

Tackling Long-Tailed Category Distribution Under Domain Shifts

Xiao Gu¹, Yao Guo², Zeju Li¹, Jianing Qiu¹, Qi Dou³, Yuxuan Liu²,
Benny Lo^{1*}, and Guang-Zhong Yang^{2*}

¹ Imperial College London

{xiao.gu17, zeju.li18, jianing.qiu17, benny.lo}@imperial.ac.uk

² Shanghai Jiao Tong University

{yao.guo, 200009051yx, gzyang}@sjtu.edu.cn

³ The Chinese University of Hong Kong

qidou@cuhk.edu.hk

Abstract. Machine learning models fail to perform well on real-world applications when 1) the category distribution $P(Y)$ of the training dataset suffers from long-tailed distribution and 2) the test data is drawn from different conditional distributions $P(X|Y)$. Existing approaches cannot handle the scenario where both issues exist, which however is common for real-world applications. In this study, we took a step forward and looked into the problem of long-tailed classification under domain shifts. We designed three novel core functional blocks including Distribution Calibrated Classification Loss, Visual-Semantic Mapping and Semantic-Similarity Guided Augmentation. Furthermore, we adopted a meta-learning framework which integrates these three blocks to improve domain generalization on unseen target domains. Two new datasets were proposed for this problem, named AWA2-LTS and ImageNet-LTS. We evaluated our method on the two datasets and extensive experimental results demonstrate that our proposed method can achieve superior performance over state-of-the-art long-tailed/domain generalization approaches and the combinations. Source codes and datasets can be found at our project page <https://xiaogu.site/LTDS>.

Keywords: Long Tail, Domain Generalization, Cross-Modal Representation Learning, Meta Learning

1 Introduction

Deep learning has made unprecedented achievements on various applications ranging from self-driving [2], service robots [8], to health and wellbeing [26]. The model would perform well with the assumption that training and testing data are independent identically distributed (*i.i.d.*); however it seldomly holds for real-world applications. The violation of *i.i.d.* assumption could hinder the performance of deep learning models upon practical deployment. Without loss of generality, we denote the data and label as X and Y , the joint distribution as $P(X, Y)$, the latter of which can be formulated by $P(Y)P(X|Y)$. We argue

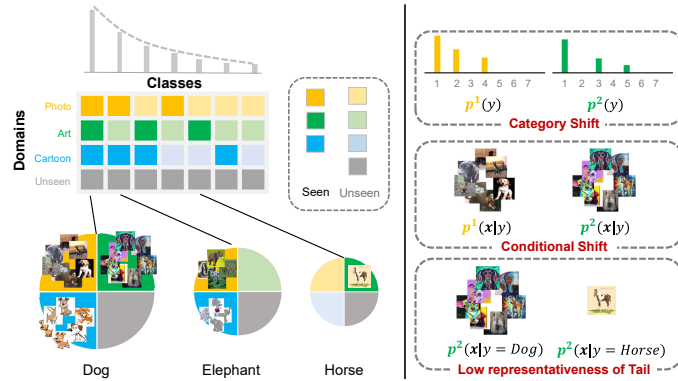


Fig. 1. Visual illustrations of the issues complicated with long-tailed category distribution and conditional distribution shifts across domains. The overall dataset¹ is long-tailed distributed over classes. Meanwhile, with data collected from multiple domains (e.g., different styles), head classes (e.g., *Dog*) can be observed in most domains, whereas tail classes (e.g., *Horse*) only contain a few samples in certain specific domains.

that the reason why current models fail to generalize well in real-world applications is rooted in both the categorical distribution $P(Y)$ and class conditional distribution $P(X|Y)$.

On one hand, real-world data exhibits long-tailed distribution over categories, with only a few classes (*head*) accounting for the major proportions, whilst many more classes (*tail*) presenting extremely limited samples [43]. For instance, in action recognition, the case of “*open door*” is common in daily activities, whereas some actions such as “*repair door*” occur much less frequently. This leads to a long-tailed label distribution of $P(Y)$, where conventional training strategies that apply common classification losses (mostly cross-entropy) on instance-level sampled batches would fail. In this case, the trained model would gain high performance on the head classes but behave poorly on tail classes, failing to achieve consistently good performance across all categories. On the other hand, the conditional distribution $P(X|Y)$ is also prone to changes in the real world [35]. Different styles of image recognition data, camera viewpoints of action recognition data, acquisition protocols of medical images, etc., would alter the distribution of $P(X|Y)$, leading to diversified distributions, a.k.a. domain shifts.

In this regard, long-tailed categorical distribution (**LT**) [43] and domain shifts (**DS**) [35] have been two major issues concerned with real-world datasets. Although increasing research efforts have been made, these two issues are so far tackled individually, with their complex co-existence situation not being considered yet. Existing solutions cannot deal with the entanglement of **LT** and **DS**, since a balanced distribution $P(Y)$ or identical $P(X|Y)$ are their prior assump-

¹Images are adopted from PACS [15]. Its distribution originally is not long-tailed. Here just for intuitive explanation of our focused problem.

tions of those **DS** [6,35] and **LT** [27,13,17] solutions, respectively. As we know, in real-world scenarios, these two issues often come together. Take medical image data as an example, the conditional distribution varies across different hospitals, and, there are a large population of patients with common diseases whilst some patients with rare diseases. In addition, the low prevalence of those rare diseases may lead to the inclusion of corresponding patients only by certain hospitals. This also similarly applies to many other applications [4] where the head classes are common in most domains, whilst tail classes only appear in certain domains due to the low-frequency. Such combination of **LT** and **DS** leads to a more challenging, yet more practical scenario where $P(Y)$ of each individual domain is not only imbalanced, but also partial. Ideally a reasonable model should be robust across classes and generalize across domains, simultaneously.

We argue that there are three main challenges posed by the problem of **LT-DS** (cf. Fig. 1). **1)** Because of the existence of multiple domains, the categorical distributions $P(Y)$ are different across domains. Given the relatively low frequency of non-head classes, their corresponding samples may be collected only in certain domains. As a result, the spurious correlation between non-head classes and domain-specific characteristics might be learned as biased shortcuts. **2)** The conditional distributions $P(X|Y)$ are varied significantly across domains. It is expected that the model can handle such shifts, with domains aligned and unbiased representations learned. **3)** It is hard to explicitly model the distribution of tail classes $P(X|Y = \textit{tail})$, as only a paucity of domain-specific samples exist. This poses challenges to avoiding overfitting on the tails. Hence, research is desperately needed to solve the co-occurrence of these issues in **LT-DS**.

In this work, we propose an effective solution to tackle all of the aforementioned challenges. First, a novel domain-specific distribution calibrated loss is introduced to address the infinite imbalance ratio of each domain. Subsequently, we leverage class distributional embeddings as unbiased semantic features, to align the derived visual representations to unbiased semantic space via the alignment between domain-specific visual prototypes and semantic embeddings. Furthermore, we propose a semantic-similarity guided module by leveraging the knowledge learned from head classes, for implicit augmentation of tail classes. In addition, to ensure the model is capable of handling out-of-distribution data in unseen domains, a meta-learning framework integrating the above three core modules is proposed to boost the generalization capability. To evaluate the effectiveness of the proposed method, we developed two datasets with **LT-DS** problems, namely **AWA2-LTS** and **ImageNet-LTS**, and conducted extensive comparison experiments on both datasets. Results demonstrate that our proposed method exceeds state-of-the-art **LT** or **DS** methods by a large margin.

2 Related Works

Long-Tailed Category Distribution. To learn from class imbalanced training data, one line of existing works aims to manipulate class-wise contributions by resampling [1,9], reweighting [11,38,19], logits adjustment [27,10], and two-

stage training [14,32]. Another emerging line has made attempts at ensemble learning under long-tailed settings, such as contrastive learning [36], knowledge distillation [12], variance-bias calibration [37]. In particular, RIDE [37] indicates that the predictions of head classes would be of larger intra-class variances, whereas tail classes would exhibit larger biases. This becomes more serious in the **LT-DS** scenario, since the intra-class variances are related to not only semantics but also domains shifts; whereas the representations of non-head classes may easily be biased by domain-specific characteristics. Unfortunately, most of the existing methods do not take into account conditional-distribution-shift introduced biases, instead assuming identical in their work. Similar issues exist in recent meta-learning based approaches. To ensure good performance across all classes, recent works investigated the category shift between long-tailed and balanced distributions, and introduced meta-learning strategies to optimize parameters on a held-out balanced meta-test subset [13,17]. This is not applicable for **LT-DS**, since it is impossible to sample a held-out meta-test proportion with balanced category distribution without conditional distribution shifts.

Model Generalization at Domain Shifts. Domain generalization (DG) aims to develop computational models that are capable of handling data from unseen domains. Existing domain generalization solutions are varied, including aligning intra-class representations across domains [18], factoring out domain-specific information [34,7], simulating domain gaps via sophisticated training strategies [6,16], or performing data augmentation [22,44]. The shortcomings of most solutions become apparent when faced with **LT-DS**, as **LT-DS** poses imbalanced distribution over a large number of classes. For the methods benefiting from explicit categorical distribution alignment [18], it is computationally prohibitive to design class-specific aligning models, and impossible to align domain-specific tail classes. Furthermore, the large class number makes sampling classes from multiple domains intractable, thus being difficult to cover relatively large portions of the label set in a mini batch [6]. Even worse, for those tail classes, since they are only available in certain domains due to the low-frequency, some shortcuts of the classifier may be learned due to the spurious correlation between the domain-specific information and the occurrence of associated tail classes.

On the other hand, most current domain generalization approaches assume similar categorical distribution across domains, yet this can hardly hold true in the real world [20,29,42]. A similar issue has been raised in [29] referred to as open domain generalization, where the distribution and label sets of each source domain and target domain can be different. Shu *et al.* [29] introduced the domain-specific model in each individual source domain and applied a meta learning strategy to generalize each domain-specific model to other domains, by knowledge distillation from other domains. However, this framework cannot well apply to **LT-DS** settings. To be specific, the knowledge derived from each domain is easy to get biased due to the specific long-tailed and incomplete categorical distribution in each domain. Even worse, there is no guarantee that under such bias, the spurious correlation between non-semantic domain-specific characteristics and domain-specific classes can be avoided.

Cross-Modal Representation Learning. Leveraging information from multiple modalities is popular for related multi-modality applications [30,25] to facilitate effective representation learning of each individual modality. One of the related applications is few/zero-shot learning, where one line of research aims to establish the relationship between semantic space and visual space [41]. Maniayar *et al.* [23] leveraged the semantic space to enable zero-shot domain generalization, and Mancini *et al.* [22] used the semantic embeddings as the classifier for zero-shot learning in unseen domains. These inspired our work; however as a different task setting, our goal is to derive unbiased predictor under long-tailed settings such that the missing classes in seen domains can be recognized as well.

Recently, Samuel *et al.* [28] leveraged class descriptors to facilitate long-tailed classification. It developed a dual network to derive both visual features and semantic features from the input image, and then fused these two together to boost the performance of long-tailed classification. Although it applied semantic embeddings similar to our work, our task aims to address a more challenging problem, where both imbalance and conditional distribution shifts exist.

3 Methodologies

3.1 Problem Setup and Preliminaries




Fig. 2. Illustrations of domain splits, open/non-open classes, and seen/unseen classes.

We denote the input and label spaces as \mathcal{X} and \mathcal{Y} , and the domain space as \mathcal{D} . \mathcal{D} consists of totally K domains $\{D_k\}_{k=1}^K$ and there are totally C categories in the label space. Each sample is denoted as $\{\mathbf{x}_i, y_i, d_i\}$, where i indicates the sample index, \mathbf{x}_i the input sample, y_i the ground truth label, and d_i the domain index; $1 \leq d_i \leq K$. The training domains and testing domains are denoted as \mathcal{D}_{tr} and \mathcal{D}_{te} , respectively, where $\mathcal{D}_{tr} \subset \mathcal{D}$ and $\mathcal{D}_{te} = \mathcal{D}$. The categorical distribution $p^k(y)$ of each training domain k follows a long-tailed distribution, and the low prevalence of tail classes may lead to the failure of collecting training samples from rare classes, i.e., $\mathcal{Y}^k \subset \mathcal{Y}$. We denote the label set of all training data as $\mathcal{Y}^{tr} = \bigcup_{k=1}^{|\mathcal{D}_{tr}|} \mathcal{Y}^k$. To test the overall performance across all classes, test data is sampled under balanced distribution over all classes. Since there might be domain-specific non-head classes in each domain, open classes exist in the testing domains, namely $\mathcal{Y}^{tr} \subset \mathcal{Y}$, $\mathcal{Y}^{te} = \mathcal{Y}$. A visual illustration is presented in Fig. 2.

Our ultimate goal is to build a computational model that is able to recognize all the non-open classes across domains, as well as open classes belonging to $\mathcal{Y} \setminus \mathcal{Y}^{tr}$.

The computational model $g : \mathcal{X} \rightarrow \mathcal{Y}$ maps raw input to the final prediction. Following previous domain generalization works [6], it can be decoupled into a feature extractor f and a head classifier h , where $f : \mathcal{X} \rightarrow \mathcal{Z}$, and $h : \mathcal{Z} \rightarrow \mathbb{R}^C$. The final prediction $\hat{y} = g(\mathbf{x}) = h \circ f(\mathbf{x})$. With the loss function denoted as $\mathcal{L}(h \circ f(\mathbf{x}), y)$, we derive the estimated error ϵ on test data as:

$$\begin{aligned} \epsilon &= \mathbb{E}_{m \sim P_{\mathcal{D}_{te}}} \mathbb{E}_{(\mathbf{x}, y) \sim p^m(\mathbf{x}, y)} \mathcal{L}(h \circ f(\mathbf{x}), y) \\ &= \mathbb{E}_{m \sim P_{\mathcal{D}_{te}}} \mathbb{E}_{(\mathbf{x}, y) \sim p^n(\mathbf{x}, y)} \mathcal{L}(h \circ f(\mathbf{x}), y) \frac{p^m(f(\mathbf{x}), y)}{p^n(f(\mathbf{x}), y)} \\ &= \mathbb{E}_{m \sim P_{\mathcal{D}_{te}}} \mathbb{E}_{(\mathbf{x}, y) \sim p^n(\mathbf{x}, y)} \mathcal{L}(h \circ f(\mathbf{x}), y) \frac{p^m(y) p^m(f(\mathbf{x})|y)}{p^n(y) p^n(f(\mathbf{x})|y)}, \end{aligned} \quad (1)$$

where $P_{\mathcal{D}_{tr}} \& P_{\mathcal{D}_{te}}$ denotes the probability of sampling data from training or testing domains.

To minimize ϵ as in Equation (1), it is of paramount significance to model the term $\frac{p^{te}(y)p^{te}(f(\mathbf{x})|y)}{p^{tr}(y)p^{tr}(f(\mathbf{x})|y)}$ to ensure the robustness under **LT-DS**. However, there exist several issues that are challenging to resolve:

- 1) $p^n(y)$ of each individual training domain is imbalanced. Even through a balanced resampling on seen classes, those unseen classes of each individual domain still leads to “imbalance” with an infinite ratio.
- 2) It comes with challenges in aligning the distribution of $p^n(f(\mathbf{x})|y)$ to an unbiased and semantically meaningful space, since there are some classes unseen in each individual domain, especially for those tail classes.
- 3) The distributions of tail classes $p(f(\mathbf{x})|y)$ are difficult to model compared to head classes, as caused by the limited sample number in certain domains.
- 4) Since we aim to model and align $p^n(f(\mathbf{x})|y)$ rather than $p^n(\mathbf{x}|y)$, it is important to make sure that f is able to handle out-of-distribution data itself, thus enabling extracting domain-invariant and discriminative features by $f(\mathbf{x})$.

It should be noted that Equation (1) gets some inspirations from the recent work [13], while they are conceptually different. To be specific, Jamal *et al.* [13] considers the distribution shifts across long-tailed and balanced distributions; whereas the shifts of $P(X|Y)$ like style changes are not taken into consideration.

In the following, we first go through the core functional blocks to address aforementioned issues, followed by a meta-learning based framework to integrate these functional blocks.

3.2 Core Functional Blocks

Distribution Calibrated Classification Loss-Model $\frac{p^m(y)}{p^n(y)}$. Considering the term $\frac{p^m(y)}{p^n(y)}$, we aim to tackle the imbalance of training data $p^n(y)$ so as to work on balanced distribution $p^m(y)$ with distribution calibrated classification loss. In [27], Ren *et al.* proposed a variant version of softmax function to approximate




Fig. 3. Transformation between **semantic space** \mathbf{s} and **visual feature prototypes** \mathbf{v}^n of each domain D_n . \mathbf{s} denotes semantic features based on word embeddings from class names or sentence embeddings from class descriptors. In each iteration, based on sampled batch from domain D_n , its visual prototype \mathbf{v}^n is updated by exponential moving average (EMA). After transforming it to the semantic space \mathbf{s}^n by e , the missing entries in \mathbf{v}^n are filled by the corresponding semantic embeddings in \mathbf{s} to derive the complete $\hat{\mathbf{s}}^n$ and then converted back to the visual space as $\hat{\mathbf{v}}^n$.

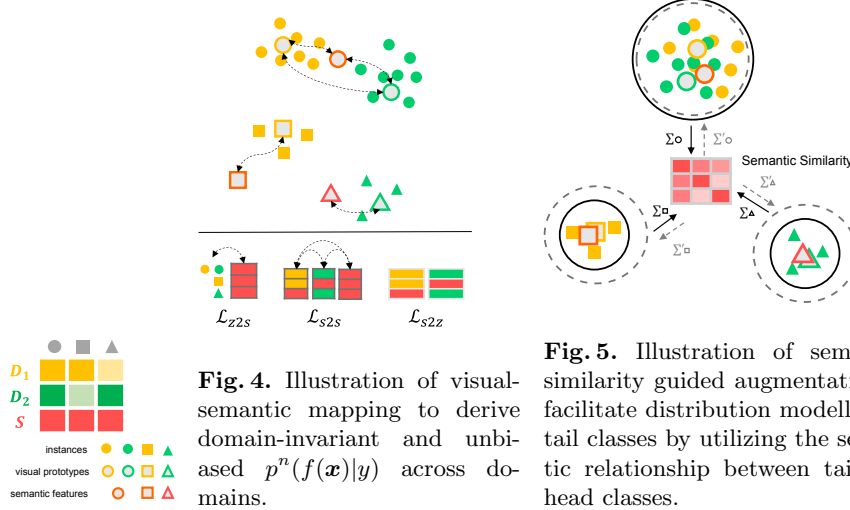


Fig. 4. Illustration of visual-semantic mapping to derive domain-invariant and unbiased $p^n(f(\mathbf{x})|y)$ across domains.

Fig. 5. Illustration of semantic-similarity guided augmentation to facilitate distribution modelling of tail classes by utilizing the semantic relationship between tail and head classes.

the discrepancy of the posterior distributions between training and testing data. Similar ideas have also been introduced in [10]. Based on [27,10], the distribution calibrated classification loss is formulated as:

$$\mathcal{L}_{dc}(\mathbf{x}_i, y_i, d_i; f, h) = -\log \frac{n_{y_i}^{d_i} \exp([h \circ f(\mathbf{x}_i)]_{y_i})}{\sum_{c=1}^C n_c^{d_i} \exp([h \circ f(\mathbf{x}_i)]_c)}, \quad (2)$$

where $n_c^{d_i}$ denotes the sample number of class c in the training domain D_{d_i} . Please see supplementary material for proof.

Visual-Semantic Mapping-Align $p^n(f(\mathbf{x})|y)$. To ensure unbiased and semantically meaningful representations, we leverage the semantic embeddings based on word embeddings from class names or sentence embeddings from class descriptors, inspired by existing zero-shot learning works [23,22]. With totally C classes and feature dim as d_s , the semantic embedding is denoted as $\mathbf{s} \in \mathbb{R}^{C \times d_s}$, with its c element \mathbf{s}_c corresponding to the embedding of class c .

On the other hand, for each domain n with classes C in total and feature dim as d_v , we have a visual prototype $\mathbf{v}^n \in \mathbb{R}^{C \times d_v}$, which is derived in an online manner by exponential moving average (EMA). Some entries of \mathbf{v}^n are probably empty caused by the missing categories in each individual domain n . The index mask of valid entries is denoted as \mathbf{M}^n for convenience.

To achieve the alignment between visual feature prototype $\{\mathbf{v}^n\}$ and \mathbf{s} , we introduce another two functions e and d , where $e: \mathcal{Z} \rightarrow \mathcal{S}$ and $d: \mathcal{S} \rightarrow \mathcal{Z}$. $\{\mathbf{v}^n\}$ is firstly transformed to the semantic space by e as \mathbf{s}^n and subsequently, the missing entries of \mathbf{s}^n are filled by the corresponding semantic features in \mathbf{s} to $\hat{\mathbf{s}}^n$, by $e(\mathbf{v}^n) \cdot \mathbf{M}^n \oplus \mathbf{s} \cdot \tilde{\mathbf{M}}^n$. By the function d , $\hat{\mathbf{s}}^n$ is transformed back to visual space as $\hat{\mathbf{v}}^n$. The data flow is visualized in Fig. 3.

We introduce three typical losses to fulfill the goal of visual-semantic alignment, which are visually illustrated in Fig. 4. First of all, \mathcal{L}_{z2s} of Equation (3) is utilized to align each training sample to its corresponding semantic feature. In Equation (3), we adopt margin contrastive loss on a unit-normalized embedding space, where the margin α is inspired by [33] as to encourage more tight distribution of intra-class embeddings, and τ is the temperature constant.

Furthermore, to avoid shifts across domains, we enforce a cross-prototype contrastive loss as in Equation (4), which aims to decrease intra-class inter-domain discrepancies and enlarge inter-class distances. Additionally, to enforce the manifold constraint of our learned feature representation, we convert $\hat{\mathbf{s}}_v^n$ back to visual space $\hat{\mathbf{v}}^n$ and apply \mathcal{L}_{s2s} as in Equation (5). It applied the classification loss based on the visual classifier h , and also aims to align $\hat{\mathbf{v}}^n$ to semantic embeddings by $\mathcal{L}_{s2s}([e(\hat{\mathbf{v}}^n)], \mathbf{s})$ similar to a cycle loss [45]. To deal with the class imbalance in each batch, \mathcal{L}_{z2s} is further integrated with the last module when calculating the loss.

$$\mathcal{L}_{z2s}(\mathbf{x}_i, y_i, \mathbf{s}; e, f) = -\log \frac{\exp(([\mathbf{e} \circ f(\mathbf{x}_i)]^\top \mathbf{s}_{y_i} - \alpha)/\tau)}{\exp(([\mathbf{e} \circ f(\mathbf{x}_i)]^\top \mathbf{s}_{y_i} - \alpha)/\tau) + \sum_{j \neq y_i} \exp(([\mathbf{e} \circ f(\mathbf{x}_i)]^\top \mathbf{s}_j)/\tau)}, \quad (3)$$

$$\mathcal{L}_{s2s}(\hat{\mathbf{s}}^m, \hat{\mathbf{s}}^n) = \mathbb{E}_c \left[-\log \frac{\exp((\hat{\mathbf{s}}_c^m \hat{\mathbf{s}}_c^n - \alpha)/\tau)}{\exp((\hat{\mathbf{s}}_c^m \hat{\mathbf{s}}_c^n - \alpha)/\tau) + \underbrace{\sum_{j \neq c} \exp(\hat{\mathbf{s}}_c^m \hat{\mathbf{s}}_j^n / \tau) + \sum_{j \neq c} \exp(\hat{\mathbf{s}}_c^m \hat{\mathbf{s}}_j^m / \tau)}_{\text{enlarge inter-class inter-/intra-domain distance}} \right], \quad (4)$$

$$\mathcal{L}_{s2z}(\hat{\mathbf{v}}^n; e, h) = \mathbb{E}_i \left[-\log \frac{\exp([h(\hat{\mathbf{v}}_i^n)]_i)}{\sum_{c=1}^C \exp([h(\hat{\mathbf{v}}_i^n)]_c)} \right] + \mathcal{L}_{s2s}([e(\hat{\mathbf{v}}^n)], \mathbf{s}). \quad (5)$$

Semantic-Similarity Guided Augmentation-Model $p(f(\mathbf{x})|y)$. Another troubling issue lies in the poor diversity of tail classes. In addition to achieving semantically meaningful and unbiased representations, it is also expected that

overfitting on the tail classes can be avoided. It emphasizes the importance of adding to the diversity and richness of tail classes. Therefore, a semantic similarity guided feature augmentation method is proposed as below.

We define the conditional feature distribution (assumed as multi-variate Gaussian distribution, aggregated from all domains) as $p(f(\mathbf{x})|c) \sim \mathcal{N}(\boldsymbol{\mu}_c, \boldsymbol{\Sigma}_c)$. The classifier h is composed of a weight matrix $[w_1, \dots, w_C]^\top$ and biases $[b_1, \dots, b_C]^\top$. Without loss of generality and for simplicity, we only consider the weight matrix in the following. The upper bound of softmax cross entropy loss [39] can therefore be derived as in Equation (6), with proof in supplementary material.

$$\begin{aligned} \mathbb{E}_{f(\mathbf{x}_i)} \left[-\log \frac{\exp(w_{y_i}^\top f(\mathbf{x}_i))}{\sum_{c=1}^C \exp(w_c^\top f(\mathbf{x}_i))} \right] \\ \leq \log \left[\sum_{c=1}^C \exp((w_c^\top - w_{y_i}^\top) \boldsymbol{\mu}_{y_i} + \frac{\lambda}{2} (w_c^\top - w_{y_i}^\top) \boldsymbol{\Sigma}_{y_i} (w_c - w_{y_i})) \right]. \end{aligned} \quad (6)$$

This indicates that by adding the penalty of $\frac{\lambda}{2} (w_c^\top - w_{y_i}^\top) \boldsymbol{\Sigma}_{y_i} (w_c - w_{y_i})$, the up-boundary of classification loss can be approximated by implicit augmentation, where λ can be considered as a term to control the augmentation degree [39].

Thus far, we assume a nearly-identical visual space (i.e., similar $f(\mathbf{x}|y=c) \approx \boldsymbol{\mu}_c$ across domains) after visual-semantic mapping; however the estimation of $\boldsymbol{\Sigma}_c$ is hardly possible for tail classes. Guided by the semantic inter-class relationship from \mathbf{s} , we select the top k classes that are most similar to the corresponding class c (including c itself). This stems from the observation that similar classes are supposed to have similar semantic variances. For example, *deer* and *antelope* may share similar characteristics of the variations of shape, color, etc. Motivated by this, we introduce a weighted covariance estimation strategy to leverage the knowledge learned from head classes,

$$\begin{aligned} Sim_c = \{\mathbf{s}_c^\top \mathbf{s}_i | i = 1, 2, \dots, C\}; \mathbf{k}_c = \{i | \mathbf{s}_c^\top \mathbf{s}_i \in topk(Sim_c)\}, \\ \boldsymbol{\Sigma}'_c = \frac{\sum_{k \in \mathbf{k}_c} n_k \boldsymbol{\Sigma}_k}{\sum_{k \in \mathbf{k}_c} n_k}. \end{aligned} \quad (7)$$

Afterwards, we applied an surrogate loss introduced in Equation (6) to optimize the boundary of classification loss by adding implicit augmentation terms:

$$\mathcal{L}_{aug}(\mathbf{x}_i, y_i; f, h) = -\log \frac{\exp(w_{y_i}^\top f(\mathbf{x}_i))}{\sum_{c=1}^C \exp(w_{y_i}^\top f(\mathbf{x}_i) + \frac{\lambda}{2} (w_c - w_{y_i})^\top \boldsymbol{\Sigma}'_{y_i} (w_c - w_{y_i}))}. \quad (8)$$

In practice, the covariance $\boldsymbol{\Sigma}$ is online calculated from T_Σ steps onwards to avoid the effect of inter-domain variances.

3.3 Meta-Learning Based Generalization

The objective of meta-learning is to ensure that the trained models are robust against domain shifts and perform well on all seen and unseen classes. If combin-

ing the functional blocks introduced above, we can obtain relatively good results by conventional training strategies. However, the generalization capability on unforeseen domains is not guaranteed. Thus, we apply meta-learning to simulate the domain distribution gaps in an episodic manner, inspired by previous DG works [6]. The optimization process is illustrated in Algorithm 1. For each iteration, the training domains are randomly divided into two splits, \mathcal{D}_{mtr} and \mathcal{D}_{mte} . We make sure that data in \mathcal{D}_{mte} always come from different domains from \mathcal{D}_{mtr} . After training on \mathcal{D}_{mtr} , the model is expected to perform well on unseen domains (especially for domain-unique tail classes) in \mathcal{D}_{mte} .

Meta Train. Over the course of meta-training, we make the model trained on \mathcal{D}_{mtr} able to derive semantically meaningful and unbiased representations. With a batch data of size B sampled from each domain in \mathcal{D}_{mtr} , i.e., $\{\mathbf{x}_i, y_i, d_i\}_{i=1}^{B \times |\mathcal{D}_{mtr}|}$, we exert the following typical loss functions.

First of all, in order to calibrate the loss from imbalanced distributions to balanced ones, we apply the domain calibrated softmax loss with $\mathcal{L}_{Cls} = \frac{1}{B \times |\mathcal{D}_{mtr}|} \sum_i \mathcal{L}_{dc}(\mathbf{x}_i, y_i, d_i; f, h)$. In this way, we can improve the performance over all classes and not propagating discouraging gradients to unseen classes. Subsequently, to build unbiased representations, the visual-semantic alignment loss $\mathcal{L}_{Z2S} = \frac{1}{B \times |\mathcal{D}_{mtr}|} \sum_i \mathcal{L}_{z2s}(f(\mathbf{x}_i), y_i, \mathbf{s}; e)$ is adopted to align the embeddings in the semantic space. Furthermore, to enable domain-invariant feature learning as well as to avoid the prohibitive costs of sampling all classes in all domains when the class number is huge, a prototype alignment loss is utilized, $\mathcal{L}_{S2S} = \mathbb{E}_{m,n \in \mathcal{D}_{mtr}} \mathcal{L}_{s2s}(\hat{\mathbf{s}}^m, \hat{\mathbf{s}}^n) + \mathbb{E}_{n \in \mathcal{D}_{mtr}} \mathcal{L}_{s2s}(\hat{\mathbf{s}}^n, \mathbf{s})$. In addition, we apply $\mathcal{L}_{S2Z} = \mathbb{E}_{n \in \mathcal{D}_{mtr}} \mathcal{L}_{z2s}(\hat{\mathbf{v}}^n; e, h)$ to further constrain the semantic manifold.

With intra-class inter-domain distribution aligned, the intra-class variances would be more semantically relevant. We then track the covariance of $f(\mathbf{x})$ feature distribution from T_Σ steps onwards, and apply the surrogate augmentation loss, $\mathcal{L}_{Aug} = \frac{1}{B \times |\mathcal{D}_{mtr}|} \sum_i \mathcal{L}_{aug}(\mathbf{x}_i, y_i; f, h)$, to increase the diversity of feature distributions, especially for the tail classes. \mathcal{L}_{Aug} is set to 0 before T_Σ .

The overall meta-training loss is formulated as $\mathcal{L}_{mtr} = \mathcal{L}_{Cls} + w_1 \mathcal{L}_{Z2S} + w_2 \mathcal{L}_{S2S} + w_3 \mathcal{L}_{S2Z} + w_4 \mathcal{L}_{Aug}$, where w_1, w_2, w_3, w_4 are weight hyperparameters. The model parameters $\theta\{f, h, e, d\}$ at step t are firstly updated based on \mathcal{L}_{mtr} with an optimization step with learning rate β_1 :

$$\theta^t\{f', h', e', d'\} = \theta^t\{f, h, e, d\} - \beta_1 \nabla \mathcal{L}_{mtr}. \quad (9)$$

Meta Test. The model optimized on \mathcal{D}_{mtr} is expected to perform well on held-out domains \mathcal{D}_{mte} . In other words, the optimized representations $\theta\{f', h', e', d'\}$ should be unbiased and semantically meaningful, generalizing well when faced with label distribution shifts and novel domain-specific classes. With samples $\{\mathbf{x}_i, y_i, d_i\}_{i=1}^{B \times |\mathcal{D}_{mte}|}$ from \mathcal{D}_{mte} , the following losses are utilized. First comes the calibrated classification loss $\mathcal{L}_{MCls} = \frac{1}{B \times |\mathcal{D}_{mte}|} \sum_i \mathcal{L}_{dc}(\mathbf{x}_i, y_i, d_i; f', h')$. Subsequently, with $\hat{\mathbf{s}}^n$ of meta training domains updated to $\hat{\mathbf{s}}^{n'}$, another loss $\mathcal{L}_{MZ2S} = \frac{1}{B \times |\mathcal{D}_{mte}|} \sum_i \mathcal{L}_{z2s}(f'(\mathbf{x}_i), y_i, \mathbf{s}; e') + \mathbb{E}_{n \in \mathcal{D}_{mtr}} \mathcal{L}_{z2s}(f'(\mathbf{x}_i), y_i, \hat{\mathbf{s}}^{n'}; e')$ aims to align

Algorithm 1 Meta-learning for long-tailed domain generalization.

Input: Training set \mathcal{D}^{tr} ; semantic embeddings \mathbf{s} ; initialized visual prototype \mathbf{v}^n .
Hyperparameters: Steps T_Σ, T_{max} ; Weights $w_1, w_2, w_3, w_4, w_{mte}$; LR β_1, β_2 .
Initialized model parameters: Feature extractor f ; classifier h ; models e and d .
Output: Optimized models f and h .

- 1: **for** $t \leq T_{max}$ **do**
- 2: Randomly split \mathcal{D}_{tr} into \mathcal{D}_{mtr} and \mathcal{D}_{mte} .
- 3:
- 4: Sample $\{\mathbf{x}_i, y_i, d_i\}_{i=1}^{B \times |\mathcal{D}_{mtr}|}$ from \mathcal{D}_{mtr} . ▷ Meta Train.
- 5: Calculate losses: $\mathcal{L}_{Cls}, \mathcal{L}_{Z2S}$.
- 6: Update $\{\mathbf{v}^n | n \in \mathcal{D}_{mtr}\}$ based on $f(\mathbf{x}_i)$.
- 7: Calculate new $\{\hat{\mathbf{s}}^n, \hat{\mathbf{v}}^n | n \in \mathcal{D}_{mtr}\}$ based on e, d, \mathbf{v} .
- 8: Update Σ, Σ' when $t \geq T_\Sigma$.
- 9: Calculate losses: $\mathcal{L}_{S2S}, \mathcal{L}_{S2Z}$, and \mathcal{L}_{Aug} .
- 10: Calculate meta-training loss: $\mathcal{L}_{mtr} = \mathcal{L}_{Cls} + w_1 \mathcal{L}_{Z2S} + w_2 \mathcal{L}_{S2S} + w_3 \mathcal{L}_{S2Z} + w_4 \mathcal{L}_{Aug}$.
- 11: Calculate new $\theta^t\{f', h', e', d'\} = \theta^t\{f, h, e, d\} - \beta_1 \nabla \mathcal{L}_{mtr}$.
- 12:
- 13: Sample $\{\mathbf{x}_i, y_i, d_i\}_{i=1}^{B \times |\mathcal{D}_{mte}|}$ from \mathcal{D}_{mte} . ▷ Meta Test.
- 14: Calculate losses: $\mathcal{L}_{MCls}, \mathcal{L}_{MZ2S}, \mathcal{L}_{MAug}$.
- 15: Calculate new $\{\hat{\mathbf{s}}^{n'}, \hat{\mathbf{v}}^{n'} | n \in \mathcal{D}_{mtr}\}$ based on $e', d', \{\mathbf{v}^n\}$.
- 16: Calculate meta-testing loss: $\mathcal{L}_{mte} = \mathcal{L}_{MCls} + w_1 \mathcal{L}_{MZ2S} + w_4 \mathcal{L}_{MAug}$.
- 17: Update $\theta^{t+1}\{f, h, e, d\} = \theta^t\{f, h, e, d\} - \beta_2 \nabla (\mathcal{L}_{mtr} + w_{mte} \mathcal{L}_{mte})$.
- 18: **end for**

the visual features to both the semantic embeddings \mathbf{s} and the visual prototypes of \mathcal{D}_{mtr} . This ensures that the knowledge extracted from meta-training steps are domain-invariant and semantically meaningful across all classes seen in \mathcal{D}_{mte} . In addition, the surrogate augmentation loss is enforced from T_Σ onwards, to increase the feature diversity, $\mathcal{L}_{MAug} = \frac{1}{B \times |\mathcal{D}_{mte}|} \sum \mathcal{L}_{aug}(\mathbf{x}_i, y_i; f', h')$.

The overall meta-test loss is $\mathcal{L}_{mte} = \mathcal{L}_{MCls} + w_1 \mathcal{L}_{MZ2S} + w_4 \mathcal{L}_{MAug}$, and we finally update the $\theta\{f, h, e, d\}$ based on $\mathcal{L}_{mtr} + w_{mte} \mathcal{L}_{mte}$ by learning rate β_2 :

$$\theta^{t+1}\{f, h, e, d\} = \theta^t\{f, h, e, d\} - \beta_2 \nabla (\mathcal{L}_{mtr} + w_{mte} \mathcal{L}_{mte}). \quad (10)$$

4 Experiments

4.1 Experimental Settings

Datasets. Two datasets were adopted to evaluate the effectiveness of our proposed methods. **AWA2** [40] and **ImageNet-LT** [21]. To benchmark our proposed task, **AWA2** was firstly randomly resampled to convert to a long-tailed version. Regarding **AWA2-LT** and **ImageNet-LT**, we applied off-the-shelf style transfer models (Hayao & Shinkai²; Vangogh & Ukiyoe³) to simulate domain shifts as shown in Fig. 6 and the generated new datasets are referred to as **AWA2-LTS** and **ImageNet-LTS**, respectively. Some classes from certain domains were deliberately dropped out to simulate a more realistic settings with



Fig. 6. Examples of simulated domain shifts by off-the-shelf style transfer models. Each original sample was randomly assigned a style to make up for the final dataset. Some samples of non-head classes in certain domains were randomly dropped out to simulate the practical **LT-DS** problem studied in this work.




Fig. 7. Data distribution of **AWA2-LTS**.

Table 1. Dataset Details.

Dataset	$ D $	$ C $	$ x $	Ratio*
AWA2-LTS	5	50	~8k	~78
ImageNet-LTS	5	1000	~100k	~256

entangled long-tailed and shifted distributions. A visual explanation of **AWA2-LTS** training data distribution is given in Fig. 7. Please see Table 1 and refer to supplementary material for more details.

Model Training Setup. We deployed all models using Pytorch with NVIDIA RTX3090. For **AWA2-LTS** and **ImageNet-LTS**, ResNet10 was adopted as the backbone with random initialization. Regarding semantic features, Word2Vec embeddings [24] were adopted for **ImageNet-LT**. BERT embeddings [5] of class descriptors were utilized for **AWA2-LTS**. More details can be found in supplementary material and our project page <https://xiaogu.site/LTDS>.

Evaluation Setup. Following [29,42], a threshold was empirically set up to apply on the prediction confidence. Classes below the given threshold are defined as open classes, belonging to $\mathcal{Y} \setminus \mathcal{Y}^{tr}$. We applied leave-one-domain-out protocol for evaluation. We reported the **Acc-U** for the non-open classes in the held-out unseen domain. In addition, the domain-averaged accuracy of non-open classes in all the domains **Acc**, and harmonic score **H** of all the classes in all the domains are reported. The score **Acc** and **H** take into account those classes belonging to $\mathcal{Y}^{tr} \setminus \mathcal{Y}^i$ of each individual domain i . Therefore both metrics can to some extent reflect whether those domain-specific tail classes could lead to spurious correlations during training.

4.2 Long-Tail with Conditional Distribution Shifts

We evaluated the performance under **LT-DS** based on two benchmarks **AWA2-LTS** and **ImageNet-LTS** proposed in this work. For comparison, Agg Baseline, **LT** solutions (cRT [14], BSCE [27], Equal [31], Remix [3]), **DG** solutions (Epic-FCR [16], MixStyle [44], CuMix [22], DAML [29]) were implemented.

²<https://github.com/Yijunmaverick/CartoonGAN-Test-Pytorch-Torch>

³<https://github.com/junyanz/pytorch-CycleGAN-and-pix2pix>

Table 2. Results on **AWA2-LTS** based on leave-one-domain-out evaluation.

		Original			Hayao			Shinkai			Vangogh			Ukiyoe		
Method		Acc-U	Acc	H	Acc-U	Acc	H	Acc-U	Acc	H	Acc-U	Acc	H	Acc-U	Acc	H
LT	Agg	29.4	27.0	34.5	20.4	29.9	36.2	30.8	33.5	38.6	27.1	34.2	41.0	25.5	34.2	38.6
	cRT[14]	30.4	29.1	34.8	23.5	33.6	38.9	34.7	35.8	39.6	28.6	35.8	43.0	28.4	36.7	35.7
	BSCE[27]	41.8	35.9	41.7	24.7	36.1	41.7	30.2	35.8	40.7	29.0	37.7	43.2	25.9	33.6	35.1
	Equal[31]	34.1	32.9	36.6	24.3	35.3	42.7	33.5	36.2	40.5	28.8	35.8	42.4	27.3	34.7	34.0
	Remix[3]	32.7	30.3	35.9	16.9	30.7	33.3	27.6	32.0	37.5	26.9	31.8	41.2	26.5	32.0	34.9
DG	Epi-FCR[16]	34.0	33.1	40.5	23.3	34.0	40.6	29.7	35.5	39.1	27.5	36.1	42.0	27.0	35.7	38.0
	MixStyle[44]	36.7	34.0	41.2	27.1	36.2	41.7	32.0	36.2	40.6	28.4	36.0	41.9	28.8	36.2	38.3
	CuMix[22]	36.1	33.8	38.6	24.7	35.3	41.0	30.2	35.1	41.4	28.2	35.1	40.9	26.5	34.7	34.6
	DAML[29]	13.9	10.7	16.2	14.7	22.5	29.8	17.3	24.9	30.4	14.3	19.5	25.6	22.9	28.9	36.0
	DAML[29]-Warmup	<u>42.2</u>	35.3	<u>42.5</u>	25.7	35.2	39.5	31.2	36.8	44.0	29.4	37.5	<u>45.3</u>	28.6	36.0	41.4
	MixStyle+BSCE	40.0	36.8	41.8	<u>28.8</u>	<u>39.7</u>	<u>43.5</u>	32.4	38.3	<u>44.2</u>	<u>30.8</u>	38.2	43.3	<u>29.8</u>	<u>38.9</u>	<u>39.3</u>
	Epi-FCR+BSCE	41.3	<u>36.9</u>	42.0	24.0	35.9	41.2	32.0	<u>39.2</u>	<u>42.5</u>	30.1	<u>38.5</u>	41.7	26.6	35.9	38.7
Ours	49.4	42.1	45.8	29.8	42.4	46.3	<u>34.3</u>	42.6	45.3	32.7	40.3	46.3	32.9	42.4	<u>39.3</u>	

Table 3. Results on **ImageNet-LTS** based on leave-one-domain-out evaluation.

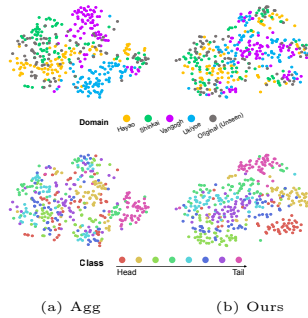
		Original			Hayao			Shinkai			Vangogh			Ukiyoe		
Method		Acc-U	Acc	H	Acc-U	Acc	H	Acc-U	Acc	H	Acc-U	Acc	H	Acc-U	Acc	H
	Agg	19.5	18.1	23.5	13.0	18.2	24.4	15.2	18.0	24.2	13.2	17.5	23.7	12.5	17.9	24.1
LT	cRT[14]	20.7	18.7	24.8	13.8	19.0	24.3	16.2	18.8	24.8	14.0	18.4	<u>25.5</u>	13.6	18.7	25.3
	BSCE[27]	<u>20.8</u>	<u>19.1</u>	<u>25.6</u>	<u>14.3</u>	<u>19.4</u>	<u>25.2</u>	<u>16.5</u>	19.1	<u>25.2</u>	14.7	<u>18.8</u>	<u>25.5</u>	<u>14.1</u>	<u>19.1</u>	<u>26.4</u>
	Equal[31]	16.3	15.4	20.6	10.7	15.2	20.6	13.2	16.4	22.2	10.5	14.8	21.4	10.9	15.8	21.3
	Remix[3]	14.8	13.8	18.6	10.1	14.1	20.0	11.3	14.1	22.0	11.1	13.5	18.2	10.5	14.8	20.8
DG	Epi-FCR[16]	19.2	18.8	23.1	13.5	19.0	23.3	15.0	<u>20.0</u>	25.1	12.2	18.0	23.8	13.0	17.5	25.2
	MixStyle[44]	17.7	16.4	22.0	12.1	16.7	22.1	13.6	16.5	22.9	11.8	16.0	22.9	11.5	16.2	21.4
	CuMix[22]	18.2	17.2	23.7	13.2	17.5	22.7	14.2	17.1	22.4	12.1	16.8	22.3	12.1	17.1	23.7
	DAML[29]-Warmup	14.7	12.7	18.2	10.5	13.0	18.2	11.5	13.2	18.0	9.7	12.8	16.9	10.1	13.0	18.6
Ours		24.3	20.8	25.9	16.3	21.3	25.4	17.4	20.9	25.8	<u>14.3</u>	20.3	26.6	15.4	20.3	26.7

Quantitative Results. As shown in Tables 2 & 3, our method achieves overall superior performance compared to other methods. In addition, two combinations MixStyle+BSCE and Epi-FCR+BSCE were applied on **AWA2-LTS** dataset for comparison. We noticed the extremely low performance of DAML [29] under **LT-DS**. Since it is based on the knowledge distillation from each domain-specific model, the knowledge learned from individual long-tailed distribution would be significantly biased towards the head classes. We simply moderated such bias by a warm-up pretraining, and the variant is referred to as DAML-Warmup. Although DAML-Warmup can alleviate the class imbalance issue when the imbalance ratio and class number is small, we observed its failure on **ImageNet-LTS**. It may indicate that it cannot handle the semantic information when a large proportion of classes are missing. Overall, our results demonstrate superior performance compared to existing methods.

Qualitative Results. We present the t-SNE visualizations of features from the test split of **AWA2-LTS** in Fig. 8 along with the results from the baseline Agg. Different colors indicate different domains on the upper row, whereas different categories ranging from head to tail on the bottom. It can be observed that the samples which are sampled from the same classes but from different domains

Table 4. Ablation studies on **AWA2-LTS**. CE indicates cross-entropy loss.

Index	\mathcal{L}_{dc}	CE	\mathcal{L}_{Z2S}	\mathcal{L}_{S2S}	\mathcal{L}_{S2Z}	\mathcal{L}_{Aug}	Meta	Avg		
								Acc-U	Acc	H
a	-	✓	-	-	-	-	-	26.6	31.8	37.8
b	✓	-	-	-	-	-	-	30.6	36.7	39.2
c	-	✓	-	-	-	-	✓	28.9	34.3	33.8
d	✓	-	-	-	-	-	✓	34.3	40.0	40.2
e	✓	-	✓	-	-	-	-	31.5	<u>41.2</u>	41.0
f	✓	-	✓	✓	-	-	-	32.0	40.8	41.2
g	✓	-	✓	✓	✓	-	-	32.5	40.8	42.8
h	✓	-	-	-	-	✓	-	26.2	30.0	36.8
i	✓	-	✓	✓	✓	✓	-	34.0	40.9	42.8
j	✓	-	✓	✓	✓	✓	✓	35.7	42.0	44.6
k	Single Prototype Alignment							32.5	41.0	41.6
l	No-Weight Augmentation							<u>34.6</u>	40.7	<u>43.0</u>

**Fig. 8.** t-SNE of **AWA2-LTS** test set. Left: Agg. Right: Ours.

(including the unseen domain) are better clustered. More qualitative analysis can be found in supplementary material.

Ablation Studies. We validate the performance of each proposed module and the whole meta learning framework by ablation studies as shown in Table 4 indicates the effectiveness of each individual modules. In particular, the ablated model indexed by d also presents relatively good performance, demonstrating that even without additional semantic features for alignment, it is still comparable to existing solutions. We also did two another ablation studies, with just single prototype for alignment and without weighted term in Equation (8) for updating covariance matrix. Please see more detailed discussions in our supplementary material. Additional experiments on sample complexity, choices of embeddings, as well as results on open domain generalization datasets can also be found in supplementary materials.

5 Conclusions

Long-tailed category distribution and domain shifts have been two major issues concerned with real-world datasets, leading to degraded performance upon practical deployment. The combination of these two problems poses a significantly challenging scenario, where not only these two problems should be addressed, but the domain-specific non-head classes should also be paid attention to, to avoid shortcuts. We proposed a meta-learning framework to ensure that the model can perform well over all classes and all domains, including unseen novel domains. We evaluated two benchmarks proposed in this paper. The experimental results demonstrate that the proposed method can achieve superior performance, when compared to either long-tailed/domain-generalization solutions or the combinations. In the future, we are going to apply our method to more specific applications like behavioural analysis and health care.

References

1. Buda, M., Maki, A., Mazurowski, M.A.: A systematic study of the class imbalance problem in convolutional neural networks. *Neural Networks* **106**, 249–259 (2018)
2. Chen, Y., Rong, F., Duggal, S., Wang, S., Yan, X., Manivasagam, S., Xue, S., Yumer, E., Urtasun, R.: Geosim: Realistic video simulation via geometry-aware composition for self-driving. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 7230–7240 (2021)
3. Chou, H.P., Chang, S.C., Pan, J.Y., Wei, W., Juan, D.C.: Remix: Rebalanced mixup. In: *European Conference on Computer Vision*. pp. 95–110. Springer (2020)
4. Damen, D., Doughty, H., Farinella, G.M., Furnari, A., Kazakos, E., Ma, J., Moltisanti, D., Munro, J., Perrett, T., Price, W., et al.: Rescaling egocentric vision. *arXiv preprint arXiv:2006.13256* (2020)
5. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018)
6. Dou, Q., Coelho de Castro, D., Kamnitsas, K., Glocker, B.: Domain generalization via model-agnostic learning of semantic features. *Advances in Neural Information Processing Systems* **32**, 6450–6461 (2019)
7. Gu, X., Guo, Y., Deligianni, F., Lo, B., Yang, G.Z.: Cross-subject and cross-modal transfer for generalized abnormal gait pattern recognition. *IEEE Transactions on Neural Networks and Learning Systems* **32**(2), 546–560 (2020)
8. Gupta, A., Savarese, S., Ganguli, S., Fei-Fei, L.: Embodied intelligence via learning and evolution. *Nature communications* **12**(1), 1–12 (2021)
9. Han, H., Wang, W.Y., Mao, B.H.: Borderline-smote: a new over-sampling method in imbalanced data sets learning. In: *International conference on intelligent computing*. pp. 878–887. Springer (2005)
10. Hong, Y., Han, S., Choi, K., Seo, S., Kim, B., Chang, B.: Disentangling label distribution for long-tailed visual recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6626–6636 (2021)
11. Huang, C., Li, Y., Loy, C.C., Tang, X.: Learning deep representation for imbalanced classification. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 5375–5384 (2016)
12. Iscen, A., Araujo, A., Gong, B., Schmid, C.: Class-balanced distillation for long-tailed visual recognition. In: *Proceedings of the British Machine Vision Conference (BMVC)*. BMVA Press (2021)
13. Jamal, M.A., Brown, M., Yang, M.H., Wang, L., Gong, B.: Rethinking class-balanced methods for long-tailed visual recognition from a domain adaptation perspective. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 7610–7619 (2020)
14. Kang, B., Xie, S., Rohrbach, M., Yan, Z., Gordo, A., Feng, J., Kalantidis, Y.: Decoupling representation and classifier for long-tailed recognition. *arXiv preprint arXiv:1910.09217* (2019)
15. Li, D., Yang, Y., Song, Y.Z., Hospedales, T.M.: Deeper, broader and artier domain generalization. In: *Proceedings of the IEEE international conference on computer vision*. pp. 5542–5550 (2017)
16. Li, D., Zhang, J., Yang, Y., Liu, C., Song, Y.Z., Hospedales, T.M.: Episodic training for domain generalization. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 1446–1455 (2019)

17. Li, S., Gong, K., Liu, C.H., Wang, Y., Qiao, F., Cheng, X.: Metasaug: Meta semantic augmentation for long-tailed visual recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5212–5221 (2021)
18. Li, Y., Tian, X., Gong, M., Liu, Y., Liu, T., Zhang, K., Tao, D.: Deep domain generalization via conditional invariant adversarial networks. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 624–639 (2018)
19. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE international conference on computer vision. pp. 2980–2988 (2017)
20. Liu, X., Hu, B., Jin, L., Han, X., Xing, F., Ouyang, J., Lu, J., El Fakhri, G., Woo, J.: Domain generalization under conditional and label shifts via variational bayesian inference. In: Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21. pp. 881–887. International Joint Conferences on Artificial Intelligence Organization (8 2021)
21. Liu, Z., Miao, Z., Zhan, X., Wang, J., Gong, B., Yu, S.X.: Large-scale long-tailed recognition in an open world. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2537–2546 (2019)
22. Mancini, M., Akata, Z., Ricci, E., Caputo, B.: Towards recognizing unseen categories in unseen domains. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIII 16. pp. 466–483. Springer (2020)
23. Maniayar, U., Deshmukh, A.A., Dogan, U., Balasubramanian, V.N., et al.: Zero shot domain generalization. arXiv preprint arXiv:2008.07443 (2020)
24. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems* **26** (2013)
25. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. arXiv preprint arXiv:2103.00020 (2021)
26. Ravì, D., Wong, C., Deligianni, F., Berthelot, M., Andreu-Perez, J., Lo, B., Yang, G.Z.: Deep learning for health informatics. *IEEE journal of biomedical and health informatics* **21**(1), 4–21 (2016)
27. Ren, J., Yu, C., Sheng, S., Ma, X., Zhao, H., Yi, S., Li, H.: Balanced meta-softmax for long-tailed visual recognition. arXiv preprint arXiv:2007.10740 (2020)
28. Samuel, D., Atzmon, Y., Chechik, G.: From generalized zero-shot learning to long-tail with class descriptors. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 286–295 (2021)
29. Shu, Y., Cao, Z., Wang, C., Wang, J., Long, M.: Open domain generalization with domain-augmented meta-learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9624–9633 (2021)
30. Socher, R., Ganjoo, M., Manning, C.D., Ng, A.: Zero-shot learning through cross-modal transfer. In: *Advances in Neural Information Processing Systems*. pp. 935–943 (2013)
31. Tan, J., Wang, C., Li, B., Li, Q., Ouyang, W., Yin, C., Yan, J.: Equalization loss for long-tailed object recognition. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11662–11671 (2020)
32. Tang, K., Huang, J., Zhang, H.: Long-tailed classification by keeping the good and removing the bad momentum causal effect. arXiv preprint arXiv:2009.12991 (2020)
33. Wang, F., Cheng, J., Liu, W., Liu, H.: Additive margin softmax for face verification. *IEEE Signal Processing Letters* **25**(7), 926–930 (2018)

34. Wang, G., Han, H., Shan, S., Chen, X.: Cross-domain face presentation attack detection via multi-domain disentangled representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6678–6687 (2020)
35. Wang, J., Lan, C., Liu, C., Ouyang, Y., Zeng, W., Qin, T.: Generalizing to unseen domains: A survey on domain generalization. arXiv preprint arXiv:2103.03097 (2021)
36. Wang, P., Han, K., Wei, X.S., Zhang, L., Wang, L.: Contrastive learning based hybrid networks for long-tailed image classification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 943–952 (2021)
37. Wang, X., Lian, L., Miao, Z., Liu, Z., Yu, S.: Long-tailed recognition by routing diverse distribution-aware experts. In: International Conference on Learning Representations (2021), <https://openreview.net/forum?id=D9I3drBz4UC>
38. Wang, Y.X., Ramanan, D., Hebert, M.: Learning to model the tail. In: Proceedings of the 31st International Conference on Neural Information Processing Systems. pp. 7032–7042 (2017)
39. Wang, Y., Pan, X., Song, S., Zhang, H., Huang, G., Wu, C.: Implicit semantic data augmentation for deep networks. *Advances in Neural Information Processing Systems* **32**, 12635–12644 (2019)
40. Xian, Y., Lampert, C.H., Schiele, B., Akata, Z.: Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly. *IEEE transactions on pattern analysis and machine intelligence* **41**(9), 2251–2265 (2018)
41. Xian, Y., Sharma, S., Schiele, B., Akata, Z.: f-vaegan-d2: A feature generating framework for any-shot learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10275–10284 (2019)
42. You, K., Long, M., Cao, Z., Wang, J., Jordan, M.I.: Universal domain adaptation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2720–2729 (2019)
43. Zhang, Y., Kang, B., Hooi, B., Yan, S., Feng, J.: Deep long-tailed learning: A survey. arXiv preprint arXiv:2110.04596 (2021)
44. Zhou, K., Yang, Y., Qiao, Y., Xiang, T.: Mixstyle neural networks for domain generalization and adaptation. arXiv preprint arXiv:2107.02053 (2021)
45. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE international conference on computer vision. pp. 2223–2232 (2017)