

Bidirectional Logits Tree: Pursuing Granularity Reconciliation in Fine-Grained Classification

Zhiguang Lu^{1,2}, Qianqian Xu^{1*}, Shilong Bao², Zhiyong Yang², Qingming Huang^{1,2,3*}

¹Key Laboratory of Intelligent Information Processing, Institute of Computing Technology, Chinese Academy of Sciences

²School of Computer Science and Technology, University of Chinese Academy of Sciences

³Key Laboratory of Big Data Mining and Knowledge Management, University of Chinese Academy of Sciences
{luzhiguang23s, xuqianqian}@ict.ac.cn, {baoshilong, yangzhiyong21, qmhuang}@ucas.ac.cn

Abstract

This paper addresses the challenge of *Granularity Competition* in fine-grained classification tasks, which arises due to the semantic gap between multi-granularity labels. Existing approaches typically develop independent hierarchy-aware models based on shared features extracted from a common base encoder. However, because coarse-grained levels are inherently easier to learn than finer ones, the base encoder tends to prioritize coarse feature abstractions, which impedes the learning of fine-grained features. To overcome this challenge, we propose a novel framework called the Bidirectional Logits Tree (BiLT) for *Granularity Reconciliation*. The key idea is to develop classifiers sequentially from the finest to the coarsest granularities, rather than parallelly constructing a set of classifiers based on the same input features. In this setup, the outputs of finer-grained classifiers serve as inputs for coarser-grained ones, facilitating the flow of hierarchical semantic information across different granularities. On top of this, we further introduce an Adaptive Intra-Granularity Difference Learning (AIGDL) approach to uncover subtle semantic differences between classes within the same granularity. Extensive experiments demonstrate the effectiveness of our proposed method.

Code — <https://github.com/ZhiguangLuu/BiLT>

Extended version — <https://arxiv.org/abs/2412.12782>

Introduction

Fine-grained visual classification (FGVC) has been a long-standing and challenging research focus in the deep learning community (Xu et al. 2023; Jiang et al. 2024; Pu et al. 2024). Unlike traditional image classification tasks, FGVC demands models to capture subtle distinctions between categories for accurate predictions. Recently, several studies (Zhang et al. 2024; Chen et al. 2022; Wang et al. 2021) have recognized that the hierarchical label structure (HLS) inherent in class names (Deng et al. 2009; Krizhevsky, Hinton et al. 2009; Zhao, Li, and Xing 2011; Maji et al. 2013) can significantly enhance the understanding and performance of FGVC tasks. Consequently, substantial efforts have been made to explore methods for effectively integrating hierarchical information into FGVC tasks (Jain, Karthik, and

Gandhi 2024; Wang et al. 2023). Generally, once hierarchical semantic information is incorporated, the primary goal is to ensure that the model’s predictions consistently align with the inherent hierarchical structure of the data. In other words, these methods aim to make precise predictions or at least keep the results as close as possible to semantically similar categories even when errors occur, thereby minimizing the severity of mistake (Bertinetto et al. 2020). For example, suppose a model misclassifies a “Husky”. In that case, it is more acceptable within the hierarchy for the output to be another dog breed, such as a “Corgi”, rather than an unrelated category like a “Siamese Cat”.

In pursuit of this goal, one effective paradigm (Silla and Freitas 2011; Zhang and Zhou 2013) usually adopts a common shared encoder to extract the semantic features of images, and then **constructs hierarchy-aware models independently** for the classification tasks at different levels. Typically, (Bertinetto et al. 2020) introduces a multi-task learning method that employs hierarchical cross-entropy loss and a soft label technique to train these hierarchy-aware models. On top of this, (Karthik et al. 2021) develops a post hoc approach using Conditional Risk Minimization (CRM) to calibrate likelihoods during the test phase. Besides, borrowing the idea from neural collapse (Papayan, Han, and Donoho 2020), (Liang and Davis 2023) utilizes an Equiangular Tight Frame to achieve feature alignment across all classes.

Despite significant progress, a critical challenge remains: the current FGVC paradigm struggles with the issue of **Granularity competition** in feature learning due to the semantic gap between multi-granularity labels (as shown in Fig.2). Specifically, coarse labels are generally much easier to distinguish than finer ones. As a result, if we pay equal attention to all levels, the learning process of the underlying encoder is likely to be dominated by the coarse branches, leading to insufficient learning of detailed information at finer granularities. To address this issue, (Chang et al. 2021) disentangles coarse-level features from fine-grained ones via level-specific classification heads, and also uses the finer-grained features to enhance the performance of coarser-grained predictions. (Garg, Sani, and Anand 2022) proposes a soft balance strategy to align the learning process of multi-granularity classifications. However, these methods primarily emphasize the negative impacts of coarse-grained learning on finer details, while overlooking the crucial fact

*Corresponding Authors

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

that **coarse-grained information can also enhance fine-grained learning**. Consequently, the issue of granularity competition remains **inadequately considered**.

In this paper, we propose a generic framework called *Bidirectional Logits Tree* (BiLT), designed to fully utilize hierarchical label information. Specifically, BiLT builds hierarchy-aware models in a sequential manner rather than a parallel one, where the inputs to the coarse classifier depend on the outputs of the preceding finer-grained model. This paradigm improves the priority of learning finer features and also facilitates semantic information flow among multi-granularity labels during training. Meanwhile, classification errors at coarser levels could also serve as an auxiliary supervision signal for upstream finer models' updates, reconciling the granularity competition issue end-to-end. Last but not least, we recognize that accurately identifying subtle semantic differences between sub-classes within the same granularity is also crucial for achieving promising performance. To this end, an Adaptive Intra-Granularity Difference Learning (AIGDL) approach is developed to serve our purpose better.

In summary, the contributions of this paper are as follows:

- We propose a novel paradigm called BiLT to construct hierarchy-aware predictors for fine-grained classifications, which can effectively alleviate the granularity competition issue by facilitating the flow of hierarchical semantic information across all granularities.
- An Adaptive Intra-Granularity Difference Learning (AIGDL) method is further developed to serve our strategy better. It empowers BiLT to learn the fine-grained semantic differences of classes within the same granularity, boosting the final performance sharply.
- Empirical studies over three widely used benchmark datasets consistently demonstrate the effectiveness of our proposed approach.

Related Work

Hierarchical Architecture Methods Hierarchical architecture methods aim to design a hierarchy-aware model architecture that leverages the hierarchical relationship. (Silla and Freitas 2011) first introduced hierarchy-aware classifiers that incorporated a well-defined hierarchy, empirically outperforming flat classifiers across various application scenarios. Building on this idea, (Wu et al. 2016; Bertinetto et al. 2020) proposed connecting classifiers at all levels to a shared feature vector, framing the overall optimization as a multi-task learning problem. However, (Chang et al. 2021) later highlighted a potential issue with this method that sharing the same feature vector across multiple granularities can cause granularity competition problems, where coarse-level predictions impede fine-grained feature learning, whereas fine-grained feature learning can facilitate coarse-grained classification. To avoid this issue, (Liang and Davis 2023) from the perspective of neural collapse, aligned features with class distances by fixing classifier weights to a pre-computed Equiangular Tight Frame based on the distance matrix. However, this method is limited by the requirement that the feature dimension must equal the number of classes,

making it difficult to apply to large datasets and certain downstream tasks. More recently, (Jain, Karthik, and Gandhi 2024) proposed a multi-granularity ensemble method, training multiple neural networks at different levels separately, and adjusting fine-grained outputs based on coarse-grained outputs during testing. This approach improved Top-1 Accuracy and reduced Mistake Severity in fine-grained classification. However, this approach requires training multiple networks separately, which is computationally expensive and difficult to scale to datasets with many levels.

Hierarchical Loss/Cost Methods Hierarchical loss methods aim to design loss functions that can effectively leverage the hierarchical relationships between labels. (Bertinetto et al. 2020) proposed a hierarchical cross-entropy (HXE) loss, a conditional probabilistic loss that conditions class probabilities on those of their ancestors. This approach jointly optimizes the HXE loss across all granularities, framing it as a multi-task learning problem. In contrast, (Karthik et al. 2021) introduced a cost-sensitive method based on conditional risk minimization, calibrating outputs according to the class-relationship matrix during the test phase to minimize the risk. As noted earlier, using hard labels for coarse-grained classification can impede fine-grained feature learning. To address this, (Garg, Sani, and Anand 2022) proposed summing the soft labels of subclasses and aligning them with coarse-grained hard labels using Jensen-Shannon divergence, while also aligning subclass features with those of their superclass using a geometric consistency loss.

Label Embedding Methods This approach focuses on uncovering relationships in a unified label semantic space and therefore describes differences between labels by the distance in this space. For instance, (Frome et al. 2013; Xian et al. 2016) started early attempts to explore a generic algorithm for learning a joint embedding space for images and labels simultaneously. Inspired by the strengths of hyperbolic space in modeling hierarchical structures, (Liu et al. 2020) proposed learning a hyperbolic label space for fine-grained classifications, leading to promising performance. In addition, to learn a favorable space, (Bertinetto et al. 2020) first initialized relationships between labels based on their lowest common ancestor (LCA) height, and proposed a soft label method for precise fine-tuning. Taking a step further, (Zhang et al. 2021) developed an online label smoothing strategy to adaptively learn the label embedding during training, while (Collins, Bhatt, and Weller 2022) proposed a crowdsourced soft label method for individual images. Additionally, label smoothing plays a significant role in various tasks, including but not limited to knowledge distillation (Yuan et al. 2020; Park et al. 2023; Han et al. 2024), image generation (Zhang et al. 2023), weakly supervised learning (Gong, Bisht, and Xu 2024; Wei et al. 2022), and trustworthy machine learning (Qin et al. 2021). Overall, label smoothing is a widely used technique to enhance model generalizability and performance (Müller, Kornblith, and Hinton 2019). Nevertheless, these methods often fail to utilize prior knowledge of inter-label distances. They heavily rely on pre-defined class relationships, or even solely depend on heuristic strategies for label embedding adjustment, which results

in suboptimal performance and slow convergence.

Problem Definition

Let \mathcal{X} and \mathcal{Y} be the input space and label space, respectively. Under the context of fine-grained visual classifications (FGVC) (Liu et al. 2024; Wang et al. 2024; Du et al. 2021), the label space \mathcal{Y} could be formulated as a hierarchy label tree with $H + 1$ levels, where each node corresponds to a specific class, and each edge contains high-level semantic relationships between classes. Typically, for $h \in \{0, 1, 2, \dots, H\}$, the root node ($h = 0$) represents the superclass of all classes, and the label becomes finer and finer as the level h increases in the tree. Subsequently, let C^h be the number of classes at the h -th level, the label space could be expressed as $\mathcal{Y} = \bigcup_{h=0}^H \mathcal{Y}_h$, where $\mathcal{Y}_h = \{0, 1, \dots, C^h - 1\}$ and we have $C^0 = 1 < C^1 \leq C^2 \leq \dots \leq C^H$. Finally, suppose that there are N training samples, denoted as $\mathcal{D} = \{(\mathbf{x}_i, \{y_i^h\}_{h=1}^H)\}_{i=1}^N$, where $\mathbf{x}_i \in \mathcal{X}$ is an input image and $y_i^h \in \mathcal{Y}_h$ corresponds to its ground-truth label at the h -th level. The primary concern of this paper is to train a well-performed classifier that can accurately predict the fine-grained class for each image while ensuring that the predictions are close to the ground truth when it makes mistake.

Methodology

In this section, we will first reveal the fundamental limitations of current fine-grained classification methods, i.e., *the granularity competition issue*. To address this, we propose a generic framework called Bidirectional Logits Tree (BiLT). Fig. 1 presents the overall pipeline of our proposed BiLT. On top of this, a novel *Adaptive Intra-Granularity Difference Learning* is developed to exploit the fine-grained semantics among all classes sufficiently. Further details on BiLT will be discussed in the following.

Granularity Competition Impedes Fine-grained Learning

We start with our discussions by reviewing the traditional FGVC paradigms (Wang et al. 2015) using the hierarchy label tree. Briefly speaking, given an image x with its label sets $\{y^h\}_{h=1}^H$, conventional FGVC methods using the hierarchy label tree first adopt a base encoder Φ to extract the semantic features of each sample, i.e., $\Phi(x)$. On this basis, an independent classifier (denoted as f^h) will be constructed for each label level h whose goal is to make an accurate prediction \hat{y}^h as much as possible at its own level. In order to learn these classifiers effectively, current studies (Wang et al. 2024; Garg, Sani, and Anand 2022; Chang et al. 2021) usually consider the following multi-task optimization problem:

$$\min_{\Phi, f^1, \dots, f^H} \mathbb{E}_{\mathcal{D}} \left[\sum_{h=1}^H \lambda_h \cdot \mathcal{L}^h(\mathcal{Y}^h, f^h(\Phi(x))) \right], \quad (1)$$

where λ_h is a tunable parameter, and \mathcal{L}^h represents the classification error (such as cross-entropy loss) at h -th level in the hierarchy label tree. Fig. 3 presents a toy example

($H = 3$) to instantiate the mainstream learning paradigm of FGVC.

Despite great success, in this paper, we argue that current FGVC paradigms using the hierarchy label tree, suffer from *Granularity competition* issues, leading to limited performance. According to Eq.1, it is apparent that coarse-grained levels are inherently simpler to learn than fine-grains, so the learning process of the base encoder Φ will be almost dominated by the shallow level (i.e., coarse labels). Fig. 2 examines this phenomenon on the real-world benchmark dataset FGVC-Aircraft, showing that **methods sharing the same features at fine-grained levels converge more slowly and yield lower accuracy** compared to only fine-grained level trained exclusively with standard cross-entropy loss (Only CE). As a result, the unique features extracted by Φ inevitably collapse, where the detailed information for fine-grained feature learns insufficiently (Chang et al. 2021). In addition, current paradigms have not fully explored the semantic relationships between classes at different granularities, which further impedes learning fine-grained features for accurate predictions.

Bidirectional Logits Tree for Granularity Reconciliation

According to the above discussions, the fundamental limitations of previous FGVC literature (Bertinetto et al. 2020; Chang et al. 2021; Garg, Sani, and Anand 2022) come from the fact that all classifiers at different granularities built on unique base encoders. We propose a novel generic framework called Bidirectional Logits Tree (BiLT) to address this. The principle of BiLT is to develop coarse-level classifiers on top of its previous finer classifier outputs instead of constructing a set of classifiers separately, as shown in Fig. 1.

Specifically, let $\phi_i := \Phi(x_i)$ be the feature vector for simplicity. BiLT first constructs the finest-level classifier f^H , which directly takes ϕ_i as the input, i.e.,

$$\mathbf{z}_i^H := f^H(\phi_i),$$

where $\mathbf{z}_i^H \in \mathbb{R}^{C^H}$ is the output logits at level H , and $z_{ij}^H, j \in [C^H]$ indicates the probability of image x_i belonging to the class j . Subsequently, for the class-level $h \in \{1, \dots, H-1\}$, its corresponding classifier f^h is developed following the previous $h - 1$ outputs:

$$\mathbf{z}_i^h := f^h(g^h(\mathbf{z}_i^{h+1})),$$

where $\mathbf{z}_i^h \in \mathbb{R}^{C^h}$ and g^h is a trainable transformation layer to explore nonlinear relationships between different levels. In this paper, g^h is implemented by a sequential module, including batch normalization, linear transformation, batch normalization, and ELU activation (Liang and Davis 2023).

We still adopt a similar manner as Eq.1 to learn BiLT. Here the loss for each sample x_i at h -th level is defined as:

$$\mathcal{L}_{\text{BiLT}}^h(\mathcal{Y}^h, f^h) = \mathbf{y}_i^h \log(\text{softmax}(\mathbf{z}_i^h)),$$

where \mathbf{y}_i^h is a one-hot encoding with the y_i^h -th indices as 1.

Intuitively, the learning process of our proposed BiLT *boosts a bidirectional semantic information flow across all*

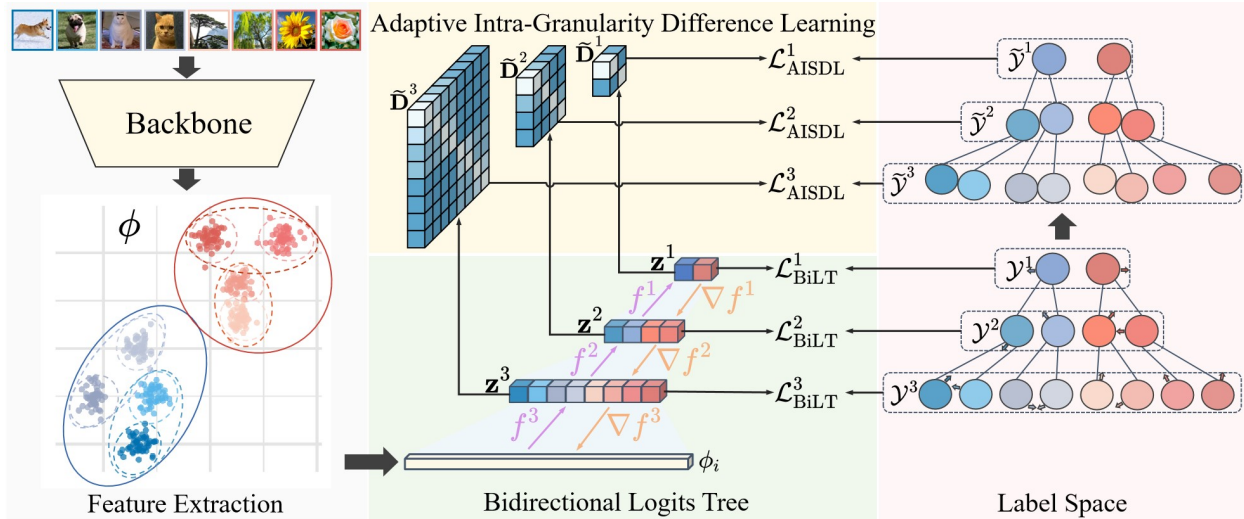


Figure 1: The overall framework of our method. In the **forward** phase of the Bidirectional Logits Tree (BiLT), coarse-grained logits are derived from fine-grained counterparts, while the gradients from coarse-grained classifiers influence fine-grained classifiers and feature learning in the **backward** phase. Simultaneously, Adaptive Intra-Granularity Difference Learning (AIGDL) adjusts the output of BiLT and supervision by learning differences between categories within the same granularity.

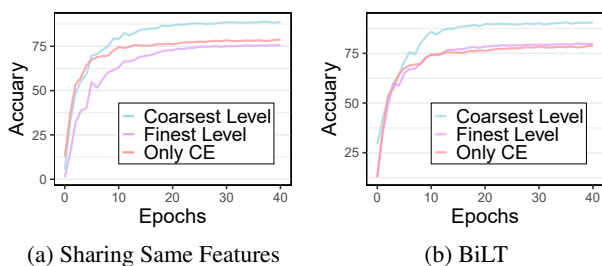


Figure 2: Comparison of convergence speed of Sharing Same Features and BiLT

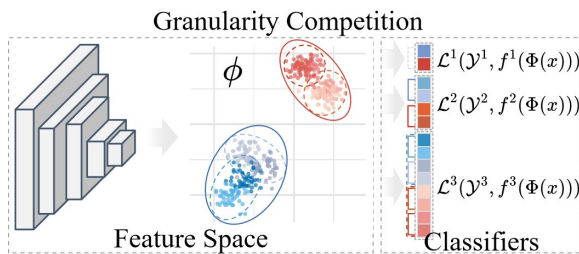


Figure 3: An illustrative example demonstrates the granularity competition problem, where the model prioritizes coarse-grained learning at Level 1 and Level 2, thereby rendering the fine-grained features at Level 3 difficult to distinguish.

granularities properly, as shown in Fig. 1. First of all, during the **forward propagation**, the finer outputs serve as the input features for the coarse-level model recursively, enabling cross-level information explorations. This to some extent **promotes the learning priority of finer features** in the base encoder Φ such that more discriminative features

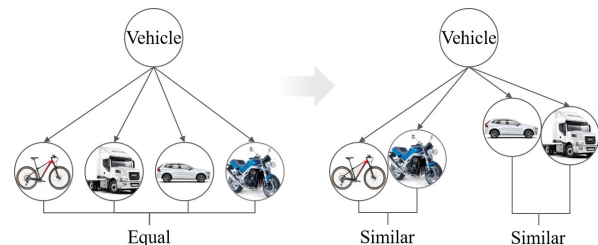


Figure 4: Predefined label trees struggle to articulate differences amongst classes at the same hierarchical level, and our method aims to learn disparities among classes and apply relevant corrections accordingly.

will be abstracted for finer classifications. Meanwhile, during the **backward propagation**, each classifier is not only optimized by the classification risk at the current level but also supervised by those coarser levels. In this way, classifiers at different granularities will interact with each other, leverage hierarchical information effectively, and thus reconcile the granularity competition problem.

Adaptive Intra-Granularity Difference Learning

So far, we have explored how to excavate fine-grained class relationships effectively at different levels. However, mining the semantic information between labels involved **at the same granularity** is also essential for promising performance. Current studies (Garg, Sani, and Anand 2022; Liang and Davis 2023) generally introduce a cost matrix defined on LCA (Least Common Ancestor) to calibrate the final predictions. Based on the fact that there are often obscure semantic differences between sub-classes at the same granularity, **such a strategy** considering intra-granularity classes

equally might **struggle to** articulate them. Fig. 4 provides an example to illustrate it. We know that both “bicycles”, “motorcycles”, “trucks” and “car” belong to “Vehicle”. Yet, in the feature space, it is obvious that the distance between “bicycles” and “motorcycles” should be closer than “bicycles” and “trucks” since their appearance is highly similar. Apparently, this cannot be achieved through simple LCA-based methods.

We propose a novel Adaptive Intra-Granularity Difference Learning (AIGDL) method to remedy this. Concretely, let $\mathbf{p}_i^h = \text{softmax}(\mathbf{z}_i^h)$ be the prediction probability and $\mathbf{D}^h \in \mathbb{R}^{C^h \times C^h}$ be the intra-granularity class distance matrix at level h , where D_{ij}^h represents the LCA distance between class y_i^h and class y_j^h . We first follow the widely adopted Condition Risk Minimization (CRM) paradigm (Garg, Sani, and Anand 2022) to calibrate:

$$\underset{k}{\operatorname{argmin}} R(\mathbf{y}_i^h = k | \mathbf{x}_i) = \underset{k}{\operatorname{argmin}} \sum_{j=1}^{C^h} \mathbf{D}_{k,j}^h \cdot p^h(\mathbf{y}_i^h = k | \mathbf{x}_i).$$

Compared with the typical likelihood maximization, this CRM adjustment guarantees us to select the Bayes optimal prediction $R(\mathbf{y}_i^h = k | \mathbf{x}_i)$ resulting in the lowest possible overall cost (Duda and Hart 1974).

On top of \mathbf{D}^h , a learnable intra-granularity distance matrix $\Delta^h \in \mathbb{R}^{C^h \times C^h}$ is further introduced to capture the nuanced semantic relationships between classes. Thereafter, the decision rule is augmented as follows:

$$\begin{aligned} & \underset{k}{\operatorname{argmin}} R(\mathbf{y}_i^h | \mathbf{x}_i) \\ &= \underset{k}{\operatorname{argmin}} \sum_{j=1}^{C^h} (D_{k,j}^h - \beta \cdot \Delta_{k,j}^h) \cdot p^h(\mathbf{y}_i^h = k | \mathbf{x}_i) \\ &= \underset{k}{\operatorname{argmax}} \sum_{j=1}^{C^h} (\beta \cdot \Delta_{k,j}^h - D_{k,j}^h) \cdot p^h(\mathbf{y}_i^h = k | \mathbf{x}_i) \\ &= \underset{k}{\operatorname{argmax}} \sum_{j=1}^{C^h} \tilde{D}_{k,j}^h \cdot p^h(\mathbf{y}_i^h = k | \mathbf{x}_i), \end{aligned}$$

where $\tilde{\mathbf{D}}^h = \beta \cdot \Delta^h - \mathbf{D}^h$ and β signifies the weight of the learnable matrix Δ^h . Meanwhile, to avoid self-cost, the diagonal entries of Δ^h (i.e., Δ_{ii}^h) are masked as zeros and then each row of Δ^h undergoes an L2 normalization.

To effectively learn intra-granularity differences, we thus introduce a label smoothing strategy (Müller, Kornblith, and Hinton 2019) for training. Different from previous studies (Bertinetto et al. 2020), as the goal is to pursue a model with better mistake, we propose to employ the above class-wise matrix to label smoothing explicitly. Specifically, the label smoothing technique for AIGDL could be expressed as:

$$\tilde{\mathbf{y}}_i^h = (1 - \epsilon) \mathbf{y}_i^h + \epsilon \cdot \frac{\exp(\gamma(\beta \cdot \Delta_{y_i^h}^h - \mathbf{D}_{y_i^h}^h))}{\sum_{j=1}^{C^h} \exp(\gamma(\beta \cdot \Delta_{y_i^h, j}^h - \mathbf{D}_{y_i^h, j}^h))},$$

where ϵ is the smoothing factor, γ is the temperature parameter, and \mathbf{D}_k^h and Δ_k^h are the k -th row of \mathbf{D}^h and Δ^h respectively. In this way, the semantic similarities among these

subclasses can be implicitly reflected at the label level, ensuring consistency in predictions. Moreover, it is to be noted that the above smoothing method is a general version of the Soft-Label method (Bertinetto et al. 2020) by setting $\epsilon = 1$ and $\beta = 0$.

Similarly, the loss function of AIGDL for each level is defined as:

$$\mathcal{L}_{\text{AIGDL}}^h = \frac{1}{N} \sum_{i=1}^N \tilde{\mathbf{y}}_i^h \log(\text{softmax}(\tilde{\mathbf{D}}^h \mathbf{p}_i^h)).$$

Overall Optimization Objective

Combining all the above components, the overall optimization objective is defined as:

$$\begin{aligned} \mathcal{L} &= \sum_h \lambda_h \cdot (\mathcal{L}_{\text{BiLT}}^h + \mathcal{L}_{\text{AIGDL}}^h) \\ &= \sum_h \frac{\lambda_h}{N} \sum_{i=1}^N (\mathbf{y}_i^h \log(\mathbf{p}_i^h) + \tilde{\mathbf{y}}_i^h \log(\text{softmax}(\tilde{\mathbf{D}}^h \mathbf{p}_i^h))), \end{aligned}$$

where $\lambda_h = \exp(\alpha \cdot (h - H))$ is the weight for each level. As the hierarchical level becomes coarser, its contribution to fine-grained feature learning diminishes. So the corresponding loss weight λ_h decreases as the hierarchical level becomes coarser.

Experiments

Experiments Setup

Datasets. We evaluate methods on four datasets: FGVC-Aircraft (Maji et al. 2013), CIFAR-100 (Krizhevsky, Hinton et al. 2009), iNaturalist2019 (Van Horn et al. 2018) and tieredImageNet-H(Ren et al. 2018), all of which have been used in previous studies(Liang and Davis 2023; Garg, Sani, and Anand 2022; Bertinetto et al. 2020). Dataset statistics and split settings are provided in the Appendix.

Competitors. We compare our method against the following competitors: **HXE**(Bertinetto et al. 2020), **Soft-Labels**(Bertinetto et al. 2020), **Flamingo**(Chang et al. 2021), **CRM**(Karthik et al. 2021), **HAF**(Garg, Sani, and Anand 2022), **HAFrime**(Liang and Davis 2023), and **HiE**(Jain, Karthik, and Gandhi 2024). We also include the vanilla cross-entropy loss as a baseline, denoted as **Cross-Entropy**. Note that, all competitors are implemented following the guidelines suggested by the corresponding paper.

Performance Comparisons

Overall Performance. Tab. 1 and Tab. 2 present the detailed performance of the FGVC-Aircraft and CIFAR-100 datasets, respectively. Detailed performance results for the iNaturalist2019 dataset are provided in the Appendix. Notably, all metrics are correlated to log-scaled LCA distance except for Top-1 Accuracy (Karthik et al. 2021). According to the reported results, we can draw the following remarks: In most cases, our proposed BiLT could outperform all competitors at all metrics, except for the Top-1 Accuracy and Hier Dist@1 on the CIFAR-100 and iNaturalist2019 datasets. Even in failure cases, the performance of BiLT is still comparable. This speaks to the efficacy of our approach.

Method	Mistake Severity(\downarrow)	Hier Dist@1(\downarrow)	Hier Dist@5(\downarrow)	Hier Dist@20(\downarrow)	Top-1 Accuracy(\uparrow)
Cross-Entropy	2.12 +/- 0.0288	0.44 +/- 0.0188	2.10 +/- 0.0053	2.67 +/- 0.0022	79.35 +/- 0.7021
HXE	2.04 +/- 0.0074	0.43 +/- 0.0283	1.96 +/- 0.0119	2.60 +/- 0.0085	78.75 +/- 0.9481
Soft-Labels	2.10 +/- 0.0124	0.49 +/- 0.0223	2.07 +/- 0.0149	2.65 +/- 0.0058	77.59 +/- 0.9698
Flamingo	2.10 +/- 0.0352	0.40 +/- 0.0116	2.06 +/- 0.0099	2.65 +/- 0.0043	<u>80.72 +/- 0.5849</u>
CRM	2.08 +/- 0.0366	0.42 +/- 0.0163	<u>1.74 +/- 0.0053</u>	<u>2.44 +/- 0.0018</u>	79.57 +/- 0.5880
HAF	2.53 +/- 0.0610	0.67 +/- 0.0295	2.10 +/- 0.0063	2.61 +/- 0.0028	73.68 +/- 1.2166
HAFrame	<u>2.01 +/- 0.0103</u>	<u>0.39 +/- 0.0162</u>	1.75 +/- 0.0060	2.45 +/- 0.0015	80.52 +/- 0.8375
HiE	2.06 +/- 0.0279	0.44 +/- 0.0146	1.82 +/- 0.0028	2.46 +/- 0.0013	78.66 +/- 0.9660
Ours	2.00 +/- 0.0220	0.38 +/- 0.0097	1.72 +/- 0.0032	2.44 +/- 0.0016	81.23 +/- 0.5820

Table 1: Performance comparisons on the FGVC-Aircraft dataset with different metrics. The first and second best results are highlighted with **bold** text and underline, respectively.

Method	Mistake Severity(\downarrow)	Hier Dist@1(\downarrow)	Hier Dist@5(\downarrow)	Hier Dist@20(\downarrow)	Top-1 Accuracy(\uparrow)
Cross-Entropy	2.35 +/- 0.0225	0.52 +/- 0.0076	2.25 +/- 0.0146	3.18 +/- 0.0079	77.85 +/- 0.2699
HXE	2.40 +/- 0.0137	0.62 +/- 0.0205	2.05 +/- 0.0082	3.02 +/- 0.0171	72.76 +/- 0.6816
Soft-Labels	2.33 +/- 0.0270	0.62 +/- 0.0144	1.29 +/- 0.0056	2.73 +/- 0.0122	74.34 +/- 0.4588
Flamingo	2.32 +/- 0.0186	0.51 +/- 0.0094	2.07 +/- 0.0214	3.08 +/- 0.0077	77.83 +/- 0.2942
CRM	2.31 +/- 0.0229	0.51 +/- 0.0019	1.15 +/- 0.0037	2.20 +/- 0.0034	77.80 +/- 0.2462
HAF	2.24 +/- 0.0177	0.50 +/- 0.0075	1.42 +/- 0.0057	2.64 +/- 0.0037	77.51 +/- 0.3091
HAFrame	2.21 +/- 0.0108	0.49 +/- 0.0066	<u>1.11 +/- 0.0018</u>	<u>2.18 +/- 0.0013</u>	77.71 +/- 0.2319
HiE	<u>2.20 +/- 0.0232</u>	0.47 +/- 0.0056	1.19 +/- 0.0031	2.86 +/- 0.0097	78.63 +/- 0.3101
Ours	2.17 +/- 0.0197	0.48 +/- 0.0106	1.08 +/- 0.0057	2.17 +/- 0.0030	77.77 +/- 0.4140

Table 2: Performance comparisons on the CIFAR-100 dataset with different metrics. The first and second best results are highlighted with **bold** text and underline, respectively.

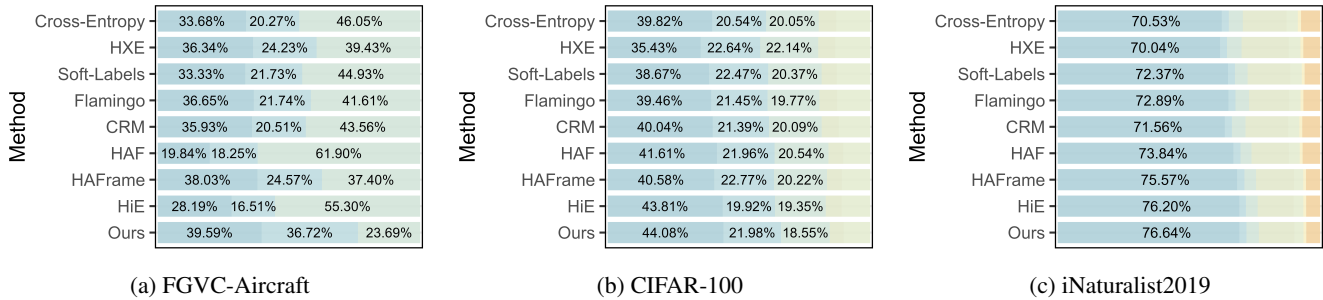


Figure 5: The probability of mistakes made by our method and competitors across the three datasets. In each subfigure, different colors represent varying LCA distances, increasing from left to right. The length of each bar indicates the severity of the mistake, while the values on the bars show the probability of the model making a mistake at a specific LCA distance.

Mistake Severity Distribution. Fig. 5 shows the mistake severity distribution for our method and the competitors. Specifically, here we report the performance comparisons at the finest granularity by examining the LCA distances between wrong and ground truth classes. It is obvious that, when misclassifying an example, our proposed BiLT could make a more reasonable prediction. As shown in Fig. 5a, when the model makes a mistake, our method has a 39.59% probability that the LCA distance between the prediction and the ground truth is 1. Additionally, there is a 76.31% probability (39.59% + 36.72%) that the LCA distance is less than or equal to 2, resulting in a significantly lower proportion of severe mistakes compared to those of competitors.

Intra-Granularity Difference Visualization. Fig. 6 shows the heat map of the learnable intra-granularity difference matrix Δ^H , illustrating relationships between the finest-level classes on CIFAR-100. Heat map of Δ^H for FGVC-Aircraft are provided in the Appendix. In a specific coarse class (zoomed portion **A**), the colors of the 3rd and the 4th classes are greener, indicating a weaker correlation compared to other fine-grained classes. At the coarse-grained level, square **(B)** is more yellow than square **(A)**, indicating finer-grained classes in **(B)** share closer relations with each other. This illustrates that the label tree fails to accurately describe the difference between classes, necessitating the use of Δ^H to represent their relationships.

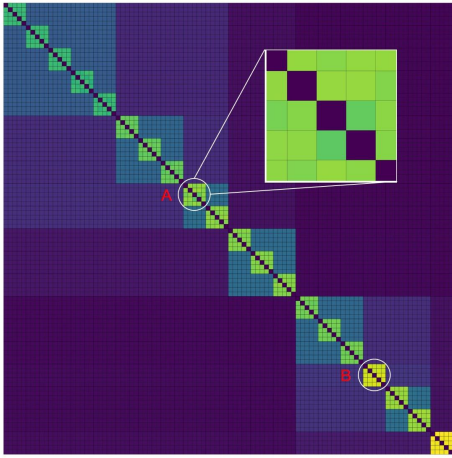


Figure 6: Heat maps of Δ^H for CIFAR-100, with colors transitioning from purple to yellow as values increase. This heat map is composed of 100×100 small squares, which exactly indicates 100 classes (5th level) within CIFAR-100. 20 green or yellow squares, each composed of 5 smaller squares, represent coarser labels for 20 classes (4th level), and larger squares represent coarser classes.

Ablation Study

Sensitivity analysis of α . Fig.7 presents the sensitivity analysis of α on the FGVC-Aircraft dataset, while the analysis for the CIFAR-100 dataset is provided in the Appendix. The weight assigned to level h in the overall optimization objective is given by $\lambda_h = \exp(\alpha \cdot (h - H))$. A smaller α increases the weights of the coarser levels. When α is set to 0, all levels receive equal weight, resulting in a uniform optimization objective. Within the α range of 0.5 to 1.0, the model reduces Mistake Severity while maintaining Accuracy. However, a large α significantly decreases the weight of the coarse level, leading to a notable increase in Mistake Severity. This phenomenon illustrates the importance of coarse-grained classification for fine-grained learning.

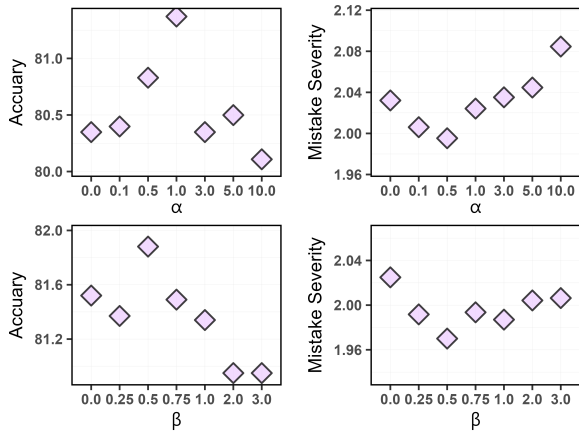


Figure 7: Sensitivity analysis about α and β on FGVC-Aircraft.

Sensitivity analysis of β . Fig.7 shows the sensitivity analysis of β on the FGVC-Aircraft dataset. For the sensitivity analysis of β on the CIFAR-100 dataset please see Appendix. β is the weight for the learnable adjustment matrix Δ^h . Too large or too small β leads to decreased performance. When $\beta = 0$, the learnable adjustment matrix Δ^h has no effect. Conversely, too large β may cause hierarchical level spanning, where some levels deviate significantly from the originally defined distance matrix after adjustment. The model achieved the best performance at $\beta = 0.5$.

Sensitivity analysis of ϵ and γ . Fig.8 shows the sensitivity analysis of ϵ and γ on the FGVC-Aircraft and CIFAR-100 datasets. γ is the temperature coefficient in the soft label generation process, and ϵ is the proportion of the soft label. When $\epsilon = 1$, it means only soft label is used. ϵ and γ regulate the label encoding together. In the FGVC-Aircraft dataset, Mistake Severity is low when ϵ is 0.3 and γ is around 0.7. In the CIFAR-100 dataset, Mistake Severity is low when ϵ and γ are both around 0.3. It can be observed that using only hard label or only soft label is always not the optimal solution.

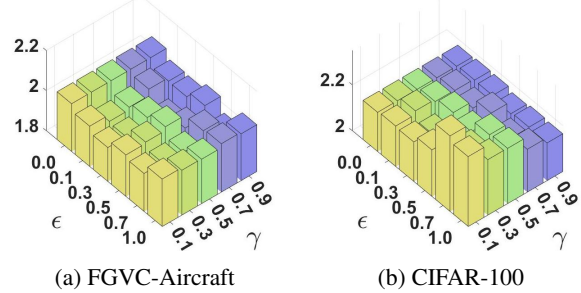


Figure 8: Sensitivity analysis about ϵ and γ on FGVC-Aircraft and CIFAR-100.

Conclusion

In this paper, we focus on the granularity competition problem between different granularities in fine-grained visual classification (FGVC) tasks. However, coarse-grained features are naturally easy to learn, leading feature-shared-based methods focus only on coarse features instead of fine-grained ones. To address this issue, we propose a novel method, the Bidirectional Logits Tree (BiLT), which organizes classifiers from the finest to the coarsest levels, rather than a parallel framework where all classifiers share the same feature vector. Additionally, we observe that predefined label trees cannot accurately capture semantic differences between labels at the same granularity, and therefore propose Adaptive Intra-Granularity Difference Learning (AIGDL). This method empowers BiLT to learn the fine-grained semantic differences of classes within the same granularity. Finally, extensive experiments and visualizations justify the effectiveness of our method.

Acknowledgments

This work was supported in part by the National Key R&D Program of China under Grant 2018AAA0102000, in part by National Natural Science Foundation of China: 62236008, U21B2038, U23B2051, 61931008, 62122075, 62206264, and 92370102, in part by Youth Innovation Promotion Association CAS, in part by the Strategic Priority Research Program of the Chinese Academy of Sciences, Grant No. XDB0680000, in part by the Innovation Funding of ICT, CAS under Grant No.E0000000, and in part by the Postdoctoral Fellowship Program of CPSF under Grant GZB20240729.

References

- Bertinetto, L.; Mueller, R.; Tertikas, K.; Samangoei, S.; and Lord, N. A. 2020. Making better mistakes: Leveraging class hierarchies with deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12506–12515.
- Bukchin, G.; Schwartz, E.; Saenko, K.; Shahar, O.; Feris, R.; Giryas, R.; and Karlinsky, L. 2021. Fine-grained angular contrastive learning with coarse labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8730–8740.
- Chang, D.; Pang, K.; Zheng, Y.; Ma, Z.; Song, Y.-Z.; and Guo, J. 2021. Your” flamingo” is my” bird”: Fine-grained, or not. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11476–11485.
- Chen, J.; Wang, P.; Liu, J.; and Qian, Y. 2022. Label relation graphs enhanced hierarchical residual network for hierarchical multi-granularity classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4858–4867.
- Collins, K. M.; Bhatt, U.; and Weller, A. 2022. Eliciting and learning with soft labels from every annotator. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, 40–52.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*, 248–255.
- Du, R.; Xie, J.; Ma, Z.; Chang, D.; Song, Y.-Z.; and Guo, J. 2021. Progressive learning of category-consistent multi-granularity features for fine-grained visual classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12): 9521–9535.
- Duda, R. O.; and Hart, P. E. 1974. Pattern classification and scene analysis. In *A Wiley-Interscience publication*.
- Frome, A.; Corrado, G. S.; Shlens, J.; Bengio, S.; Dean, J.; Ranzato, M.; and Mikolov, T. 2013. Devise: A deep visual-semantic embedding model. In *Advances in Neural Information Processing Systems*.
- Garg, A.; Sani, D.; and Anand, S. 2022. Learning hierarchy aware features for reducing mistake severity. In *European Conference on Computer Vision*, 252–267.
- Gong, X.; Bisht, N.; and Xu, G. 2024. Does Label Smoothing Help Deep Partial Label Learning? In *International Conference on Machine Learning*.
- Grcic, M.; Gadetsky, A.; and Brbic, M. 2024. Fine-grained Classes and How to Find Them. In *International Conference on Machine Learning*, 16275–16294.
- Han, B.; Xu, Q.; Yang, Z.; Bao, S.; Wen, P.; Jiang, Y.; and Huang, Q. 2024. AUCSeg: AUC-oriented Pixel-level Long-tail Semantic Segmentation. In *Advances in Neural Information Processing Systems*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- Jain, K.; Karthik, S.; and Gandhi, V. 2024. Test-time amendment with a coarse classifier for fine-grained classification. In *Advances in Neural Information Processing Systems*.
- Jiang, X.; Tang, H.; Gao, J.; Du, X.; He, S.; and Li, Z. 2024. Delving into Multimodal Prompting for Fine-Grained Visual Classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2570–2578.
- Karthik, S.; Prabhu, A.; Dokania, P. K.; and Gandhi, V. 2021. No Cost Likelihood Manipulation at Test Time for Making Better Mistakes in Deep Networks. In *International Conference on Learning Representations*.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations*.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Landrieu, L.; and Garnot, V. S. F. 2021. Leveraging class hierarchies with metric-guided prototype learning. In *British Machine Vision Conference*.
- Liang, T.; and Davis, J. 2023. Inducing neural collapse to a fixed hierarchy-aware frame for reducing mistake severity. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1443–1452.
- Liu, M.; Roy, S.; Li, W.; Zhong, Z.; Sebe, N.; and Ricci, E. 2024. Democratizing Fine-grained Visual Recognition with Large Language Models. In *International Conference on Learning Representations*.
- Liu, S.; Chen, J.; Pan, L.; Ngo, C.-W.; Chua, T.-S.; and Jiang, Y.-G. 2020. Hyperbolic visual embedding learning for zero-shot recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9273–9281.
- Maji, S.; Rahtu, E.; Kannala, J.; Blaschko, M.; and Vedaldi, A. 2013. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*.
- Müller, R.; Kornblith, S.; and Hinton, G. E. 2019. When does label smoothing help? In *Advances in Neural Information Processing Systems*.
- Ni, J.; Cheng, W.; Chen, Z.; Asakura, T.; Soma, T.; Kato, S.; and Chen, H. 2021. Superclass-conditional gaussian mixture model for learning fine-grained embeddings. In *International Conference on Learning Representations*.

- Papayan, V.; Han, X.; and Donoho, D. L. 2020. Prevalence of neural collapse during the terminal phase of deep learning training. *Proceedings of the National Academy of Sciences*, 117(40): 24652–24663.
- Park, H.; Noh, J.; Oh, Y.; Baek, D.; and Ham, B. 2023. Acls: Adaptive and conditional label smoothing for network calibration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3936–3945.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*.
- Pu, Y.; Han, Y.; Wang, Y.; Feng, J.; Deng, C.; and Huang, G. 2024. Fine-grained recognition with learnable semantic data augmentation. *IEEE Transactions on Image Processing*, 33: 3130–3144.
- Qin, Y.; Wang, X.; Beutel, A.; and Chi, E. 2021. Improving calibration through the relationship with adversarial robustness. In *Advances in Neural Information Processing Systems*, 14358–14369.
- Ren, M.; Triantafillou, E.; Ravi, S.; Snell, J.; Swersky, K.; Tenenbaum, J. B.; Larochelle, H.; and Zemel, R. S. 2018. Meta-learning for semi-supervised few-shot classification. In *International Conference on Learning Representations*.
- Silla, C. N.; and Freitas, A. A. 2011. A survey of hierarchical classification across different application domains. *Data Mining and Knowledge Discovery*, 22: 31–72.
- Touvron, H.; Sablayrolles, A.; Douze, M.; Cord, M.; and Jégou, H. 2021. Graft: Learning fine-grained image representations with coarse labels. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 874–884.
- Van Horn, G.; Mac Aodha, O.; Song, Y.; Cui, Y.; Sun, C.; Shepard, A.; Adam, H.; Perona, P.; and Belongie, S. 2018. The inaturalist species classification and detection dataset. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8769–8778.
- Wang, D.; Shen, Z.; Shao, J.; Zhang, W.; Xue, X.; and Zhang, Z. 2015. Multiple granularity descriptors for fine-grained categorization. In *Proceedings of the IEEE International Conference on Computer Vision*, 2399–2406.
- Wang, R.; Zou, C.; Zhang, W.; Zhu, Z.; and Jing, L. 2023. Consistency-aware Feature Learning for Hierarchical Fine-grained Visual Classification. In *Proceedings of the 31st ACM International Conference on Multimedia*, 2326–2334.
- Wang, W.; Sun, Y.; Li, W.; and Yang, Y. 2024. Transhp: Image classification with hierarchical prompting. In *Advances in Neural Information Processing Systems*.
- Wang, Y.; Wang, Z.; Hu, Q.; Zhou, Y.; and Su, H. 2021. Hierarchical semantic risk minimization for large-scale classification. *IEEE Transactions on Cybernetics*, 52(9): 9546–9558.
- Wei, J.; Liu, H.; Liu, T.; Niu, G.; Sugiyama, M.; and Liu, Y. 2022. To Smooth or Not? When Label Smoothing Meets Noisy Labels. In *International Conference on Machine Learning*, 23589–23614.
- Wu, H.; Merler, M.; Uceda-Sosa, R.; and Smith, J. R. 2016. Learning to make better mistakes: Semantics-aware visual food recognition. In *Proceedings of the 24th ACM International Conference on Multimedia*, 172–176.
- Xian, Y.; Akata, Z.; Sharma, G.; Nguyen, Q.; Hein, M.; and Schiele, B. 2016. Latent embeddings for zero-shot classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 69–77.
- Xu, S.; Sun, Y.; Zhang, F.; Xu, A.; Wei, X.; and Yang, Y. 2023. Hyperbolic Space with Hierarchical Margin Boosts Fine-Grained Learning from Coarse Labels. In *Advances in Neural Information Processing Systems*.
- Yuan, L.; Tay, F. E.; Li, G.; Wang, T.; and Feng, J. 2020. Revisiting knowledge distillation via label smoothing regularization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3903–3911.
- Zagoruyko, S.; and Komodakis, N. 2016. Wide Residual Networks. In *British Machine Vision Conference*.
- Zhang, C.-B.; Jiang, P.-T.; Hou, Q.; Wei, Y.; Han, Q.; Li, Z.; and Cheng, M.-M. 2021. Delving deep into label smoothing. *IEEE Transactions on Image Processing*, 30: 5984–5996.
- Zhang, J.; Chen, Z.; Zhang, H.; Xiao, C.; and Li, B. 2023. {DiffSmooth}: Certifiably robust learning via diffusion models and local smoothing. In *32nd USENIX Security Symposium (USENIX Security 23)*, 4787–4804.
- Zhang, M.-L.; and Zhou, Z.-H. 2013. A review on multi-label learning algorithms. *IEEE Transactions on Knowledge and Data Engineering*, 26(8): 1819–1837.
- Zhang, S.; Zheng, S.; Shui, Z.; and Yang, L. 2024. HLS-FGVC: Hierarchical Label Semantics Enhanced Fine-Grained Visual Classification. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, 7370–7374.
- Zhao, B.; Li, F.; and Xing, E. 2011. Large-scale category structure aware image categorization. In *Advances in Neural Information Processing Systems*.