 25 4. Section **D** proves Lemma. 1 on the consistency of the parametric cross-entropy with multi-  
26 plicative adjustments.
- 27 5. Finally, Section **E** establishes the benefit of data augmentation by relating “data-augmented  
28 cross-entropy” to parametric cross-entropy loss.

## 29 1 Introduction

30 Recently, deep learning, large datasets, and the evolution of computing power have led to unprece-  
31 dented success in computer vision, and natural language processing [15, 41, 66]. This success is  
32 partially driven by the availability of high-quality datasets, built by carefully collecting a sufficient  
33 number of samples for each class. In practice, real-world datasets are frequently imbalanced and  
34 exhibit long-tailed behavior, necessitating a careful treatment of the minorities [21, 56, 24]. Indeed,

35 modern classification tasks can involve thousands of classes, so it is perhaps intuitive that some  
 36 classes should be over/under-represented compared to others. Besides class imbalance, minorities can  
 37 also appear at the feature-level; for instance, the specific values of the features of an example can vary  
 38 depending on that example’s membership in certain sensitive or protected groups, e.g. race, gender,  
 39 disabilities (see also Figure 1a). In scenarios where imbalances are induced by heterogeneous client  
 40 datasets (e.g., in the context of federated learning), addressing these imbalances can help ensure that a  
 41 machine learning model works well for all clients, rather than just those that generate the majority of  
 42 the training data. This rich set of applications motivate the careful treatment of imbalanced datasets.

43 In the imbalanced classification literature, the recurring theme is maximizing a fairness-seeking  
 44 objective, such as balanced accuracy. Unlike standard accuracy, which can be dominated by the  
 45 majorities, a fairness-seeking objective seeks to promote examples from minorities, and downweigh  
 46 examples from majorities. Here, note that there is a distinction between the test and training  
 47 objectives. While the overall goal is typically to maximize a non-differentiable objective such as  
 48 balanced accuracy on the test set, during training, we use a differentiable proxy for this, such as  
 49 weighted cross-entropy. Thus, the fundamental question of interest is:

How to design a training loss to maximize a fairness-seeking objective on the test set?

50 A classical answer to this question is to use a Bayes-consistent training loss function, so that  
 51 as the sample size grows, the training process returns the Bayes-optimal decision rule. Thus,  
 52 weighted cross-entropy (e.g., each class gets a different weight, see §4) is traditionally a good  
 53 choice for optimizing weighted accuracy objectives. Unfortunately, this intuition starts to break  
 54 down when the training problem is overparameterized, which is a common practice in deep learn-  
 55 ing: in essence, for large capacity deep nets, the training process can perfectly fit to the data, and  
 56 training loss is no longer indicative of test error. In fact, recent works [8, 40] show that weighted  
 57 cross-entropy has minimal benefit to balanced accuracy, and instead alternative methods based on  
 58 margin adjustment can be effective (namely, by ensuring that minority classes are further away from  
 59 decision boundary). These ideas led to the development of a parametric cross-entropy function  
 60  $\ell(y, f(\mathbf{x})) = w_y \log \left( 1 + \sum_{k \neq y} e^{l_k - l_y} \cdot e^{\Delta_k f_k(\mathbf{x}) - \Delta_y f_y(\mathbf{x})} \right)$ , which allows for a personalized treat-  
 61 ment of the individual classes via the design parameters  $(w_k, l_k, \Delta_k)_{k=1}^K$  [8, 40, 56, 37, 71]. Here,  $w_k$   
 62 is the classical weighting term whereas  $l_k$  and  $\Delta_k$  are additive and multiplicative logit adjustments.  
 63 However, despite these developments, it is unclear how such parametric cross-entropy functions can  
 64 be tuned for use for different fairness objectives, for example to tackle class or group imbalances. The  
 65 works by [8, 56] provide theoretically-motivated choices for  $(w_k, l_k)$ , while [40] argues that  $(w_k, l_k)$   
 66 is not as effective as  $\Delta_k$  in the interpolating regime of zero training error. However, these works  
 67 do not provide an optimized loss function that can be systematically tailored for different fairness  
 68 objectives, such as balanced accuracy common in class imbalanced scenarios, or equal opportunity  
 69 [24, 18] which is relevant in group-sensitive settings.

70 In this work, we address these shortcomings by designing the loss function within the optimization *in*  
 71 *a principled fashion*, to handle different fairness-seeking objectives. Our main idea is to use bi-level  
 72 optimization, where the model weights are optimized over the training data, and the loss function is  
 73 automatically tuned by monitoring the validation loss. Our core intuition is that unlike training data,  
 74 the validation data is difficult to fit and will provide a consistent estimator of the test objective.

75 **Contributions.** Based on this high-level idea, this paper takes a step towards a systematic treatment  
 76 of imbalanced learning problems with contributions along several fronts: state-of-the-art performance,  
 77 data augmentation, applications to different imbalance types, and theoretical intuitions. Specifically:

- 78 • We introduce *AutoBalance* —a bilevel optimization framework— that designs a fairness-seeking  
 79 loss function by jointly training the model and the loss function hyperparameters in a systematic  
 80 way (Figure 1b, Section 4). To further improve the test performance, our design also incorporates  
 81 *data augmentation policies* personalized to subpopulations (classes or groups). We demonstrate the  
 82 benefits of AutoBalance when optimizing various fairness-seeking objectives over the state-of-the-art,  
 83 such as logit-adjustment (LA) [56] and label-distribution-aware margin (LDAM) [8] losses.
- 84 • Extensive experiments provide several takeaways (Section 5). First, the impact of individual design  
 85 parameters in the loss function is revealed, with the additive adjustment  $l_k$  and multiplicative adjust-  
 86 ment  $\Delta_k$  synergistically improving the fairness objective. Second, personalized data augmentation  
 87 can further improve the performance over a single generic augmentation policy. Third, the loss  
 88 functions designed by AutoBalance are transferable across datasets.

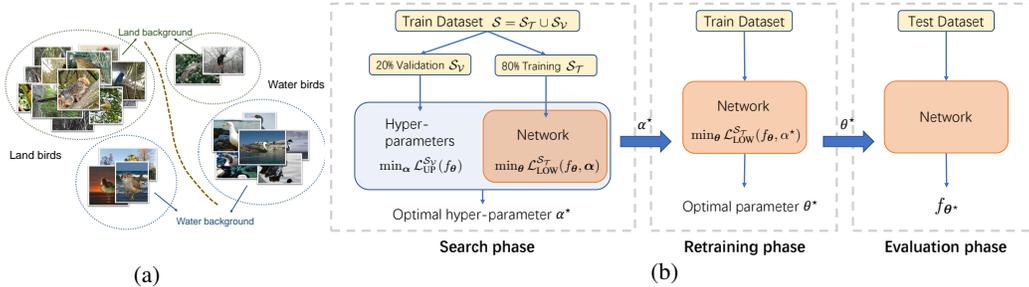


Figure 1: (a) Example group-imbalance on the Waterbirds dataset [67, 78]. Groups correspond to the distinct background types, while classes are distinct bird types. (b) Framework overview. The search phase conducts a bilevel optimization to design the optimal training loss function parameterized by  $\alpha^*$  by minimizing the validation loss, using a train-validation split (e.g. 80%-20%). The retrain phase uses the original training data and  $\alpha^*$  to obtain the optimal model parameters  $\theta^*$ . The evaluation phase predicts the test data using  $\theta^*$ .

89 • Beyond class imbalance, we consider applications of loss function design to the group-sensitive  
 90 setting (Section 6). Our experiments show that AutoBalance consistently outperforms various  
 91 baselines, leading to a more efficient Pareto-frontier of accuracy-fairness tradeoffs.

## 92 2 Related Work

93 Our work relates to imbalanced classification, fairness, bilevel optimization, and data augmentation.  
 94 Below we focus on the former three and defer the extended discussion to the supplementary.

95 **Long-tailed learning.** Learning with long-tailed data has received substantial interest historically,  
 96 with classical methods focusing on designing sampling strategies, such as over- or under-sampling [42,  
 97 63, 9, 43, 1, 59, 72, 5, 53]. Several loss re-weighting schemes [57, 55, 28, 13, 6, 36] have been  
 98 proposed to adjust weights of different classes or samples during training. Another line of work [76,  
 99 39, 56] focuses on post-hoc correction. More recently, several works [47, 37, 17, 36, 8, 13, 56, 71]  
 100 develop more refined class-balanced loss functions (e.g. (4.1)) that better adapt to the training data.  
 101 In addition, several works [32, 79] point out that separating the representation learning and class  
 102 balancing can lead to improvements. In this work, our approach is in the vein of class-balanced  
 103 losses; however, rather than fixing a balanced loss function (e.g. based on the class probabilities in  
 104 the training dataset), we employ our Algorithm 1 to automatically guide the loss design.

105 **Group-sensitive and Fair Learning.** The group-sensitive learning aims to ensure the fairness of the  
 106 classifier under setups where there exists under-represented groups (e.g., gender, race). [7, 24, 74, 68]  
 107 propose several fairness metrics as well as insightful methodologies. A line of research [4, 20]  
 108 optimize the worst-case loss over the test distribution and further applications motivate (label, group)  
 109 metrics such as equality of opportunity [24, 18] (also recall DEO (3.1)). [61] discusses group-sensitive  
 110 learning in an over-parameterized regime and proposes that strong regularization ensures fairness.  
 111 Closer to our work, [61, 40] also study Waterbirds dataset. Compared to the regularization-based  
 112 approach of [61], we explore a parametric loss design (which is inspired from [40]) to optimize  
 113 fairness-risk over validation. [18] proposes methods and statistical guarantees for fair empirical  
 114 risk minimization. A key observation of our work is that, such guarantees based on training-only  
 115 optimization can be vacuous in the overparameterized regime. Thus, using train-validation split  
 116 (e.g. our Algo 1) is critical for optimizing fairness metrics more reliably. This is verified by the  
 117 effectiveness of our approach in the evaluations of Section 6.

118 **Bilevel Optimization.** Classical approaches [65] for hyper-parameter optimization are typically based  
 119 on derivative-free schemes, including random search [69] and reinforcement learning [81, 3, 70, 77].  
 120 Recently, a growing line of works focus on differentiable algorithms that are often faster and can  
 121 scale up to millions of parameters [49, 52, 62, 30, 51]. These techniques [48, 38, 80, 50] with  
 122 continuous relaxations have shown significant success in neural architecture search, learning rate  
 123 scheduling, regularization, etc. These methods are typically formulated as a bilevel optimization  
 124 problem: the upper and lower optimizations minimize the validation and training losses, respectively.  
 125 Some theoretical guarantees (albeit restrictive) are also available [11, 22, 2, 58]. Different from  
 126 these, our work focuses on principled design of training loss function to optimize fairness-seeking  
 127 objectives for imbalanced data. Here, a key algorithmic distinction (e.g. compared to architecture  
 128 search) is that, our loss function design is only used during optimization and not during inference.

129 This leads to a more sophisticated hyper-gradient and necessitates additional measures to ensure  
 130 stability of our approach (see Algo 1).

### 131 3 Problem Setup

132 Let  $[K]$  denote the set  $\{1, \dots, K\}$ . Let  $1_E$  be the indicator function of event  $E$ . Suppose we have a  
 133 dataset  $\mathcal{S} = (\mathbf{x}_i, y_i, g_i)_{i=1}^n$  sampled i.i.d. from a distribution  $\mathcal{D}$  with input space  $\mathcal{X}$ ,  $K$  classes and  $G$   
 134 groups. For an example  $(\mathbf{x}, y, g)$ ,  $\mathbf{x} \in \mathcal{X}$  is the input features,  $y \in [K]$  is the output label, and  $g \in [G]$   
 135 is the group membership. Let  $f : \mathcal{X} \rightarrow \mathbb{R}^K$  be a model that outputs a distribution over classes and let  
 136  $\hat{y}_f(\mathbf{x}) = \arg \max_{i \in [K]} f(\mathbf{x})_i$ . Standard classification error is defined as  $\mathcal{E}(f) = \mathbb{P}_{\mathcal{D}}[y \neq \hat{y}_f(\mathbf{x})]$ . Let  
 137  $\ell(y, \hat{y})$  be a differentiable proxy for 0-1 loss (specifically cross-entropy). We similarly denote

$$\text{Population risk: } \mathcal{L}(f) = \mathbb{E}_{\mathcal{D}}[\ell(y, \hat{y}_f(\mathbf{x}))], \quad \text{Empirical risk: } \mathcal{L}^{\mathcal{S}}(f) = \frac{1}{n} \sum_{i=1}^n \ell(y_i, \hat{y}_f(\mathbf{x}_i)).$$

138 For  $(\mathbf{x}, y, g) \sim \mathcal{D}$ , define the class, group, and (class, group) frequencies as follows

$$\pi_k = \mathbb{P}_{\mathcal{D}}(y = k), \quad \bar{\pi}_j = \mathbb{P}_{\mathcal{D}}(g = j), \quad \text{and} \quad \pi_{k,j} = \mathbb{P}_{\mathcal{D}}(y = k, g = j), \quad \text{for } (k, j) \in [K] \times [G].$$

139 Similarly, given sample  $(\mathbf{x}, y, g)$  from  $\mathcal{D}$ , let  $\mathcal{D}_k$  be the distribution conditioned on class  $y = k$ , let  $\bar{\mathcal{D}}_j$   
 140 be the distribution conditioned on group  $g = j$ , let  $\mathcal{D}_{k,j}$  be the distribution of  $(\mathbf{x}, y, g)$  conditioned  
 141 on  $y = k$  and  $g = j$ . We say that a problem has class (or group) imbalance if majority class (or  
 142 group) is substantially more frequent than minority. We introduce

$$\text{Class-conditional risk: } \mathcal{L}_k(f) = \mathbb{E}_{\mathcal{D}_k}[\ell(y, \hat{y}_f(\mathbf{x}))], \quad \text{Balanced risk: } \mathcal{L}_{\text{bal}}(f) = \frac{1}{K} \sum_{k=1}^K \mathcal{L}_k(f).$$

143 Similarly, group and (class, group)-conditional risks are denoted via  $\bar{\mathcal{L}}_j(f)$  and  $\mathcal{L}_{k,j}(f)$  respectively.  
 144 We restrict our attention to the following applications that can benefit from our approach.

145 ✓ **Setting A: Imbalanced classes.** This occurs when class frequencies differ, i.e.,  $\max_{i \in [K]} \pi_i \gg$   
 146  $\min_{i \in [K]} \pi_i$ . In this setting, we ignore group membership and focus on classes. Instead of standard  
 147 accuracy, we will optimize a class-balanced error  $\mathcal{E}_{\text{bal}}(f)$  by designing a class-personalized loss.

148 ✓ **Setting B: Imbalanced groups.** This occurs when group or (class, group) frequencies differ,  
 149 i.e.,  $\max_{j \in [G]} \bar{\pi}_j \gg \min_{j \in [G]} \bar{\pi}_j$  or  $\max_{(k,j)} \pi_{k,j} \gg \min_{(k,j)} \pi_{k,j}$ . Specifically, in the fairness  
 150 literature, groups represent sensitive or protected attributes. A typical goal is ensuring that the  
 151 prediction of the model is independent of these attributes. While many fairness metrics exist, in this  
 152 work, we focus on the Difference of Equal Opportunity (DEO) [24, 18]. Our evaluations also focus  
 153 on binary classification (with labels denoted via  $\pm$ ) and two groups ( $K = G = 2$ ). With this setup,  
 154 the DEO risk is defined as  $\mathcal{L}_{\text{DEO}}(f) = |\mathcal{L}_{+,1}(f) - \mathcal{L}_{+,2}(f)|$ . When both classes are equally relevant  
 155 (rather than  $y = +1$  implying a semantically positive outcome), we use the following definition:

$$\mathcal{L}_{\text{DEO}}(f) = |\mathcal{L}_{+,1}(f) - \mathcal{L}_{+,2}(f)| + |\mathcal{L}_{-,1}(f) - \mathcal{L}_{-,2}(f)|. \quad (3.1)$$

### 156 4 Loss Function Design and Proposed Method

157 Our main goal in this paper is automatically designing loss functions to optimize target objectives  
 158 for imbalanced learning (e.g., Settings A and B). We will employ a parametrizable family of loss  
 159 functions that can be tailored to the needs of different classes or groups. Cross-entropy variations  
 160 have been proposed by [47, 37, 17] to optimize balanced objectives. Our design space will utilize  
 161 recent works by [40, 71], who introduce Vector Scaling (VS)-loss and Class-Dependent Temperatures  
 162 (CDT). Specifically, for Setting A<sup>1</sup> we build on the following loss function parametrized by three  
 163 vectors  $\mathbf{w}, \mathbf{l}, \mathbf{\Delta} \in \mathbb{R}^K$ :

$$\ell(y, f(\mathbf{x})) = w_y \log \left( 1 + \sum_{k \neq y} e^{l_k - l_y} \cdot e^{\Delta_k f_k(\mathbf{x}) - \Delta_y f_y(\mathbf{x})} \right). \quad (4.1)$$

164 This loss is same as VS-loss of [40], which also contains  $\mathbf{w}, \mathbf{l}, \mathbf{\Delta} \in \mathbb{R}^K$ . However, the denominator  
 165 of VS-loss uses  $\Delta_y$  whereas we use  $\Delta_k$  similar to CDT of [71]. Here,  $w_y$  enables conventional  
 166 weighted CE and  $l_y$  is the additive adjustment to the logits. Recently [56] proposed a Fisher consistent  
 167 logit adjustment (LA) loss parameterized by  $w_y$  and  $l_y$ ; we make the following related observation.

<sup>1</sup>Discussion of Setting B (group-imbalance) is deferred to Section 6, however the main ideas are similar.

---

**Algorithm 1:** Bilevel Optimization for AutoBalance

---

**Input:** Model  $f_\theta$  with weights  $\theta$ , dataset  $\mathcal{S} = \mathcal{S}_\mathcal{T} \cup \mathcal{S}_\mathcal{V}$ , step sizes  $\eta_\alpha$  &  $\eta_\theta$ , # iterations  $t_2 > t_1$

- 1 Initialize  $\alpha$  with  $\ell_{\text{UP}}(\cdot) = \ell_{\text{LOW}}(\cdot; \alpha)$  // Consistent initialization
- 2 Train  $\theta$  for  $t_1$  iterations ( $\alpha$  is fixed) // Warm-up training
- 3 **for**  $i \leftarrow t_1$  **to**  $t_2$  **do**
- 4     Sample training batch  $\mathcal{B}_\mathcal{T}$  from  $\mathcal{S}_\mathcal{T}$ ;  
      // Apply class-personalized augmentation  $\mathcal{B}_\mathcal{T} \leftarrow \mathcal{A}(\mathcal{B}_\mathcal{T})$  (implicitly)
- 5      $\theta \leftarrow \theta - \eta_\theta \nabla_\theta \mathcal{L}_{\text{LOW}}^{\mathcal{B}_\mathcal{T}}(f_\theta; \alpha)$
- 6     Sample validation batch  $\mathcal{B}_\mathcal{V}$  from  $\mathcal{S}_\mathcal{V}$ ;
- 7     Compute hyper-gradient  $\nabla_\alpha \mathcal{L}_{\text{UP}}^{\mathcal{B}_\mathcal{V}}(f_\theta)$  //via Implicit Differentiation e.g. [49]
- 8      $\alpha \leftarrow \alpha - \eta_\alpha \nabla_\alpha \mathcal{L}_{\text{UP}}^{\mathcal{B}_\mathcal{V}}(f_\theta)$  // Update loss function hyper-parameters
- 9 **end**
- 10 Set  $\alpha_* \leftarrow \alpha$ ,  $\mathcal{S}_\mathcal{T} \leftarrow \mathcal{S}$ , reset weights  $\theta$
- 11 Train  $\theta$  for  $t_2$  iterations using  $\alpha_*$  // Final training over full  $\mathcal{S}$  using  $\alpha_*$

**Result:** The final model  $\theta_* \leftarrow \theta$  and hyper-parameters  $\alpha_*$

---

168 **Lemma 1** Parametric loss function (4.1) is not consistent for standard or balanced errors if there  
169 are distinct multiplicative adjustments i.e.  $\Delta_i \neq \Delta_j$  for some  $i, j \in [K]$ .

170 While consistency is a desirable property, it is intuitively more critical during the earlier phase of  
171 the training where the training risk is more indicative of the test risk. In the interpolating regime of  
172 zero-training error, [40] shows that  $w, l$  can be ineffective and multiplicative  $\Delta$ -adjustment can be  
173 more favorable. Our algorithm will be initialized with a consistent weighted-CE; however, we will  
174 allow the algorithm to automatically adapt to the interpolating regime by tuning  $l$  and  $\Delta$ .

175 **Proposed training loss function.** For our algorithm, we will augment (4.1) with *data augmentation*  
176 that can be *personalized to distinct classes*. Let us denote the data augmentation policies by  $\mathcal{A} =$   
177  $(\mathcal{A}_y)_{y=1}^K$  where each  $\mathcal{A}_y$  stochastically augments an input example with label  $y$ . Additionally, we  
178 clamp  $\Delta_i$  with the sigmoid function  $\sigma$  to limit its range to (0,1) to ensure non-negativity. To this end,  
179 our loss function for the lower-level optimization (over training data) is as follows:

$$\ell_{\text{LOW}}(y, \mathbf{x}, f; \alpha) = -\mathbb{E}_{\mathcal{A}} \left[ w_y \log \left( \frac{e^{\sigma(\Delta_y) f_y(\mathcal{A}_y(\mathbf{x})) + l_y}}{\sum_{i \in [K]} e^{\sigma(\Delta_i) f_i(\mathcal{A}_y(\mathbf{x})) + l_i}} \right) \right]. \quad (4.2)$$

180 Here,  $\alpha$  is the set of hyperparameters of the loss function that we wish to optimize, specifically  
181  $\alpha = [w, l, \Delta, \text{param}(\mathcal{A})]$ .  $\text{param}(\mathcal{A})$  is the parameterization of the augmentation policies  $(\mathcal{A}_y)_{y \in [K]}$ .

182 Personalized data augmentation (PDA). Remark-  
183 able benefits of data augmentation techniques  
184 provide a natural motivation to investigate whether  
185 one can benefit from learning class-personalized  
186 augmentation policies. To explain our intuition,  
187 consider a spherical augmentation strategy where  
188  $\mathcal{A}_y(\mathbf{x})$  samples a vector uniformly from an  $\ell_2$ -ball of  
189 radius  $\varepsilon_y$  around  $\mathbf{x}$ . As visualized in Figure 2 for a  
190 linear classifier, if the augmentation strengths of both  
191 classes are equal, the max-margin classifier is not  
192 affected by the application of the data augmentation  
193 and remains identical. Thus, augmentation has no  
194 benefit. However by applying a stronger augmenta-  
195 tion on minority, the decision boundary is shifted to protect minority which can provably benefit the  
196 balanced accuracy [40, 8]. The following intuitive observation links the PDA to parametric loss (4.1).

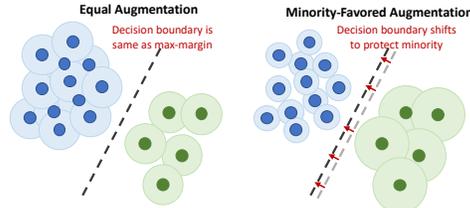


Figure 2: Data augmentation can shift the decision boundary to benefit the minority class by providing a larger margin. Lemma 2 establishes an equivalence between spherical data augmentation and parametric cross-entropy loss.

197 **Lemma 2** Consider a binary classification task with labels 0 and 1 and a linearly separable training  
198 dataset. For any parametric loss (4.1) choices of  $(l_i, \Delta_i, w_i)_{i=0}^1$ , there exists spherical augmentation  
199 strengths for minority/majority classes so that, without regularization, **optimizing the logistic loss**  
200 **with personalized augmentations returns the same classifier as optimizing (4.1).**

201 This Lemma is similar in flavor to the result of [33], which considers a larger uncertainty set  
202 around the minority class. As discussed in the supplementary materials, Lemma 2 is relevant for

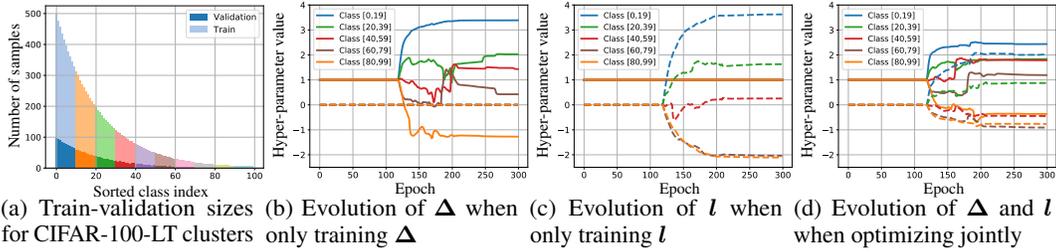


Figure 3: (a) Visualizing class clustering and train-validation split. (b), (c), (d) Evolution of loss function parameters  $l$ ,  $\Delta$  over epochs for CIFAR-100-LT where solid curves and dashed curves corresponds to  $\Delta$  and  $l$  respectively. We display average value of 20 classes for better visualization. Based on theory, the minority classes should be assigned a larger margin. During the initial 120 epochs, we use weighted cross-entropy training and AutoBalance kicks in after epoch 120. Observe that, AutoBalance does indeed learn larger parameters ( $l_y$ ,  $\Delta_y$ ) for minority class clusters (each containing 20 classes) consistent with theoretical intuition. In all Figures (b), (c), (d), by the end of training, the colors are ordered according to the class frequency. However, when  $\Delta_y$  is trained jointly with  $l_y$  (Fig (d)), the training is more stable compared to training  $\Delta_y$  alone (Fig (b)). Thus, besides its accuracy benefits in Table 1,  $l_y$  also seems to have optimization benefits.

203 the overparameterized regime whereas the approach of [33] is ineffective for separable data [54].  
 204 Algorithmically, the augmentations that we consider are much more flexible than the  $\ell_p$ -balls of  
 205 [33] and our experiments showcase the value of our approach in state-of-the-art multiclass settings.  
 206 Finally, we note that, the (theoretical) benefits of PDA can go well-beyond Lemma 2 by leveraging  
 207 the invariances [10, 14] (via rotation, translation).

#### 208 4.1 Proposed Bilevel Optimization Method

209 We formulate the loss function design as a bilevel optimization over  $\alpha$  and a hypothesis set  $\mathcal{F}$ .  
 210 Split the dataset  $\mathcal{S}$  into training  $\mathcal{S}_{\mathcal{T}}$  and validation  $\mathcal{S}_{\mathcal{V}}$  sets with  $n_{\mathcal{T}}$  and  $n_{\mathcal{V}}$  examples respectively.  
 211 The upper-level variable  $\alpha$  aims to minimize a desired fairness-seeking objective  $\ell_{\text{UP}}(y, \hat{y})$ , and the  
 212 lower-level hypothesis  $f \in \mathcal{F}$  aims to minimize the training loss (4.2) as follows:

$$\min_{\alpha} \mathcal{L}_{\text{UP}}^{\mathcal{S}_{\mathcal{V}}}(f_{\alpha}) \quad \text{WHERE} \quad f_{\alpha} = \arg \min_{f \in \mathcal{F}} \mathcal{L}_{\text{LOW}}^{\mathcal{S}_{\mathcal{T}}}(f; \alpha) := \frac{1}{n_{\mathcal{T}}} \sum_{i=1}^{n_{\mathcal{T}}} \ell_{\text{LOW}}(y_i, \mathbf{x}_i, f; \alpha). \quad (4.3)$$

213 Here,  $\mathcal{L}_{\text{UP}}$  is obtained as the empirical average of  $\ell_{\text{UP}}$ . The function  $\ell_{\text{UP}}$  is weighted cross-entropy,  
 214 chosen to be a consistent proxy for the desired accuracy objective  $\mathcal{E}_{\text{UP}}$ . For instance, if  $\mathcal{E}_{\text{UP}}$  is  
 215 a superposition of standard and balanced errors  $\mathcal{E}_{\text{UP}} = (1 - \lambda)\mathcal{E} + \lambda\mathcal{E}_{\text{bal}}$ , then similarly  $\mathcal{L}_{\text{UP}} =$   
 216  $(1 - \lambda)\mathcal{L} + \lambda\mathcal{L}_{\text{bal}}$ . Algorithm 1 summarizes our approach and highlights the key components. The  
 217 training loss  $\ell_{\text{LOW}}(\cdot; \alpha)$  is also initialized as a consistent proxy for  $\mathcal{E}_{\text{UP}}$  (e.g. same as  $\ell_{\text{UP}}$ ).

218 **Implicit Differentiation and Warm-up training.** For a loss function parameter  $\alpha$ , the hyper-  
 219 gradient can be written via the chain-rule  $\frac{\partial \mathcal{L}_{\text{UP}}(\theta^*)}{\partial \alpha} = \frac{\partial \mathcal{L}_{\text{UP}}(\theta^*)}{\partial \theta^*} \frac{\partial \theta^*}{\partial \alpha}$ . Here,  $\theta^*$  is the solution of the  
 220 lower-level problem. We note that  $\partial \mathcal{L}_{\text{UP}} / \partial \alpha = 0$  since  $\alpha$  does not appear within the upper-level  
 221 loss. Also observe that  $\partial \mathcal{L}_{\text{UP}}(\theta^*) / \partial \theta^*$  can be directly computed by taking the gradient. To compute  
 222  $\partial \theta^* / \partial \alpha$ , we follow the recent work by [49] and employ the Implicit Function Theorem (IFT). If  
 223 there exists a fixed point  $(\theta^*, \alpha^*)$  that satisfies  $\partial \mathcal{L}_{\text{LOW}}(\theta^*, \alpha^*) / \partial \theta = 0$  and regularity conditions  
 224 are satisfied, then around  $\alpha^*$ , there exists a function  $\theta(\alpha)$  such that  $\theta(\alpha^*) = \theta^*$  and we also  
 225 have  $\frac{\partial \theta}{\partial \alpha} = (\frac{\partial^2 \mathcal{L}_{\text{LOW}}}{\partial \theta^2})^{-1} \frac{\partial^2 \mathcal{L}_{\text{LOW}}}{\partial \theta \partial \alpha}$ . However, directly computing inverse Hessian  $(\frac{\partial^2 \mathcal{L}_{\text{LOW}}}{\partial \theta^2})^{-1}$  is usually  
 226 time consuming or even impossible for modern neural networks which have millions of parameters.  
 227 To compute the hyper-gradient while avoiding extensive computation, we approximate the inverse  
 228 Hessian via the Neumann series, which is widely used for inverse Hessian estimation [45, 49]. Finally,  
 229 the warm-up phase of our method (Line 2 of Algo. 1) is essential to guarantee that the IFT assumption  
 230  $\frac{\partial \mathcal{L}_{\text{LOW}}(\theta^*, \alpha^*)}{\partial \theta} = 0$  is approximately satisfied.

231 **Why is train-validation split critical?** It is well-understood that large capacity neural networks  
 232 can perfectly fit and achieve 100% training accuracy [19, 82, 31]. This also implies that over the  
 233 training data, different accuracy metrics or fairness constraints can be perfectly satisfied (e.g. 100%  
 234 balanced accuracy, and 0% DEO). To truly find a model that lies on the Pareto-front of the (accuracy,

Method	CIFAR-10-LT	CIFAR100-LT	ImageNet-LT	iNaturalist
Cross-Entropy	30.33	62.69	60.81	39.76
LDAM loss [8]	26.45	59.40	58.14	35.42
LA loss ( $\tau = 1$ ) [56]	23.32	58.92	55.60	34.37
Algo. 1: $\alpha \leftarrow \tau$ of LA loss	21.75	58.76	55.16	34.18
Algo. 1: $\alpha \leftarrow \mathbf{l}$	22.86	58.73	55.39	34.42
Algo. 1: $\alpha \leftarrow \mathbf{\Delta}$	22.52	58.49	54.72	34.14
Algo. 1: $\alpha \leftarrow \mathbf{\Delta \& l}$	21.39	56.77	52.90	33.16
Algo. 1: $\alpha \leftarrow \mathbf{\Delta \& l}$ , LA init	<b>21.16</b>	<b>56.71</b>	<b>51.96</b>	<b>33.09</b>

Table 1: Evaluations of balanced accuracy on long-tailed data. Algo. 1 with  $\mathbf{\Delta \& l}$  design space and LA initialization (bottom row) outperforms other baselines, across various datasets.

235 fairness) tradeoff, the optimization procedure should (approximately) evaluate on the population loss.  
236 To this end, the validation phase provides this crucial population loss-proxy in the overparameterized  
237 setting where training error can be vacuous. Following the literature [34, 35], the intuition is that,  
238 as the dimensionality of the hyper-parameter  $\alpha$  is typically smaller than the validation size  $n_V$ ,  
239 validation loss will not overfit and will be indicative of the test even if the training loss is zero. In  
240 the supplementary materials, we formalize this intuition with emphasis on multi-objective tradeoffs;  
241 namely, under intuitive stability conditions, we show that a small amount of validation data is  
242 sufficient to ensure that the Pareto-front of the validation risk uniformly approximates that of the test-  
243 risk. Concretely, for two objectives  $(\mathcal{L}_1, \mathcal{L}_2)$ , uniformly over all  $\lambda$ , the  $\alpha$  minimizing the validation  
244 risk  $\mathcal{L}_{UP}^{S_V}(f) = (1 - \lambda)\mathcal{L}_1^{S_V} + \lambda\mathcal{L}_2^{S_V}$  in (4.3) also approximately minimizes the test risk  $\mathcal{L}_{UP}(f)$ .

## 245 5 Evaluations for Imbalanced Classes

246 In this section, we present our experiments on various datasets (CIFAR-10, CIFAR-100, iNaturalist-  
247 2018 and ImageNet) when the classes are imbalanced. The goal is to understand whether our bilevel  
248 optimization can design effective loss functions that improve balanced error  $\mathcal{E}_{\text{bal}}$  on the test set. The  
249 setup is as follows.  $\mathcal{E}_{\text{bal}}$  is the test objective. The validation loss  $\mathcal{L}_{UP}$  is the balanced cross-entropy  
250  $\text{CE}_{\text{bal}}$ . We consider various designs for  $\ell_{\text{LOW}}$  such as individually tuning  $w, l, \mathbf{\Delta}$  and augmentation.  
251 We report the average result of 3 random experiments under this setup.

252 **Datasets.** We follow previous works [56, 13, 8] to construct long-tailed versions of the datasets.  
253 Specifically, for a  $K$ -class dataset, we create a long-tailed dataset by reducing the number of  
254 examples per class according to the exponential function  $n'_i = n_i \mu^i$ , where  $n_i$  is the original number  
255 of examples for class  $i$ ,  $n'_i$  is the new number of examples per class, and  $\mu < 1$  is a scaling factor.  
256 Then, we define the imbalance factor  $\rho = n'_0/n'_K$ , which is the ratio of the number of examples in the  
257 largest class ( $n'_0$ ) to the smallest class ( $n'_K$ ). For the **CIFAR-10-LT** and **CIFAR-100-LT** dataset, we  
258 construct long-tailed versions of the datasets with imbalance factor  $\rho = 100$ . **ImageNet-LT** contains  
259 115,846 training examples and 1,000 classes, with imbalance factor  $\rho = 256$ . **iNaturalist-2018**  
260 contains 435,713 images from 8,142 classes, and the imbalance factor is  $\rho = 500$ . These choices  
261 follow that of [56]. For all datasets, we split the long-tailed training set into 80% training and 20%  
262 validation during the search phase (Figure 1b).

263 **Implementation.** In both CIFAR datasets, the lower-level optimization trains a ResNet-32 model  
264 with standard mini-batch stochastic gradient descent (SGD) using learning rate 0.1, momentum 0.9,  
265 and weight decay  $1e - 4$ , over 300 epochs. The learning rate decays at epochs 220 and 260 with a  
266 factor 0.1. The upper-level hyper-parameter optimization computes the hyper-gradients via implicit  
267 differentiation. Because the hyper-gradient is mostly meaningful when the network achieves near zero  
268 loss (Thm 1 of [49]), we start the upper optimization after 120 epochs of the lower-level optimization,  
269 using SGD with initial learning rate 0.01, momentum 0.9, and weight decay  $1e - 4$ . For CIFAR-LT,  
270 20 hyper-parameters are trained, corresponding to  $\mathbf{l}$  and  $\mathbf{\Delta}$  for all classes. For CIFAR-100, because  
271 the size of validation data in minority classes can be too small (e.g., only one example in a tail class),  
272 as visualized in Figure 3(a), we gather classes of similar frequencies into clusters of size 10 to share  
273 the same hyper-parameters (i.e., the values and updates of  $(l_y, \Delta_y)$ 's). Similarly, for ImageNet-LT,  
274 we gather classes into clusters of size 10, and for iNaturalist, we gather classes into clusters of size  
275 40. For ImageNet-LT and iNaturalist, following previous work [56], we use ResNet-50 and SGD for  
276 the lower and upper optimizations, and the same learning rate as CIFAR-LT over 150 epochs. We

Method	CIFAR-10-LT	CIFAR100-LT	ImageNet-LT
MADAO [26]	24.39	59.10	59.92
Algo. 1: $\alpha \leftarrow$ PDA	22.53	58.55	58.77
Algo. 1: $\alpha \leftarrow$ PDA, $\Delta \& l$	20.76	56.49	52.11

Table 2: The result of personalized data optimization.

Source\Target	CIFAR-10-LT	CIFAR100-LT	ImageNet-LT
CIFAR-10-LT	21.39	58.12	56.13
CIFAR100-LT	22.32	56.77	55.89
ImageNet-LT	23.85	59.17	52.90

Table 3: The results of hyper-parameter transfer.

277 train the network for 40 epochs to warm up before the loss function design starts. The learning rate  
 278 decays by a factor 0.1 at epochs 80 and 120.

279 **Personalized Data Augmentation (PDA).** For PDA, we utilize the AutoAugment [12] policy space  
 280 and apply a bilevel search for the augmentation policy. Our approach follows existing differentiable  
 281 augmentation strategies (e.g, [26]); however, we train separate policies for each class cluster to  
 282 ensure that the resulting policies can adjust to class frequencies. Due to space limitations, please see  
 283 supplementary materials for further details.

284 **Results and discussion.** We compared our methods with the state-of-the-art long-tail learning  
 285 methods. Table 1 shows the results of our experiments where the design space is parametric CE  
 286 (4.1). In the first part of the table, we conduct experiments for three baseline methods: normal CE,  
 287 LDAM [8] and Logit Adjustment loss with temperature parameter  $\tau = 1$  [56]. The latter choice  
 288 guarantees Fisher consistency. In the second part of the Table 1, we study Algo. 1 with design spaces  
 289  $l$ ,  $\Delta$ , and  $l \& \Delta$ . The first version of Algo. 1 in Table 1 tunes the LA loss parameter  $\tau$  where  $l$  is  
 290 parameterized by a single scalar  $\tau$  as  $l_y = \tau \log(\pi_y)$ . The next three versions of Algo. 1 consider  
 291 tuning  $l$ ,  $\Delta$ ,  $l \& \Delta$  respectively (Figure 3b-d shows the evolution of the  $l$  and  $\Delta$  parameters during the  
 292 optimization). Finally, in last version of Algo. 1, the loss design is initialized with LA loss with  $\tau = 1$   
 293 (rather than balanced CE). The takeaway from these results is that our approach consistently leads to  
 294 a superior balanced accuracy objective. That said, tuning the LA loss alone is highly competitive with  
 295 optimizing  $\Delta$  and  $l$  alone (in fact, strictly better for CIFAR-10-LT, indicating Algo. 1 does not always  
 296 converge to the optimal design). Importantly, when combining  $l \& \Delta$ , our algorithm is able to design  
 297 a better loss function and outperform all rows across all benchmarks. Finally, when the algorithm  
 298 further is initialized with LA loss, the performance further improves accuracy, demonstrating that  
 299 warm-starting with good designs improves performance.

300 In Table 2, we study the benefits of data augmentation, following our intuitions from Lemma 2. We  
 301 compare to the differentiable augmentation baseline of MADAO [26] which trains a single policy for  
 302 the full dataset. PDA is a personalized variation of MADAO and leads to noticeable improvement  
 303 across all benchmarks (most noticeably in CIFAR-10-LT). More importantly, the last line of the table  
 304 demonstrates that PDA can be synergistically combined with the parametric CE (4.1) which leads  
 305 to further improvements. Finally, in Table 3, we investigate the transferability of our loss design  
 306 (including augmentation). Here, a class within the new dataset uses the loss function designed for the  
 307 closest class from the old dataset, where closest means the most similar class frequency in terms of  
 308 percentile (when all classes of both datasets are sorted). These results demonstrate that AutoBalance  
 309 designs transferable loss functions; e.g., loss functions transferred from CIFAR-LT to ImageNet have  
 310 only slight performance degradation, compared to training the loss from scratch (the diagonals).

## 311 6 Approaches and Evaluations for Imbalanced Groups

312 While Section 5 focuses on the fundamental challenge of balanced error minimization, a more ambi-  
 313 tious goal is optimizing generic fairness-seeking objectives. In this section, we study accuracy-fairness  
 314 tradeoffs by examining the pareto-frontiers of the DEO (3.1), group-balanced error, and standard  
 315 error. We also investigate group-balanced risk defined as  $\mathcal{L}_{\text{bal}}^G(f) = \frac{1}{KG} \sum_{k=1}^K \sum_{j=1}^G \mathcal{L}_{k,j}(f)$ . Note  
 316 that this definition treats each (class, group) pair as its own (sub)group. Throughout, we explicitly set  
 317 the loss to cross-entropy for clarity, thus we use CE,  $\text{CE}_{\text{bal}}^G$ ,  $\text{CE}_{\text{DEO}}$  to refer to  $\mathcal{L}$ ,  $\mathcal{L}_{\text{bal}}^G$ ,  $\mathcal{L}_{\text{DEO}}$ .

318 **Validation (upper-level) loss function.** In Algo. 1, we set  $\mathcal{L}_{\text{UP}} = (1 - \lambda) \cdot \text{CE} + \lambda \cdot \text{CE}_{\text{DEO}}$  for  
 319 varying  $0 \leq \lambda \leq 1$ . The parameter  $\lambda$  enables a trade-off between accuracy and fairness.

320 **Group-sensitive training loss design.** The parametric cross-entropy (CE) can be extended to (class,  
 321 group) imbalance by extending hyper-parameter  $\alpha$  to  $[K] \times [G]$  variables  $w, l, \Delta \in \mathbb{R}^{[K] \times [G]}$   
 322 generalizing (4.1), (4.2). This leads us to the following design:

$$\ell_{\text{low}}(y, g, f(\mathbf{x}); \alpha) = -w_{yg} \log \left( \frac{e^{\sigma(\Delta_{yg})f_y(\mathbf{x}) + l_{yg}}}{\sum_{k \in [K]} e^{\sigma(\Delta_{kg})f_k(\mathbf{x}) + l_{kg}}} \right). \quad (6.1)$$

Loss function	Balanced Error	Worst (class, group) error	DEO
Cross entropy (CE)	23.38	43.25	33.75
CE <sub>bal</sub>	20.83	36.67	20.25
Group-LA loss	22.83	40.50	29.33
CE <sub>DEO</sub>	19.29	35.17	25.25
0.1 · CE + 0.9 · CE <sub>DEO</sub> (λ = 0.1)	20.06	31.67	26.25
Algo. 1: with λ = 0.1	<b>15.13</b>	<b>30.33</b>	<b>4.25</b>

Table 4: Comparison of fairness metrics for group-imbalanced experiments. The first five rows are different training loss choices, where CE<sub>bal</sub>, Group-LA, and CE<sub>DEO</sub> promote group fairness. The last row is Algo. 1, which designs training loss for the validation loss choice of 0.1 · CE + 0.9 · CE<sub>DEO</sub>. We note that, DEO can be trivially minimized by always predicting the same class. To avoid this, we use a mild amount of CE loss with λ = 0.1 in Algo. 1.

323 Here,  $w_{yg}$  applies weighted CE, while  $\Delta_{yg}$  and  $t_{yg}$  are logit adjustments for different (class, groups).  
324 Note that a similar group-sensitive loss is proposed in [40] for binary classification.

325 **Baselines.** We will compare Algo. 1 with training loss functions parameterized via  $(1-\lambda) \cdot \text{CE} + \lambda \cdot \mathcal{L}_{\text{reg}}$ .  
326 Here  $\mathcal{L}_{\text{reg}}$  is a fairness-promoting regularization. Specifically, as displayed in Table 4 and Figure 4,  
327 we will set  $\mathcal{L}_{\text{reg}}$  to be CE<sub>bal</sub><sup>G</sup>, CE<sub>DEO</sub> and Group LA. ‘‘Group LA’’ is a natural generalization of the  
328 LA loss to group-sensitive setting; it chooses weights  $w_g = 1/\bar{\pi}_g$  to balance group frequencies and  
329 then applies logit-adjustment with  $\tau = 1$  over the classes conditioned on the group-membership:

$$330 \ell(y, g, f(\mathbf{x})) = -\frac{1}{G\bar{\pi}_g} \log \left( \frac{e^{f y(\mathbf{x}) + \tau \log \pi_{yg}}}{\sum_{k \in [K]} e^{f k(\mathbf{x}) + \tau \log \pi_{kg}}} \right).$$

331 **Datasets.** We experiment with the modified Waterbird dataset [61]. The goal is to correctly classify  
332 the bird type despite the spurious correlations due to the image background. The distribution of the  
333 original data is as follows. The binary classes  $k \in \{-, +\}$  correspond to {waterbird, landbird}, and  
334 the groups  $[G] = \{1, 2\}$  correspond to {land background, water background}. The fraction of data  
335 in each (class, group) pair is  $\pi_{-,2} = 0.22$ ,  $\pi_{-,1} = 0.012$ ,  $\pi_{+,2} = 0.038$ , and  $\pi_{+,1} = 0.73$ . The  
336 landbird on the water background ( $\{+, 2\}$ ) and the waterbird on the land background ( $\{-, 1\}$ ) are  
337 minority sub-groups within their respective classes. The test set, following [61], has equally allocated  
338 bird types on different backgrounds, i.e.,  $\pi_{\pm,j} = 0.25$ . As the test dataset is balanced, the standard  
339 classification error  $\mathcal{E}(f)$  is defined to be the weighted error  $\mathcal{E}(f) = \pi_{y,g} \mathcal{E}_{y,g}(f)$ .

340 **Implementation.** We follow the feature extraction method from [61], where  $x_i$  are 512-dimensional  
341 ResNet18 features. When using Algo. 1, we split the original training data into 50% training and  
342 50% validation. The search phase uses 150 epochs of warm up followed by 350 epochs of bilevel  
343 optimization. The remaining implementation details are similar to Section 5.

344 **Results and discussion.** We consider various  
345 fairness-related metrics, including the worst  
346 (class, group) error, DEO  $\mathcal{E}_{\text{DEO}}(f)$  and the  
347 balanced error  $\mathcal{E}_{\text{bal}}^G(f)$ . We seek to understand  
348 whether AutoBalance algorithm can improve  
349 performance on the test set compared to the  
350 baseline training loss functions of the form  
351  $(1-\lambda) \cdot \text{CE} + \lambda \cdot \mathcal{L}_{\text{reg}}$ . In Figure 4, we show  
352 the influence of the parameter  $\lambda$  where  $\mathcal{L}_{\text{reg}}$   
353 is chosen to be CE<sub>DEO</sub>, CE<sub>bal</sub>, or Group-LA  
354 (each point on the plot represents a different  
355  $\lambda$  value). As we sweep across values of  $\lambda$ ,  
356 there arises a tradeoff between standard clas-  
357 sification error  $\mathcal{E}(f)$  and the fairness metrics.  
358 We observe that Algo. 1 significantly Pareto-  
359 dominates alternative approaches, for example  
360 achieving lower DEO or balanced error for the same standard error. This demonstrates the value of  
361 automatic loss function design for a rich class of fairness-seeking objectives.

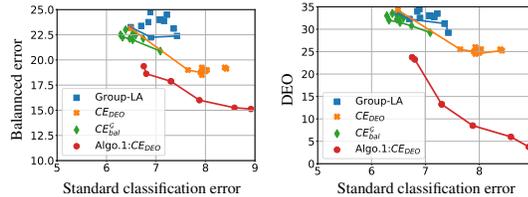


Figure 4: Waterbirds fairness-accuracy tradeoffs for parametrized loss designs  $(1-\lambda) \cdot \text{CE} + \lambda \cdot \mathcal{L}_{\text{reg}}$ , for different  $\mathcal{L}_{\text{reg}}$  choices. Group-balanced error  $\mathcal{E}^G(f)$  (left) and DEO  $\mathcal{E}_{\text{DEO}}(f)$  (right) are plotted as a function of the misclassification error  $\mathcal{E}(f)$ . Algo. 1 exhibits a noticeably better tradeoff curve as it uses a DEO-based validation objective to design an optimized training loss function.

362 Next, in Table 4, we solely focus on optimizing the fairness objectives (rather than standard error).  
363 Thus, we simply set  $\lambda = 0$  and compare against CE<sub>DEO</sub>, CE<sub>bal</sub>, Group-LA as the baseline approaches  
364 (without blending CE). We also display the outcome of Algo. 1 with  $\lambda = 0.1$ . Similar too Figure 4,

365 our approach outperforms all baselines for all metrics. The performance gap is particularly noticeable  
366 when it comes to DEO (4.25% for Algo. 1, vs 20.25% for the best baseline).

## 367 7 Conclusions and Future Directions

368 This work provides an optimization-based approach to automatically design loss functions to address  
369 imbalanced learning problems. Our algorithm consistently outperforms, or is at least competitive  
370 with, the state-of-the-art approaches for optimizing balanced accuracy. Importantly, our approach is  
371 not restricted to imbalanced classes or specific objectives, and can achieve good tradeoffs between  
372 (accuracy, fairness) on the Pareto frontier. We also provide theoretical insights on certain algorithmic  
373 aspects including loss function design, data augmentation, and train-validation split.

374 **Potential Limitations, Negative Societal Impacts, & Precautions:** Our algorithmic approach can  
375 be considered within the realm of automated machine learning literature (AutoML) [29]. AutoML  
376 algorithms often optimize the model performance, thus reducing the need for engineering expertise at  
377 the expense of increased computational cost and increased carbon footprint. For instance, our procedure  
378 is computationally more intensive compared to the theory-inspired loss function prescriptions of  
379 [56, 8]. A related limitation is that Algo. 1 can be brittle in extremely imbalanced scenarios with very  
380 few samples per class. We took the several steps to help mitigate such issues: first, our algorithm is  
381 initialized with a Bayes consistent loss function to provide a warm-start (such as the proposal of [56]).  
382 Second, we reduce the hyper-parameter search space by grouping classes with similar frequencies (to  
383 help with brittleness). Finally, evaluations show that the designed loss functions are transferable, and  
384 hence our algorithm does not have to train from scratch for a new dataset.

## 385 References

- 386 [1] Shin Ando and Chun Yuan Huang. Deep over-sampling framework for classifying imbalanced  
387 data. In *Joint European Conference on Machine Learning and Knowledge Discovery in*  
388 *Databases*, pages 770–785. Springer, 2017.
- 389 [2] Yu Bai, Minshuo Chen, Pan Zhou, Tuo Zhao, Jason D Lee, Sham Kakade, Huan Wang, and  
390 Caiming Xiong. How important is the train-validation split in meta-learning? *arXiv preprint*  
391 *arXiv:2010.05843*, 2020.
- 392 [3] Bowen Baker, Otkrist Gupta, Nikhil Naik, and Ramesh Raskar. Designing neural network  
393 architectures using reinforcement learning. *arXiv preprint arXiv:1611.02167*, 2016.
- 394 [4] Aharon Ben-Tal, Dick den Hertog, Anja De Waegenare, Bertrand Melenberg, and Gijs Rennen.  
395 Robust solutions of optimization problems affected by uncertain probabilities. *Management*  
396 *Science*, 59(2):341–357, 2013.
- 397 [5] Mateusz Buda, Atsuto Maki, and Maciej A Mazurowski. A systematic study of the class  
398 imbalance problem in convolutional neural networks. *Neural Networks*, 106:249–259, 2018.
- 399 [6] Jonathon Byrd and Zachary Lipton. What is the effect of importance weighting in deep learning?  
400 In *International Conference on Machine Learning*, pages 872–881. PMLR, 2019.
- 401 [7] Toon Calders, Faisal Kamiran, and Mykola Pechenizkiy. Building classifiers with independency  
402 constraints. In *2009 IEEE International Conference on Data Mining Workshops*, pages 13–18.  
403 IEEE, 2009.
- 404 [8] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arechiga, and Tengyu Ma. Learning imbalanced  
405 datasets with label-distribution-aware margin loss. *arXiv preprint arXiv:1906.07413*, 2019.
- 406 [9] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. Smote:  
407 synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–  
408 357, 2002.
- 409 [10] Shuxiao Chen, Edgar Dobriban, and Jane H Lee. A group-theoretic framework for data  
410 augmentation. *Journal of Machine Learning Research*, 21(245):1–71, 2020.
- 411 [11] Nicolas Couellan and Wenjuan Wang. On the convergence of stochastic bi-level gradient  
412 methods. *Optimization*, 2016.

- 413 [12] ED Cubuk, B Zoph, D Mane, V Vasudevan, and QV Le. Autoaugment: Learning augmentation  
414 policies from data. arxiv 2018. *arXiv preprint arXiv:1805.09501*.
- 415 [13] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based  
416 on effective number of samples. In *Proceedings of the IEEE/CVF Conference on Computer  
417 Vision and Pattern Recognition*, pages 9268–9277, 2019.
- 418 [14] Tri Dao, Albert Gu, Alexander Ratner, Virginia Smith, Chris De Sa, and Christopher Ré. A  
419 kernel theory of modern data augmentation. In *International Conference on Machine Learning*,  
420 pages 1528–1537. PMLR, 2019.
- 421 [15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-  
422 scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern  
423 recognition*, pages 248–255. Ieee, 2009.
- 424 [16] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural  
425 networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017.
- 426 [17] Qi Dong, Shaogang Gong, and Xiatian Zhu. Imbalanced deep learning by minority class  
427 incremental rectification. *IEEE transactions on pattern analysis and machine intelligence*,  
428 41(6):1367–1381, 2018.
- 429 [18] Michele Donini, Luca Oneto, Shai Ben-David, John Shawe-Taylor, and Massimiliano Pontil.  
430 Empirical risk minimization under fairness constraints. *arXiv preprint arXiv:1802.08626*, 2018.
- 431 [19] Simon Du, Jason Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global  
432 minima of deep neural networks. In *International Conference on Machine Learning*, pages  
433 1675–1685. PMLR, 2019.
- 434 [20] John C Duchi, Peter W Glynn, and Hongseok Namkoong. Statistics of robust optimization: A  
435 generalized empirical likelihood approach. *Mathematics of Operations Research*, 2021.
- 436 [21] Vitaly Feldman. Does learning require memorization? a short tale about a long tail. In  
437 *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages  
438 954–959, 2020.
- 439 [22] Luca Franceschi, Paolo Frasconi, Saverio Salzo, Riccardo Grazi, and Massimiliano Pontil.  
440 Bilevel programming for hyperparameter optimization and meta-learning. In *International  
441 Conference on Machine Learning*, pages 1568–1577. PMLR, 2018.
- 442 [23] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural  
443 networks. In *International Conference on Machine Learning*, pages 1321–1330. PMLR, 2017.
- 444 [24] Moritz Hardt, Eric Price, and Nathan Srebro. Equality of opportunity in supervised learning.  
445 *arXiv preprint arXiv:1610.02413*, 2016.
- 446 [25] Ryuichiro Hataya, Jan Zdenek, Kazuki Yoshizoe, and Hideki Nakayama. Faster autoaugment:  
447 Learning augmentation strategies using backpropagation. In *European Conference on Computer  
448 Vision*, pages 1–16. Springer, 2020.
- 449 [26] Ryuichiro Hataya, Jan Zdenek, Kazuki Yoshizoe, and Hideki Nakayama. Meta approach to data  
450 augmentation optimization. *arXiv preprint arXiv:2006.07965*, 2020.
- 451 [27] Daniel Ho, Eric Liang, Xi Chen, Ion Stoica, and Pieter Abbeel. Population based augmentation:  
452 Efficient learning of augmentation policy schedules. In *International Conference on Machine  
453 Learning*, pages 2731–2741. PMLR, 2019.
- 454 [28] Chen Huang, Yining Li, Chen Change Loy, and Xiaoou Tang. Learning deep representation  
455 for imbalanced classification. In *Proceedings of the IEEE conference on computer vision and  
456 pattern recognition*, pages 5375–5384, 2016.
- 457 [29] Frank Hutter, Lars Kotthoff, and Joaquin Vanschoren. *Automated machine learning: methods,  
458 systems, challenges*. Springer Nature, 2019.

- 459 [30] Simon Jenni and Paolo Favaro. Deep bilevel learning. In *Proceedings of the European*  
460 *conference on computer vision (ECCV)*, pages 618–633, 2018.
- 461 [31] Ziwei Ji and Matus Telgarsky. Polylogarithmic width suffices for gradient descent to achieve  
462 arbitrarily small test error with shallow relu networks. *arXiv preprint arXiv:1909.12292*, 2019.
- 463 [32] Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and  
464 Yannis Kalantidis. Decoupling representation and classifier for long-tailed recognition. *arXiv*  
465 *preprint arXiv:1910.09217*, 2019.
- 466 [33] Shuichi Katsumata and Akiko Takeda. Robust cost sensitive support vector machine. In  
467 *Artificial intelligence and statistics*, pages 434–443. PMLR, 2015.
- 468 [34] Michael Kearns. A bound on the error of cross validation using the approximation and estimation  
469 rates, with consequences for the training-test split. *Advances in Neural Information Processing*  
470 *Systems*, pages 183–189, 1996.
- 471 [35] Michael Kearns and Dana Ron. Algorithmic stability and sanity-check bounds for leave-one-out  
472 cross-validation. *Neural computation*, 11(6):1427–1453, 1999.
- 473 [36] Salman Khan, Munawar Hayat, Syed Waqas Zamir, Jianbing Shen, and Ling Shao. Striking  
474 the right balance with uncertainty. In *Proceedings of the IEEE/CVF Conference on Computer*  
475 *Vision and Pattern Recognition*, pages 103–112, 2019.
- 476 [37] Salman H Khan, Munawar Hayat, Mohammed Bennamoun, Ferdous A Sohel, and Roberto  
477 Togneri. Cost-sensitive learning of deep feature representations from imbalanced data. *IEEE*  
478 *transactions on neural networks and learning systems*, 29(8):3573–3587, 2017.
- 479 [38] Mikhail Khodak, Liam Li, Maria-Florina Balcan, and Ameet Talwalkar. On weight-sharing and  
480 bilevel optimization in architecture search. 2019.
- 481 [39] Byungju Kim and Junmo Kim. Adjusting decision boundary for class imbalanced learning.  
482 *IEEE Access*, 8:81674–81685, 2020.
- 483 [40] Ganesh Ramachandra Kini, Orestis Paraskevas, Samet Oymak, and Christos Thrampoulidis.  
484 Label-imbalanced and group-sensitive classification under overparameterization. *arXiv preprint*  
485 *arXiv:2103.01550*, 2021.
- 486 [41] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep  
487 convolutional neural networks. In *Advances in neural information processing systems*, pages  
488 1097–1105, 2012.
- 489 [42] Miroslav Kubat, Stan Matwin, et al. Addressing the curse of imbalanced training sets: one-sided  
490 selection. In *Icml*, volume 97, pages 179–186. Citeseer, 1997.
- 491 [43] Hansang Lee, Minseok Park, and Junmo Kim. Plankton classification on imbalanced large scale  
492 database via convolutional neural networks with transfer learning. In *2016 IEEE international*  
493 *conference on image processing (ICIP)*, pages 3713–3717. IEEE, 2016.
- 494 [44] Joseph Lemley, Shabab Bazrafkan, and Peter Corcoran. Smart augmentation learning an optimal  
495 data augmentation strategy. *Ieee Access*, 5:5858–5869, 2017.
- 496 [45] Renjie Liao, Yuwen Xiong, Ethan Fetaya, Lisa Zhang, KiJung Yoon, Xaq Pitkow, Raquel Urta-  
497 sun, and Richard Zemel. Reviving and improving recurrent back-propagation. In *International*  
498 *Conference on Machine Learning*, pages 3082–3091. PMLR, 2018.
- 499 [46] Sungbin Lim, Ildoo Kim, Taesup Kim, Chiheon Kim, and Sungwoong Kim. Fast autoaugment.  
500 *arXiv preprint arXiv:1905.00397*, 2019.
- 501 [47] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense  
502 object detection. In *Proceedings of the IEEE international conference on computer vision*,  
503 pages 2980–2988, 2017.
- 504 [48] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search.  
505 *arXiv preprint arXiv:1806.09055*, 2018.

- 506 [49] Jonathan Lorraine, Paul Vicol, and David Duvenaud. Optimizing millions of hyperparameters  
507 by implicit differentiation. In *International Conference on Artificial Intelligence and Statistics*,  
508 pages 1540–1552. PMLR, 2020.
- 509 [50] Jelena Luketina, Mathias Berglund, Klaus Greff, and Tapani Raiko. Scalable gradient-based  
510 tuning of continuous regularization hyperparameters. In *International conference on machine*  
511 *learning*, pages 2952–2960. PMLR, 2016.
- 512 [51] Matthew MacKay, Paul Vicol, Jon Lorraine, David Duvenaud, and Roger Grosse. Self-tuning  
513 networks: Bilevel optimization of hyperparameters using structured best-response functions.  
514 *arXiv preprint arXiv:1903.03088*, 2019.
- 515 [52] Dougal Maclaurin, David Duvenaud, and Ryan Adams. Gradient-based hyperparameter opti-  
516 mization through reversible learning. In *International conference on machine learning*, pages  
517 2113–2122. PMLR, 2015.
- 518 [53] Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li,  
519 Ashwin Bharambe, and Laurens Van Der Maaten. Exploring the limits of weakly supervised  
520 pretraining. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages  
521 181–196, 2018.
- 522 [54] Hamed Masnadi-Shirazi, Nuno Vasconcelos, and Arya Iranmehr. Cost-sensitive support vector  
523 machines. *arXiv preprint arXiv:1212.0975*, 2012.
- 524 [55] Aditya Menon, Harikrishna Narasimhan, Shivani Agarwal, and Sanjay Chawla. On the statistical  
525 consistency of algorithms for binary classification under class imbalance. In *International*  
526 *Conference on Machine Learning*, pages 603–611. PMLR, 2013.
- 527 [56] Aditya Krishna Menon, Sadeep Jayasumana, Ankit Singh Rawat, Himanshu Jain, Andreas Veit,  
528 and Sanjiv Kumar. Long-tail learning via logit adjustment. *arXiv preprint arXiv:2007.07314*,  
529 2020.
- 530 [57] Katharina Morik, Peter Brockhausen, and Thorsten Joachims. Combining statistical learning  
531 with a knowledge-based approach: a case study in intensive care monitoring. Technical report,  
532 Technical Report, 1999.
- 533 [58] Samet Oymak, Mingchen Li, and Mahdi Soltanolkotabi. Generalization guarantees for neural  
534 architecture search with train-validation split. *arXiv preprint arXiv:2104.14132*, 2021.
- 535 [59] Samira Pouyanfar, Yudong Tao, Anup Mohan, Haiman Tian, Ahmed S Kaseb, Kent Gauen,  
536 Ryan Dailey, Sarah Aghajanzadeh, Yung-Hsiang Lu, Shu-Ching Chen, et al. Dynamic sampling  
537 in convolutional neural networks for imbalanced data classification. In *2018 IEEE conference*  
538 *on multimedia information processing and retrieval (MIPR)*, pages 112–117. IEEE, 2018.
- 539 [60] Saharon Rosset, Ji Zhu, and Trevor Hastie. Margin maximizing loss functions. In *NIPS*, pages  
540 1237–1244, 2003.
- 541 [61] Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. Distributionally  
542 robust neural networks for group shifts: On the importance of regularization for worst-case  
543 generalization. *arXiv preprint arXiv:1911.08731*, 2019.
- 544 [62] Amirreza Shaban, Ching-An Cheng, Nathan Hatch, and Byron Boots. Truncated back-  
545 propagation for bilevel optimization. In *The 22nd International Conference on Artificial*  
546 *Intelligence and Statistics*, pages 1723–1732. PMLR, 2019.
- 547 [63] Li Shen, Zhouchen Lin, and Qingming Huang. Relay backpropagation for effective learning  
548 of deep convolutional neural networks. In *European conference on computer vision*, pages  
549 467–482. Springer, 2016.
- 550 [64] Ashish Shrivastava, Tomas Pfister, Oncel Tuzel, Joshua Susskind, Wenda Wang, and Russell  
551 Webb. Learning from simulated and unsupervised images through adversarial training. In  
552 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2107–  
553 2116, 2017.

- 554 [65] Ankur Sinha, Pekka Malo, and Kalyanmoy Deb. A review on bilevel optimization: from  
555 classical to evolutionary approaches and applications. *IEEE Transactions on Evolutionary  
556 Computation*, 22(2):276–295, 2017.
- 557 [66] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural  
558 networks. In *International Conference on Machine Learning*, pages 6105–6114. PMLR, 2019.
- 559 [67] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The  
560 caltech-ucsd birds-200-2011 dataset. 2011.
- 561 [68] Robert Williamson and Aditya Menon. Fairness risk measures. In *International Conference on  
562 Machine Learning*, pages 6786–6797. PMLR, 2019.
- 563 [69] Sirui Xie, Hehui Zheng, Chunxiao Liu, and Liang Lin. Snas: stochastic neural architecture  
564 search. *arXiv preprint arXiv:1812.09926*, 2018.
- 565 [70] Zhen Xu, Andrew M Dai, Jonas Kemp, and Luke Metz. Learning an adaptive learning rate  
566 schedule. *arXiv preprint arXiv:1909.09712*, 2019.
- 567 [71] Han-Jia Ye, Hong-You Chen, De-Chuan Zhan, and Wei-Lun Chao. Identifying and compen-  
568 sating for feature deviation in imbalanced deep learning. *arXiv preprint arXiv:2001.01385*,  
569 2020.
- 570 [72] Xi Yin, Xiang Yu, Kihyuk Sohn, Xiaoming Liu, and Manmohan Chandraker. Feature transfer  
571 learning for deep face recognition with under-represented data. *arXiv preprint arXiv:1803.09014*,  
572 2018.
- 573 [73] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon  
574 Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In  
575 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6023–6032,  
576 2019.
- 577 [74] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P Gummadi.  
578 Fairness beyond disparate treatment & disparate impact: Learning classification without dis-  
579 parate mistreatment. In *Proceedings of the 26th international conference on world wide web*,  
580 pages 1171–1180, 2017.
- 581 [75] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond  
582 empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.
- 583 [76] Junjie Zhang, Lingqiao Liu, Peng Wang, and Chunhua Shen. To balance or not to balance:  
584 A simple-yet-effective approach for learning with long-tailed distributions. *arXiv preprint  
585 arXiv:1912.04486*, 2019.
- 586 [77] Zhao Zhong, Zichen Yang, Boyang Deng, Junjie Yan, Wei Wu, Jing Shao, and Cheng-Lin Liu.  
587 Blockqnn: Efficient block-wise neural network architecture generation. *IEEE transactions on  
588 pattern analysis and machine intelligence*, 2020.
- 589 [78] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A  
590 10 million image database for scene recognition. *IEEE transactions on pattern analysis and  
591 machine intelligence*, 40(6):1452–1464, 2017.
- 592 [79] Boyan Zhou, Quan Cui, Xiu-Shen Wei, and Zhao-Min Chen. Bbn: Bilateral-branch network  
593 with cumulative learning for long-tailed visual recognition. In *Proceedings of the IEEE/CVF  
594 Conference on Computer Vision and Pattern Recognition*, pages 9719–9728, 2020.
- 595 [80] Pan Zhou, Caiming Xiong, Richard Socher, and Steven CH Hoi. Theory-inspired path-  
596 regularized differential network architecture search. *arXiv preprint arXiv:2006.16537*, 2020.
- 597 [81] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. *arXiv  
598 preprint arXiv:1611.01578*, 2016.
- 599 [82] Difan Zou, Yuan Cao, Dongruo Zhou, and Quanquan Gu. Stochastic gradient descent optimizes  
600 over-parameterized deep relu networks. *arXiv preprint arXiv:1811.08888*, 2018.

601 **A Extended related works**

602 Below we include the related work on data augmentation which was omitted from Section 2 due to  
 603 space considerations.

604 **Data Augmentation.** Data augmentation techniques have been studied for decades, and many  
 605 approaches such as random crop, flip, rotation, Mixup [75], Cutout [16], CutMix [73] have been  
 606 applied in model training. Recently, researchers focus on automatically finding data augmentation  
 607 policies to achieve better performance. Some methods [44, 64] obtain augmentation policies through  
 608 an additional network or GAN. Inspired by neural architecture search, AutoAugment [12] and its  
 609 followup works [27, 46, 25, 26] formulate data augmentation as a hyper-parameter search problem.  
 610 [25, 26] propose optimizing the augmentation policy using bi-level optimization by conducting  
 611 differentiable relaxation on policies. Different from the above works that perform a bi-level search for  
 612 an augmentation policy on a balanced dataset, our approach employs personalized data augmentation  
 613 for different classes on long-tailed datasets, which leads to a better result in long-tailed learning  
 614 problems.

---

**Algorithm 2:** Hyper gradient computation [49]

---

**Input:** Model  $f_\theta$  with weights  $\theta$ , hyper-parameter  $\alpha$ , dataset  $\mathcal{S} = \mathcal{S}_T \cup \mathcal{S}_V$ , step sizes  $\eta$ , order  
 of Neumann approximation  $i$

```

1  $v_1 = \frac{\partial \mathcal{L}_{UP}^{S_V}}{\partial \theta}$ 
2  $p = v_1$ 
3 for  $j \leftarrow 1$  to  $i$  do
  | // Compute approximate Hessian inverse using Neumann series
4    $v_1 = v_1 (I - \eta \frac{\partial^2 \mathcal{L}_{LOW}^{S_T}}{\partial \theta \partial \theta^T})$ 
5    $p \leftarrow v_1$ 
6 end
7  $v_2 = -p \frac{\partial^2 \mathcal{L}_{LOW}^{S_T}}{\partial \theta \partial \alpha^T}$ 

```

**Result:** The hyper gradient  $v_2$                       //  $v_2 = \frac{\partial \mathcal{L}_{UP}}{\partial \theta} \left[ \frac{\partial^2 \mathcal{L}_{LOW}}{\partial \theta \partial \theta^T} \right]^{-1} \frac{\partial^2 \mathcal{L}_{LOW}}{\partial \theta \partial \alpha^T}$

---

615 **B Extended experiments**

616 In this section, we conduct additional experiments to extend the results from Sections 5 and 6.  
 617 Importantly, for the experiments in Section 5 (imbalanced classes) and Section 6 (imbalanced groups),  
 618 we conduct more trials (for a total of 5 trials per method) and we also provide standard error bounds  
 619 in the tables. We also included additional baselines for Section 6.

620 Tables 5, 6, and 7 display the updated results of Section 5 for the balanced accuracy, personalized  
 621 data optimization, and hyper-parameter transfer evaluation scenarios, respectively. These correspond  
 622 to the original Tables 1, 2, and 3, respectively. The additional trials are generally consistent with the  
 623 results and insights in the main body of the paper and demonstrates the validity of our approach.

624 For Section 6, we evaluate two additional baselines: Distributionally Robust Optimization (DRO) [61],  
 625 and a post-hoc model – described below – that tries to address group imbalance. The results are shown  
 626 in Table 8. Overall, the results show that our Algo. 1 with  $\lambda = 0.1$  consistently outperforms the other  
 627 baselines. Below, we describe these two baselines in more detail.

628 *DRO baseline:* We follow the work of [61] where we optimize the DRO loss as an additional baseline.  
 629 For consistency, we use a slight variation where we change the network to ResNet-18. Table 8 shows  
 630 that although DRO achieves a lot better performance compared to the other baselines, our Algo. 1  
 631 with  $\lambda = 0.1$  still performs noticeably better across all performance metrics.

632 *Post-hoc baseline:* We first train a ResNet-18 model with training dataset using simple CE loss. As  
 633 the training dataset is imbalanced, the error of worst group is quite high at this intermediate point,  
 634 more than 45%. As our posthoc model, we use vector scaling [23] which adjusts the logits. Vector  
 635 scaling is essentially a generalization of Platt scaling where each logit gets its own weights in a

Method	CIFAR-10-LT	CIFAR100-LT	ImageNet-LT	iNaturalist
Cross-Entropy	30.45 ± 0.45	62.69 ± 0.16	60.82 ± 0.24	39.72 ± 0.22
LDAM loss [8]	26.37 ± 0.33	59.47 ± 0.38	58.14 ± 0.20	35.63 ± 0.21
LA loss ( $\tau = 1$ ) [56]	23.13 ± 0.35	58.96 ± 0.20	55.57 ± 0.23	34.37 ± 0.18
Algo. 1: $\alpha \leftarrow \tau$ of LA loss	21.82 ± 0.13	58.68 ± 0.20	55.15 ± 0.25	34.19 ± 0.19
Algo. 1: $\alpha \leftarrow \mathbf{l}$	23.02 ± 0.43	58.71 ± 0.25	55.30 ± 0.29	34.35 ± 0.21
Algo. 1: $\alpha \leftarrow \Delta$	22.59 ± 0.26	58.40 ± 0.22	54.55 ± 0.17	34.37 ± 0.22
Algo. 1: $\alpha \leftarrow \Delta \& \mathbf{l}$	21.39 ± 0.18	56.84 ± 0.17	53.16 ± 0.17	33.41 ± 0.30
Algo. 1: $\alpha \leftarrow \Delta \& \mathbf{l}$ , LA init	21.15 ± 0.22	56.70 ± 0.18	52.11 ± 0.12	33.16 ± 0.13

Table 5: Evaluations of balanced accuracy on long-tailed data with 5 trials. Algo. 1 with  $\Delta \& \mathbf{l}$  design space and LA initialization (bottom row) outperforms other baselines, across various datasets.

Method	CIFAR-10-LT	CIFAR100-LT	ImageNet-LT
MADAO [26]	24.42 ± 0.28	59.10 ± 0.21	59.56 ± 0.41
Algo. 1: $\alpha \leftarrow$ PDA	22.55 ± 0.38	58.52 ± 0.56	58.73 ± 0.51
Algo. 1: $\alpha \leftarrow$ PDA, $\Delta \& \mathbf{l}$	20.69 ± 0.17	56.47 ± 0.23	52.18 ± 0.22

Table 6: Personalized data optimization with 5 trials.

636 similar fashion to the parametric cross-entropy loss. The inputs to the vector scaling post-hoc model  
637 are the output logits ( $2 \times N$ ) from ResNet-18 where  $N$  is the sample size. We use  $w$  ( $2 \times 1$ ) and  $b$   
638 ( $2 \times 1$ ) to adjust it  $f_{\text{posthoc}}(x) = wx + b$ . As there are only 4 parameters to tune, we use grid search to  
639 find the parameters that can minimize loss (DEO or balanced error  $\mathcal{E}_{\text{bal}}(f)$ ) on the test dataset<sup>2</sup>. Note  
640 that, optimizing over the test data, intuitively, makes this baseline stronger than it actually is. The  
641 posthoc model uses just 4 parameters, which can aid in balancing the result; however, as it can only  
642 adjust per class instead of per (class, group), the performance is limited compared to our approach.  
643 We note that, intuitively, our approach (or in general choosing an intelligent loss function) can be  
644 perceived as applying a posthoc adjustment during training rather than after training.

## 645 C Theoretical Insights into Pareto-Efficiency with Validation Data

646 In Section 4.1, we provided theoretical intuitions on why validation is necessary to build models  
647 that optimize multiple learning objectives. Within the context of this work, these objectives can be a  
648 blend of accuracy and fairness. To recap, our main intuition is that large capacity neural networks can  
649 perfectly maximize different accuracy metrics or satisfy fairness constraints such as DOE by simply  
650 fitting training data perfectly and achieving 100% training accuracy. To truly find a model that lie on  
651 the multi-objective pareto-front, the optimization procedure should (approximately) evaluate on the  
652 population landscape. Validation phase enables this as the dimensionality of the validation parameter  
653 is typically much smaller than the sample size and prevents overfitting. Below, we formalize this in  
654 a general constrained multiobjective learning setting. Suppose there are  $R$  objectives to optimize.  
655 Let  $(\ell_i)_{i=1}^R$  be  $R$  loss functions and set the corresponding  $\mathcal{L}_i(f) = \mathbb{E}[\ell_i(y, f(x))]$ . These can be  
656 accuracy or fairness objectives or it can even be class- or group-conditional risks (i.e.  $R = K$ , every  
657 class gets its own loss function).

658 Split  $\mathcal{S} = \mathcal{T} \cup \mathcal{V}$  where  $\mathcal{T}$  and  $\mathcal{V}$  are training and validation respectively. During the training  
659 phase, we assume that, there is an algorithm (e.g. SGD, Adam, convex optimization, etc)  $A$  that  
660 optimizes over the training data  $\mathcal{T}$  (e.g. by minimizing ERM with gradient descent).  $A$  admits the  
661 hyper-parameters  $\alpha$  (e.g., parameterization of the loss function) as input and returns a hypothesis

$$f_\alpha = A(\mathcal{T}, \alpha)$$

662 For the discussion in this section, we use  $\mathcal{H}$  to denote the hyperparameter search space i.e. the values  
663 the hyperparameter  $\alpha$  can take. Let  $\mathcal{L}_i^\mathcal{V}(f)$  be the empirical version of  $\mathcal{L}_i(f)$  computed over  $\mathcal{V}$ . Fix  
664 penalties  $\lambda = (\lambda_i)_{i=1}^R$  which govern the combination of the loss functions (e.g. blending accuracy and  
665 fairness or weighing individual classes). The validation phase then optimizes  $\theta$  via a Multi-objective

<sup>2</sup>This is in contrast to using differentiable proxies based on cross-entropy.

Source\Target	CIFAR-10-LT	CIFAR100-LT	ImageNet-LT
CIFAR-10-LT	21.39 ± 0.18	58.10 ± 0.28	56.13 ± 0.29
CIFAR100-LT	22.40 ± 0.23	56.84 ± 0.17	55.72 ± 0.35
ImageNet-LT	23.70 ± 0.29	59.24 ± 0.23	53.16 ± 0.17

Table 7: Hyper-parameter transfer with 5 trials.

Loss function	Balanced Error	Worst (class, group) error	DEO
Cross entropy (CE)	25.37(±0.31)	46.69(±4.18)	33.75(±1.86)
CE <sub>bal</sub>	21.09(±0.27)	36.63(±4.82)	20.61(±1.52)
Group-LA loss	22.91(±0.36)	40.27(±5.23)	29.34(±1.46)
CE <sub>DEO</sub>	19.32(±0.31)	33.04(±5.46)	25.33(±1.35)
0.9 · CE + 0.1 · CE <sub>DEO</sub> (λ = 0.1)	20.38(±0.27)	33.36(±6.00)	26.42(±1.39)
DRO [61]	16.47(±0.23)	32.67(±3.06)	6.91(±1.30)
Posthoc: $\mathcal{E}_{\text{bal}}(f)$	21.15(±0.39)	42.83(±6.43)	32.30(±1.60)
Posthoc: $0.9 \cdot \mathcal{E}_{\text{bal}}(f) + 0.1 \cdot \mathcal{E}_{\text{DEO}}$	21.56(±0.53)	44.13(±9.39)	29.37(±2.45)
Algo. 1: with λ = 0.1	<b>15.50(±0.18)</b>	<b>30.33(±2.47)</b>	<b>4.25(±0.94)</b>

Table 8: Comparison of fairness metrics for group-imbalanced experiments, with 5 trials. Two additional baselines, DRO and post-hoc model, are evaluated.

666 ERM problem<sup>3</sup>

$$\min_{\alpha \in \mathcal{H}} \mathcal{L}_{\lambda}^{\mathcal{V}}(f_{\alpha}) \quad \text{WHERE} \quad \mathcal{L}_{\lambda}^{\mathcal{V}}(f) = \sum_{i=1}^R \lambda_i \mathcal{L}_i^{\mathcal{V}}(f_{\alpha}). \quad (\text{M-ERM})$$

667 Our theoretical analysis (provided below) of this setting is similar to the model-selection and cross-  
668 validation literature [34, 35, 58]. However, these works focus on a single objective. Unlike these,  
669 we will show that, small amount of validation data is enough to guarantee the pareto-efficiency of  
670 the train-validation split for all choices of  $\lambda \in \Lambda$ . The following assumption is used by earlier work  
671 and useful for studying continuous hyperparameter spaces. The basic idea is ensuring stability of the  
672 training algorithms. This has been verified for different hyperparameter types (e.g., ridge regression  
673 parameter, continuous parameterization of the neural architecture) under proper settings [58]. We  
674 remind that, our setting is also continuous as we use differentiable optimization to determine the best  
675 loss function.

676 **Assumption 1 (Training algorithm is stable)** Suppose  $\mathcal{H} \subset \mathbb{R}^h$ . There exists a partitioning of  $\mathcal{H}$   
677 into at most  $2^h$  sets  $(\mathcal{H}_i)_{i \geq 1}$  such that, over each set, the training algorithm  $A$  is locally-Lipschitz.  
678 That is, for all  $i$  and some  $L > 0$ , all pairs  $\alpha_1, \alpha_2 \in \mathcal{H}_i$  and inputs  $\mathbf{x}$  (over the support of  $\mathcal{D}$ ) satisfies  
679 that  $|f_{\alpha_1}(\mathbf{x}) - f_{\alpha_2}(\mathbf{x})| \leq L \|\alpha_1 - \alpha_2\|_{\ell_2}$ .

680 Here the Lipschitz constant  $L$  governs the stability level of the training algorithm. Note that, the  
681 partitioning is optional and it is included to account for discrete and discontinuous hyperparameter  
682 spaces. The following theorem shows that, if the training algorithm  $A$  satisfies stability conditions  
683 over  $\mathcal{H}$  and if the validation sample size  $n_{\mathcal{V}}$  is larger than  $R$  and the effective dimension of  $\mathcal{H}$ ,  
684 then (M-ERM) does return an approximately pareto-optimal solution uniformly over all choices of  
685  $(\lambda_i, \gamma_i)_{i=1}^R$ .

686 **Theorem 1 (Multi-objective generalization)** Suppose Assumption 1 holds. Let penalties  $\lambda$  take  
687 values over the sets  $\Lambda \subset \mathbb{R}^R$ . Assume the elements of the sets  $\mathcal{H}, \Lambda$  have bounded  $\ell_2$  norm. Suppose  
688 the loss functions have bounded derivatives (in absolute value) and, for some  $\Xi > 0$ , they are bounded  
689 as follows

$$\sup_{\lambda \in \Lambda} \left| \sum_{i=1}^R \lambda_i \ell_i(y, \hat{y}) \right| \leq \Xi.$$

<sup>3</sup>We believe the results can be stated for a mixture of regularizations and constraints (e.g. enforcing the condition  $\mathcal{L}_i^{\mathcal{V}}(f_{\alpha}) \leq \tau_i$ ). We opted to restrict our attention to regularization in consistence with the general setting of the paper.

690 Given  $\lambda$ , define the corresponding  $\hat{\alpha} = \arg \min_{\alpha \in \mathcal{H}} \mathcal{L}_\lambda^\vee(f_\alpha)$  solving (M-ERM). Then, with proba-  
 691 bility  $1 - 2e^{-t}$ , for all penalties  $\lambda \in \Lambda$ , the associated  $\hat{\alpha}_\lambda$  achieves the population multi-objective  
 692 risk

$$\mathcal{L}_\lambda(f_{\hat{\alpha}}) \leq \arg \min_{\alpha \in \mathcal{H}} \mathcal{L}_\lambda(f_\alpha) + \Xi \sqrt{\frac{\tilde{\mathcal{O}}(h + R + t)}{n_\vee}}. \quad (\text{C.1})$$

693 Here  $\tilde{\mathcal{O}}(\cdot)$  hides logarithmic terms. Specifically, the sample size grows only logarithmically in the  
 694 stability parameter of Assumption 1 (i.e.  $\log L$  factor).

695 **Interpretation:** In words, this result shows that, as soon as the validation sample size is larger  
 696 than  $\mathcal{O}(h + R)$ , train-validation split selects a model that is as good as the optimal model whose  
 697 hyperparameter is tuned over the test data. That is, as  $n_\vee$  grows, the multi-objective risk of  $f_{\hat{\alpha}}$   
 698 converges to risk of training with the optimal hyperparameter. In our context, it means the ability to  
 699 select the optimal loss function via validation. An important remark is that, this selection happens  
 700 regardless of the training phase and whether training risk overfits or not. That is, even if training  
 701 returns poor models, validation phase selects the best one (out of poor options). Importantly,  $h + R$  is  
 702 a small number in practice. For instance, for imbalanced loss function design,  $h$  is at most  $\mathcal{O}(K)$  as  
 703 we use three parameters for each class. If we cluster the classes, then  $h$  is in the order of clusters.  $R$   
 704 is typically 1 (e.g. balanced accuracy in Sec 5) or 2 (e.g. DEO/accuracy tradeoffs in Sec 6). However,  
 705 in the extreme case of optimizing a general combination of class-conditional risks,  $R$  can be as  
 706 large as number of classes  $K$ . A remarkable aspect of this result is that, we get multi-objective  
 707 pareto-efficiency of the validation-based optimization by using an extra  $\mathcal{O}(R)$  samples (compared to  
 708 single-loss scenario which requires  $\mathcal{O}(h)$  samples [58]).

709 **Proof** The strategy is based on applying a covering argument over all variables namely  $\alpha, \lambda$  and  
 710 can be seen as a multi-objective generalization of Theorem 1 of [58]. Let  $\mathcal{H}_\varepsilon, \Lambda_\varepsilon$  be  $\varepsilon$ -covers with  
 711 respect to  $\ell_2$ -norm of the corresponding sets  $\mathcal{H}, \Lambda$ . The size of these sets obey  $\log |\mathcal{H}_\varepsilon| \leq N_h(\varepsilon) =$   
 712  $h \log(B/\varepsilon)$  and  $\log |\Lambda_\varepsilon| \leq N_R(\varepsilon) = R \log(B/\varepsilon)$  where  $B > 0$  depends on the radius of  $\Lambda, \mathcal{H}$ .

713 To proceed, pick a pair  $\alpha, \lambda$  from the cover  $\mathcal{H}_\varepsilon, \Lambda_\varepsilon$ . Define the loss function  $\ell_\lambda = \sum_{i=1}^R \lambda_i \ell_i(y, \hat{y})$ .  
 714 Since this is bounded by  $\Xi$ , we can apply Hoeffding bound for the individual cover elements. Union  
 715 bounding these Hoeffding (or  $\Xi$ -sub-gaussian) concentration bounds over all cover elements, with  
 716 probability  $1 - 2e^{-t}$ , we have that

$$|\mathcal{L}_\lambda^\vee(f_\alpha) - \mathcal{L}_\lambda(f_\alpha)| \lesssim \Xi \sqrt{\frac{(h + R) \log(B/\varepsilon) + t}{n_\vee}}. \quad (\text{C.2})$$

717 **Perturbation analysis:** To proceed, given  $(\alpha, \lambda)$  from  $\mathcal{H}, \Lambda$ , choose an  $\varepsilon$ -neighboring point  $(\alpha', \lambda')$   
 718 from the cover  $\mathcal{H}_\varepsilon, \Lambda_\varepsilon$ . Let  $\Gamma > 0$  be the Lipschitz constant of the loss  $\ell_\lambda$  (specifically over worst  
 719 case  $\lambda$ ) and  $\bar{\Xi} = \sup_{1 \leq r \leq R} |\ell_i(y, \hat{y})|$ . We note that dependence on  $\Gamma, \bar{\Xi}$  will be only logarithmic.  
 720 Applying triangle inequalities, we find that

$$\begin{aligned} |\ell_\lambda(y, f_\alpha(\mathbf{x})) - \ell_{\lambda'}(y, f_{\alpha'}(\mathbf{x}))| &\leq |\ell_\lambda(y, f_\alpha(\mathbf{x})) - \ell_\lambda(y, f_{\alpha'}(\mathbf{x}))| + |\ell_\lambda(y, f_{\alpha'}(\mathbf{x})) - \ell_{\lambda'}(y, f_{\alpha'}(\mathbf{x}))| \\ &\leq \Gamma \|f_{\alpha'}(\mathbf{x}) - f_\alpha(\mathbf{x})\| + \left| \sum_{r=1}^R (\lambda_r - \lambda'_r) \ell_r(y, f_{\alpha'}(\mathbf{x})) \right| \\ &\leq \Gamma L \|\alpha' - \alpha\|_{\ell_2} + \bar{\Xi} \sqrt{R} \|\lambda - \lambda'\|_{\ell_2} \\ &\leq (\Gamma L + \bar{\Xi} \sqrt{R}) \varepsilon. \end{aligned}$$

721 This also implies that  $|\mathcal{L}_\lambda(f_\alpha) - \mathcal{L}_{\lambda'}(f'_\alpha)|, |\mathcal{L}_\lambda^\vee(f_\alpha) - \mathcal{L}_{\lambda'}^\vee(f'_\alpha)| \leq (\Gamma L + \bar{\Xi} \sqrt{R}) \varepsilon$ . Combining this  
 722 with (C.2), for all  $\lambda, \alpha$ , we obtained the uniform convergence guarantee

$$|\mathcal{L}_\lambda^\vee(f_\alpha) - \mathcal{L}_\lambda(f_\alpha)| \lesssim \Xi \sqrt{\frac{(h + R) \log(B/\varepsilon) + t}{n_\vee}} + 2(\Gamma L + \bar{\Xi} \sqrt{R}) \varepsilon.$$

723 Setting  $\varepsilon \rightarrow \frac{\Xi}{2(\Gamma L + \Xi\sqrt{R})\sqrt{n_{\mathcal{V}}}}$ , we find that

$$|\mathcal{L}_{\lambda}^{\mathcal{V}}(f_{\alpha}) - \mathcal{L}_{\lambda}(f_{\alpha})| \lesssim \Xi \sqrt{\frac{(h+R) \log(B(\Gamma L + \Xi\sqrt{R})\sqrt{n_{\mathcal{V}}}/\Xi) + t}{n_{\mathcal{V}}}} \quad (\text{C.3})$$

$$\leq \Xi \sqrt{\frac{\tilde{\mathcal{O}}(h+R+t)}{n_{\mathcal{V}}}}, \quad (\text{C.4})$$

724 where we dropped the logarithmic terms. To proceed, let  $\alpha_{*}$  be an optimal hyperparameter for the  
725 population risk of the validation phase i.e.  $\arg \min_{\alpha \in \mathcal{H}} \mathcal{L}_{\alpha}(f_{\alpha})$ . We find the generalization risk of  
726 the optimal  $\hat{\alpha}$  via

$$\mathcal{L}_{\hat{\alpha}}^{\mathcal{V}}(f_{\alpha}) \leq \mathcal{L}_{\alpha_{*}}^{\mathcal{V}}(f_{\alpha}) \leq \mathcal{L}_{\alpha_{*}}(f_{\alpha}) + \Xi \sqrt{\frac{\tilde{\mathcal{O}}(h+R+t)}{n_{\mathcal{V}}}},$$

727 concluding with the advertised result. ■

## 728 D Proof of Lemma 1

729 Let us recall the parametric cross-entropy loss function

$$\ell(y, f(\mathbf{x})) = w_y \log \left( 1 + \sum_{k \neq y} e^{l_k - l_y} \cdot e^{\Delta_k f_k(\mathbf{x}) - \Delta_y f_y(\mathbf{x})} \right) = -w_y \log \left( \frac{e^{\Delta_y f_y(\mathbf{x}) + l_y}}{\sum_{i \in [K]} e^{\Delta_i f_i(\mathbf{x}) + l_i}} \right).$$

730 Denote the labeling likelihood  $\eta_y(\mathbf{x}) = P(y | \mathbf{x})$ . When the weights  $w_y$  are not all ones, the class  
731 frequencies are effectively adjusted as  $\pi'_y \propto w_y \pi_y$ . Let  $\bar{\eta}_y(\mathbf{x})$  be the corresponding likelihood  
732 function. The optimal score function minimizing the cross-entropy loss is given by  $\Delta_y f_y^*(\mathbf{x}) +$   
733  $l_y = \log \bar{\eta}_y(\mathbf{x})$ . This choice is determined by minimizing the expected loss (given  $\mathbf{x}$ ) which sets  
734 the KL divergence between  $\bar{\eta}_y(\mathbf{x})$  and softmax output to zero. This leads to the decision rule  
735  $f_y^*(\mathbf{x}) = \bar{\Delta}_y \log \frac{\bar{\eta}_y(\mathbf{x})}{e^{l_y}}$  where  $\bar{\Delta}_y = \Delta_y^{-1}$  (to simplify the subsequent notation). Equivalently, the  
736 classification rule becomes

$$f_y^*(\mathbf{x}) = \log \left( \frac{\bar{\eta}_y(\mathbf{x})}{e^{l_y}} \right)^{\bar{\Delta}_y} \iff \text{rule}(\mathbf{x}) = \arg \max_{y \in [K]} \alpha_y \bar{\eta}_y^{\bar{\Delta}_y}(\mathbf{x}),$$

737 where  $\alpha_y = e^{-\bar{\Delta}_y l_y}$ . For standard accuracy, Bayes-optimal decision rule is  $\arg \max_{y \in [K]} \bar{\eta}_y(\mathbf{x})$   
738 and for balanced accuracy, it is  $\arg \max_{y \in [K]} \frac{\bar{\eta}_y(\mathbf{x})}{\pi_y}$ . In both cases, it can be written as  
739  $\arg \max_{y \in [K]} c_y \bar{\eta}_y(\mathbf{x})$  where  $c_y$  are adjustments.

740 We complete the proof by constructing a simple distribution that shows the rule( $\mathbf{x}$ ) is sub-optimal.  
741 Without losing generality, we may assume  $\Delta_1 \neq \Delta_2$  i.e., multiplicative adjustments of the first  
742 two classes differ. Given this multiplicative adjustment choice of  $\Delta$ , we will construct a simple  
743 distribution for which minimizing parametric CE don't result in Bayes-optimal decision. Specifically,  
744 we construct input features  $\mathbf{x}_1$  and  $\mathbf{x}_2$  so that the first two classes (labels  $y \in \{1, 2\}$ ) have the highest  
745 score  $c_y \bar{\eta}_y(\mathbf{x})$  and the top two scores are close. That is, for some arbitrarily small scalars  $\varepsilon, \varepsilon'$  we  
746 have  $c_1 \bar{\eta}_1(\mathbf{x}_1) = c_2 \bar{\eta}_2(\mathbf{x}_1) + \varepsilon$  and  $c_1 \bar{\eta}_1(\mathbf{x}_2) = c_2 \bar{\eta}_2(\mathbf{x}_2) + \varepsilon'$ . Additionally, set  $\bar{\eta}_i(\mathbf{x}_1) = \Gamma \bar{\eta}_i(\mathbf{x}_2)$   
747 for  $i = 1, 2$  and  $\Gamma \neq 1$  an arbitrary scalar.<sup>4</sup> Since  $\varepsilon \leq 0$  dictates the Bayes-optimal class decision  
748 ( $y = 1$  vs  $y = 2$ ), we need the score function  $f^*$  to satisfy

$$\frac{\bar{\eta}_1(\mathbf{x}_1)^{\bar{\Delta}_1}}{\bar{\eta}_2(\mathbf{x}_1)^{\bar{\Delta}_2}} \geq \frac{\alpha_2}{\alpha_1}, \quad \frac{\bar{\eta}_1(\mathbf{x}_2)^{\bar{\Delta}_1}}{\bar{\eta}_2(\mathbf{x}_2)^{\bar{\Delta}_2}} \geq \frac{\alpha_2}{\alpha_1}.$$

749 Letting  $\varepsilon, \varepsilon' \rightarrow 0$ , this implies that  $\frac{\bar{\eta}_1(\mathbf{x}_1)^{\bar{\Delta}_1}}{\bar{\eta}_2(\mathbf{x}_1)^{\bar{\Delta}_2}} = \frac{\bar{\eta}_1(\mathbf{x}_2)^{\bar{\Delta}_1}}{\bar{\eta}_2(\mathbf{x}_2)^{\bar{\Delta}_2}}$ . However, this contradicts with the initial  
750 assumption of  $\Gamma \neq 1$  via

$$\frac{\bar{\eta}_1(\mathbf{x}_1)^{\bar{\Delta}_1}}{\bar{\eta}_2(\mathbf{x}_1)^{\bar{\Delta}_2}} = \frac{\Gamma^{\bar{\Delta}_1} \bar{\eta}_1(\mathbf{x}_2)^{\bar{\Delta}_1}}{\Gamma^{\bar{\Delta}_2} \bar{\eta}_2(\mathbf{x}_2)^{\bar{\Delta}_2}} = \frac{\bar{\eta}_1(\mathbf{x}_2)^{\bar{\Delta}_1}}{\bar{\eta}_2(\mathbf{x}_2)^{\bar{\Delta}_2}} \iff \Gamma^{\bar{\Delta}_1 - \bar{\Delta}_2} = 1 \iff \Gamma = 1.$$

<sup>4</sup>Scaling the likelihoods by  $\Gamma$  doesn't affect  $\arg \max_y c_y \bar{\eta}_y(\mathbf{x})$  as the other classes are assigned small probabilities.

751 **E Proof of Lemma 2**

752 This lemma considers the solution of the binary parametric loss defined as

$$\ell(y, f_{\boldsymbol{\theta}}(\mathbf{x})) = w_y \cdot \log \left( 1 + e^{l_y} \cdot e^{-\Delta_y y f_{\boldsymbol{\theta}}(\mathbf{x})} \right).$$

753 Consider the ridge-constrained problem

$$\boldsymbol{\theta}_R = \arg \min_{\boldsymbol{\theta}} \sum_{i=1}^n \ell(y_i, f_{\boldsymbol{\theta}}(\mathbf{x}_i)) \quad \text{subject to} \quad \|\boldsymbol{\theta}\|_{\ell_2} \leq R. \quad (\text{E.1})$$

754 The ridgeless model described in Lemma 2 obtained by minimizing the parametric loss is given by  
 755 the limit  $\boldsymbol{\theta}_{\infty} = \lim_{R \rightarrow \infty} \boldsymbol{\theta}_R / R$ . Here, we focus on linear models  $f_{\boldsymbol{\theta}}(\mathbf{x}) = \boldsymbol{\theta}^T \mathbf{x}$ . The result will  
 756 be established by connecting the above loss to Cost Sensitive (CS)-SVM which enforces different  
 757 margins on classes. Fix  $\delta > 0$ . Define the CS-SVM problem as

$$\hat{\mathbf{w}}_{\delta} := \arg \min \|\mathbf{w}\|_2 \quad \text{subject to} \quad \begin{cases} \mathbf{w}^T \mathbf{x}_i \geq \delta & , y_i = 1 \\ \mathbf{w}^T \mathbf{x}_i \leq -1 & , y_i = -1 \end{cases}, \quad i \in [n]. \quad (\text{E.2})$$

758 Assume the spherical data-augmentation with radii  $\varepsilon_{\pm}$  for the two classes. Then, standard SVM on  
 759 the augmented data solves

$$\min \|\mathbf{w}\|_2 \quad \text{subject to} \quad \begin{cases} \mathbf{w}^T \mathbf{x}_i - \varepsilon_+ \|\mathbf{w}\|_2 \geq 1 & , y_i = +1 \\ \mathbf{w}^T \mathbf{x}_i + \varepsilon_- \|\mathbf{w}\|_2 \leq -1 & , y_i = -1 \end{cases}, \quad i \in [n]. \quad (\text{E.3})$$

760 Here, the first observation is that, since ridgeless logistic loss is equivalent to SVM<sup>5</sup> [60], ridgeless  
 761 logistic loss with the augmented data also converges to the solution of (E.3).

762 For the loss function above, fix  $\delta = \Delta_- / \Delta_+ > 0$ . Applying Proposition 1 of [40],  $\boldsymbol{\theta}_{\infty}$  coincides  
 763 with the ( $\ell_2$  normalized) solution of the CS-SVM i.e.

$$\hat{\mathbf{w}}_{\delta} := \arg \min \|\mathbf{w}\|_2 \quad \text{subject to} \quad \begin{cases} \mathbf{w}^T \mathbf{x}_i \geq 1/\Delta_+ & , y_i = 1 \\ \mathbf{w}^T \mathbf{x}_i \leq -1/\Delta_- & , y_i = -1 \end{cases}, \quad i \in [n]. \quad (\text{E.4})$$

for  $\Delta_+, \Delta_- > 0$ . Note that, without losing generality, we can assume  $\Delta_{\pm} < 1$  by preserving the  
 ratio to  $\delta$  as it doesn't change the classification rule. We will prove that  $\hat{\mathbf{w}}_{\delta}$  is optimal in (E.3) for the  
 following choice of  $\varepsilon_{\pm}$ :

$$\varepsilon_{\pm} := \frac{1/\Delta_{\pm} - 1}{\|\hat{\mathbf{w}}_{\delta}\|_2}.$$

764 This will in turn conclude that ridgeless augmented logistic regression is equivalent to ridgeless  
 765 regression with parametric cross-entropy.

766 **Proof of optimality of  $\hat{\mathbf{w}}_{\delta}$  for (E.3).** To prove the claim let  $\hat{\alpha}_i, i \in [n]$  be the dual variables associated  
 767 with (E.4) corresponding to the minimizer  $\hat{\mathbf{w}}_{\delta}$ . By KKT conditions it holds that  $(\{\hat{\alpha}_i\}_{i \in [n]}, \hat{\mathbf{w}}_{\delta})$  is a  
 768 solution to:

$$\sum_{i \in [n]} \alpha_i y_i \mathbf{x}_i = \mathbf{w} / \|\mathbf{w}\|_2, \quad \alpha_i \geq 0, \quad \alpha_i \mathbf{x}_i^T \mathbf{w} = \begin{cases} \frac{\alpha_i}{\Delta_+} & , y_i = +1 \\ -\frac{\alpha_i}{\Delta_-} & , y_i = -1 \end{cases}, \quad i \in [n] \quad (\text{E.5})$$

Set

$$\hat{\beta}_i = \left( \frac{1}{1 - \varepsilon_+ \sum_{i: y_i = +1} \hat{\alpha}_i - \varepsilon_- \sum_{i: y_i = -1} \hat{\alpha}_i} \right) \hat{\alpha}_i$$

769 With these it only takes a few algebra steps to verify that  $(\{\hat{\beta}_i\}_{i \in [n]}, \hat{\mathbf{w}}_{\delta})$  is a solution to:

$$\frac{\mathbf{w}}{\|\mathbf{w}\|_2} - \sum_i \beta_i y_i \mathbf{x}_i + \varepsilon_+ \left( \sum_{i: y_i = +1} \beta_i \right) \frac{\mathbf{w}}{\|\mathbf{w}\|_2} + \varepsilon_- \left( \sum_{i: y_i = -1} \beta_i \right) \frac{\mathbf{w}}{\|\mathbf{w}\|_2} = 0 \quad (\text{E.6})$$

$$\beta_i \geq 0, \quad \beta_i \mathbf{x}_i^T \mathbf{w} = \begin{cases} \beta_i (1 + \varepsilon_+ \|\mathbf{w}\|_2) & , y_i = +1 \\ \beta_i (-1 - \varepsilon_- \|\mathbf{w}\|_2) & , y_i = -1. \end{cases} \quad (\text{E.7})$$

<sup>5</sup>In the sense that  $\ell_2$  normalized solution of logistic regression with infinitesimal ridge is equal to the  $\ell_2$  normalized SVM solution.

In particular, to verify  $\hat{\beta}_i \geq 0$ ,  $i \in [n]$  we used that from the optimality of the primal-dual pair  $(\{\hat{\alpha}_i\}_{i \in [n]}, \hat{\mathbf{w}}_\delta)$ :

$$\sum_{i:y_i=+1} \frac{\hat{\alpha}_i}{\Delta_+} + \sum_{i:y_i=-1} \frac{\hat{\alpha}_i}{\Delta_-} = \|\hat{\mathbf{w}}_\delta\|_2,$$

770 and the definition of  $\varepsilon_+, \varepsilon_i$ . This completes the proof as it can be checked that the above corresponds  
771 exactly to the KKT conditions of (E.3).