

Rethinking Long-tailed Dataset Distillation: A Uni-Level Framework with Unbiased Recovery and Relabeling

Xiao Cui^{1,2}, Yulei Qin³, Xinyue Li^{1,2}, Wengang Zhou¹, Hongsheng Li², Houqiang Li^{1*}

¹University of Science and Technology of China

²The Chinese University of Hong Kong, MMLab

³Independent Researcher

{cuixiao2001,lrel7}@mail.ustc.edu.cn, yuleichin@126.com, {zhwg,lihq}@ustc.edu.cn, hsli@ee.cuhk.edu.hk

Abstract

Dataset distillation creates a small distilled set that enables efficient training by capturing key information from the full dataset. While existing dataset distillation methods perform well on balanced datasets, they struggle under long-tailed distributions, where imbalanced class frequencies induce biased model representations and corrupt statistical estimates such as Batch Normalization (BN) statistics. In this paper, we rethink long-tailed dataset distillation by revisiting the limitations of trajectory-based methods, and instead adopt the statistical alignment perspective to jointly mitigate model bias and restore fair supervision. To this end, we introduce three dedicated components that enable unbiased recovery of distilled images and soft relabeling: (1) enhancing expert models (an observer model for recovery and a teacher model for relabeling) to enable reliable statistics estimation and soft-label generation; (2) recalibrating BN statistics via a full forward pass with dynamically adjusted momentum to reduce representation skew; (3) initializing synthetic images by incrementally selecting high-confidence and diverse augmentations via a multi-round mechanism that promotes coverage and diversity. Extensive experiments on four long-tailed benchmarks show consistent improvements over state-of-the-art methods across varying degrees of class imbalance. Notably, our approach improves top-1 accuracy by 15.6% on CIFAR-100-LT and 11.8% on Tiny-ImageNet-LT under IPC=10 and IF=10.

Introduction

Dataset distillation (DD) is the process of synthesizing a significantly smaller yet representative dataset that retains the essential characteristics of an original, larger dataset (Wang et al. 2018; Yu, Liu, and Wang 2023; Liu and Du 2025). By drastically reducing data volume, DD facilitates efficient model training and substantially reduces computational costs, rendering it particularly valuable for resource-constrained scenarios (Jia et al. 2024; Chai et al. 2024). Beyond reducing computational burden, DD also offers a compact lens for studying how data distribution affects model learning (Zhu et al. 2023; Cheng et al. 2024).

Long-tailed dataset distillation (Zhao et al. 2025) specifically addresses scenarios characterized by class imbalance,

*Corresponding authors: Wengang Zhou, Hongsheng Li and Houqiang Li.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

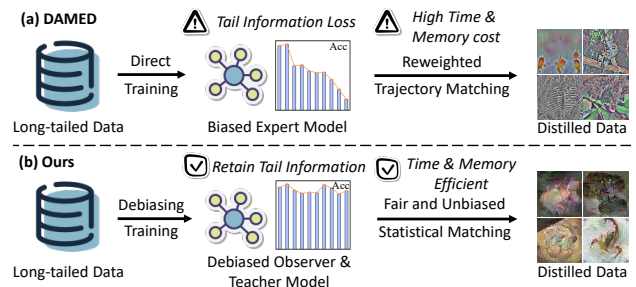


Figure 1: Comparison of DAMED and our method. (a) DAMED directly trains a biased expert on long-tailed data and applies reweighted trajectory matching, leading to tail information loss and high computational cost. (b) Our method debiases both observer and teacher models and performs unbiased statistical matching, effectively retaining tail knowledge with shorter time and lower memory cost.

where a small subset of head classes contain abundant samples, whereas the remaining tail classes are sparsely represented. Such imbalance is ubiquitous in real-world applications as the acquisition of sufficient samples for rare categories is costly or infeasible. Most existing DD approaches (Liu et al. 2023; Guo et al. 2024; Shao et al. 2024b) perform well on balanced datasets, but struggle under class-imbalanced conditions. Their assumption of uniform data density leads to head-class dominance in the synthetic set and poor representation of minority classes, ultimately degrading performance in long-tailed scenarios.

Very few studies have explicitly tackled the limitations of conventional DD methods under long-tailed distributions, mainly because the widely used benchmarks typically feature balanced class structures (Krizhevsky 2012; Le and Yang 2015; Deng et al. 2009). To our knowledge, DAMED (Zhao et al. 2025) is the only recent work that explicitly addresses this issue. It simulates long-tailed training dynamics by injecting class-frequency-aware offsets into the softmax layer, thereby inducing gradient behaviors that mimic those observed under imbalanced training. However, as illustrated in Fig. 1, DAMED still presents significant limitations. 1) **Under-representation of tail classes**. It relies on a feature extraction expert trained on long-tailed data but

without debiasing, causing tail-class representations to be poorly preserved in the distilled dataset. 2) **Unintentional trade-off in trajectory matching.** Mid-frequency classes receive unstable or insufficient gradient feedback, leading to compromised performance across the distribution. 3) **Heavy computation overhead.** Its bi-level trajectory-based optimization suffers from computational inefficiency and excessive GPU memory usage, severely restricting practical applicability (Yin, Xing, and Shen 2023).

To comprehensively resolve these shortcomings, we propose a novel uni-level optimization framework that explicitly counteracts biases stemming from class imbalance in a cost-efficient manner. This formulation is critical because effective debiasing strategies often cause the expert’s training trajectory to deviate significantly from that induced by standard training, rendering trajectory matching unstable and challenging. Our framework considers debiasing through two complementary components: unbiased synthetic image recovery and unbiased soft relabeling. To achieve effective recovery in long-tailed settings, our framework ensures a diverse and representative synthetic image initialization, and leverages a debiased expert model (observer model) to perform a fair extraction of BN statistics for precise alignment. Meanwhile, unbiased soft relabeling provides effective semantic supervision through soft labels generated by another well-trained, debiased expert model (teacher model).

To implement this design, we introduce three tailored strategies to mitigate model bias, statistical unfairness, and suboptimal initialization under long-tailed distributions. First, we propose a **mixture consistency** loss and a **class-wise debias** loss to regularize both the observer and teacher models. The former ensures robust feature learning against multiple data augmentations and the latter adopts dynamic weighting to rebalance class-wise supervision. Second, we recalibrate the estimation of BN statistics to address the challenges of class imbalance and temporal dependency. We freeze the observer model and perform an efficient forward-only pass over the entire training set. During this process, our dynamically adjusted momentum ensures equal contribution from all samples within each class, eliminating **intra-class bias**. We then average the per-class statistics to obtain a globally balanced estimate, removing **inter-class bias**. Third, we introduce a confidence-aware, class-independent strategy for **synthesis initialization**. For each real image, multiple augmentations are precomputed and scored by the teacher model using negative cross-entropy. We adopt a multi-round selection strategy where each image contributes at most one augmentation per round, progressively selecting high-confidence variants to ensure diversity. To ensure consistent batch structure, we insert zero-filled placeholders for all classes having fewer instances than the largest one.

Our main contributions are as follows:

- We rethink long-tailed dataset distillation by moving from bi-level trajectory matching to a uni-level statistical alignment framework that better supports debiasing.
- We implement unbiased recovery and soft relabeling through three key strategies: expert model debiasing; fair BN statistics recalibration; and confidence-guided,

multi-round synthetic data initialization.

- Extensive experiments on CIFAR-10-LT, CIFAR-100-LT, Tiny-ImageNet-LT, and ImageNet-LT demonstrate our consistent superiority against state-of-the-art baselines. It improves accuracy by 15.6% on CIFAR-100-LT and 11.8% on Tiny-ImageNet-LT (IPC=10, IF=10).

Related Work

Dataset Distillation

Early dataset distillation methods, such as K-Center (Sener and Savarese 2017) and GraphCut (Iyer et al. 2021), select a subset of real data directly, which limits the expressiveness of the resulting distilled dataset. Subsequent methods fall into three major categories. Gradient-matching-based approaches (Liu et al. 2023; Wang et al. 2025) align gradients between real and distilled data but scale poorly due to high memory usage. Trajectory-matching-based methods (Cazenavette et al. 2022; Zhong et al. 2025) simulate training dynamics but are computationally expensive and memory-intensive. Distribution-matching-based methods (Zhao and Bilen 2023; Cui et al. 2025a) speed up convergence by matching features but still suffer from high memory costs and degrade on larger datasets like Tiny-ImageNet or ImageNet. Recent efforts attempt to reduce memory overhead via generative-model-based approaches (Cui et al. 2025b; Chen et al. 2025) or by adopting uni-level optimization (Sun et al. 2024; Shao et al. 2024b). However, generative approaches typically rely on pretrained generators trained on balanced large-scale datasets, while existing uni-level methods operate under balanced assumptions and lack explicit debiasing strategies. DAMED (Zhao et al. 2025) is the only prior work explicitly targeting long-tailed DD. However, it inherits representation bias from long-tailed expert training and retains the inefficiencies of trajectory-matching frameworks. In contrast, our work is the first to systematically address long-tailed DD within a uni-level framework, with principled strategies for expert debiasing, image initialization, and unbiased alignment.

Long-tailed Recognition

Long-tailed recognition refers to visual tasks performed under imbalanced data distributions (Zhang et al. 2025). To mitigate the resulting representation bias, data augmentation strategies have been widely studied (Zheng et al. 2024; Wang et al. 2024; Li and Jia 2025). For instance, Mixup (Zhang et al. 2018) and its class-aware extension UniMix (Li et al. 2021) promote feature interpolation to enrich supervision for tail classes, while CMO (Shi, Dong, and Shen 2021) generates context-aware mixed samples that better preserve semantic coherence for rare categories. Beyond augmentation, other approaches mitigate long-tailed bias through network-level optimization (Zhang et al. 2023; Zhu et al. 2024; He 2024), data synthesis using generative models or instance composition (Shao et al. 2024a; Khorram et al. 2024; Zhao et al. 2024), and loss rebalancing strategies (Xiong and Yao 2024; Du, Han, and Huang 2024; Lin et al. 2017; Du et al. 2023) that amplify learning signals from underrepresented classes. Given the limited attention to

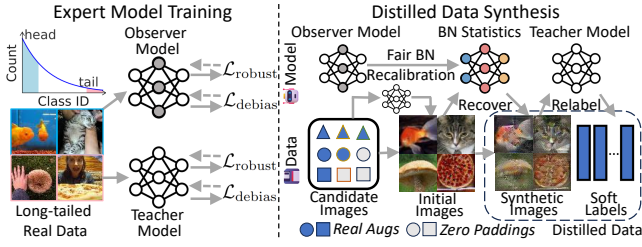


Figure 2: Illustration of our pipeline. We first train robust and debiased observer and teacher models on real data to enable reliable statistical alignment and semantic guidance. A candidate image pool tailored for long-tailed distributions is then used for teacher-guided initialization. Synthetic images are subsequently recovered via fair BN alignment and relabeled using teacher model to ensure semantic consistency.

long-tailed dataset distillation, we draw conceptual insights from the broader long-tailed recognition literature to debias both the observer and teacher models, thereby enabling effective distillation under severe class imbalance.

Methods

Problem Statement

Long-tailed dataset distillation aims to construct a small, class-balanced synthetic dataset $\mathcal{S} = \{(\mathbf{x}_s^i, y_s^i, \tilde{\mathbf{y}}_s^i)\}_{i=1}^{|\mathcal{S}|}$ from a long-tailed real dataset $\mathcal{D} = \{(\mathbf{x}^i, y^i)\}_{i=1}^{|\mathcal{D}|}$, where $y^i \in \{0, \dots, C-1\}$ and $|\mathcal{S}| \ll |\mathcal{D}|$, such that models trained on \mathcal{S} achieve strong generalization performance on a class-balanced real test set \mathcal{T} . Here, y_s^i and $\tilde{\mathbf{y}}_s^i$ denote the hard and soft labels, respectively. Let $\mathcal{D}_c = \{(\mathbf{x}^i, y^i) \in \mathcal{D} \mid y^i = c\}$ denote the subset of class c , where $|\mathcal{D}_0| > |\mathcal{D}_1| > \dots > |\mathcal{D}_{C-1}|$ and $|\mathcal{D}_0| \gg |\mathcal{D}_{C-1}|$. The synthetic (distilled) dataset is class-balanced such that $\mathcal{S}_c = \{(\mathbf{x}_s^i, y_s^i, \tilde{\mathbf{y}}_s^i) \in \mathcal{S} \mid y_s^i = c\}$ satisfies $|\mathcal{S}_0| = |\mathcal{S}_1| = \dots = |\mathcal{S}_{C-1}| = \text{IPC}$.

To address the problem, we propose a uni-level statistical alignment framework for unbiased recovery and soft relabeling. Our approach departs from traditional trajectory-matching methods, whose fundamental limitations are detailed in the subsequent section. The success of our framework relies on three core strategies: expert model debiasing, fair BN statistics recalibration, and confidence-guided data initialization. The entire pipeline is illustrated in Fig. 2.

Drawback of Trajectory-Matching-Based Methods

Trajectory-based dataset distillation methods typically begin by training an expert model θ_R on the real dataset \mathcal{D} , followed by a bi-level optimization process, where the trajectory student model θ_S is trained on synthetic data, and the synthetic dataset \mathcal{S} is updated by matching its induced trajectory to that of the expert θ_R . This objective is often formulated as minimizing the discrepancy between the training trajectories of the student and expert models:

$$\mathcal{L}(\mathcal{S}) = \|F(\mathcal{S}, \theta_S) - F(\mathcal{D}, \theta_R)\|_2, \quad (1)$$

where F denotes training trajectories. However, when the expert model θ_R is trained on a long-tailed dataset, its internal representations are inevitably prone to class imbalance

if no proper intervention is enforced. Optimizing the student model to mimic such an expert causes the distilled dataset to inherit the bias, leading to over-emphasis on head-class semantics and under-representation of minority ones.

Although DAMED (Zhao et al. 2025) attempts to simulate imbalanced training dynamics within the student to reduce trajectory mismatch, it relies on an representation expert trained on imbalanced data without debiasing. As a result, the distilled dataset inherits the expert’s representational bias. More broadly, the trajectory-based methods are difficult to balance between explicit debiasing and strict trajectory matching. Pre-distillation adjustments such as reweighting or logit correction alter the expert’s optimization path, breaking the premise of trajectory matching. Meanwhile, post-hoc debiasing is impractical, as these methods only reproduce parameter evolution and lack fine-grained control over per-class representation quality. In addition to these limitations, such methods incur substantial computational overhead due to the bi-level nature of the optimization, multi-step training trajectory simulation, and backpropagation through the unrolled computation graph.

Uni-level Long-tailed Dataset Distillation

Our uni-level distillation framework avoids trajectory simulation and instead enables explicit debiasing through unbiased synthetic image recovery and unbiased soft relabeling. To enable unbiased recovery of the synthetic images \mathcal{S} , we adopt a feature-statistics alignment objective to match their internal distribution to that of real data, without relying on repeated access to real samples during training. Specifically, we leverage BN statistics as lightweight summaries of feature distributions. We train an expert model θ_R (also referred to as the observer model) on real data and freeze it to serve as a stable source of BN statistics for both real and synthetic inputs. The corresponding alignment loss is defined as:

$$\mathcal{L}(\mathcal{S}) = \sum_{l=1}^L \mathbf{D}_l^\mu(\mathcal{S}, \mathcal{D}; \theta_R) + \mathbf{D}_l^\sigma(\mathcal{S}, \mathcal{D}; \theta_R), \quad (2)$$

where l indexes the monitored BN layers, and \mathbf{D}_l^μ , \mathbf{D}_l^σ represent the ℓ_2 distances between the mean (μ_l) and variance (σ_l) of synthetic and real features in layer l :

$$\mathbf{D}_l^\mu(\mathcal{S}, \mathcal{D}; \theta_R) = \|\mu_l(\mathcal{S}; \theta_R) - \mu_l(\mathcal{D}; \theta_R)\|_2 \quad (3)$$

$$\mathbf{D}_l^\sigma(\mathcal{S}, \mathcal{D}; \theta_R) = \|\sigma_l(\mathcal{S}; \theta_R) - \sigma_l(\mathcal{D}; \theta_R)\|_2 \quad (4)$$

The success of this alignment in long-tailed scenarios hinges on three factors: (1) a robust and debiased observer model θ_R to serve as a reliable anchor for extracting statistics; (2) fair and accurate estimation of BN statistics (μ_l , σ_l) to handle both the intra-class and inter-class skew; (3) a diverse and semantically meaningful initialization of \mathcal{S} to effectively cover tail classes. Although unbiased recovery enforces distribution-level consistency, it does not directly enforce semantic discriminability. To address this, we incorporate soft relabeling via an unbiased teacher model θ_T , which produces balanced and informative supervision for \mathcal{S} . The resulting soft labels capture inter-class relationships and provide dense guidance that is critical under class imbalance.

Debiasing the Observer and Teacher Models To ensure effective supervision in long-tailed settings, both the observer model θ_R and the teacher model θ_T must be trained to mitigate class imbalance and enhance robustness. The observer model guides distribution-level alignment via BN statistics, while the teacher model provides soft label (Cui et al. 2024b,a) supervision for semantic guidance. If either model is biased, the distilled dataset would inherit skewed feature distributions or inaccurate soft labels. We address these challenges through two complementary loss functions.

First, to enhance the robustness of both θ_R and θ_T , we use a symmetric mixture consistency loss that aligns representations across different mixed-label augmented images:

$$\mathcal{L}_{\text{robust}} = - \sum_{i=1}^2 \cos(\mathbf{z}_i, \text{sg}(\mathbf{p}_{\bar{i}})), \quad (5)$$

where $\bar{i} = 3 - i$ denotes the index of the opposite view. For the i -th augmented view, \mathbf{z}_i is obtained by projecting the encoder output through a shared linear layer that maps features into a representation space. \mathbf{p}_i is further transformed from \mathbf{z}_i via an additional prediction layer. The stop-gradient operator $\text{sg}(\cdot)$ ensures one-sided alignment for each term.

Second, to explicitly debias class imbalance during training, we adopt a dynamically-weighted rebalancing strategy that gradually emphasizes minority classes:

$$\mathcal{L}_{\text{debias}} = \alpha \sum_{k=0}^{C-1} \frac{-(r_k)^{-q} y_k \log p_k}{\sum_{j=0}^{C-1} (r_j)^{-q}} - \beta \sum_{k=0}^{C-1} y_k \log p_k, \quad (6)$$

where y_k is the mixed target probability, p_k is the predicted softmax output, r_k is the sample frequency of class k , and q controls the sharpness of reweighting. We set $\alpha = (\frac{t}{T})^2$ and $\beta = 1 - \alpha$, where t and T denote the current and total training iterations. This schedule gradually shifts focus toward minority classes while preserving early-stage stability.

Both θ_R and θ_T are trained with the combined loss:

$$\mathcal{L} = \gamma_1 \mathcal{L}_{\text{robust}} + \gamma_2 \mathcal{L}_{\text{debias}}, \quad (7)$$

where γ_1 and γ_2 are scalar weights for balancing.

Fair Recalibration of BN Statistics Accurate and fair BN statistics are critical in our framework, as they serve as alignment targets for image recovery. However, under the standard exponential moving average update with fixed momentum, the running estimates become biased due to unequal sample contributions: recent batches dominate the statistics, while earlier ones are quickly forgotten. This effect is particularly problematic in long-tailed settings, where each tail-class sample carries high representational value and must contribute equally to the accumulated statistics. These limitations motivate our post-hoc recalibration strategy.

To address this, we recalibrate BN statistics for the observer model θ_R after training, using a fair, sample-balanced and class-balanced accumulation strategy. Specifically, we freeze model parameters and perform a full forward pass over the real dataset, processing one batch at a time. For each monitored BN layer l , the class- c mean is updated at batch index t using a dynamically adjusted momentum α_t^c :

$$\mu_{l,t}^c = (1 - \alpha_t^c) \cdot \mu_{l,t-1}^c + \alpha_t^c \cdot \hat{\mu}_{l,t}^c, \quad \alpha_t^c = \frac{B_t^c}{N_{t-1}^c + B_t^c}, \quad (8)$$

where $\mu_{l,t}^c$ denotes the cumulative mean of class- c samples up to step t , $\hat{\mu}_{l,t}^c$ is their batch mean, B_t^c is the number of class- c samples in that batch, and N_{t-1}^c is the cumulative count before step t . Unlike standard BN with fixed momentum, our formulation introduces a dynamic momentum coefficient α_t^c , ensuring that each sample contributes equally to the final statistics regardless of its temporal order.

After the full pass, we aggregate global BN statistics by uniformly averaging across all class-specific means to ensure class-balanced statistics. Let T denote the total number of batches in the forward pass. The global mean is given by:

$$\mu_l(\mathcal{D}; \theta_R) = \frac{1}{C} \sum_{c=0}^{C-1} \mu_{l,T}^c(\mathcal{D}; \theta_R), \quad (9)$$

and similarly for variance σ . This two-stage strategy mitigates both intra-class bias and inter-class bias, thereby serving as reliable supervision signals for statistical alignment.

Confidence-guided Multi-round Initialization Initialization primarily governs the diversity of final synthetic images, while also playing a supportive role in promoting stable optimization under long-tailed distributions. Traditional initialization strategies typically rely on either sampling real images or random noise. However, random initialization often leads to poor convergence and degraded downstream performance. In highly imbalanced settings, directly sampling real images becomes infeasible, as tail classes often contain too few samples to provide adequate initialization.

To overcome this limitation, we propose a confidence-guided, multi-round initialization strategy tailored for long-tailed distributions. Specifically, we generate multiple augmentations (e.g., crops) per real image and score them via the teacher model θ_T using negative cross-entropy loss. These augmentations are stored in a class-wise candidate pool. In each round, every real image contributes its most confident unused augmentation to a temporary selection pool. If the total number of candidates exceeds the remaining slots for that class, we select the top-scoring augmentations; otherwise, we retain all. This process repeats until each class reaches its target IPC. This strategy ensures high-confidence selection while maintaining sample-level diversity across varying class sizes. To maintain structural consistency across classes, we insert zero-initialized placeholders for classes with fewer real samples than the largest class. These placeholders are excluded from the augmentation and selection process, ensuring all synthetic samples are semantically meaningful.

Recovery, Relabeling and Evaluation Given an initialized distilled set $\mathcal{S}_{\text{init}}$, an observer model θ_R , a teacher model θ_T , and precomputed BN statistics from the real dataset, we optimize the distilled set \mathcal{S} (initialized from $\mathcal{S}_{\text{init}}$) through statistical recovery using θ_R and real-data BN statistics ($\mu_l(\mathcal{D}; \theta_R)$, $\sigma_l(\mathcal{D}, \theta_R)$), followed by soft relabeling via θ_T .

For statistical recovery, we pass all images in \mathcal{S} through the observer model θ_R to compute per-layer BN statistics $\mu_l(\mathcal{S}; \theta_R)$ and $\sigma_l(\mathcal{S}; \theta_R)$, and align them with the precomputed real-data statistics $\mu_l(\mathcal{D}; \theta_R)$ and $\sigma_l(\mathcal{D}; \theta_R)$ as defined in Eq. (2). In practice, we also perform class-wise alignment

Dataset	CIFAR-10-LT								
	10			50			100		
Imbalance Factor	10	20	50	10	20	50	10	20	50
Random	32.5±2.2	39.6±0.9	51.9±1.5	33.2±0.4	42.0±1.3	51.6±1.3	34.4±2.0	41.4±0.7	52.6±0.5
K-Center (ICLR'18)	21.9±0.8	24.2±0.8	31.7±0.9	17.8±0.2	20.8±0.5	26.1±0.2	16.2±0.5	19.0±1.0	24.2±1.2
Graph-Cut (ALT'21)	28.7±0.9	34.2±1.0	40.6±1.0	24.2±0.7	28.6±0.8	33.9±0.4	22.9±0.9	26.0±0.5	33.3±1.0
DC (ICLR'21)	37.9±0.9	38.5±0.9	37.4±1.4	37.3±0.9	38.8±1.0	35.8±1.2	36.7±0.8	38.1±1.0	35.3±1.4
MTT (CVPR'22)	58.0±0.8	59.5±0.4	62.0±0.9	45.8±1.4	49.9±0.8	53.6±0.5	37.7±0.6	41.6±1.1	47.8±1.1
DREAM (ICCV'23)	34.6±0.6	42.2±1.5	50.5±0.7	30.8±0.6	38.4±0.3	45.5±0.9	30.8±1.7	34.9±0.8	42.2±0.8
IDM (CVPR'23)	54.8±0.4	57.1±0.3	60.1±0.3	51.9±0.7	53.3±0.6	56.1±0.4	49.8±0.6	50.9±0.5	53.1±0.4
Minimax (CVPR'24)	29.2±0.5	28.5±0.6	39.9±0.1	18.4±0.3	22.5±0.2	25.2±0.2	19.9±0.4	23.3±0.2	28.0±0.6
DATM (ICLR'24)	57.2±0.4	60.4±0.2	66.7±0.6	41.6±0.2	43.4±0.3	50.3±0.2	37.3±0.2	38.9±0.1	44.3±0.1
RDED* (CVPR'24)	46.2±0.2	55.5±0.3	62.0±0.3	43.3±0.3	46.1±0.3	53.2±0.3	41.2±0.3	45.3±0.2	49.9±0.3
EDC* (NeurIPS'24)	56.4±0.2	62.7±0.3	68.5±0.1	55.2±0.2	60.2±0.2	64.1±0.3	53.2±0.5	57.4±0.3	60.5±0.4
DAMED (CVPR'25)	58.1±0.3	63.0±1.0	70.5±0.4	54.2±1.0	59.4±0.7	65.8±0.2	53.4±0.1	58.2±0.6	64.0±0.9
Ours	63.6±0.5	68.3±0.2	74.1±0.2	62.9±0.2	67.3±0.3	70.6±0.1	62.7±0.1	66.4±0.2	68.8±0.4

Table 1: Quantitative comparisons on CIFAR-10-LT. Rows marked with * indicate experiments conducted using open-source implementations with adaptations. Other results are taken from DAMED (Zhao et al. 2025).

Dataset	CIFAR-100-LT								
	10			20			50		
Imbalance Factor	10	20	50	10	20	50	10	20	50
Random	14.2±0.6	21.7±0.6	32.1±0.6	15.0±0.3	21.6±0.5	30.5±0.5	13.4±0.5	20.6±0.6	26.9±0.5
K-Center (ICLR'18)	10.7±0.9	15.9±1.0	24.8±0.2	10.0±0.5	15.1±0.6	23.8±0.3	8.7±0.6	12.4±0.6	19.8±0.5
Graph-Cut (ALT'21)	16.9±0.3	22.2±0.4	29.9±0.4	16.0±0.5	20.7±0.5	28.7±0.3	13.1±0.5	17.2±0.6	24.8±0.5
DC (ICLR'21)	24.0±0.3	27.4±0.3	27.4±0.3	23.2±0.3	26.2±0.3	27.4±0.3	19.8±0.4	22.7±0.4	25.9±0.3
MTT (CVPR'22)	14.3±0.1	16.7±0.2	13.8±0.2	12.6±0.3	15.0±0.2	10.6±0.5	8.2±0.5	11.2±0.5	6.3±0.5
DREAM (ICCV'23)	10.1±0.4	12.0±1.0	13.1±0.4	9.4±0.4	10.3±0.6	12.3±0.3	6.8±0.5	7.6±0.5	9.8±0.5
DATM (ICLR'24)	28.2±0.4	34.1±0.2	31.6±0.1	25.3±0.3	27.2±0.1	27.1±0.1	22.2±0.4	19.8±0.4	22.9±0.4
RDED* (CVPR'24)	30.5±0.4	36.0±0.2	37.6±0.1	32.1±0.3	33.9±0.2	-	28.0±0.2	-	-
EDC* (NeurIPS'24)	31.5±0.3	36.6±0.1	38.5±0.2	32.7±0.3	34.2±0.2	-	30.2±0.2	-	-
DAMED (CVPR'25)	31.5±0.2	37.5±0.4	40.0±0.1	31.4±0.5	35.1±0.4	37.0±0.7	29.8±0.3	31.9±0.5	33.2±0.5
Ours	47.1±0.1	48.8±0.1	49.9±0.3	45.5±0.4	46.9±0.2	48.1±0.2	42.1±0.1	43.4±0.2	44.2±0.1

Table 2: Quantitative comparisons on CIFAR-100-LT. Rows marked with * indicate experiments conducted using open-source code with adaptations. Other results are taken from DAMED (Zhao et al. 2025). ‘-’ means distillation could not be executed.

by computing statistics separately for each class in \mathcal{S} and matching them to their corresponding real-data counterparts.

Once recovery is complete, each synthetic image is relabeled by passing it through θ_T . Specifically, for each synthetic image \mathbf{x}_s^i , we obtain a soft label $\tilde{\mathbf{y}}_s^i = \theta_T(\mathbf{x}_s^i)$, while retaining the original one-hot label \mathbf{y}_s^i for supervision.

To evaluate the quality of \mathcal{S} , we follow (Shao et al. 2024b) and train a student model s from scratch on \mathcal{S} using a dual-objective loss that combines one-hot labels and soft targets:

$$\mathcal{L}_{\text{match}} = \kappa_1 \cdot \mathcal{L}_{\text{CE}}(s(\mathbf{x}_s^i), \mathbf{y}_s^i) + \kappa_2 \cdot \|\tilde{\mathbf{y}}_s^i - s(\mathbf{x}_s^i)\|_2^2, \quad (10)$$

where $\mathcal{L}_{\text{CE}}(\cdot, \cdot)$ denotes the cross-entropy loss, $\|\cdot\|_2$ represents l_2 distance, κ_1 and κ_2 are weights for balancing.

Experiments

Experimental Settings

Datasets To comprehensively evaluate the effectiveness of our method under varying degrees of class imbalance, we conduct experiments on four long-tailed benchmarks across different scales: CIFAR-10-LT, CIFAR-100-LT (Cui et al. 2019), Tiny-ImageNet-LT (Le and Yang 2015), and

ImageNet-LT (Liu et al. 2019). These datasets are constructed by applying an exponential decay sampling strategy to the balanced datasets, following the protocols of prior works (Zhao et al. 2025; Cui et al. 2019). Specifically, the number of samples for class c is determined by $|\mathcal{D}_c| = |\mathcal{D}_0|\phi^c$, where $\phi^c = \beta^{-(c/(C-1))}$, with β representing the imbalance factor (IF). A larger value of β indicates more severe class imbalance. We evaluate under multiple imbalance ratios and varying images-per-class (IPC) budgets to reflect both mild and extreme long-tailed settings.

Network Architectures Following the settings established in DAMED (Zhao et al. 2025), we adopt a depth-3 ConvNet as the student model for CIFAR-10-LT and CIFAR-100-LT, and a depth-4 ConvNet for Tiny-ImageNet-LT and ImageNet-LT. Given the superior representational capacity of ResNet architectures for large-scale datasets, we additionally evaluate ResNet-50 on ImageNet-LT under highly imbalanced scenarios. During evaluation, all student models are trained for 1000 epochs on the distilled dataset. All experiments were repeated five times for fairness, and conducted primarily on a single NVIDIA RTX 3090 GPU.

Dataset	Tiny-ImageNet-LT			
	10		20	
IPC	10	20	10	20
Random	7.4±0.2	13.5±0.4	7.6±0.1	13.2±0.4
K-Center	10.4±1.5	11.3±1.2	12.6±1.4	9.9±2.2
Graph-Cut	9.8±0.7	5.6±0.6	3.4±0.8	5.0±1.0
MTT	11.1±0.2	18.1±0.2	7.7±0.1	14.7±0.2
DREAM	5.4±0.3	6.8±0.1	4.8±0.1	6.0±0.2
DATM	21.3±0.1	14.5±0.1	14.0±0.6	19.0±0.3
RDED*	22.7±0.3	23.3±0.2	21.0±0.3	21.8±0.3
EDC*	23.9±0.6	25.2±0.2	22.1±0.5	23.7±0.4
DAMED	26.0±0.3	27.9±0.2	23.6±0.3	25.5±0.3
Ours	37.8±0.4	38.9±0.2	36.1±0.1	37.0±0.1

Table 3: Quantitative comparisons on Tiny-ImageNet-LT. Rows marked with * indicate experiments conducted using open-source implementations with adaptations. Other baseline results are taken from DAMED (Zhao et al. 2025).

Dataset	ImageNet-LT			
	5		10	
IPC	10	20	10	20
Random	3.9±0.1	7.0±0.1	3.9±0.1	6.8±0.1
G-VBSM	-	1.0±0.1	-	1.0±0.1
TESLA	3.0±0.1	-	2.7±0.1	-
DATM	7.4±0.1	8.1±0.2	7.9±0.1	8.2±0.1
SRe ² L	6.7±0.1	10.1±0.1	7.7±0.3	10.9±1.0
DAMED	20.8±0.2	21.0±0.1	20.3±0.1	20.7±0.1
Ours	24.7±0.4	25.0±0.3	23.5±0.2	24.1±0.4

Table 4: Quantitative comparisons on ImageNet-LT. Baseline results are taken from DAMED (Zhao et al. 2025).

Baselines We compare our method with a diverse set of representative baselines, including coreset selection methods such as Random, K-Center Greedy (Sener and Savarese 2017), and Graph Cut (Iyer et al. 2021); gradient-matching-based methods including DC (Zhao, Mopuri, and Bilen 2020) and DREAM (Liu et al. 2023); distribution-matching-based methods such as CAFE (Wang et al. 2022) and IDM (Zhao et al. 2023); trajectory-matching-based methods including MTT (Cazenavette et al. 2022), DATM (Guo et al. 2024), TESLA (Cui et al. 2023), and DAMED (Zhao et al. 2025); uni-level optimization methods including SRe²L (Yin, Xing, and Shen 2023), RDED (Sun et al. 2024) and EDC (Shao et al. 2024b); and generative-model-based methods such as Minimax (Gu et al. 2024).

Results and Discussions

Main Results We conduct comprehensive evaluations covering a broad spectrum of IF and IPC configurations, covering datasets of varying complexity. As shown in Figs. 1, 2, 3, and 4, our method consistently outperforms strong baselines under all evaluated settings. While DAMED (Zhao et al. 2025) yields student performance closely matching those of its biased expert—effectively saturating its performance ceiling—our method explicitly mitigates expert bias, enabling the distilled data to supervise more accurate and generalizable student models, thereby raising the achievable upper bound. By aligning

Dataset	Tiny-ImageNet-LT		ImageNet-LT (ResNet)	
	100		256	
IPC	10	20	10	20
MTT	5.6±0.2	7.7±0.3	-	-
DATM	10.1±0.3	11.9±0.1	6.2±0.3	6.7±0.2
DAMED	17.1±0.8	18.0±0.5	17.2±0.2	17.9±0.2
Ours	28.5±0.2	29.8±0.1	48.2±0.7	48.9±0.2

Table 5: Quantitative comparisons under severe class imbalance, with ResNet-50 used for ImageNet-LT evaluation.

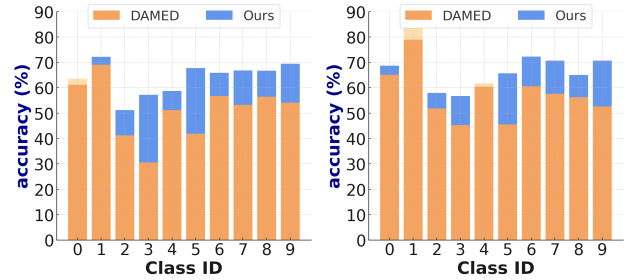


Figure 3: Class-wise accuracy comparison on CIFAR-10-LT (IF = 100). The left figure corresponds to IPC = 10, and the right figure to IPC = 20. For both settings, we compare our method with DAMED (Zhao et al. 2025).

model representations in a class-balanced manner and recalibrating balanced BN statistics, our approach avoids the typical overfitting to head classes and promotes equitable learning across classes and training samples. Our debiasing mechanisms allow the distilled dataset to preserve both structural fidelity and semantic completeness, making our method broadly effective across datasets of different scales.

Results under Highly Imbalanced Settings Table 5 summarizes the results under highly imbalanced scenarios. These settings pose significant challenges for dataset distillation, particularly when the number of available real images falls below the target IPC for certain classes. Under such constraints, some prior methods become inapplicable; for example, EDC’s initialization and RDED’s sampling mechanism fail due to insufficient tail-class examples. Other baselines also struggle to achieve competitive performance, often due to biased representation learning or optimization instability. In contrast, our method consistently achieves stronger performance across all tested configurations. Notably, under

Dataset	CIFAR-10-LT	CIFAR-100-LT	Tiny-IN-LT
IF	100	50	100
CAFE	15.2±0.8	11.3±0.9	-
DREAM	15.3±0.7	11.7±0.7	1.9±0.2
DATM	20.4±0.5	12.8±0.4	5.5±0.2
G-VBSM	4.1±0.2	1.8±0.2	2.3±0.1
RDED	21.8±0.4	15.6±0.2	9.1±0.3
DAMED	24.1±0.5	7.8±0.3	6.0±0.3
Ours	44.8±0.3	31.8±0.1	20.1±0.2

Table 6: Performance comparison under extreme settings with only 1 image per class (IPC = 1).

Method	ConvNet-3	VGG-11	ResNet-18	AlexNet
MTT	37.7	27.4	29.6	26.3
DATM	37.3	34.4	35.1	36.8
RDED	41.2	41.6	37.5	28.2
DAMED	53.4	29.7	43.6	37.9
Ours	62.7	64.6	58.7	56.5

Table 7: Cross-architecture evaluation on CIFAR-10-LT under an imbalance factor of 100, using 10 images per class.

Method	IPC=10	IPC=20	IPC=50
w/o Model Debiasing	31.7	32.3	32.8
w/o Statistics Recalibration	40.9	41.8	42.1
w/o Adapted Initialization	40.8	-	-
Full	42.1	43.4	44.2

Table 8: Ablation study on CIFAR-100-LT (IF=50).

IF = 256 with ResNet-50 as the evaluation model, our distilled set for ImageNet-LT not only outperforms those generated by all competing methods under the same imbalanced setting, but also surpasses several methods whose distilled sets are obtained using the full, balanced ImageNet-1K.

Results under Extremely Low IPC Settings We further evaluate our method under severely compressed distillation regimes, where only one synthetic image per class is retained. As shown in Table 6, our method achieves over $2\times$ accuracy improvement than all baselines on most datasets. This strong performance stems from two key factors. First, fair BN statistic recalibration ensures that even a single synthetic image per class reflects accurate distribution-level information, providing reliable supervision despite minimal data capacity. Additionally, unbiased soft labels provide semantic guidance that compensates for the limited representational expressiveness of low-IPC synthetic samples. Together, these mechanisms enable our method to remain robust under extreme distillation constraints.

Cross-Architecture Performance To evaluate its generalization across architectures, we train multiple student models with different structures on the same distilled dataset. As shown in Table 7, our method consistently outperforms existing approaches across four representative evaluation backbones. Notably, baseline methods often show significant accuracy variation across architectures, while our distilled data supports uniformly strong performance. These results suggest that our method captures semantically meaningful and transferable patterns, facilitating generalizable learning across diverse student architectures.

Class-wise Accuracy for Long-tailed Dataset Fig. 3 compares class-wise accuracy between DAMED (Zhao et al. 2025) and our method. DAMED underperforms on tail classes due to its biased expert training, which fails to preserve rare-class semantics. Moreover, the frequency-adjusted loss used during trajectory matching overlooks mid-frequency classes, resulting in suppressed performance. In contrast, our method avoids these issues by first training a debiased expert model and then aligning fair BN statistics.

Dataset Stage	CIFAR-10-LT (IF=100)		CIFAR-100-LT (IF=50)	
	EMT	DDS	EMT	DDS
DAMED	31388s	30141s	26269s	29328s
Ours	2395s	118s	2183s	273s

Table 9: Runtime comparison on expert model training (EMT) and distilled data synthesis (DDS) between DAMED (Zhao et al. 2025) and our approach (IPC=1).

Dataset	CIFAR-100-LT			Tiny-ImageNet-LT	
	IPC=10	IPC=20	IPC=50	IPC=1	IPC=10
DAMED	10.2	12.1	15.8	10.2	>24.0
Ours	3.1	3.1	3.1	6.1	6.1

Table 10: Peak GPU memory (GB) comparison between DAMED (Zhao et al. 2025) and our approach.

Ablation on Different Components As shown in Table 8, each component contributes critically to the success of our approach. The model debiasing strategy preserves tail-class semantics without degrading head or mid-frequency performance, thereby raising the overall performance ceiling. The recalibrated BN statistics ensure that each training sample, particularly those from low-shot categories, contributes fairly to the accumulated representation. The initialization strategy offers diverse, class-representative starting points for synthetic images, even when real data per class is scarce.

Computational Efficiency We assess the computational efficiency of our method by comparing both runtime and peak GPU memory usage with DAMED (Zhao et al. 2025), the only existing approach explicitly designed for long-tailed dataset distillation. As shown in Table 9, our method substantially reduces the computation time required for both expert model training and distilled data synthesis. Specifically, across both datasets, the total runtime of our pipeline is less than one-twentieth that of DAMED. In addition to faster execution, our method also exhibits more favorable memory behavior. As shown in Table 10, DAMED’s GPU memory usage grows rapidly with IPC, constraining its applicability at higher values. In contrast, our method maintains constant memory usage regardless of IPC, allowing for stable and efficient execution across a wide range of settings.

Conclusion

We present a uni-level framework for long-tailed dataset distillation, explicitly designed to address the representation bias and inefficiency inherent in prior methods. We enhance distillation effectiveness under class imbalance via three key components: expert model debiasing, BN statistics recalibration, and confidence-aware initialization. Extensive experiments demonstrate that our method consistently outperforms existing baselines across a wide spectrum of IF and IPC settings, including highly imbalanced and low-sample regimes, showcasing strong robustness and generalizability.

Broader Impact Our method can potentially be extended to multi-domain or federated dataset distillation, where data imbalance naturally occurs across clients or domains.

Acknowledgments

This work was supported by the GPU cluster of the MCC Lab, Information Science and Technology Institution, and the Supercomputing Center of USTC.

References

- Cazenavette, G.; Wang, T.; Torralba, A.; Efros, A. A.; and Zhu, J.-Y. 2022. Dataset distillation by matching training trajectories. In *CVPR*.
- Chai, Z.; Gao, Z.; Lin, Y.; Zhao, C.; Yu, X.; and Xie, Z. 2024. Adaptive Backdoor Attacks Against Dataset Distillation for Federated Learning. In *ICC*, 4614–4619. IEEE.
- Chen, M.; Du, J.; Huang, B.; Wang, Y.; Zhang, X.; and Wang, W. 2025. Influence-Guided Diffusion for Dataset Distillation. In *ICLR*.
- Cheng, L.; Chen, K.; Li, J.; Tang, S.; Zhang, S.; and Wang, M. 2024. Dataset Distillers Are Good Label Denoisers In the Wild. *arXiv preprint arXiv:2411.11924*.
- Cui, J.; Wang, R.; Si, S.; and Hsieh, C.-J. 2023. Scaling up dataset distillation to imagenet-1k with constant memory. 6565–6590.
- Cui, X.; Qin, Y.; Gao, Y.; Zhang, E.; Xu, Z.; Wu, T.; Li, K.; Sun, X.; Zhou, W.; and Li, H. 2024a. Sinkhorn Distance Minimization for Knowledge Distillation. In *LREC-COLING*, 14846–14858.
- Cui, X.; Qin, Y.; Zhou, W.; Li, H.; and Li, H. 2025a. OPTICAL: Leveraging Optimal Transport for Contribution Allocation in Dataset Distillation. In *CVPR*.
- Cui, X.; Qin, Y.; Zhou, W.; Li, H.; and Li, H. 2025b. Optimizing Distributional Geometry Alignment with Optimal Transport for Generative Dataset Distillation. In *NeurIPS*.
- Cui, X.; Zhu, M.; Qin, Y.; Xie, L.; Zhou, W.; and Li, H. 2024b. Multi-Level Optimal Transport for Universal Cross-Tokenizer Knowledge Distillation on Language Models. *arXiv preprint arXiv:2412.14528*.
- Cui, Y.; Jia, M.; Lin, T.-Y.; Song, Y.; and Belongie, S. 2019. Class-balanced loss based on effective number of samples. In *CVPR*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *CVPR*.
- Du, C.; Han, Y.; and Huang, G. 2024. SimPro: A Simple Probabilistic Framework Towards Realistic Long-Tailed Semi-Supervised Learning. In *ICML*.
- Du, F.; Yang, P.; Jia, Q.; Nan, F.; Chen, X.; and Yang, Y. 2023. Global and Local Mixture Consistency Cumulative Learning for Long-tailed Visual Recognitions. In *CVPR*.
- Gu, J.; Vahidian, S.; Kungurtsev, V.; Wang, H.; Jiang, W.; You, Y.; and Chen, Y. 2024. Efficient dataset distillation via minimax diffusion. In *CVPR*.
- Guo, Z.; Wang, K.; Cazenavette, G.; Li, H.; Zhang, K.; and You, Y. 2024. Towards lossless dataset distillation via difficulty-aligned trajectory matching. In *ICLR*.
- He, J. 2024. Gradient Reweighting: Towards Imbalanced Class-Incremental Learning. In *CVPR*, 16668–16677.
- Iyer, R.; Khargoankar, N.; Bilmes, J.; and Asanani, H. 2021. Submodular combinatorial information measures with applications in machine learning. In *International Conference on Algorithmic Learning Theory*.
- Jia, Y.; Vahidian, S.; Sun, J.; Zhang, J.; Kungurtsev, V.; Gong, N. Z.; and Chen, Y. 2024. Unlocking the potential of federated learning: The symphony of dataset distillation via deep generative latents. In *ECCV*, 18–33. Springer.
- Khorram, S.; Jiang, M.; Shahbazi, M.; Danesh, M. H.; and Fuxin, L. 2024. Taming the Tail in Class-Conditional GANs: Knowledge Sharing via Unconditional Training at Lower Resolutions. In *CVPR*, 7580–7590.
- Krizhevsky, A. 2012. Learning Multiple Layers of Features from Tiny Images. *University of Toronto*.
- Le, Y.; and Yang, X. 2015. Tiny imagenet visual recognition challenge. *CS 231N*.
- Li, X.; Yang, J.; Wang, Y.; Lin, W.; and Song, Y. 2021. UniMix: Tail-Aware Mixup for Long-Tailed Visual Recognition. In *ICCV*.
- Li, Z.; and Jia, Y. 2025. ConMix: Contrastive Mixup at Representation Level for Long-tailed Deep Clustering. In *ICLR*.
- Lin, T.; Goyal, P.; Girshick, R. B.; He, K.; and Dollár, P. 2017. Focal Loss for Dense Object Detection. In *ICCV*, 2999–3007.
- Liu, P.; and Du, J. 2025. The Evolution of Dataset Distillation: Toward Scalable and Generalizable Solutions. *arXiv preprint arXiv:2502.05673*.
- Liu, Y.; Gu, J.; Wang, K.; Zhu, Z.; Jiang, W.; and You, Y. 2023. Dream: Efficient dataset distillation by representative matching. In *ICCV*.
- Liu, Z.; Miao, Z.; Zhan, X.; Wang, J.; Gong, B.; and Yu, S. X. 2019. Large-scale long-tailed recognition in an open world. In *CVPR*.
- Sener, O.; and Savarese, S. 2017. Active learning for convolutional neural networks: A core-set approach. *arXiv preprint arXiv:1708.00489*.
- Shao, J.; Zhu, K.; Zhang, H.; and Wu, J. 2024a. DiffuLT: Diffusion for Long-Tail Recognition Without External Knowledge. In *NeurIPS*.
- Shao, S.; Zhou, Z.; Chen, H.; and Shen, Z. 2024b. Elucidating the Design Space of Dataset Condensation. In *NeurIPS*.
- Shi, X.; Dong, Y.; and Shen, X. 2021. CMO: Contextual Mixup for Long-Tailed Recognition. In *CVPR*.
- Sun, P.; Shi, B.; Yu, D.; and Lin, T. 2024. On the diversity and realism of distilled dataset: An efficient dataset distillation paradigm. In *CVPR*.
- Wang, B.; Wang, P.; Xu, W.; Wang, X.; Zhang, Y.; Wang, K.; and Wang, Y. 2024. Kill Two Birds with One Stone: Rethinking Data Augmentation for Deep Long-tailed Learning. In *CVPR*.
- Wang, K.; Li, Z.; Cheng, Z.-Q.; Khaki, S.; Sajedi, A.; Vedantam, R.; Plataniotis, K. N.; Hauptmann, A.; and You, Y. 2025. Emphasizing Discriminative Features for Dataset Distillation in Complex Scenarios. In *CVPR*.

Wang, K.; Zhao, B.; Peng, X.; Zhu, Z.; Yang, S.; Wang, S.; Huang, G.; Bilen, H.; Wang, X.; and You, Y. 2022. Cafe: Learning to condense dataset by aligning features. In *CVPR*.

Wang, T.; Zhu, J.-Y.; Torralba, A.; and Efros, A. A. 2018. Dataset distillation. *arXiv preprint arXiv:1811.10959*.

Xiong, H.; and Yao, A. 2024. Deep Imbalanced Regression via Hierarchical Classification Adjustment. In *CVPR*, 23721–23730.

Yin, Z.; Xing, E.; and Shen, Z. 2023. Squeeze, recover and relabel: Dataset condensation at imagenet scale from a new perspective. In *NeurIPS*.

Yu, R.; Liu, S.; and Wang, X. 2023. Dataset Distillation: A Comprehensive Review. *arXiv preprint arXiv:2301.07014*.

Zhang, C.; Almpanidis, G.; Fan, G.; Deng, B.; Zhang, Y.; Liu, J.; Kamel, A.; Soda, P.; and Gama, J. 2025. A systematic review on long-tailed learning. *TNNLS*.

Zhang, H.; Cisse, M.; Dauphin, Y. N.; and Lopez-Paz, D. 2018. mixup: Beyond Empirical Risk Minimization. In *ICLR*.

Zhang, H.; Zhu, L.; Wang, X.; and Yang, Y. 2023. Divide and Retain: A Dual-Phase Modeling for Long-Tailed Visual Recognition. *TNNLS*.

Zhao, B.; and Bilen, H. 2023. Dataset condensation with distribution matching. In *WACV*, 6514–6523.

Zhao, B.; Mopuri, K. R.; and Bilen, H. 2020. Dataset condensation with gradient matching. *arXiv preprint arXiv:2006.05929*.

Zhao, G.; Li, G.; Qin, Y.; and Yu, Y. 2023. Improved distribution matching for dataset condensation. In *CVPR*.

Zhao, Q.; Dai, Y.; Li, H.; Hu, W.; Zhang, F.; and Liu, J. 2024. LTGC: Long-Tail Recognition via Leveraging LLMs-driven Generated Content. In *CVPR Workshop*.

Zhao, Z.; Wang, H.; Shang, Y.; Wang, K.; and Yan, Y. 2025. Distilling long-tailed datasets. In *CVPR*, 30609–30618.

Zheng, H.; Zhou, L.; Li, H.; Su, J.; Wei, X.; and Xu, X. 2024. BEM: Balanced and Entropy-based Mix for Long-Tailed Semi-Supervised Learning. In *CVPR*, 22893–22903.

Zhong, W.; Tang, H.; Zheng, Q.; Xu, M.; Hu, Y.; and Nie, L. 2025. Towards Stable and Storage-efficient Dataset Distillation: Matching Convexified Trajectory. In *CVPR*.

Zhu, D.; Lei, B.; Zhang, J.; Fang, Y.; Xie, Y.; Zhang, R.; and Xu, D. 2023. Rethinking data distillation: Do not overlook calibration. In *CVPR*, 4935–4945.

Zhu, M.; Fan, C.; Chen, H.; Liu, Y.; Mao, W.; Xu, X.; and Shen, C. 2024. Generative Active Learning for Long-tailed Instance Segmentation. In *ICML*.