

# Distilling Balanced Knowledge from a Biased Teacher

Seonghak Kim

Agency for Defense Development (ADD), Republic of Korea

seonghak35@gmail.com

## Abstract

Conventional knowledge distillation, designed for model compression, fails on long-tailed distributions because the teacher model tends to be biased toward head classes and provides limited supervision for tail classes. We propose Long-Tailed Knowledge Distillation (LTKD), a novel framework that reformulates the conventional objective into two components: a cross-group loss, capturing mismatches in prediction distributions across class groups (head, medium, and tail), and a within-group loss, capturing discrepancies within each group’s distribution. This decomposition reveals the specific sources of the teacher’s bias. To mitigate the inherited bias, LTKD introduces (1) a rebalanced cross-group loss that calibrates the teacher’s group-level predictions and (2) a reweighted within-group loss that ensures equal contribution from all groups. Extensive experiments on CIFAR-100-LT, TinyImageNet-LT, and ImageNet-LT demonstrate that LTKD significantly outperforms existing methods in both overall and tail-class accuracy, thereby showing its ability to distill balanced knowledge from a biased teacher for real-world applications.

## 1. Introduction

Knowledge Distillation (KD) is a well-established model compression technique designed to transfer knowledge from a large and powerful teacher model to a compact student model [16]. Its effectiveness has led to successful applications across diverse domains, including computer vision [25, 37, 40, 44], NLP [13, 19, 34, 43], and LLM [1, 11, 39, 41]. This approach aims to retain the high accuracy of computationally expensive and large-scale models within a lightweight architecture, enabling deployment in resource-constrained environments [2, 5, 54].

Conventional KD methods have evolved along two primary branches: logit-based [16, 17, 28, 49, 52] and feature-based [6, 12, 18, 29, 30, 35, 45, 51] approaches. Among them, logit-based methods are particularly prevalent, typically formulated by minimizing the Kullback–Leibler (KL) divergence between the teacher’s and student’s softened

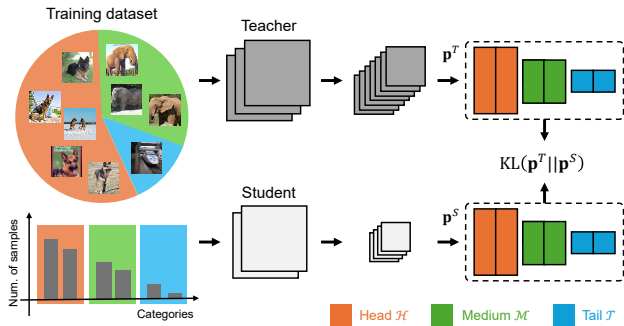


Figure 1. Overview of standard KD on long-tailed distributions. The training data is highly imbalanced (pie chart and bar graph), with most samples belonging to head classes (orange) and few to tail classes (blue). This creates a biased teacher whose predictions ( $\mathbf{p}^T$ ), visually represented by varying bar heights, are skewed toward head classes. Standard KD forces the student to mimic this bias ( $\mathbf{p}^S$ ), resulting in poor generalization on tail classes.

predictions [16]. This framework operates under a critical assumption: the training data is balanced [22, 23], allowing the teacher to offer reliable guidance across all classes. However, *what happens when this assumption breaks down?* More importantly, *can a teacher trained on imbalanced data still offer trustworthy supervision?*

As shown in Fig. 1, real-world datasets often follow a long-tailed distribution [3, 36, 42, 50]. When trained on such data, the teacher model becomes biased toward head classes, performing well on frequent classes but poorly on rare ones due to insufficient exposure. Consequently, applying standard KD is not only ineffective but can be detrimental. The student inherits the teacher’s bias, overfitting to head-class predictions while receiving little meaningful guidance on tail-class examples. This ultimately results in poor generalization. This highlights a critical need for new KD frameworks designed to distill balanced knowledge even from a biased teacher.

To address this, we propose *Long-Tailed Knowledge Distillation (LTKD)*, a novel knowledge distillation approach tailored for class-imbalanced scenarios, as illustrated in Fig. 2. Our approach reformulates the conventional KL-

based KD objective into two components: a cross-group loss and a within-group loss. This decomposition enables a theoretical analysis of how knowledge is transferred under class imbalance. By defining appropriate class groups (*e.g.* head, medium, and tail), we can elucidate the distinct contributions of these components: the cross-group term reflects mismatches in aggregate probability across groups, while the within-group term captures discrepancies within each group. Our analysis reveals that both terms are distorted by the teacher’s bias in distinct ways. The cross-group term leads to an overestimation of head-class probabilities and an underestimation of tail-class ones. Meanwhile, because the within-group term is weighted by each group’s aggregate probability, it tends to disproportionately favor the head group while neglecting the tail group.

Building on this insight, we introduce two core components to counteract the teacher’s bias: (1) a rebalanced cross-group loss that calibrates the skewed group-level predictions, and (2) a reweighted within-group loss that ensures equal learning focus across all groups. Conceptually, these adjustments offset the distortions caused by class imbalance, enabling the student to learn a more balanced representation.

We comprehensively validate the effectiveness of LTKD across a wide range of model architectures on standard long-tailed benchmarks. Experimental results demonstrate substantial improvements in both overall and, crucially, tail-class accuracy, consistently achieving state-of-the-art performance. Remarkably, in nearly all cases, our method surpasses the teacher’s own performance. These results underscore the ability of LTKD to distill balanced knowledge from a biased teacher, marking a significant step toward deploying robust and high-performance models in real-world and imbalanced scenarios.

In summary, our contributions are threefold:

- We reformulate the KL-based objective into cross-group and within-group components, enabling analysis of teacher bias under long-tailed distributions.
- We propose rebalancing and reweighting strategies that equalize the influence of all class groups during distillation, mitigating biased supervision.
- We achieve state-of-the-art performance on long-tailed benchmarks, improving both overall and tail-class accuracy—even surpassing the teacher in most cases.

## 2. Related work

### 2.1. Knowledge distillation on balanced datasets

Knowledge Distillation (KD) transfers rich information—often referred to as dark knowledge [21]—from a large teacher model to a compact student, aiming to preserve performance while reducing computational cost [9]. KD methods are broadly categorized into feature-based [6, 12, 18,

29, 30, 35, 45, 51] and logit-based [16, 17, 28, 49, 52] approaches. Feature-based distillation focuses on transferring internal knowledge by aligning hidden representations [30]. It has been extended to leverage spatial attention maps [45], sample-wise relational information [29], and contrastive objectives [35]. ReviewKD [6] introduced a residual-style fusion of earlier feature layers to guide deeper ones, while CAT-KD [12] transfers class activation maps (CAMs) to help the student focus on class-discriminative regions. Although effective, these methods typically require complex feature transformations and often incur higher computational overhead. In contrast, logit-based methods, first introduced by Hinton [16], are more cost-effective as they aim to match the softened logits using KL divergence. DKD [52] decouples the KL divergence loss into target and non-target components to better reflect their distinct contributions, while DIST [17] further relaxes the KL objective through correlation-based loss, preserving inter-class and intra-class prediction relations. Despite their effectiveness, a key limitation of existing methods is the assumption of a balanced class distribution [10, 12, 18, 51], which is rarely met in real-world datasets. A teacher trained on imbalanced data tends to be biased toward head classes, and this bias is directly inherited by the student. As a result, the student may overfit to frequent classes while receiving inadequate supervision for tail classes, leading to degraded generalization and reduced robustness.

### 2.2. Knowledge distillation on LT datasets

A specific line of research has explored KD under long-tailed distributions [4, 15, 20, 24, 38, 47, 53], where the imbalance in training data leads to biased teacher predictions. One prominent approach is to modify the distillation objective to account for class imbalance. BKD [47] proposes a dual-loss framework that combines an instance-balanced classification loss with a class-balanced distillation loss to enhance performance on underrepresented classes. Another direction involves multi-teacher distillation. LFME [38] introduces a self-paced framework that aggregates knowledge from multiple expert models, each trained on less imbalanced subsets of the data. By adaptively selecting both experts and samples, LFME enhances generalization in long-tailed scenarios. However, these methods generally assume that the student mirrors the teacher in architecture, placing emphasis primarily on classification accuracy while largely overlooking the compression aspect.

In contrast, our work is, to the best of our knowledge, the first to address knowledge distillation under long-tailed distributions from the perspective of model compression. We begin by analyzing the role of KL divergence in long-tailed scenarios and demonstrate that the performance degradation of existing logit-based KD methods stems from the direct transfer of teacher bias to the student.

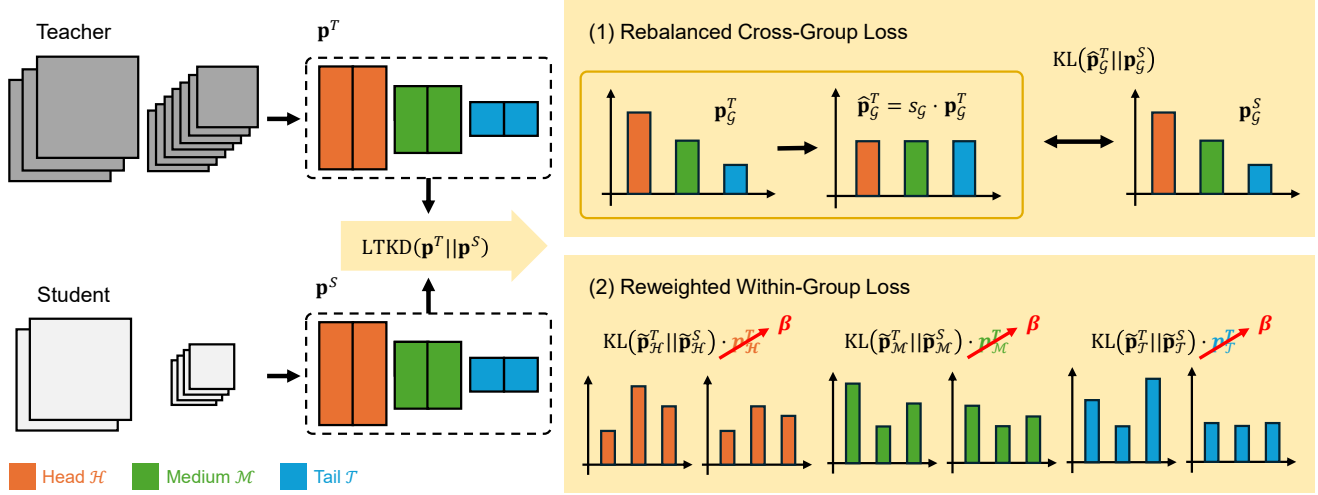


Figure 2. Overview of the proposed Long-Tailed Knowledge Distillation (LTKD). Our method first decomposes the standard KL-based KD loss into a cross-group component (capturing mismatches in aggregated group-level predictions) and a within-group component (capturing internal discrepancies). To correct the teacher’s class bias, LTKD then applies a rebalanced cross-group loss and reweighted within-group loss, ensuring a balanced knowledge transfer to the student.

### 3. Method

In this section, we begin by revisiting the KL divergence to diagnose the sources of teacher bias under long-tailed distributions. Based on this analysis, we propose *Long-Tailed Knowledge Distillation* (LTKD), a framework that mitigates bias by rebalancing cross-group predictions and equalizing within-group contributions during distillation.

#### 3.1. Revisiting KL divergence

In a classification task where a sample  $\mathbf{x}$  is categorized into one of  $C$  classes, the predictive probability vector is represented as  $\mathbf{p} = [p_1, p_2, \dots, p_C] \in \mathbb{R}^C$ , which is obtained by applying the softmax function  $\sigma(\cdot)$  to the logit vector  $\mathbf{z} = [z_1, z_2, \dots, z_C] \in \mathbb{R}^C$  as follows:

$$p_i = \sigma(z_i) = \frac{\exp(z_i)}{\sum_{j=1}^C \exp(z_j)}, \quad (1)$$

where  $p_i$  and  $z_i$  denote the probability and logit value corresponding to the  $i$ -th class, respectively.

Given our focus on long-tailed distributions, we partition the  $C$  classes into three mutually disjoint groups,  $\mathcal{G} \in \{\mathcal{H}, \mathcal{M}, \mathcal{T}\}$ . These correspond to a head group ( $\mathcal{H}$ ) with many samples, a medium group ( $\mathcal{M}$ ), and a tail group ( $\mathcal{T}$ ) with few samples. Accordingly, the KD loss function, defined using KL divergence, can be formulated as follows:

$$\begin{aligned} \text{KD} &= \text{KL}(\mathbf{p}^T \parallel \mathbf{p}^S) \\ &= \sum_{\mathcal{G}} \sum_{i \in \mathcal{G}} p_i^T \log \left( \frac{p_i^T}{p_i^S} \right), \end{aligned} \quad (2)$$

where the teacher and student models are denoted by  $T$  and  $S$ , respectively.

To proceed with the formulation, we introduce two new notations. The first is the cross-group probability distribution,  $\mathbf{p}_{\mathcal{G}} = [p_{\mathcal{H}}, p_{\mathcal{M}}, p_{\mathcal{T}}] \in \mathbb{R}^3$ , where each element represents the aggregated probability over the head, medium, and tail class groups, respectively, and is defined as:

$$p_{\mathcal{G}} = \frac{\sum_{i \in \mathcal{G}} \exp(z_i)}{\sum_{j=1}^C \exp(z_j)}. \quad (3)$$

The second is the within-group probability distribution,  $\tilde{\mathbf{p}}_{\mathcal{G}} = [\tilde{p}_{\mathcal{G}_1}, \tilde{p}_{\mathcal{G}_2}, \dots, \tilde{p}_{\mathcal{G}_i}]_{i \in \mathcal{G}} \in \mathbb{R}^{|\mathcal{G}|}$ , which defines the probability distribution within each class group and is calculated as:

$$\tilde{p}_{\mathcal{G}_i} = \frac{\exp(z_i)}{\sum_{j \in \mathcal{G}} \exp(z_j)}. \quad (4)$$

Using the relationship between  $p_i$ ,  $p_{\mathcal{G}}$ , and  $\tilde{p}_{\mathcal{G}_i}$ , we can rewrite  $p_i$  (for  $i \in \mathcal{G}$ ) as a product of cross-group and within-group probabilities:  $p_i = p_{\mathcal{G}} \cdot \tilde{p}_{\mathcal{G}_i}$ . Based on this, Eq. (2) can be reformulated as:

$$\text{KD} = \sum_{\mathcal{G}} \sum_{i \in \mathcal{G}} p_i^T \log \left( \frac{p_{\mathcal{G}}^T}{p_{\mathcal{G}}^S} \right) + \sum_{\mathcal{G}} \sum_{i \in \mathcal{G}} p_i^T \log \left( \frac{\tilde{p}_{\mathcal{G}_i}^T}{\tilde{p}_{\mathcal{G}_i}^S} \right). \quad (5)$$

Since both  $p_{\mathcal{G}}^T$  and  $p_{\mathcal{G}}^S$  are independent of the class index  $i$ , and  $\sum_{i \in \mathcal{G}} p_i^T = p_{\mathcal{G}}^T$ , the first term in Eq. (5) can be simplified as follows:

$$\begin{aligned} \sum_{\mathcal{G}} \sum_{i \in \mathcal{G}} p_i^T \log \left( \frac{p_{\mathcal{G}}^T}{p_{\mathcal{G}}^S} \right) &= \sum_{\mathcal{G}} p_{\mathcal{G}}^T \log \left( \frac{p_{\mathcal{G}}^T}{p_{\mathcal{G}}^S} \right) \\ &= \text{KL}(\mathbf{p}_{\mathcal{G}}^T \parallel \mathbf{p}_{\mathcal{G}}^S), \end{aligned} \quad (6)$$

where  $\mathbf{p}_G^T = [p_{\mathcal{H}}^T, p_{\mathcal{M}}^T, p_{\mathcal{T}}^T]$  and  $\mathbf{p}_G^S = [p_{\mathcal{H}}^S, p_{\mathcal{M}}^S, p_{\mathcal{T}}^S]$  denote the cross-group probability distributions of the teacher and student, respectively. Analogously, the second term in Eq. (5) can also be simplified as follows:

$$\begin{aligned} \sum_G \sum_{i \in G} p_i^T \log \left( \frac{\tilde{p}_{G_i}^T}{\tilde{p}_{G_i}^S} \right) &= \sum_G p_G^T \sum_{i \in G} \tilde{p}_{G_i}^T \log \left( \frac{\tilde{p}_{G_i}^T}{\tilde{p}_{G_i}^S} \right) \\ &= \sum_G p_G^T \cdot \text{KL}(\tilde{\mathbf{p}}_G^T \| \tilde{\mathbf{p}}_G^S), \end{aligned} \quad (7)$$

where  $\tilde{\mathbf{p}}_G^T$  and  $\tilde{\mathbf{p}}_G^S$  denote the within-group probability distributions of the teacher and student for group  $G$ , respectively.

Finally, Eq. (2) can be expressed as:

$$\begin{aligned} \text{KD} &= \text{KL}(\mathbf{p}_G^T \| \mathbf{p}_G^S) + \sum_G p_G^T \cdot \text{KL}(\tilde{\mathbf{p}}_G^T \| \tilde{\mathbf{p}}_G^S) \\ &= \text{KL}(\mathbf{p}_G^T \| \mathbf{p}_G^S) + p_{\mathcal{H}}^T \cdot \text{KL}(\tilde{\mathbf{p}}_{\mathcal{H}}^T \| \tilde{\mathbf{p}}_{\mathcal{H}}^S) \\ &\quad + p_{\mathcal{M}}^T \cdot \text{KL}(\tilde{\mathbf{p}}_{\mathcal{M}}^T \| \tilde{\mathbf{p}}_{\mathcal{M}}^S) + p_{\mathcal{T}}^T \cdot \text{KL}(\tilde{\mathbf{p}}_{\mathcal{T}}^T \| \tilde{\mathbf{p}}_{\mathcal{T}}^S). \end{aligned} \quad (8)$$

This formulation reveals that the overall KD loss consists of two key components: (1) cross-group loss, which aligns the aggregate probability distributions across the class groups (head, medium, and tail); (2) weighted sum of within-group losses, which aligns the probability distributions within each group. Notably, the weight assigned to each within-group loss corresponds to the teacher’s cross-group probability  $p_G^T$ . This decomposition enables analysis of the distinct contributions from cross- and within-group learning, revealing the inherent limitations of existing KD in long-tailed settings.

### 3.2. Rebalanced cross-group loss

In long-tailed classification scenarios, a teacher model trained on an imbalanced dataset naturally exhibits a prediction bias toward the data-rich head classes. To examine this phenomenon, we analyze the cross-group output probabilities of a pre-trained teacher model  $\mathbf{p}_G^T = [p_{\mathcal{H}}^T, p_{\mathcal{M}}^T, p_{\mathcal{T}}^T]$  on both the balanced and long-tailed versions of CIFAR-100.

Specifically, we measure the aggregate probability for each class group—head, medium, and tail—within a single batch. On a representative batch from the balanced dataset, the teacher model produces nearly uniform group-wise predictions, with values of [22.54, 20.76, 20.70]. In contrast, on the long-tailed dataset, the corresponding values become [27.88, 19.28, 16.83], indicating a clear bias toward the head group. This trend remains consistent across batches within a single epoch (see Fig. 3), indicating that the teacher model systematically assigns higher confidence to head classes.

This imbalance poses a critical challenge in knowledge distillation. The KL-based loss incentivizes the student to




Figure 3. Group-wise prediction trends of the teacher model on the CIFAR-100 dataset across training batches. On the balanced version (solid lines), the teacher produces nearly uniform group-wise outputs. In contrast, on the long-tailed version (dotted lines), the teacher assigns higher probabilities to head classes and lower probabilities to tail classes.

mimic the teacher’s biased outputs, thereby forcing it to inherit the same head-class bias. Consequently, the student’s predictions become skewed toward dominant classes, which is detrimental to its tail-class performance.

To mitigate this issue, we introduce a group-wise rebalancing mechanism that adjusts the teacher’s cross-group probability distribution before distillation. The goal is to align the group-wise probabilities to a uniform distribution—*e.g.* [21, 21, 21]—instead of the skewed values observed under long-tailed settings. This is achieved by calculating scaling factors for each group such that their rebalanced predictions become equal across a batch.

Let the sum of cross-group probabilities in a batch  $B$  be  $\mathbf{p}_{\text{batch}} = [p_{\mathcal{H}}^B, p_{\mathcal{M}}^B, p_{\mathcal{T}}^B]$ . We define the target average as:

$$p_{\text{avg}}^B = \text{Mean}(p_{\mathcal{H}}^B, p_{\mathcal{M}}^B, p_{\mathcal{T}}^B), \quad (9)$$

and set the desired balanced vector as  $[p_{\text{avg}}^B, p_{\text{avg}}^B, p_{\text{avg}}^B]$ . The scaling factors for each group are then given by:

$$s_{\mathcal{H}} = \frac{p_{\text{avg}}^B}{p_{\mathcal{H}}^B}, s_{\mathcal{M}} = \frac{p_{\text{avg}}^B}{p_{\mathcal{M}}^B}, s_{\mathcal{T}} = \frac{p_{\text{avg}}^B}{p_{\mathcal{T}}^B}. \quad (10)$$

However, simply applying these scaling factors for each sample,  $\hat{\mathbf{p}}_G^T = [s_{\mathcal{H}}p_{\mathcal{H}}^T, s_{\mathcal{M}}p_{\mathcal{M}}^T, s_{\mathcal{T}}p_{\mathcal{T}}^T]$ , does not guarantee a valid probability distribution, as the results will not sum to one ( $\sum_G s_G p_G^T \neq 1$ ). Therefore, to ensure proper normalization, we apply the following correction:

$$\hat{\mathbf{p}}_G^T = \left[ \frac{s_{\mathcal{H}}p_{\mathcal{H}}^T}{\sum_G s_G p_G^T}, \frac{s_{\mathcal{M}}p_{\mathcal{M}}^T}{\sum_G s_G p_G^T}, \frac{s_{\mathcal{T}}p_{\mathcal{T}}^T}{\sum_G s_G p_G^T} \right], \quad (11)$$




Figure 4. Distillation loss curves for KD, DKD, and LTKD for Head (left), Medium (middle), and Tail (right) class groups. LTKD achieves a consistently lower loss across all three groups compared to the baselines, overcoming the suboptimal convergence caused by teacher bias.

where  $\hat{\mathbf{p}}_{\mathcal{G}}^T$  is the rebalanced cross-group probability vector. This operation rebalances the group-level contributions while ensuring the output remains a valid probability distribution. Through this, the student is guided by a rebalanced teacher distribution that no longer disproportionately emphasizes the head group.

### 3.3. Reweighted within-group loss

We now focus on the within-group loss component, which is a weighted sum of KL divergences within each group:

$$\begin{aligned} \sum_{\mathcal{G}} p_{\mathcal{G}}^T \cdot \text{KL}(\tilde{\mathbf{p}}_{\mathcal{G}}^T \| \tilde{\mathbf{p}}_{\mathcal{G}}^S) &= p_{\mathcal{H}}^T \cdot \text{KL}(\tilde{\mathbf{p}}_{\mathcal{H}}^T \| \tilde{\mathbf{p}}_{\mathcal{H}}^S) \\ &+ p_{\mathcal{M}}^T \cdot \text{KL}(\tilde{\mathbf{p}}_{\mathcal{M}}^T \| \tilde{\mathbf{p}}_{\mathcal{M}}^S) + p_{\mathcal{T}}^T \cdot \text{KL}(\tilde{\mathbf{p}}_{\mathcal{T}}^T \| \tilde{\mathbf{p}}_{\mathcal{T}}^S), \end{aligned} \quad (12)$$

where  $\tilde{\mathbf{p}}_{\mathcal{G}}^T$  and  $\tilde{\mathbf{p}}_{\mathcal{G}}^S$  are the within-group distributions of the teacher and student, respectively, and  $p_{\mathcal{G}}^T$  denotes the aggregated probability that the teacher assigns to group  $\mathcal{G} \in \{\mathcal{H}, \mathcal{M}, \mathcal{T}\}$ .

In long-tailed classification, due to the imbalanced nature of the training data, the teacher model tends to allocate higher probability to head classes ( $p_{\mathcal{H}}^T > p_{\mathcal{M}}^T > p_{\mathcal{T}}^T$ ). As a result, the within-group KL loss becomes heavily biased toward the head group, while the contribution from the tail group is substantially diminished. This bias causes the student to receive weaker supervision signals for underrepresented classes, resulting in suboptimal convergence across the full label space (see Fig. 4).

To address this issue, we propose a reweighting strategy that enforces equal importance across all three groups, regardless of the teacher’s confidence. Specifically, we replace the teacher-derived weights  $p_{\mathcal{G}}^T$  with a uniform constant  $\beta$ , yielding the modified within-group loss:

$$\begin{aligned} &\beta \sum_{\mathcal{G}} \text{KL}(\tilde{\mathbf{p}}_{\mathcal{G}}^T \| \tilde{\mathbf{p}}_{\mathcal{G}}^S) \\ &= \beta (\text{KL}(\tilde{\mathbf{p}}_{\mathcal{H}}^T \| \tilde{\mathbf{p}}_{\mathcal{H}}^S) + \text{KL}(\tilde{\mathbf{p}}_{\mathcal{M}}^T \| \tilde{\mathbf{p}}_{\mathcal{M}}^S) + \text{KL}(\tilde{\mathbf{p}}_{\mathcal{T}}^T \| \tilde{\mathbf{p}}_{\mathcal{T}}^S)). \end{aligned} \quad (13)$$

This simple modification ensures that each group contributes equally to the within-group distillation loss, preventing the overrepresented head group from dominating the gradient flow during distillation.

### 3.4. Long-tailed knowledge distillation

By combining these two complementary strategies, we propose the Long-Tailed Knowledge Distillation (LTKD):

$$\text{LTKD} = \alpha \cdot \text{KL}(\hat{\mathbf{p}}_{\mathcal{G}}^T \| \mathbf{p}_{\mathcal{G}}^S) + \beta \cdot \sum_{\mathcal{G}} \text{KL}(\tilde{\mathbf{p}}_{\mathcal{G}}^T \| \tilde{\mathbf{p}}_{\mathcal{G}}^S). \quad (14)$$

- $\hat{\mathbf{p}}_{\mathcal{G}}^T$  and  $\mathbf{p}_{\mathcal{G}}^S$  denote the rebalanced cross-group distributions from the teacher and the original distribution from the student, respectively,
  - $\tilde{\mathbf{p}}_{\mathcal{G}}^T$  and  $\tilde{\mathbf{p}}_{\mathcal{G}}^S$  represent the within-group normalized distributions within each group  $\mathcal{G} \in \{\mathcal{H}, \mathcal{M}, \mathcal{T}\}$ ,
  - $\alpha$  and  $\beta$  are hyperparameters that balance the contributions of the cross- and within-group distillation terms.
- See the supplementary material for the pseudo-code.

## 4. Experiment

### 4.1. Dataset

We construct CIFAR-100-LT, TinyImageNet-LT, and ImageNet-LT by introducing class imbalance into the balanced CIFAR-100 [22], TinyImageNet [23], and ImageNet [31] datasets, following previous works [7, 8, 26].

For each class  $c$ , the number of training samples, denoted by  $|\mathcal{D}_c|$ , is reduced using an exponential decay function:

$$|\hat{\mathcal{D}}_c| = |\mathcal{D}_c| \cdot \gamma^{-c/C}, \quad (15)$$

where  $C$  is the total number of classes and  $\gamma$  is the imbalance factor, defined as the ratio between the number of training samples in the largest and smallest classes [7]. For both CIFAR-100-LT and TinyImageNet-LT, we construct three versions of the training set with imbalance factors of [10, 20, 100]. For ImageNet-LT, imbalance factors

Table 1. Accuracy (%) on tail ( $\mathcal{T}$ ) and overall (All) classes for CIFAR-100-LT with homogeneous (Top) and heterogeneous (Bottom) architectures. **Bold** and underline denote the best and second-best results, respectively.  $\Delta$  is the gap between them.

T-S Pairs	ResNet32×4 – ResNet8×4						VGG13 – VGG8					
$\gamma$	10		20		100		10		20		100	
Group	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All
Teacher	50.72	64.95	39.19	58.82	15.28	45.35	45.67	60.77	36.43	55.10	14.01	43.11
Student	47.32	60.59	36.99	55.44	13.38	42.48	43.67	57.43	33.89	52.29	13.13	40.70
DKD [52]	49.86	64.55	37.87	58.78	13.25	<u>46.11</u>	48.00	61.84	37.65	56.68	14.42	44.22
ReviewKD [6]	<u>52.08</u>	64.71	<u>40.12</u>	<u>59.17</u>	<u>15.09</u>	45.91	47.75	61.43	37.69	56.51	<u>14.76</u>	44.19
DIST [17]	50.28	63.74	38.69	58.28	13.86	45.21	45.57	60.53	34.36	54.68	12.46	42.12
CAT-KD [12]	49.83	<u>64.74</u>	37.67	58.73	12.83	45.33	<u>48.53</u>	<u>62.01</u>	<u>37.95</u>	<u>56.78</u>	14.22	<u>44.33</u>
<b>LTKD</b>	<b>58.66</b>	<b>66.76</b>	<b>49.70</b>	<b>62.54</b>	<b>27.21</b>	<b>51.08</b>	<b>53.95</b>	<b>63.04</b>	<b>45.77</b>	<b>58.86</b>	<b>23.30</b>	<b>47.66</b>
$\Delta$	<b>+6.58</b>	<b>+2.02</b>	<b>+9.58</b>	<b>+3.37</b>	<b>+12.12</b>	<b>+4.97</b>	<b>+5.42</b>	<b>+1.03</b>	<b>+7.82</b>	<b>+2.08</b>	<b>+8.54</b>	<b>+3.33</b>

T-S Pairs	WRN-40-2 – ShuffleNetV1						ResNet50 – MobileNetV2					
$\gamma$	10		20		100		10		20		100	
Group	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All
Teacher	49.77	63.05	39.88	58.27	14.88	44.78	49.74	63.51	37.74	56.70	14.42	42.26
Student	40.04	54.06	29.91	48.22	10.74	36.21	30.47	44.58	22.32	39.25	7.04	27.56
DKD [52]	50.86	63.65	39.94	58.28	15.04	45.24	<u>43.29</u>	57.20	<u>33.23</u>	<u>52.20</u>	<u>12.45</u>	<u>39.21</u>
ReviewKD [6]	<u>51.24</u>	<u>63.90</u>	<u>40.44</u>	<u>58.63</u>	<u>15.81</u>	<u>45.40</u>	33.68	47.75	24.80	42.08	9.75	31.86
DIST [17]	48.40	62.47	37.48	56.92	12.23	41.95	37.86	52.36	27.11	46.50	9.81	34.96
CAT-KD [12]	51.02	63.68	40.23	58.26	14.68	44.84	43.18	<u>57.23</u>	33.17	51.90	11.61	38.45
<b>LTKD</b>	<b>57.40</b>	<b>65.42</b>	<b>48.42</b>	<b>60.94</b>	<b>23.99</b>	<b>48.60</b>	<b>48.43</b>	<b>57.79</b>	<b>40.82</b>	<b>53.70</b>	<b>21.04</b>	<b>42.45</b>
$\Delta$	<b>+6.16</b>	<b>+1.52</b>	<b>+7.98</b>	<b>+2.31</b>	<b>+8.18</b>	<b>+3.20</b>	<b>+5.14</b>	<b>+0.56</b>	<b>+7.59</b>	<b>+1.50</b>	<b>+8.59</b>	<b>+3.24</b>

of [5, 10, 20] are used. The classes are divided into three groups based on the number of training samples: Head (the top 33% of classes), Medium (the next 34%), and Tail (the bottom 33%). The test set remains balanced to ensure fair evaluation.

## 4.2. Implementation details

We conduct experiments using widely used CNN architectures, including ResNet [14], VGG [33], WideResNet (WRN) [46], ShuffleNet [48], and MobileNet [32]. Each experiment is conducted three times, and the average performance is reported. Additional details and results are provided in the supplementary material.

## 4.3. Main results

**CIFAR-100-LT.** In all settings, LTKD consistently shows remarkable improvements in both overall accuracy (All) and tail-class accuracy ( $\mathcal{T}$ ), as shown in Tab. 1. For instance, with a ResNet32×4–ResNet8×4 pair ( $\gamma = 100$ ), LTKD boosts the tail-class accuracy from 15.09% to 27.21%, and overall accuracy from 46.11% to 51.08%. Similarly, for the ResNet50–MobileNetV2 ( $\gamma = 100$ ) pair, LTKD delivers a substantial gain of +8.59% in tail accuracy and +3.24% in overall accuracy. These gains are consistent across all imbalance levels and architecture combinations, underscoring the importance of our rebalanced and reweighted supervision.

**TinyImageNet-LT.** Even on the more challenging TinyImageNet-LT, LTKD proves its effectiveness with consistent performance improvements, as shown in Tab. 2. For example, with the ResNet32×4–ResNet8×4 pair ( $\gamma = 100$ ), LTKD improves tail accuracy from 9.09% to 10.48% and overall accuracy from 34.61% to 36.21%. Similarly, for VGG13–MobileNetV2 pair ( $\gamma = 100$ ), LTKD boosts tail-class accuracy from 7.02% to 9.52% and overall accuracy from 29.97% to 32.30%. These results highlight the robustness of LTKD on a more complex benchmark.

**ImageNet-LT.** On the large-scale ImageNet-LT, LTKD demonstrates its scalability by consistently outperforming all baselines. For instance, with a ResNet34–ResNet18 pair, LTKD improves tail accuracy by +1.06%, +1.75%, and +2.40% for  $\gamma = 5, 10, 20$ , respectively. This strong performance extends to the ResNet50–MobileNetV1 setting, with gains of up to +3.20% on tail classes, highlighting the method’s ability to perform effective knowledge transfer even on large-scale, severely imbalanced datasets.

## 4.4. Ablation study

**Rebalance effect.** To isolate the effect of our rebalanced cross-group loss, we compare it against a non-rebalanced baseline, disabling the within-group loss for both settings (Tab. 4). Our rebalanced loss consistently improves tail accuracy, improving it from 38.30% to 40.51% for the ResNet32×4–ResNet8×4 pair. Similar positive trends are

Table 2. Accuracy (%) on tail ( $\mathcal{T}$ ) and overall (All) classes for TinyImageNet-LT with homogeneous (Top) and heterogeneous (Bottom) architectures. **Bold** and underline denote the best and second-best results, respectively.  $\Delta$  is the gap between them.

T-S Pairs	ResNet32 $\times$ 4 – ResNet8 $\times$ 4						VGG13 – VGG8					
$\gamma$	10		20		100		10		20		100	
Group	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All
Teacher	38.47	52.64	28.74	47.49	9.53	35.37	32.53	45.23	21.85	39.75	6.29	29.71
Student	29.72	44.60	21.46	40.25	4.73	30.62	31.12	43.76	22.19	39.00	6.88	29.96
KD [16]	27.34	45.38	18.11	41.35	3.38	31.42	31.92	47.16	20.60	41.49	3.99	30.95
DKD [52]	34.70	48.93	<u>26.58</u>	44.84	<u>9.09</u>	<u>34.61</u>	33.20	<u>48.01</u>	<u>22.65</u>	<u>42.44</u>	5.88	31.82
ReviewKD [6]	32.85	49.13	23.43	44.62	5.39	33.51	<u>34.39</u>	47.66	<u>24.84</u>	42.36	<u>7.61</u>	<u>32.18</u>
DIST [17]	<u>34.81</u>	<u>50.14</u>	25.71	<u>45.52</u>	7.30	33.98	33.48	47.22	23.19	41.46	5.98	31.01
<b>LTKD</b>	<b>40.66</b>	<b>51.33</b>	<b>31.33</b>	<b>47.05</b>	<b>10.48</b>	<b>36.21</b>	<b>38.90</b>	<b>49.43</b>	<b>29.30</b>	<b>44.22</b>	<b>9.73</b>	<b>33.78</b>
$\Delta$	<b>+5.85</b>	<b>+1.19</b>	<b>+4.75</b>	<b>+1.53</b>	<b>+1.39</b>	<b>+1.60</b>	<b>+4.51</b>	<b>+1.42</b>	<b>+4.46</b>	<b>+1.78</b>	<b>+2.12</b>	<b>+1.60</b>

T-S Pairs	ResNet32 $\times$ 4 – ShuffleNetV1						VGG13 – MobileNetV2					
$\gamma$	10		20		100		10		20		100	
Group	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All
Teacher	38.47	52.64	28.74	47.49	9.53	35.37	32.53	45.23	21.85	39.75	6.29	29.71
Student	24.43	37.10	16.80	31.85	4.71	22.77	26.98	39.73	17.38	33.14	3.97	22.99
KD [16]	34.67	49.05	24.12	42.81	5.62	30.74	31.24	45.60	19.53	39.50	3.27	28.44
DKD [52]	<u>36.83</u>	<u>50.22</u>	26.64	44.38	8.34	<u>33.23</u>	<u>33.07</u>	<u>46.79</u>	22.06	<u>40.87</u>	5.70	<u>29.97</u>
ReviewKD [6]	36.30	49.27	<u>27.09</u>	<u>44.72</u>	<u>8.49</u>	33.12	32.37	44.99	<u>22.77</u>	39.36	<u>7.02</u>	29.20
DIST [17]	36.49	50.07	25.74	43.59	7.23	31.19	32.92	46.09	21.77	40.05	5.52	29.08
<b>LTKD</b>	<b>42.12</b>	<b>51.64</b>	<b>33.06</b>	<b>46.41</b>	<b>12.85</b>	<b>35.09</b>	<b>39.04</b>	<b>48.71</b>	<b>28.28</b>	<b>43.22</b>	<b>9.52</b>	<b>32.30</b>
$\Delta$	<b>+5.29</b>	<b>+1.42</b>	<b>+5.97</b>	<b>+1.69</b>	<b>+4.36</b>	<b>+1.86</b>	<b>+5.97</b>	<b>+1.92</b>	<b>+5.51</b>	<b>+2.35</b>	<b>+2.50</b>	<b>+2.33</b>

Table 3. Accuracy (%) on tail ( $\mathcal{T}$ ) and overall (All) classes for ImageNet-LT. **Bold** and underline denote the best and second-best results, respectively.  $\Delta$  is the gap between them.

T-S Pairs	ResNet34 – ResNet18						ResNet50 – MobileNetV1					
$\gamma$	5		10		20		5		10		20	
Group	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All
Teacher	57.75	67.61	50.61	63.85	43.00	59.91	59.62	69.06	52.69	65.57	44.82	61.46
Student	54.48	64.73	47.45	61.18	39.98	57.25	55.42	65.46	49.49	62.55	41.68	58.94
KD [16]	56.14	66.27	49.35	63.03	41.72	59.15	56.58	66.51	50.42	63.70	42.45	59.91
DKD [52]	56.83	66.84	49.95	63.50	42.65	59.82	58.54	68.09	52.39	65.04	45.04	61.46
ReviewKD [6]	<u>57.27</u>	<u>66.96</u>	<u>50.80</u>	<u>63.72</u>	<u>43.48</u>	<u>60.15</u>	<u>58.59</u>	<u>68.10</u>	<u>53.06</u>	<u>65.31</u>	<u>45.35</u>	<u>61.84</u>
DIST [17]	56.79	66.66	50.28	63.56	42.78	59.94	57.29	66.94	50.89	64.09	43.70	60.71
CAT-KD [12]	55.92	66.14	49.58	63.20	42.20	59.67	57.08	66.96	51.00	64.07	43.55	60.54
<b>LTKD</b>	<b>58.33</b>	<b>67.23</b>	<b>52.55</b>	<b>64.29</b>	<b>45.88</b>	<b>60.80</b>	<b>60.17</b>	<b>68.48</b>	<b>54.40</b>	<b>65.52</b>	<b>48.55</b>	<b>62.22</b>
$\Delta$	<b>+1.06</b>	<b>+0.27</b>	<b>+1.75</b>	<b>+0.57</b>	<b>+2.40</b>	<b>+0.65</b>	<b>+1.58</b>	<b>+0.38</b>	<b>+1.34</b>	<b>+0.21</b>	<b>+3.20</b>	<b>+0.38</b>

observed in other architecture pairs, confirming the effectiveness of the rebalance process, particularly for underrepresented classes. Interestingly, we observed a minor trade-off in the ResNet32 $\times$ 4–ShuffleNetV1 pair, where tail accuracy improved (30.01% to 30.68%) but overall accuracy slightly dropped (-0.01%), likely due to suppressing head-class signals during the rebalance process. However, this is resolved by incorporating the within-group loss (Tab. 5), demonstrating the complementary strength of our LTKD.

**Impact of each component.** We conduct an ablation study to disentangle the contributions of our cross-group and within-group components (Tab. 5). Our findings show that each component independently boosts both tail and

overall performance. For the ResNet32 $\times$ 4–ResNet8 $\times$ 4 pair, the cross-group loss alone yields a +3.70% improvement in tail accuracy, while the within-group loss alone provides an even larger gain of +5.53%. This highlights that uniformly weighting within-group losses is particularly effective in mitigating teacher bias. The best performance is consistently achieved when both strategies are combined. In the VGG13–MobileNetV2 setting, our LTKD improves tail accuracy by +9.28% and overall accuracy by +4.34% over the baseline. These results confirm that the two components are complementary, and jointly addressing cross- and within-group biases is key to effective long-tailed knowledge distillation.

Table 4. Performance comparison of cross-group loss on CIFAR-100-LT ( $\gamma = 20$ ). “Biased” refers to the unbalanced cross-group loss, while “Ours” applies our balanced loss. The within-group loss is disabled for both methods.  $\mathcal{T}$  and All denote tail and overall accuracy, respectively.  $\Delta$  shows the improvement of “Ours” over the “Biased” baseline.

Models	Teacher Student	ResNet32×4				VGG13				WRN-40-2				
		ResNet8×4		ShuffleNetV1		VGG8		MobileNetV2		WRN-40-1		ShuffleNetV1		
Loss	Cross	Within	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All		
Biased	✓	✗	38.30	55.77	30.01	<b>48.82</b>	35.55	53.30	23.42	40.29	34.64	53.97	30.23	49.31
Ours	✓	✗	<b>40.51</b>	<b>56.55</b>	<b>30.68</b>	48.81	<b>37.26</b>	<b>53.70</b>	<b>23.97</b>	<b>40.30</b>	<b>35.23</b>	<b>54.35</b>	<b>32.44</b>	<b>49.91</b>
		$\Delta$	<b>+2.21</b>	<b>+0.78</b>	<b>+0.67</b>	-0.01	<b>+1.71</b>	<b>+0.40</b>	<b>+0.55</b>	<b>+0.01</b>	<b>+0.59</b>	<b>+0.38</b>	<b>+2.21</b>	<b>+0.60</b>

Table 5. Ablation study of cross-group and within-group losses on CIFAR-100-LT ( $\gamma = 20$ ). “Baseline” is vanilla KD. “Ours (Cross ✓, Within ✓)” combines both rebalanced cross-group and reweighted within-group loss.

Models	Teacher Student	ResNet32×4				VGG13				WRN-40-2				
		ResNet8×4		ShuffleNetV1		VGG8		MobileNetV2		WRN-40-1		ShuffleNetV1		
Loss	Cross	Within	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All	$\mathcal{T}$	All		
Baseline	✗	✗	36.81	57.41	36.13	55.03	37.29	56.13	29.05	47.69	35.63	56.71	38.30	56.61
	✓	✗	40.51	56.55	30.68	48.81	37.26	53.70	23.97	40.30	35.23	54.35	32.44	49.91
Ours	✗	✓	42.34	59.78	39.10	56.84	38.54	56.83	31.99	49.20	39.67	58.11	40.45	57.65
	✓	✓	<b>49.70</b>	<b>62.54</b>	<b>45.94</b>	<b>59.62</b>	<b>45.77</b>	<b>58.86</b>	<b>38.33</b>	<b>52.03</b>	<b>45.74</b>	<b>59.91</b>	<b>48.42</b>	<b>60.94</b>




Figure 5. Hyperparameter sensitivity of LTKD on CIFAR-100-LT ( $\gamma = 20$ ), showing overall accuracy (Left) and tail-class accuracy (Right). The blue line shows the effect of varying the cross-group weight  $\alpha$  (while  $\beta = 1$ ). The red line shows the effect of varying the within-group weight  $\beta$  (while  $\alpha = 6$ ).

**Hyperparameters.** We analyze the sensitivity to hyperparameters  $\alpha$  (cross-group) and  $\beta$  (within-group) on CIFAR-100-LT ( $\gamma = 20$ ) using a ResNet32×4 and ResNet8×4 pair (Fig. 5). Varying  $\alpha$  (with  $\beta = 1$ ) shows overall accuracy peaking at  $\alpha = 6$  (61.45%). When fixing  $\alpha = 6$  and varying  $\beta$ , performance peaks at  $\beta = 6$  (62.54%). Tail accuracy follows a similar trend, reaching its peak of 49.72% ( $\alpha = 6, \beta = 8$ ). Importantly, LTKD consistently outperforms the strong ReviewKD baseline (Overall: 59.17%, Tail: 40.12%) across a wide range of  $\alpha$  and  $\beta$  values, demonstrating both its robustness and its significant benefit for underrepresented classes.

**Number of groups.** To evaluate sensitivity to group definitions, we vary the number of groups  $n(\mathcal{G})$  from 3 to 100 (Tab. 6). Even at the simple  $n(\mathcal{G}) = 3$  ( $\mathcal{H}, \mathcal{M}, \mathcal{T}$ ), LTKD already achieves significant improvements over previous

Table 6. Overall accuracy on CIFAR-100-LT ( $\gamma = 100$ ) with varying number of groups, including the continuous reweighting.

$n(\mathcal{G})$	3	4	5	10	20	25	50	100
R32×4-R8×4	51.08	51.08	<u>51.10</u>	<b>51.14</b>	50.99	50.34	50.06	50.41
VGG13-VGG8	47.66	47.85	48.06	48.26	<b>48.69</b>	<u>48.58</u>	48.19	47.82
WRN402-SV1	48.60	48.98	49.03	49.27	49.54	<b>49.90</b>	<u>49.64</u>	47.95
R50-MV2	42.45	43.01	42.99	43.40	<b>43.43</b>	<u>43.42</u>	42.51	41.18

methods, and increasing  $n(\mathcal{G})$  enables more fine-grained bias correction, yielding consistent performance gains.

The  $n(\mathcal{G}) = 100$  setting corresponds to a continuous reweighting scheme, where LTKD continues to outperform existing methods (see Tab. 1), demonstrating that the framework naturally extends beyond discrete grouping. Nevertheless, the discrete grouping design enables the explicit formulation of the within-group loss, leading to additional performance gains.

## 5. Conclusion

We present Long-Tailed Knowledge Distillation (LTKD), a novel framework addressing the ineffective transfer of biased knowledge from teachers trained on long-tailed datasets. LTKD introduces a rebalanced cross-group loss and a reweighted within-group loss, designed based on our decomposition of the conventional KL divergence, which revealed its inherent bias. This enables the effective distillation of balanced knowledge from a biased teacher. Through extensive experiments, we demonstrate that LTKD consistently enhances both overall and tail-class performance in long-tailed scenarios. We plan to extend this framework to other domains where long-tail issues are pervasive, such as object detection and semantic segmentation.

## Acknowledgment

Special thanks to Young Hwa Sung and Chong Hui Kim for providing an excellent research environment. This work was supported by the Korean Government.

## References

- [1] Kamal Acharya, Alvaro Velasquez, and Houbing Herbert Song. A survey on symbolic knowledge distillation of large language models. *IEEE Transactions on Artificial Intelligence*, 2024. 1
- [2] Pedram Agand. Knowledge distillation from single-task teachers to multi-task student for end-to-end autonomous driving. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 23375–23376, 2024. 1
- [3] Shaden Alshammari, Yu-Xiong Wang, Deva Ramanan, and Shu Kong. Long-tailed recognition via weight balancing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6897–6907, 2022. 1
- [4] Yaping Bai, Jinghua Li, Dehui Kong, Suqiao Yang, and Bao-cai Yin. Ekds: Long-tailed recognition based on expert knowledge distillation for specific categories. *Neural Networks*, page 108099, 2025. 2
- [5] Justin Chih-Yao Chen, Swarnadeep Saha, Elias Stengel-Eskin, and Mohit Bansal. Magdi: Structured distillation of multi-agent interaction graphs improves reasoning in smaller language models. *arXiv preprint arXiv:2402.01620*, 2024. 1
- [6] Pengguang Chen, Shu Liu, Hengshuang Zhao, and Jiaya Jia. Distilling knowledge via knowledge review. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5008–5017, 2021. 1, 2, 6, 7, 3, 4
- [7] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9268–9277, 2019. 5
- [8] Fei Du, Peng Yang, Qi Jia, Fengtao Nan, Xiaoting Chen, and Yun Yang. Global and local mixture consistency cumulative learning for long-tailed visual recognitions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15814–15823, 2023. 5
- [9] Jianping Gou, Baosheng Yu, Stephen J Maybank, and Dacheng Tao. Knowledge distillation: A survey. *International Journal of Computer Vision*, 129:1789–1819, 2021. 2
- [10] Jianping Gou, Liyuan Sun, Baosheng Yu, Lan Du, Kotagiri Ramamohanarao, and Dacheng Tao. Collaborative knowledge distillation via multiknowledge transfer. *IEEE Transactions on Neural Networks and Learning Systems*, 35(5): 6718–6730, 2022. 2
- [11] Yuxian Gu, Li Dong, Furu Wei, and Minlie Huang. Minillm: Knowledge distillation of large language models. *arXiv preprint arXiv:2306.08543*, 2023. 1
- [12] Ziyao Guo, Haonan Yan, Hui Li, and Xiaodong Lin. Class attention transfer based knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11868–11877, 2023. 1, 2, 6, 7, 3, 4
- [13] Sangchul Hahn and Heeyoul Choi. Self-knowledge distillation in natural language processing. *arXiv preprint arXiv:1908.01851*, 2019. 1
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6, 2
- [15] Yin-Yin He, Jianxin Wu, and Xiu-Shen Wei. Distilling virtual examples for long-tailed recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 235–244, 2021. 2
- [16] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 1, 2, 7, 3, 4
- [17] Tao Huang, Shan You, Fei Wang, Chen Qian, and Chang Xu. Knowledge distillation from a stronger teacher. *Advances in Neural Information Processing Systems*, 35:33716–33727, 2022. 1, 2, 6, 7, 3, 4
- [18] Tao Huang, Yuan Zhang, Mingkai Zheng, Shan You, Fei Wang, Chen Qian, and Chang Xu. Knowledge diffusion for distillation. *Advances in Neural Information Processing Systems*, 36:65299–65316, 2023. 1, 2
- [19] Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. Tinybert: Distilling bert for natural language understanding. *arXiv preprint arXiv:1909.10351*, 2019. 1
- [20] Lie Ju, Xin Wang, Lin Wang, Tongliang Liu, Xin Zhao, Tom Drummond, Dwarikanath Mahapatra, and Zongyuan Ge. Relational subsets knowledge distillation for long-tailed retinal diseases recognition. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 3–12. Springer, 2021. 2
- [21] Seonghak Kim, Gyeongdo Ham, Yucheol Cho, and Daeshik Kim. Robustness-reinforced knowledge distillation with correlation distance and network pruning. *IEEE Transactions on Knowledge and Data Engineering*, 2024. 2
- [22] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 1, 5, 2
- [23] Yann Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015. 1, 5, 2
- [24] Tianhao Li, Limin Wang, and Gangshan Wu. Self supervision to distillation for long-tailed visual recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 630–639, 2021. 2
- [25] Zhihui Li, Pengfei Xu, Xiaojun Chang, Luyao Yang, Yuanyuan Zhang, Lina Yao, and Xiaojiang Chen. When object detection meets knowledge distillation: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(8):10555–10579, 2023. 1
- [26] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2537–2546, 2019. 5
- [27] Aditya Krishna Menon, Sadeep Jayasumana, Ankit Singh Rawat, Himanshu Jain, Andreas Veit, and Sanjiv Kumar. Long-tail learning via logit adjustment. *arXiv preprint arXiv:2007.07314*, 2020. 2

- [28] Seyed Iman Mirzadeh, Mehrdad Farajtabar, Ang Li, Nir Levine, Akihiro Matsukawa, and Hassan Ghasemzadeh. Improved knowledge distillation via teacher assistant. In *AAAI*, 2020. 1, 2
- [29] Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho. Relational knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3967–3976, 2019. 1, 2
- [30] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*, 2014. 1, 2, 3, 4
- [31] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115:211–252, 2015. 5, 2
- [32] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. MobilenetV2: Inverted residuals and linear bottlenecks. In *CVPR*, 2018. 6, 2
- [33] K. Simonyan and A Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 6, 2
- [34] Siqi Sun, Yu Cheng, Zhe Gan, and Jingjing Liu. Patient knowledge distillation for bert model compression. *arXiv preprint arXiv:1908.09355*, 2019. 1
- [35] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive representation distillation. *arXiv preprint arXiv:1910.10699*, 2019. 1, 2
- [36] Chaozheng Wang, Shuzheng Gao, Pengyun Wang, Cuiyun Gao, Wenjie Pei, Lujia Pan, and Zenglin Xu. Label-aware distribution calibration for long-tailed classification. *IEEE transactions on neural networks and learning systems*, 35(5):6963–6975, 2022. 1
- [37] Jiabao Wang, Yuming Chen, Zhaohui Zheng, Xiang Li, Ming-Ming Cheng, and Qibin Hou. Crosskd: Cross-head knowledge distillation for object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16520–16530, 2024. 1
- [38] Liuyu Xiang, Guiguang Ding, and Jungong Han. Learning from multiple experts: Self-paced knowledge distillation for long-tailed classification. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*, pages 247–263. Springer, 2020. 2
- [39] Xiaohan Xu, Ming Li, Chongyang Tao, Tao Shen, Reynold Cheng, Jinyang Li, Can Xu, Dacheng Tao, and Tianyi Zhou. A survey on knowledge distillation of large language models. *arXiv preprint arXiv:2402.13116*, 2024. 1
- [40] Chuanguang Yang, Helong Zhou, Zhulin An, Xue Jiang, Yongjun Xu, and Qian Zhang. Cross-image relational knowledge distillation for semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12319–12328, 2022. 1
- [41] Chuanpeng Yang, Yao Zhu, Wang Lu, Yidong Wang, Qian Chen, Chenlong Gao, Bingjie Yan, and Yiqiang Chen. Survey on knowledge distillation for large language models: methods, evaluation, and application. *ACM Transactions on Intelligent Systems and Technology*, 2024. 1
- [42] Lu Yang, He Jiang, Qing Song, and Jun Guo. A survey on long-tailed visual recognition. *International Journal of Computer Vision*, 130(7):1837–1872, 2022. 1
- [43] Ziqing Yang, Yiming Cui, Zhipeng Chen, Wanxiang Che, Ting Liu, Shijin Wang, and Guoping Hu. Textbrewer: An open-source knowledge distillation toolkit for natural language processing. *arXiv preprint arXiv:2002.12620*, 2020. 1
- [44] Zhendong Yang, Zhe Li, Ailing Zeng, Zexian Li, Chun Yuan, and Yu Li. Vitkd: Feature-based knowledge distillation for vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1379–1388, 2024. 1
- [45] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. *arXiv preprint arXiv:1612.03928*, 2016. 1, 2
- [46] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *BMVC*, 2016. 6, 2
- [47] Shaoyu Zhang, Chen Chen, Xiyuan Hu, and Silong Peng. Balanced knowledge distillation for long-tailed learning. *Neurocomputing*, 527:36–46, 2023. 2
- [48] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6848–6856, 2018. 6, 2
- [49] Ying Zhang, Tao Xiang, Timothy M Hospedales, and Huchuan Lu. Deep mutual learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4320–4328, 2018. 1, 2
- [50] Yifan Zhang, Bingyi Kang, Bryan Hooi, Shuicheng Yan, and Jiashi Feng. Deep long-tailed learning: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 45(9):10795–10816, 2023. 1
- [51] Yuan Zhang, Tao Huang, Jiaming Liu, Tao Jiang, Kuan Cheng, and Shanghang Zhang. Freekd: Knowledge distillation via semantic frequency prompt. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15931–15940, 2024. 1, 2
- [52] Borui Zhao, Quan Cui, Renjie Song, Yiyu Qiu, and Jiajun Liang. Decoupled knowledge distillation. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pages 11953–11962, 2022. 1, 2, 6, 7, 3, 4
- [53] Qihao Zhao, Chen Jiang, Wei Hu, Fan Zhang, and Jun Liu. Mdcs: More diverse experts with consistency self-distillation for long-tailed recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11597–11608, 2023. 2
- [54] Zhonghan Zhao, Ke Ma, Wenhao Chai, Xuan Wang, Kewei Chen, Dongxu Guo, Yanting Zhang, Hongwei Wang, and Gaoang Wang. Do we really need a complex agent system? distill embodied agent into a single model. *arXiv preprint arXiv:2404.04619*, 2024. 1