Invariant Feature Learning for Generalized Long-Tailed Classification

Kaihua Tang¹, Mingyuan Tao², Jiaxin Qi¹, Zhenguang Liu³, Hanwang Zhang¹

¹ Nanyang Technological University, Singapore
² Damo Academy, Alibaba Group, Hangzhou, China
³ Zhejiang University, Hangzhou, China
kaihua.tang@ntu.edu.sg; juchen.tmy@alibaba-inc.com; jiaxin003@e.ntu.edu.sg
liuzhenguang2008@gmail.com; hanwangzhang@ntu.edu.sg

Abstract. Existing long-tailed classification (LT) methods only focus on tackling the class-wise imbalance that head classes have more samples than tail classes, but overlook the attribute-wise imbalance. In fact, even if the class is balanced, samples within each class may still be long-tailed due to the varying attributes. Note that the latter is fundamentally more ubiquitous and challenging than the former because attributes are not just implicit for most datasets, but also combinatorially complex, thus prohibitively expensive to be balanced. Therefore, we introduce a novel research problem: Generalized Long-Tailed classification (GLT), to jointly consider both kinds of imbalances. By "generalized", we mean that a GLT method should naturally solve the traditional LT, but not vice versa. Not surprisingly, we find that most class-wise LT methods degenerate in our proposed two benchmarks: ImageNet-GLT and MSCOCO-GLT. We argue that it is because they over-emphasize the adjustment of class distribution while neglecting to learn attributeinvariant features. To this end, we propose an Invariant Feature Learning (IFL) method as the first strong baseline for GLT. IFL first discovers environments with divergent intra-class distributions from the imperfect predictions, and then learns invariant features across them. Promisingly, as an improved feature backbone, IFL boosts all the LT lineup: one/two-stage re-balance, augmentation, and ensemble. Codes and benchmarks are available on Github: https://github.com/KaihuaTang/ Generalized-Long-Tailed-Benchmarks.pytorch

Keywords: Data Imbalance, Generalized Long-Tailed Classification

1 Introduction

Long-Tailed classification (LT) [68] is inevitable in real-world training, as longtailed distribution ubiquitously exists in data at scale [40,39] and it is often prohibitively expensive to balance against such nature [40]. For example in Fig. 1 (a), a frequent class such as "dog" has significantly more samples than a rare one such as "panda". Prevailing LT methods are essentially based on adjusting the **class-wise** imbalance ratio: given a biased classifier that tends to classify tail



Fig. 1. (a) The real-world long-tailed distribution is both class-wise and attribute-wise imbalanced; (b) even if we balance the class distribution of MSCOCO-Attribute [38], the attributes are still long-tailed

as head, we curb the head confidence while lifting the tail confidence, so the resultant classifier is expected to be fair to both head and tail during inference [19,71,6,33,41].

However, we'd like to point out that LT challenge cannot be simply characterized by class-wise imbalance. If we take a closer look at samples inside each class in Fig. 1 (a), we can find that attributes within each class are also long-tailed. This **attribute-wise** imbalance⁴ undermines the robustness of the classifier in two ways: **I**) it hurts the accuracy of images with tail attributes, *e.g*, members of minority groups in the human class are easier to be mis-classified than their majority counterparts [36], despite the fact that they both come from the same head class; **II**) it results in some attributes being mistakenly correlated to certain classes, *e.g*, images of a head class "tractor" are often captured on "field", so when a tail object like "harvester" is also captured on "field", its risk of being mis-classified as "tractor" is much higher, which is supported by our formulation Eq. (1) and visualization in Fig. 2 (b). Therefore, the **attribute-wise** imbalance further explains the cause of inconsistent performances within the same class and the existence of spurious correlations.

In fact, even if the class-wise imbalance is entirely eliminated like Fig. 1 (b), its attribute-wise imbalance still persists and hurts the generalization. Besides, strictly balancing attributes is not only prohibitive but also impossible due to the innumerable multi-label combinations of attributes, making the attribute-wise imbalance fundamentally different from the class-wise imbalance. To this end, we present a new task: **Generalized Long-Tailed** classification (GLT), to unify the challenges from both class-wise and attribute-wise imbalances. For rigorous and reproducible evaluations in the community, as detailed in Section 3, we introduce two benchmarks, ImageNet-GLT and MSCOCO-GLT, together with

 $^{^{4}}$ In this paper, the attribute represents all the factors causing the intra-class variations, including object-level attributes (colors, textures, postures, *etc*) and imagelevel attributes (lighting, contexts, *etc*).

three protocols to evaluate the robustness of models against class-wise long tail, attribute-wise long tail and their joint effect.

Not surprisingly, we find that nearly all existing LT methods [72,33,19,41] fail to tackle the attribute-wise imbalance in GLT (cf. Section 5). The reasons are two-fold: **I**) They rely on class-wise adjustment, which requires the access to the class statistical traits to re-balance [49,6,19,71]. Unfortunately, the attributewise traits are hidden in GLT, whose discovery *per se* is a challenging open problem [26]. **II**) As illustrated in Fig. 2, the cross-entropy baseline (biased classifier) tends to predict tail samples as the head class with similar attributes, resulting in low precision on the head and low accuracy on the tail. So, the success of LT methods is mainly based on lifting the tail class boundary to welcome more samples to increase the tail accuracy. However, such adjustment is only playing with the precision-accuracy trade-off [73], leaving the confused region of similar attributes unchanged in the feature space, while the proposed GLT requires algorithms to ignore those confusing attributes.

To this end, in Section 4, we introduce a framework called Invariant Feature Learning (IFL) to address the attribute-wise imbalance and serve as the first strong baseline for GLT. Our motivation is based on the reasonable assumption that: since the class feature is invariant to its attributes, e.q, a "dog" is always a dog regardless of its varying attributes, the variation of attributes is the main cause of lower prediction confidence within each class. Note that corrupted images are beyoud the scope of this paper, as they don't have any valid labels. Therefore, we use the current classification confidence of each training sample as an imbalance indicator of attributes inside the class: the lower the confidence is, the rarer the attributes are. Then, we sample a



Fig. 2. In ImageNet-GLT, a typical LT method, (b) LWS [19], is playing a precision-accuracy trade-off with the (a) biased cross-entropy baseline, while the proposed (c) IFL improves both metrics at the same time. We follow [31,19] to stratify classes into Many, Medium, and Few by the class frequency. The t-SNE [14] of image features further illustrates that IFL features reduce the confusing region between "tractor" and "harvester" caused by the shared attribute "field"

new environment based on the reversed confidence. Together with the original one, we obtain two environments with diverse attribute distributions for each class. Finally, to remove the imbalance impact of attributes, we design a metric learning loss, extending the center loss [59] to its Invariant Risk Minimization

(IRM) [3] version, which equips the model with the ability to learn class features that are invariant to attributes.

As shown in Fig. 2 (c), IFL improves both precision and accuracy under GLT by reducing the confusing region of attributes in the feature space. Besides, IFL, as an improved feature backbone, can be seamlessly incorporated into other LT methods. In particular, we find that by only using sample augmentation such as MixUp [66] or RandAug [10], IFL can surpass most of the LT methods. We also notice that the recent progress of LT methods [56,67,73], who claim to improve both head and tail, is indeed attributed to tackle the attribute-wise imbalance—they deserve to be more fairly evaluated by GLT.

Our contributions can be summarized as follows:

- We present a new challenging task: Generalized Long-Tailed classification (GLT), together with two benchmarks: ImageNet-GLT and MSCOCO-GLT. To solve GLT, one need to address both the conventional class-wise imbalance and the ever-overlooked attribute-wise imbalance.
- We develop Invariant Feature Learning (IFL) as the first strong GLT baseline. Its effectiveness demonstrates that learning attribute-invariant features is a promising direction.
- By extensive experiments, we show that IFL improves all the prevalent LT line-up on GLT benchmarks: one/two-stage re-balancing [19,50,41,33,71,6], augmentation [66,10], and ensemble [56,67].

2 Related Work

Long-Tailed Classification [72,31,11,68] aims to improve the performance under class-wise balanced evaluation given the class-wise long-tailed training data. Previous methods can be categorized into three types: 1) one/two-stage re-balancing algorithms [19,16,48,25,49,54,15] apply statistical adjustment based on the explicit class distribution to correct the tail class bias; 2) data augmentation either independently augments all samples [63,20,66,10] or transfer head information to the tail [30,13,24,58]; 3) ensemble [56,5,61,67] is recently explored as a strategy to improve head and tail categories at the same time. The conventional LT classification is essentially a special case of the proposed GLT, as solving GLT will naturally improve LT, but not vice versa.

Domain Adaptation (DA) and Out-of-distribution Generalization (O-ODG) are two other related tasks. DA [60,52,70,65] seeks to transfer models from source domains to the target domain. The difference between DA and GLT is that we don't need a subset of samples from the target domain. Recent papers [18,74] also notice the intrinsic correlation between LT and DA. In OODG [69,2,22,3], we desire a machine trained in one domain to work well in any domain. Recent studies [2] show that it cannot be addressed without any assumption. Therefore, regarding the type of assumptions, OODG can be divided into DA [52], domain generalization [23], long-tailed classification [68], zero-/few-shot learning [47,57,55], and even adversarial robustness [7,8]. Our GLT can also be viewed as a special case of OODG that is more general than LT.

Attribute-wise Imbalance itself is also a long-standing research field in the name of hard example mining [28,35], sub-population shift [43,27,21], spurious correlation [45,1], etc. Meanwhile, the corresponding methods in these fields may also inspire the future research in GLT, *e.g*, EIIL [9], GEORGE [44] and SDB [17] can be used to construct better environments or baselines. Compared with these fields, the proposed GLT benchmarks provide a unified formulation and benchmarks for both class-wise and attribute-wise imbalances.

3 Generalized Long-Tailed Classification

Previous LT methods [41,33] formulate the classification model as p(Y|X), predicting the label Y from the input image X, which can be further decomposed into $p(Y|X) \propto p(X|Y) \cdot p(Y)$ [33,41]. This formulation identifies the cause of class-wise bias as p(Y), so it can be elegantly solved by Logit Adjustment [33]. However, such a formulation is based on a strong assumption that the distribution of p(X|Y) won't change in different domains, *i.e.*, $p_{\text{train}}(X|Y) = p_{\text{test}}(X|Y)$, which cannot be guaranteed in real-world applications. Next, we will provide an in-depth analysis of why this over-simplified view fails to explain all biases.

3.1 **Problem Formulation**

Recent studies [34,4,53] demonstrate that an image of object X can be fully described or generated by its class and a list of attributes. That is to say, each X is generated by a set of underlying $(z_c, z_a)^5$, where the class-specific components z_c are the invariant factors that enable the existence of the robust classification and the attribute-related variables z_a are domain-specific knowledge that have inconsistent distributions. This formulation only assumes the invariance of a subset features z_c rather than the entire X, *i.e.*, $p_{\text{train}}(z_c|Y) = p_{\text{test}}(z_c|Y)$. Therefore, we can follow the Bayes theorem [46] to convert the classification model $p(Y|X) = p(Y|z_c, z_a)$ into the following formula:

$$p(Y = k|z_c, z_a) = \frac{p(z_c|Y = k)}{p(z_c)} \cdot \underbrace{\frac{p(z_a|Y = k, z_c)}{p(z_a|z_c)}}_{attribute \ bias} \cdot \underbrace{\frac{p(Y = k)}{class \ bias}}_{class \ bias},$$
(1)

where invariant components z_c only depend on Y; descriptive attributes z_a that vary across instances may depend on both Y and z_c . We generally consider $p(z_c, z_a) = p(z_a|z_c) \cdot p(z_c)$ WITHOUT introducing any independence assumption. Note that we also DO NOT impose the disentanglement assumption that a perfect feature vector $\mathbf{z} = [z_c; z_a]$ with separated z_c and z_a can be obtained, as the disentanglement is a challenging task on its own [32]. Otherwise, we only need to conduct a simple feature selection to obtain the ideal classification model.

⁵ In this paper, z_c and z_a stand for **all** class-specific components and variant attributes, respectively, but we use a single variable to represent them in the following examples, $e.g, z_c = feather$ and $z_a = brown$, for simplicity.

6

The reason why we need Eq. (1) to replace the simple p(Y|X) is because unlike those tasks that indeed require both z_c and z_a , *e.g*, the image captioning [64] or the segmentation task [12], the classification task merely relies on the class-specific components z_c of an image X, regardless of its varying attributes z_a . Therefore, the former formulation p(Y|X) over-simplifies the problem by ignoring the different roles between z_c and z_a during classification.

Class Bias: in class-wise LT [68], the distribution of p(Y) is considered as the main cause of the performance degradation. As p(Y) can be explicitly calculated from the training data, the majority of previous LT methods directly alleviate its effect by class-wise adjustment [41,33]or re-balancing [19,71]. However, they fail to answer **I**) why the performance is also longtailed within each class, and **II**) why tail images tend to be misclassified as certain head classes with similar attributes.

Attribute Bias: the above Eq.(1) extends the previous LT formulation by introducing the attribute bias caused by long-





Fig. 3. Examples of how to balance the attribute distribution for the Test-GBL evaluation environment in the proposed two benchmarks

tailed z_a , which not only explains the cause of inconsistent performances within each class but also demonstrates how spurious correlations hurt the prediction. Intuitively, **I**) for the intra-class variation, if an attribute "white"/"brown" of the class-specific component "feather" is more frequent/rarer in "dove" than in other classes, e.g, $\frac{p(z_a=brown|Y=dove,z_c=feather)}{p(z_a=brown|z_c=feather)} < \frac{p(z_a=white|Y=dove,z_c=feather)}{p(z_a=white|z_c=feather)}$, a dove with brown feather will have lower confidence than doves with white feather, following Eq. (1), *i.e.*, $p(Y=dove|z_c=feather, z_a=brown) < p(Y=dove|z_c=feather, z_a=white)$. **II**) Similarly, for spurious correlations, if a "field" background is more frequent in the class "tractor" than other vehicles with the same class-specific component "wheel", it will create false sense of correlation between "field" and "tractor" class, *i.e.*, $\frac{p(z_a=field|Y=tractor,z_c=wheel)}{p(z_a=field|z_c=wheel)} >>$ 1, resulting images of other vehicles with "wheel" blindly have larger $p(Y=tractor|z_c=wheel, z_a=field)$ in the "field" background, e.g. predicting a "harvester in field" as the "tractor".

Generalized Long-Tailed Distribution: the proposed GLT asserts that both the conventional class distribution and the ever-overlooked attribute distribution are long-tailed in real-world dataset at scale. However, most of the previous LT benchmarks, ImageNet-LT [31], Long-Tailed CIFAR-10/-100 [71], or iNaturalist [51], are only capable of evaluating the class bias, underestimating the role

of the attribute bias in the long tail challenge. To better study unbiased models that address both of biases at the same time, we introduce the following two GLT benchmarks and three evaluation protocols.

3.2 GLT Benchmarks and Evaluation Protocols

In this section, we design two benchmarks, ImageNet-GLT and MSCOCO-GLT, for the proposed GLT challenge, where ImageNet-GLT is a long-tailed version of ImageNet [42] and MSCOCO-GLT is constructed from MSCOCO-Attribute [38]. Although there are explicit attribute annotations in MSCOCO-Attribute, we forbid the access of them during training to make the algorithm more general. After all, attributes are not exhaustively annotated. Meanwhile, ImageNet, like most of the other datasets, doesn't have any attribute annotation, so we cluster image features within each class into multiple "pretext attributes" using a pre-trained model [37]. These clusters can thus serve as annotations for implicit attributes. To systematically diagnose two kinds of biases in Eq. (1), each benchmark is further organized into three evaluation protocols as follows:

Class-wise Long Tail (CLT) Protocol: same as the conventional LT, we first adopt a class-wise and attribute-wise LT training set, called **Train-GLT**, which can be easily sampled from ImageNet [42] and MSCOCO-Attribute [38] using a class-wise LT distribution. We don't need to intentionally ensure the attributewise imbalance as it's ubiquitous and inevitable in any real-world dataset, *e.g.*, the distribution of MSCOCO-Attribute [38] in Fig. 1 (b). The corresponding **Test-CBL**, which is i.i.d. sampled within each class, is a class-wise balanced and attribute-wise long-tailed testing set. (**Train-GLT**, **Test-CBL**) with the same attribute distributions and different class distributions can thus evaluate the robustness against the class-wise long tail.

Attribute-wise Long Tail (ALT) Protocol: the training set Train-CBL of this protocol has the same number of images for each class and keeps the original long-tailed attribute distribution by i.i.d. sampling images within each class, so its bias only comes from the attribute. Meanwhile, **Test-GBL**, as the most important evaluation environment for GLT task, has to balance both class and attribute distributions. As illustrated in Fig. 3, Test-GBL for ImageNet-GLT samples equal number of images from each "pretext attribute" (*i.e.*, feature clusters) and each class. Test-GBL for MSCOCO-GLT is a little bit tricky, because each object has multiple attributes, making strictly balancing the attribute distribution prohibitive. Hence, we select a fixed size of subset within each class that has the minimized standard deviation of attributes as the Test-GBL. As long as Test-GBL is relatively more balanced in attributes than Train-CBL, it can serve as a valid testing set for ALT protocol. In summary, (**Train-CBL, Test-GBL**) have the same class distributions and different attribute distributions.

Generalized Long Tail (GLT) Protocol: this protocol combines (Train-GLT, Test-GBL) from the above, so both class and attribute distributions are changed from training to testing. As the generalized evaluation protocol for the long-tailed challenge, an algorithm can only obtain satisfactory results when both class bias and attribute bias are well addressed by the final model.

4 Invariant Feature Learning

Prevalent LT studies [33,19] mainly focus on designing balanced classifiers, as the p(Y = k) in Eq. (1) is independent of input features. However, as we discussed in Section 3, our new attribute bias is caused by extracting the undesired z_a , so the overlooked feature learning is the key. Therefore, the proposed method aims to learn an improved feature backbone that is complementary for previous balanced classifiers [19,50,41,33].

As we discussed in Section 3, the attribute bias term in Eq. (1) raises two problems: I) the inconsistent performance of $p(Y|z_c, z_a)$ within each class, and II) the spurious correlations between a non-robust attribute and a class. To tackle these problems, we propose Invariant Feature Learning (IFL) that extends the center loss [59] from the original Empirical Risk Minimization (ERM) to its Invariant Risk Minimiza-



Fig. 4. The proposed IFL that first applies different sampling strategies according to the prediction confidence within each class, then combines them to construct environments with diverse attribute distributions, and finally learns features invariant to the environment change by the IFL metric loss

tion (IRM) [3] version, which forces the backbone to focus on learning features that are invariant across a set of environments.

The rationale behind the IFL is that $\frac{p(z_c|Y=k)}{p(z_c)}$ is consistent in all environments while $\frac{p(z_a|Y=k,z_c)}{p(z_a|z_c)}$ is not. If images of Y=k have the same feature center in all environments, it means that the extracted features are more related to invariant z_c instead of z_a . The overall IFL framework is summarized in Fig. 4.

4.1 Environment Construction

After warming-up the model with the vanilla cross-entropy loss for several epochs, we obtain an initial model with imperfect predictions. Since $\frac{p(z_c|Y=k)}{p(z_c)}$ and p(Y=k) are constants within a given class Y=k, the variation of $p(Y=k|z_c, z_a)$

is directly proportional to $\frac{p(z_a|Y=k,z_c)}{p(z_a|z_c)}$. It allows us to use the prediction confidence on the ground-truth class p(Y=k|X in k) as the indicator to sample diverse attribute distributions in the training set (See Appendix for more discussions). As illustrated in Fig. 4, samples collected from each class using the same sampling strategy will then be combined together to form different environments in \mathcal{E} . Each sampling strategy has a unique distribution of $\frac{p(z_a|Y=k,z_c)}{p(z_a|z_c)}$. The environments are periodically updated during the training phase.

In general, two environments are considered to be sufficient to learn the robustness against the environment change [3], which is also supported by our experiments on Table 4. Specifically, one environment directly samples each instance uniformly inside each class, *i.e*, the naïve i.i.d. sampling. The other environment is sampled based on the $(1 - p(Y = k|z_c, z_a))^{\beta}$, where β automatically adjusts the new attribute distribution by up-sampling 20% images from class Y=k with lowest $p(Y = k|z_c, z_a)$ to reach 80% population of the class Y=k in the new environment, following the Pareto Principle [40].

The pseudo code of the algorithm is provided in the Appendix.

4.2 IFL Metric Loss

After obtaining a set of diverse training environments \mathcal{E} , the goal of the proposed IFL can thus be formulated as the following optimization problem:

$$\min_{\theta,w} \sum_{e \in \mathcal{E}} \sum_{i \in e} L_{cls}(f(x_i^e; \theta), y_i^e; w),$$

subject to $\theta \in \arg\min_{\theta} \sum_{e \in \mathcal{E}} \sum_{i \in e} ||f(x_i^e; \theta) - C_{y_i^e}||_2,$ (2)

where θ and w are learnable parameters for the backbone and classifier, respectively; (x_i^e, y_i^e) are *i*-th (image, label) pair in the training environment $e \in \mathcal{E}$; $f(x_i^e; \theta)$ is the backbone extracting feature from x_i^e ; $L_{cls}(f(x_i^e; \theta), y_i^e; w)$ is the cross-entropy loss; $C_{y_i^e}$ is the mean feature of corresponding class y_i^e across all environments in $\mathcal{E}^{.6}$

The above Eq. (2) aims to optimize the model by the classification loss L_{cls} under the constraint that the intra-class variation of features across all environments is also minimized, suppressing the learning of z_a that causes nonrobust $\frac{p(z_a|Y=k,z_c)}{p(z_a|z_c)}$. So the overall training objective is thus defined as $L = L_{cls} + \alpha \cdot L_{IFL}$, where $L_{IFL} = ||f(x_i^e; \theta) - C_{y_i^e}||_2$ is a metric loss that ensures the above constraint and α is the trade-off parameter.

As the generalized IRM version of the center loss [59], the proposed IFL significantly boosts the GLT performance from its ERM counterpart using single environment. It's because class centers are also biased under the long-tailed attributes, so the original center loss would inevitably under-represent images with rare attributes. More experimental analyses are given in Section 5.4.

 $^{^{6}}$ We follow the center loss [59] to implement $C_{y_{i}^{e}}$ as the moving average for efficiency.

Table 1. Evaluation of CLT and GLT Protocols on ImageNet-GLT: Accuracy (*left in each cell*) and Precision (*right in each cell*) are reported. All methods are reimplemented under the same codebase with ResNext-50 backbone

| _ | | | | | | | | | | | |
|------|------------------------|-------------------------------------|-----------------|------------------------|----------------------|--------------------------------------|--------------------------------------|----------------------|--------------------|--|--|
| _ | Methods | Class-Wise Long Tail (CLT) Protocol | | | | Generalized Long Tail (GLT) Protocol | | | | | |
| < | Accuracy Precision > | $Many_C$ $Medium_C$ Few_C Overall | | Overall | $Many_C$ | $Medium_C$ | fedium _C Few _C | | | | |
| | Baseline | 59.34 39.08 | 36.95 52.87 | 14.39 56.65 | 42.52 47.92 | 50.98 32.90 | 28.49 44.72 | 10.28 49.11 | 34.75 40.65 | | |
| | cRT [19] | 56.55 45.79 | 42.89 46.23 | 26.67 41.47 | 45.92 45.34 | 48.02 38.40 | 34.16 38.07 | 19.92 33.50 | 37.57 37.51 | | |
| | LWS [19] | 55.38 46.67 | 43.91 46.87 | 30.11 40.92 | 46.43 45.90 | 47.15 39.16 | 34.88 38.68 | 22.56 32.88 | 37.94 38.01 | | |
| | Deconfound-TDE [50] | 54.94 49.27 | 43.18 43.91 | 28.64 33.40 | 45.70 44.48 | 46.87 42.39 | 34.43 35.77 | 22.11 26.30 | 37.56 37.00 | | |
| e | BLSoftmax [41] | 55.60 48.19 | 42.74 47.27 | 28.79 38.14 | 45.79 46.27 | 47.15 40.89 | 33.48 39.11 | 21.10 27.50 | 37.09 38.08 | | |
| an | Logit-Adj [33] | 54.55 49.70 | 44.40 45.05 | 31.53 36.04 | 46.53 45.56 | $45.94 \mid 41.97$ | 35.15 36.63 | 24.07 28.59 | 37.80 37.56 | | |
| al | BBN [71] | 61.64 42.74 | 43.80 54.44 | 13.94 55.12 | 46.46 49.86 | $52.41 \mid 35.58$ | 34.31 46.38 | 10.06 44.43 | 37.91 41.77 | | |
| 7 | LDAM [6] | 59.05 45.39 | 43.23 48.80 | 24.44 44.99 | 46.74 46.86 | 51.02 38.78 | 34.13 40.39 | 18.46 35.91 | 38.54 39.08 | | |
| R | (ours) Baseline + IFL | 62.71 42.9 | 3 40.10 56.83 | 18.92 61.92 | 45.97 52.06 | 54.09 36.74 | 31.73 49.03 | 13.62 51.42 | 37.96 44.47 | | |
| | (ours) cRT + IFL | 61.27 45.84 | 43.96 51.67 | 24.32 53.64 | 47.94 49.63 | 52.75 39.11 | 35.14 43.36 | 17.92 43.35 | 39.60 41.65 | | |
| | (ours) LWS + IFL | 61.50 45.43 | 43.79 52.85 | 23.86 55.58 | 47.89 50.29 | 53.21 38.92 | 34.99 44.44 | 17.42 45.90 | 39.64 42.45 | | |
| | (ours) BLSoftmax + IFL | 58.00 53.70 | 44.70 51.73 | 33.49 37.58 | 48.34 50.39 | 49.92 46.86 | 36.11 44.31 | 25.71 32.01 | 40.08 43.48 | | |
| | (ours) Logit-Adj + IFL | 56.96 56.2 | 2 46.54 50.10 | 36.88 33.29 | 49.26 50.02 | $48.25 \mid 49.17$ | 37.50 41.65 | 29.00 25.77 | 40.52 42.28 | | |
| nt | Mixup [66] | 59.68 37.96 | 30.83 55.74 | 7.09 34.33 | 38.81 45.41 | 51.04 31.85 | 23.10 47.25 | 4.94 22.88 | 31.55 37.44 | | |
| ne | RandAug [10] | 64.96 42.63 | 40.30 59.10 | 15.20 56.60 | 46.40 52.13 | 56.36 35.97 | 31.43 51.13 | 10.36 48.92 | 38.24 44.74 | | |
| 疑 | (ours) Mixup + IFL | 67.71 47.77 | 45.87 62.58 | 24.71 67.77 | 51.43 57.44 | 59.36 40.95 | 36.77 54.67 | 18.06 55.10 | 43.00 49.25 | | |
| Ā | (ours) RandAug + IFL | $69.35 \mid 49.4$ | 2 48.05 63.1 | 9 26.92 66.04 | $53.40 \mid 58.11$ | $60.79 \mid 42.41$ | $39.07 \mid 55.15$ | $20.04 \mid 57.90$ | $44.90 \mid 50.47$ | | |
| le | TADE [67] | 58.44 56.3 | 8 48.01 51.41 | 36.60 41.08 | 50.47 51.85 | 50.29 49.25 | 38.74 43.74 | 27.99 31.75 | 41.75 44.15 | | |
| - fe | RIDE [56] | 64.04 51.91 | 48.66 53.21 | 30.44 46.25 | 52.08 51.65 | 55.47 44.55 | 38.65 44.26 | 22.80 37.26 | 43.00 43.32 | | |
| sei | (ours) TADE + IFL | 61.71 55.59 | 48.87 53.42 | 34.02 40.93 | 51.78 52.41 | 53.75 48.73 | 39.90 45.28 | 26.77 35.34 | 43.47 45.17 | | |
| Ē | (ours) RIDE + IFL | 65.68 54.1 | 3 50.82 56.2 | 2 31.91 52.10 | 53.93 54.76 | $57.84 \mid 47.00$ | 41.80 48.65 | 24.63 42.96 | $45.64 \mid 47.14$ | | |

5 Experiments

5.1 Datasets and Metrics

ImageNet-GLT is a long-tailed version of the ImageNet [42], where CLT and GLT protocols share the same training set Train-GLT with 113k samples over 1k classes. ALT protocol adopts a class-wise balanced Train-CBL with 114k images. The evaluation splits {Val, Test-CBL, Test-GBL} have {30k, 60k, 60k} samples, respectively. The number of images for each class in Train-GLT ranges from 570 to 4, while all classes have 114 samples in Train-CBL. To collect the attribute-wise balanced Test-GBL, images from each class were clustered into 6 groups by KMeans using a pre-trained ResNet50 model [37] and we sampled 10 images per group and class. Following [31,19], all testing sets are also split into 3 subsets by the class frequency: Many_C with #sample > 100, Medium_C with $100 \geq #sample \geq 20$, and Few_C with #sample < 20. We further split them into 3 subsets by attribute groups: Many_A, Medium_A, and Few_A with images from the most/medium/least frequent 2 clusters of all classes, respectively.

MSCOCO-GLT is a long-tailed subset of MSCOCO-Attribute [38,29] with 196 different attributes. We cropped each object with multi-label attributes as independent images. Under CLT and GLT protocols, we have {Train-GLT, Val, Test-CBL, Test-GBL} with {144k, 2.9k, 5.8k, 5.8k} images over 29 classes, where the number of samples for each class ranges from 61k to 0.3k. The ALT protocol has {32k, 1.4k, 2.9k} images for {Train-CBL, Val, Test-GBL}. Since attributes usually co-occur with each other in one object, we cannot construct Many_A, Medium_A, and Few_A subsets the same as ImageNet-GLT, so we directly report the overall performance for MSCOCO-GLT. Note that attribute annotations are only used to construct Test-GBL, and they are not released in the training data. **Evaluation Metrics.** The top-1 accuracy is commonly adopted as the only metric in the conventional LT studies, yet, it cannot reveal the limitation of precision-accuracy trade-off. Therefore, in GLT classification, we report both Accuracy: $\frac{\#CorrectPredictions}{\#AllSamples}$, which is equal to Top-1 Recall in the class-wise balanced test sets [73], and Precision: $\frac{1}{\#class} \cdot \sum_{class} \frac{\#CorrectPredictions}{\#SamplesPredictedAsThisClass}$ to better evaluate the effectiveness of algorithms.

5.2 Investigated LT Algorithms

As a general feature learning method to deal with the attribute bias, the proposed IFL can be integrated into most prevalent LT methods. We followed Zhang *et al* [68] to summarize the investigated LT algorithms into three categories: 1) one-/two-stage re-balancing, 2) augmentation, and 3) module improvement.

For re-balancing approaches,

we studied two-stage re-sampling methods **cRT** [19] and **LWS** [19], post-hoc distribution adjustment **Deconfound-TDE** [50] and Logit Adjustment (**Logit-Adj**) [33], multi-branch models with diverse sampling strategies like **BBN** [71], and re-weighting loss functions like Balanced Softmax (**BLSoftmax**) [41] and **LDAM** [6].

For augmentation approaches, we empirically noticed that some common data augmentation methods are more general and effective than other longtailed transfer learning meth-

| Table | 2 . | Evaluation | of | ALT | Protocol | on |
|-------|------------|------------|----|-----|----------|----|
| Image | Ne | t-GLT | | | | |

| _ | Methods | Attri | Attribute-Wise Long Tail (ALT) Protocol | | | | | | |
|-----|------------------------|---------------------|---|--------------------|--------------------|--|--|--|--|
| < | Accuracy Precision > | > Many _A | Many _A Medium _A | | Overall | | | | |
| | Baseline | 56.95 55.83 | 40.11 39.17 | 28.12 28.16 | 41.73 41.74 | | | | |
| | cRT [19] | 57.45 56.28 | 39.72 38.65 | 27.58 27.35 | 41.59 41.43 | | | | |
| | LWS [19] | 56.95 55.85 | 40.11 39.30 | 28.03 27.98 | 41.70 41.71 | | | | |
| | Deconfound-TDE [50] | 57.10 56.58 | 39.80 40.08 | 27.29 27.96 | 41.40 42.36 | | | | |
| č | BLSoftmax [41] | $56.48 \mid 55.56$ | $39.81 \mid 38.96$ | 27.64 27.60 | 41.32 41.37 | | | | |
| la | BBN [71] | 60.90 60.17 | 41.08 40.81 | 27.79 28.26 | 43.26 43.86 | | | | |
| - Ř | LDAM [6] | $59.04 \mid 56.51$ | 40.96 39.21 | 27.96 27.22 | 42.66 41.80 | | | | |
| Å | (ours) Baseline + IFL | $61.38 \mid 60.78$ | $44.79 \mid 44.21$ | $31.49 \mid 31.98$ | $45.89 \mid 46.42$ | | | | |
| | (ours) cRT + IFL | 61.12 60.25 | 44.26 43.65 | 31.02 31.31 | 45.47 45.81 | | | | |
| | (ours) LWS + IFL | 61.19 60.45 | 44.66 44.07 | 31.43 31.91 | 45.76 46.25 | | | | |
| | (ours) BLSoftmax + IFI | 60.19 59.46 | $43.54 \mid 43.14$ | 30.85 31.46 | $44.86 \mid 45.43$ | | | | |
| nt | Mixup [66] | 58.71 58.04 | 40.09 38.99 | 27.52 27.54 | 42.11 42.42 | | | | |
| B | RandAug [10] | 62.35 61.25 | 45.04 44.27 | 31.47 31.26 | 46.29 46.32 | | | | |
| ßn | (ours) Mixup + IFL | 65.90 65.88 | 49.43 49.43 | 35.40 35.89 | 50.24 51.04 | | | | |
| × | (ours) RandAug + IFL | $67.39 \mid 66.81$ | $51.55 \mid 51.28$ | $37.47 \mid 37.97$ | $52.14 \mid 52.74$ | | | | |
| ole | TADE [67] | 62.63 61.91 | 45.84 45.21 | 32.82 32.82 | 47.10 47.32 | | | | |
| Ē | RIDE [56] | 63.48 61.42 | 45.62 44.16 | 32.59 32.26 | 47.24 46.67 | | | | |
| ISE | (ours) TADE + IFL | 63.50 62.67 | 48.03 47.32 | 34.69 34.52 | 48.74 48.78 | | | | |
| Ē | (ours) RIDE + IFL | $67.54 \mid 67.13$ | 51.92 51.72 | $37.84 \mid 38.46$ | $52.44 \mid 53.17$ | | | | |

ods [30], so we adopted **Mixup** [66] and Random Augmentation (**RandAug**) [10] in our experiments.

For module improvement, we followed the recent trend of ensemble learning [56] like **RIDE** [56] and **TADE** [67], which are proved to be state-of-the-art models in LT classification that are capable of improving both head and tail categories at the same time. Both of them used the trident version of ResNext-50 [62] developed by RIDE [56], which is denoted as RIDE-50 in Table 4.

Implementation details of the proposed IFL and the baseline methods are given in Appendix.

5.3 Comparisons with LT Line-up

We evaluate CLT and GLT protocols for ImageNet-GLT in Table 1. The corresponding ALT protocol is reported in Table 2. All three protocols for MSCOCO-GLT are shown in Table 3.

CLT Protocol (conventional long-tailed classification): although the proposed IFL is mainly designed to tackle the long-tailed intra-class attributes, it

also boost all prevalent LT methods in the conventional CLT protocol. It's because IFL also prevents tail images from being mis-classified as head classes by eliminating the spurious correlation made by attributes.

GLT Protocol: not surpris-

ingly, we observe a significant performance decline for all methods from CLT protocol to GLT, as tackling the additional attribute bias is much more challenging than class bias. Meanwhile, IFL can still successfully improve various baselines. The decline on the proposed GLT reveals that general long tail is indeed not just the pure class-wise imbalance.

ALT Protocol: as we expected, the majority of previous LT algorithms using rebalancing strategies failed to improve the robustness against the attribute-wise bias in this protocol. Therefore, their improvements on GLT protocol

| Table | 3. | Evaluati | on o | n M | SCOCO-GLI | : |
|---------|------|-------------|---------|--------|-----------|---|
| overall | perf | formances a | are rep | oorted | l | |

| | P P P | | | | | | | | | | |
|-----|--|---------|-------|-------|-------|---------|-------|--|--|--|--|
| | Protocols | CL | Л | G | LT | ALT | | | | | |
| < | ${\bf Accuracy} ~ ~ {\bf Precision} >$ | Overall | | Ove | erall | Overall | | | | | |
| | Baseline | 72.34 | 76.61 | 63.79 | 70.52 | 50.17 | 50.94 | | | | |
| | cRT [19] | 73.64 | 75.84 | 64.69 | 68.33 | 49.97 | 50.37 | | | | |
| | LWS [19] | 72.60 | 75.66 | 63.60 | 68.81 | 50.14 | 50.61 | | | | |
| | Deconfound-TDE [50] | 73.79 | 74.90 | 66.07 | 68.20 | 50.76 | 51.68 | | | | |
| e | BLSoftmax [41] | 72.64 | 75.25 | 64.07 | 68.59 | 49.72 | 50.65 | | | | |
| an | Logit-Adj [33] | 75.50 | 76.88 | 66.17 | 68.35 | 50.17 | 50.94 | | | | |
| Jac | BBN [71] | 73.69 | 77.35 | 64.48 | 70.20 | 51.83 | 51.77 | | | | |
| | LDAM [6] | 75.57 | 77.70 | 67.26 | 70.70 | 55.52 | 56.21 | | | | |
| æ | (ours) Baseline + IFL | 74.31 | 78.90 | 65.31 | 72.24 | 52.86 | 53.49 | | | | |
| | (ours) cRT + IFL | 76.21 | 79.11 | 66.90 | 71.34 | 52.07 | 52.85 | | | | |
| | (ours) LWS + IFL | 75.98 | 79.18 | 66.55 | 71.49 | 52.07 | 52.90 | | | | |
| | (ours) BLSoftmax + IFL | 73.72 | 77.08 | 64.76 | 70.00 | 52.97 | 53.52 | | | | |
| | (ours) Logit-Adj + IFL | 77.16 | 79.09 | 67.53 | 70.18 | 52.86 | 53.49 | | | | |
| nt | Mixup [66] | 74.22 | 78.61 | 64.45 | 71.13 | 48.90 | 49.53 | | | | |
| me | RandAug [10] | 76.81 | 79.88 | 67.71 | 72.73 | 53.69 | 54.71 | | | | |
| ßn | (ours) Mixup + IFL | 77.55 | 81.78 | 68.83 | 74.84 | 53.79 | 54.60 | | | | |
| A | (ours) RandAug + IFL | 77.71 | 81.10 | 68.16 | 73.97 | 56.62 | 57.12 | | | | |
| ole | TADE [67] | 76.22 | 78.84 | 66.98 | 71.22 | 54.93 | 55.48 | | | | |
| -fu | RIDE [56] | 78.29 | 80.33 | 68.59 | 72.20 | 58.90 | 59.43 | | | | |
| ıse | (ours) TADE + IFL | 76.53 | 79.15 | 67.38 | 72.42 | 56.76 | 57.43 | | | | |
| Ē | (ours) $RIDE + IFL$ | 78.86 | 80.70 | 69.09 | 72.57 | 58.93 | 59.84 | | | | |
| - | | | | | | | | | | | |

only came from the class-wise invariance. However, we noticed that the augmentation and ensemble approaches can improve all three protocols, making them good baselines for GLT as well. The main reason is that both augmentation and model ensemble also aim to improve the representation learning.

5.4 Ablation Studies and Further Analyses

We further conducted a group of ablation studies on ImageNet-GLT in Table 4 to address some common concerns. We also provide some further analyses to shed lights on the proposed GLT challenge.

Q1: how about the baselines from the attribute-wise side? A1: we also reported some popular methods like Focal loss [28] for hard-example mining and LFF (Learning-from-failure) [35] for domain generalization as the baselines solving the attribute bias. However, they didn't perform well using their default settings, proving the difficulty of the proposed GLT in real-world datasets.

Q2: does the improvement come from the center loss? **A2:** the original ERM version of center loss [59] can be considered as a special case of the proposed IFL with only one single environment. According to Table 4, using the vanilla center loss, *i.e.*, #Env=1 with IFL, as the additional constraint actually hurt all three protocols. It's because the center in a biased environment is also biased, *e.g.*, the center of "banana" may possess 90% of "yellow" attribute, which only makes the model more relying on spurious correlations.

Q3: how many environments are required? A3: although it's crucial to have more than 1 environment, IRM [3] asserts that two environments are enough to capture the invariance. We also found that additional environments only brought marginal improvement.

Q4: why not directly apply IRM loss? A4: theoretically, the original IRM loss [3] and the proposed IFL are supposed to have similar results, as they both embody the same spirit of learning invariance across envi-

Table 4. Ablation Studies on ImageNet-GLT, where overall results are reported; BLS, Focal, and IFF are balanced softmax loss [41], focal loss [28], and learning from failure [35], respectively

| Ablation Settings | | | | | | Evaluation Protocols | | | | | | |
|-------------------|-----------------------|---|----------|------------|---------|----------------------|---------|-------|---------|-------|--|--|
| #Env | #Env Loss IFL Augment | | Backbone | CLT P | rotocol | GLT F | rotocol | ALT P | rotocol | | | |
| 1 | CE | - | - | ResNext-50 | 42.52 | 47.92 | 34.75 | 40.65 | 41.73 | 41.74 | | |
| 1 | Focal | - | - | ResNext-50 | 39.93 | 46.99 | 32.52 | 39.12 | 39.58 | 39.85 | | |
| 1 | LFF | - | - | ResNext-50 | 41.07 | 45.79 | 33.84 | 38.46 | 40.14 | 40.58 | | |
| 1 | CE | ~ | - | ResNext-50 | 39.74 | 47.06 | 32.82 | 40.86 | 39.99 | 41.38 | | |
| 2 | IRM | - | - | ResNext-50 | 43.70 | 48.06 | 36.03 | 40.61 | 44.47 | 44.60 | | |
| 2 | CE | ~ | - | ResNext-50 | 45.97 | 52.06 | 37.96 | 44.47 | 45.89 | 46.42 | | |
| 3 | CE | ~ | - | ResNext-50 | 46.06 | 52.81 | 38.32 | 45.55 | 45.95 | 46.43 | | |
| 2 | BLS | ~ | - | ResNext-50 | 48.34 | 50.39 | 40.08 | 43.48 | 44.86 | 45.43 | | |
| 2 | CE | ~ | Mixup | ResNext-50 | 51.43 | 57.44 | 43.00 | 49.25 | 50.24 | 51.04 | | |
| 2 | CE | ~ | RandAug | ResNext-50 | 53.40 | 58.11 | 44.90 | 50.47 | 52.14 | 52.74 | | |
| 1 | CE | - | - | RIDE-50 | 46.14 | 52.98 | 38.25 | 45.80 | 46.32 | 46.56 | | |
| 2 | CE | ~ | - | RIDE-50 | 49.20 | 54.64 | 41.35 | 47.67 | 48.62 | 48.62 | | |
| 2 | TADE | ~ | - | RIDE-50 | 51.78 | 52.41 | 43.47 | 45.17 | 48.74 | 48.78 | | |
| 2 | LDAM | ~ | - | RIDE-50 | 53.93 | 54.76 | 45.64 | 47.14 | 52.44 | 53.17 | | |
| 2 | LDAM | ~ | Mixup | RIDE-50 | 56.48 | 57.67 | 47.54 | 49.86 | 53.25 | 54.27 | | |
| 2 | LDAM | ~ | RandAug | RIDE-50 | 58.70 | 59.61 | 49.80 | 51.62 | 55.65 | 55.81 | | |

ronments. However, we empirically noticed that the original IRM loss has the convergence issue in real-world dataset. 2 out of 5 random seeds result NaN loss at some point during training.

Q5: why do all models perform worse under ALT than GLT on MSCOCO-GLT? A5: note that the GLT protocol is always harder than ALT or CLT. This weird phenomenon in MSCOCO-GLT is caused by the severe class-wise imbalance in MSCOCO-Attributes [38]: a single class "person" possess over 40% of the training set, so Train-CBL has much less training samples than Train-GLT. It further proves the importance of long-tailed classification in real-world applications, as *large long-tailed datasets are better than small balanced counterparts*.

Q6: what is the precision-accuracy trade-off problem [73]? A6: as shown in Fig. 5 (a-b), class-wise re-balancing method like cRT, LWS, and Logit-Adj [19,33] are playing with the precision-accuracy trade-off under GLT protocol and barely work in ALT protocol. It's because the attribute bias hurts both accuracy and precision by forming spurious correlations between classes. Tackling it should improve both of metrics at the same time.

Q7: is there any other method that can be served as GLT baselines as well? A7: we also found that the recent trend of improving both head and tail categories [56,67,73], though lack a formal definition in their approaches, are essentially trying to solve the GLT challenge. Benefit from the feature learning, these ensemble learning and data augmentation approaches can also serve as good baselines for the proposed GLT as well. Meanwhile, the proposed IFL is orthogonal to them and can further boost their performances under all three protocols.

Q8: why the proposed GLT is the "generalized" version of the LT. A8: it's because GLT methods would naturally solve the conventional LT, but not vise versa. As shown in Fig. 5 (c-d), GLT baselines like augmentation, ensemble, and the proposed IFL automatically solve the class-wise LT. Yet, the majority of



Fig. 5. (a-b) The trending of precision and accuracy after applying the IFL; (c-d) GLT baselines will automatically improve class-wise LT, while conventional LT re-balancing algorithms won't improve the attribute-wise imbalance in GLT

class-wise re-balancing LT algorithms cannot tackle the attribute-wise imbalance of GLT very well.

6 Conclusion

We introduced a novel Generalized Long-Tailed (GLT) classification problem, completing the previous class-wise long-tailed classification by incorporating the attribute-wise imbalance nature in real-world dataset at scale, which deeply explains the cause of long-tailed performance within each class and the existence of spurious correlations during classification. Experiments on the proposed two benchmarks, ImageNet-GLT and MSCOCO-GLT, demonstrate the limitations of the previous LT algorithms using class-wise adjustment, and the importance of representation learning in GLT. To this end, we present invariant feature learning (IFL) as the first strong baseline for GLT. IFL adopts a metric loss to encourage the learning of invariant features across environments with diverse attribute distributions. As an improved feature backbone, IFL is orthogonal to most of the previous LT algorithms. After plugging IFL into the conventional LT line-up: one-/two-stage re-balance, augmentation, and ensemble, IFL boosts their performances under all protocols in the proposed GLT benchmarks.

Limitations: Due to the limited space, we didn't fully explore all kinds of attribute biases in this paper, *e.g*, benchmarks like BREEDS [43], MetaShift [27] or FMoW-WILDS [21] all have their unique attribute types. In the future, we are going to extend GLT to more diverse tasks and attribute settings.

Acknowledgements: This project is partially supported by Alibaba-NTU Singapore Joint Research Institute (JRI), and AI Singapore (AISG) Research Programme. We also feel grateful to the computational resources provided by Damo Academy.

References

- Agarwal, V., Shetty, R., Fritz, M.: Towards causal vqa: Revealing and reducing spurious correlations by invariant and covariant semantic editing. In: CVPR (2020)
 5
- Arjovsky, M.: Out of distribution generalization in machine learning. Ph.D. thesis, New York University (2020) 4
- Arjovsky, M., Bottou, L., Gulrajani, I., Lopez-Paz, D.: Invariant risk minimization. arXiv preprint arXiv:1907.02893 (2019) 4, 8, 9, 13
- Besserve, M., Mehrjou, A., Sun, R., Schölkopf, B.: Counterfactuals uncover the modular structure of deep generative models. ICLR (2020) 5
- Cai, J., Wang, Y., Hwang, J.N.: Ace: Ally complementary experts for solving longtailed recognition in one-shot. In: ICCV (2021) 4
- Cao, K., Wei, C., Gaidon, A., Arechiga, N., Ma, T.: Learning imbalanced datasets with label-distribution-aware margin loss. NeurIPS (2019) 2, 3, 4, 10, 11, 12
- Carlini, N., Athalye, A., Papernot, N., Brendel, W., Rauber, J., Tsipras, D., Goodfellow, I., Madry, A., Kurakin, A.: On evaluating adversarial robustness. arXiv preprint arXiv:1902.06705 (2019) 4
- Chakraborty, A., Alam, M., Dey, V., Chattopadhyay, A., Mukhopadhyay, D.: Adversarial attacks and defences: A survey. arXiv preprint arXiv:1810.00069 (2018)
 4
- Creager, E., Jacobsen, J.H., Zemel, R.: Environment inference for invariant learning. In: ICML (2021) 5
- Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: Randaugment: Practical automated data augmentation with a reduced search space. In: CVPR Workshops (2020) 4, 10, 11, 12
- Cui, Y., Jia, M., Lin, T.Y., Song, Y., Belongie, S.: Class-balanced loss based on effective number of samples. In: CVPR (2019) 4
- 12. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: ICCV (2017) 6
- 13. He, Y.Y., Wu, J., Wei, X.S.: Distilling virtual examples for long-tailed recognition. ICCV (2021) 4
- 14. Hinton, G., Roweis, S.T.: Stochastic neighbor embedding. In: NeurIPS (2002) 3
- Hong, Y., Han, S., Choi, K., Seo, S., Kim, B., Chang, B.: Disentangling label distribution for long-tailed visual recognition. In: CVPR (2021) 4
- 16. Hu, X., Jiang, Y., Tang, K., Chen, J., Miao, C., Zhang, H.: Learning to segment the tail. In: CVPR (2020) 4
- Idrissi, B.Y., Arjovsky, M., Pezeshki, M., Lopez-Paz, D.: Simple data balancing achieves competitive worst-group-accuracy. In: Conference on Causal Learning and Reasoning (2022) 5
- Jamal, M.A., Brown, M., Yang, M.H., Wang, L., Gong, B.: Rethinking classbalanced methods for long-tailed visual recognition from a domain adaptation perspective. In: CVPR (2020) 4
- Kang, B., Xie, S., Rohrbach, M., Yan, Z., Gordo, A., Feng, J., Kalantidis, Y.: Decoupling representation and classifier for long-tailed recognition. In: ICLR (2020) 2, 3, 4, 6, 8, 10, 11, 12, 13
- Kim, J., Jeong, J., Shin, J.: M2m: Imbalanced classification via major-to-minor translation. In: CVPR (2020) 4
- Koh, P.W., Sagawa, S., Marklund, H., Xie, S.M., Zhang, M., Balsubramani, A., Hu, W., Yasunaga, M., Phillips, R.L., Gao, I., et al.: Wilds: A benchmark of in-the-wild distribution shifts. In: ICML (2021) 5, 14

- 16 K. Tang et al.
- Krueger, D., Caballero, E., Jacobsen, J.H., Zhang, A., Binas, J., Zhang, D., Le Priol, R., Courville, A.: Out-of-distribution generalization via risk extrapolation (rex). In: ICML (2021) 4
- 23. Li, D., Yang, Y., Song, Y.Z., Hospedales, T.M.: Deeper, broader and artier domain generalization. In: ICCV (2017) 4
- 24. Li, T., Wang, L., Wu, G.: Self supervision to distillation for long-tailed visual recognition. In: ICCV (2021) 4
- Li, Y., Wang, T., Kang, B., Tang, S., Wang, C., Li, J., Feng, J.: Overcoming classifier imbalance for long-tail object detection with balanced group softmax. In: CVPR (2020) 4
- 26. Li, Z., Xu, C.: Discover the unknown biased attribute of an image classifier. arXiv preprint arXiv:2104.14556 (2021) 3
- 27. Liang, W., Zou, J.: Metashift: A dataset of datasets for evaluating contextual distribution shifts and training conflicts. ICLR (2022) 5, 14
- Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: ICCV. pp. 2980–2988 (2017) 5, 12, 13
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV. pp. 740–755. Springer (2014) 10
- Liu, J., Sun, Y., Han, C., Dou, Z., Li, W.: Deep representation learning on longtailed data: A learnable embedding augmentation perspective. In: CVPR (2020) 4, 11
- Liu, Z., Miao, Z., Zhan, X., Wang, J., Gong, B., Yu, S.X.: Large-scale long-tailed recognition in an open world. In: CVPR (2019) 3, 4, 6, 10
- Locatello, F., Bauer, S., Lucic, M., Raetsch, G., Gelly, S., Schölkopf, B., Bachem, O.: Challenging common assumptions in the unsupervised learning of disentangled representations. In: ICML. PMLR (2019) 5
- 33. Menon, A.K., Jayasumana, S., Rawat, A.S., Jain, H., Veit, A., Kumar, S.: Long-tail learning via logit adjustment. ICLR (2020) 2, 3, 4, 5, 6, 8, 10, 11, 12, 13
- Mirza, M., Osindero, S.: Conditional generative adversarial nets. arXiv preprint arXiv:1411.1784 (2014) 5
- Nam, J., Cha, H., Ahn, S., Lee, J., Shin, J.: Learning from failure: De-biasing classifier from biased classifier. NeurIPS (2020) 5, 12, 13
- 36. News, B.: Facebook apology as ai labels black men 'primates' (2021), https:// www.bbc.com/news/technology-58462511 2
- 37. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (eds.) Advances in Neural Information Processing Systems 32, pp. 8024-8035. Curran Associates, Inc. (2019), http://papers.neurips.cc/paper/ 9015-pytorch-an-imperative-style-high-performance-deep-learning-library. pdf 7, 10
- Patterson, G., Hays, J.: Coco attributes: Attributes for people, animals, and objects. In: ECCV. Springer (2016) 2, 7, 10, 13
- 39. Powers, D.M.: Applications and explanations of zipf's law. In: New methods in language processing and computational natural language learning (1998) 1
- Reed, W.J.: The pareto, zipf and other power laws. Economics letters 74(1), 15–19 (2001) 1, 9

- 41. Ren, J., Yu, C., Sheng, S., Ma, X., Zhao, H., Yi, S., Li, H.: Balanced meta-softmax for long-tailed visual recognition. NeurIPS (2020) 2, 3, 4, 5, 6, 8, 10, 11, 12, 13
- 42. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. International journal of computer vision (2015) 7, 10
- Santurkar, S., Tsipras, D., Madry, A.: Breeds: Benchmarks for subpopulation shift. ICLR (2021) 5, 14
- Sohoni, N., Dunnmon, J., Angus, G., Gu, A., Ré, C.: No subclass left behind: Fine-grained robustness in coarse-grained classification problems. NeurIPS (2020) 5
- 45. Srivastava, M., Hashimoto, T., Liang, P.: Robustness to spurious correlations via human annotations. In: ICML (2020) 5
- Stone, J.V.: Bayes' rule: a tutorial introduction to bayesian analysis. Sebtel Press (2013) 5
- 47. Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P.H., Hospedales, T.M.: Learning to compare: Relation network for few-shot learning. In: CVPR (2018) 4
- 48. Tan, J., Lu, X., Zhang, G., Yin, C., Li, Q.: Equalization loss v2: A new gradient balance approach for long-tailed object detection. In: CVPR (2021) 4
- Tan, J., Wang, C., Li, B., Li, Q., Ouyang, W., Yin, C., Yan, J.: Equalization loss for long-tailed object recognition. In: CVPR (2020) 3, 4
- 50. Tang, K., Huang, J., Zhang, H.: Long-tailed classification by keeping the good and removing the bad momentum causal effect. NeurIPS (2020) 4, 8, 10, 11, 12
- Van Horn, G., Mac Aodha, O., Song, Y., Cui, Y., Sun, C., Shepard, A., Adam, H., Perona, P., Belongie, S.: The inaturalist species classification and detection dataset. In: CVPR (2018) 6
- Wang, M., Deng, W.: Deep visual domain adaptation: A survey. Neurocomputing (2018) 4
- Wang, T., Yue, Z., Huang, J., Sun, Q., Zhang, H.: Self-supervised learning disentangled group representation as feature. NeurIPS (2021) 5
- Wang, T., Li, Y., Kang, B., Li, J., Liew, J., Tang, S., Hoi, S., Feng, J.: The devil is in classification: A simple framework for long-tail instance segmentation. In: ECCV (2020) 4
- 55. Wang, W., Zheng, V.W., Yu, H., Miao, C.: A survey of zero-shot learning: Settings, methods, and applications. TIST (2019) 4
- Wang, X., Lian, L., Miao, Z., Liu, Z., Yu, S.X.: Long-tailed recognition by routing diverse distribution-aware experts. ICLR (2020) 4, 10, 11, 12, 13
- 57. Wang, Y., Yao, Q.: Few-shot learning: A survey. arxiv (2019) 4
- Wang, Y.X., Ramanan, D., Hebert, M.: Learning to model the tail. In: NeurIPS (2017) 4
- Wen, Y., Zhang, K., Li, Z., Qiao, Y.: A discriminative feature learning approach for deep face recognition. In: ECCV (2016) 3, 8, 9, 12
- Wilson, G., Cook, D.J.: A survey of unsupervised deep domain adaptation. ACM Transactions on Intelligent Systems and Technology (TIST) (2020) 4
- Xiang, L., Ding, G., Han, J.: Learning from multiple experts: Self-paced knowledge distillation for long-tailed classification. In: ECCV (2020) 4
- 62. Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: CVPR (2017) 11
- Yin, X., Yu, X., Sohn, K., Liu, X., Chandraker, M.: Feature transfer learning for face recognition with under-represented data. In: CVPR (2019) 4
- 64. You, Q., Jin, H., Wang, Z., Fang, C., Luo, J.: Image captioning with semantic attention. In: CVPR (2016) 6

- 18 K. Tang et al.
- 65. Yue, Z., Sun, Q., Hua, X.S., Zhang, H.: Transporting causal mechanisms for unsupervised domain adaptation. In: ICCV (2021) 4
- Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. In: ICLR (2018) 4, 10, 11, 12
- Zhang, Y., Hooi, B., Hong, L., Feng, J.: Test-agnostic long-tailed recognition by test-time aggregating diverse experts with self-supervision. ICCV (2021) 4, 10, 11, 12, 13
- Zhang, Y., Kang, B., Hooi, B., Yan, S., Feng, J.: Deep long-tailed learning: A survey. arXiv preprint arXiv:2110.04596 (2021) 1, 4, 6, 11
- Zhao, B., Yu, S., Ma, W., Yu, M., Mei, S., Wang, A., He, J., Yuille, A., Kortylewski, A.: Robin: A benchmark for robustness to individual nuisances in real-world outof-distribution shifts. ECCV (2022) 4
- 70. Zhao, H., Des Combes, R.T., Zhang, K., Gordon, G.: On learning invariant representations for domain adaptation. In: ICML (2019) 4
- Zhou, B., Cui, Q., Wei, X.S., Chen, Z.M.: Bbn: Bilateral-branch network with cumulative learning for long-tailed visual recognition. In: CVPR (2020) 2, 3, 4, 6, 10, 11, 12
- 72. Zhu, B., Niu, Y., Hua, X.S., Zhang, H.: Cross-domain empirical risk minimization for unbiased long-tailed classification. In: AAAI (2022) 3, 4
- Zhu, B., Niu, Y., Hua, X.S., Zhang, H.: Cross-domain empirical risk minimization for unbiased long-tailed classification. AAAI (2022) 3, 4, 11, 13
- 74. Zou, Y., Yu, Z., Kumar, B., Wang, J.: Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In: ECCV (2018) 4