This CVPR paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version;

the final published version of the proceedings is available on IEEE Xplore.

Distilling Long-tailed Datasets

Zhenghao Zhao^{1*} Haoxuan Wang^{1*} Yuzhang Shang² Kai Wang³ Yan Yan^{1†} ¹University of Illinois Chicago ²Illinois Institute of Technology ³National University of Singapore {zzhao72, hwang339, yyan55}@uic.edu yshang4@hawk.iit.edu kai.wang@comp.nus.edu.sg

Abstract

Dataset distillation aims to synthesize a small, informationrich dataset from a large one for efficient model training. However, existing dataset distillation methods struggle with long-tailed datasets, which are prevalent in real-world scenarios. By investigating the reasons behind this unexpected result, we identified two main causes: 1) The distillation process on imbalanced datasets develops biased gradients, leading to the synthesis of similarly imbalanced distilled datasets. 2) The experts trained on such datasets perform suboptimally on tail classes, resulting in misguided distillation supervision and poor-quality soft-label initialization. To address these issues, we first propose Distributionagnostic Matching to avoid directly matching the biased expert trajectories. It reduces the distance between the student and the biased expert trajectories and prevents the tail class bias from being distilled to the synthetic dataset. Moreover, we improve the distillation guidance with Expert Decoupling, which jointly matches the decoupled backbone and classifier to improve the tail class performance and initialize reliable soft labels. This work pioneers the field of longtailed dataset distillation, marking the first effective effort to distill long-tailed datasets. Our code will be made public at https://github.com/ichbill/LTDD.

1. Introduction

Dataset distillation (DD) [4, 13, 28] aims to synthesize small and high-quality data summaries which contain the most important knowledge from a given target dataset [35]. Such a compact data summary offers numerous advantages, including faster model training, reduced memory consumption, and greater flexibility for various tasks [48], such as continual learning [31], neural architecture search [27], and knowledge distillation [15]. Various types of dataset distillation methods have been developed, including performance matching methods [30, 32, 50], parameter matching methods [38, 47]



Figure 1. **Performance comparison on CIFAR-10-LT.** Existing Dataset Distillation methods exhibit degraded performance when applied to imbalanced datasets, especially when the imbalance factor increases, whereas our method provides significantly better performance under different imbalanced scenarios.

and generative-based methods [12, 40]. Currently, parameter matching methods have demonstrated significant performance improvements across various datasets.

Research on dataset distillation has predominantly focused on uniformly distributed datasets, such as CIFAR10, CIFAR100 [22], and TinyImageNet [24]. However, in realworld applications such as medical image diagnosis [19], training samples typically exhibit a long-tailed class distribution [45]. In these scenarios, a small number of classes (head classes) have a large number of samples, while the majority of classes (tail classes) have only a few samples. Motivated by this discrepancy and the desire for practical application, we are interested in the novel task of longtailed dataset distillation (LTDD). LTDD aims to distill a long-tailed target dataset into a small and uniformly distributed synthetic dataset, where models trained on this distilled dataset can achieve satisfactory performance on the uniformly distributed test set.

Nevertheless, distilling long-tailed datasets presents significant challenges, as the effectiveness of current dataset distillation methods diminishes when applied to such imbalanced distributions. As shown in Figure 1, we find that the performance of different dataset distillation methods consistently decreases as the degree of imbalance increases. The degree of imbalance is quantified by the imbalance fac-

^{*}equal contribution

[†]corresponding author



Figure 2. **Relationship of classifier weights and class-wise accuracy.** We reveal that classifiers generated by existing dataset distillation methods often exhibit imbalanced distributions, resulting in poor performance on tail classes. In contrast, our method produces balanced classifiers, thereby enhancing overall accuracy.

tor, which is a commonly used metric to characterize the extent of imbalance [23]. The larger the imbalance factor is, the more imbalanced the dataset is. When the imbalance factor is large, popular dataset distillation methods may even perform worse than random selection (e.g. DATM can only achieve 40.1% accuracy when the imbalance factor $\beta = 200$, while random selection can achieve 49.9% accuracy). As shown in Figure 2, the performance degradation is concentrated in the tail classes. We attribute these performance failures to two key aspects of the learning procedure in long-tailed dataset distillation:

(1) During distillation, the weight distributions of the student and expert diverge, resulting in less informative tail class samples. Specifically, though the expert is trained on an imbalanced dataset, the student learns on a balanced one, resulting in differently distributed model weights. Such a discrepancy introduces controversy in the bi-level optimization objective, causing the distillation process to additionally focus on matching the data quantity distribution. Consequently, the tail classes of the distilled dataset contain significantly less useful information than the head classes, despite each class having the same number of samples.

(2) The expert models yield suboptimal performance on the tail classes. Concretely, the expert is biased toward the head classes when trained on imbalanced datasets [1, 49], resulting in degraded tail class accuracy. Since the student's performance is upper-bounded by that of the expert, the expressiveness of the distilled dataset is consequently restricted. Additionally, the image labels are obtained via the expert's prediction logits [13]. As the expert is biased toward the head classes, the prediction logits for the tail classes are, therefore, less reliable and confident.

To address the aforementioned issues, we propose a trajectory-matching approach named **Distributionagnostic Matching with Expert Decoupling**. Firstly, the student model is trained via a Distribution-agnostic Matching (DAM) strategy, which consists of a loss designed for matching experts trained on long-tailed data. Specifically, when training the student model on the balanced synthetic dataset, we mimic the procedure of training a model on an imbalanced dataset to reduce the distance between the student and the biased expert trajectories. Secondly, to mitigate the poor performance of the expert model and to enhance the overall distillation performance, an Expert Decoupling (ED) strategy is also proposed. In practice, we train representation experts and classification experts in a decoupled manner [20] to improve the performance of experts. We jointly match with the backbone layers of the representation experts and the classifier layers of the classification experts simultaneously. We use classification expert models to generate soft labels with high confidence for all classes.

This work pioneers the field of long-tailed dataset distillation as the first successful effort in this area. Our experiments on four long-tailed datasets, conducted under various imbalance factors and images-per-class (IPC) settings, demonstrate the effectiveness of our method. We show that our method is able to achieve state-of-the-art performance across different dataset distillation baselines.

2. Long-tailed Dataset Distillation

2.1. Problem Formulation

Assume we are given a long-tailed target training dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{|\mathcal{D}|}$ with C classes, where $y \in \{0, \dots, C-1\}$. Denote the subset of images belonging to class c as \mathcal{D}_c , then $|\mathcal{D}_0| > |\mathcal{D}_1| > \cdots > |\mathcal{D}_{C-1}|$ and $|\mathcal{D}_0| \gg |\mathcal{D}_{C-1}|$. Denote the test dataset as $\mathcal{T} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{|\mathcal{T}|}$ and is balanced. Our goal is to generate a small dataset $\mathcal{S} = \{(\hat{\mathbf{x}}_i, \hat{y}_i)\}_{i=1}^{|\mathcal{S}|}$, where $|\mathcal{S}| \ll |\mathcal{D}|$ and $|\mathcal{S}_0| = |\mathcal{S}_1| = \cdots = |\mathcal{S}_{C-1}|$, so that models trained on \mathcal{S} perform optimally on the test set \mathcal{T} .

Typical dataset distillation methods [4, 6, 13, 28] use the following procedure to optimize S: First, train an expert model $\mathcal{M}_{\mathcal{D}}$ on \mathcal{D} with acceptable performance. Then initialize S by randomly selecting images from \mathcal{D} or with Gaussian noise, and predict the corresponding (soft) labels using $\mathcal{M}_{\mathcal{D}}$. A student model \mathcal{M}_{S} is then trained based on S, and the distilled image set S is optimized by minimizing:

$$\mathcal{L}(\mathcal{S}) = ||\mathcal{F}(\mathcal{M}_{\mathcal{D}}; \mathcal{D}) - \mathcal{F}(\mathcal{M}_{\mathcal{S}}; \mathcal{S})||_{p}$$
(1)

where \mathcal{F} are model characteristics that can be used for matching, such as feature distributions [41], gradients [21], or training trajectories [4, 13]. Trajectory matching [4, 10, 13] matches the student model trained on the distilled dataset with the expert model trajectories trained on the original dataset, providing an impressive performance, and the matching loss is formulated as:

$$\mathcal{L}_{\text{match}}(\hat{\theta}_{t+N}, \theta_{t+M}^*, \theta_t^*) = \frac{\|\theta_{t+N} - \theta_{t+M}^*\|_2^2}{\|\theta_t^* - \theta_{t+M}^*\|_2^2}, \quad (2)$$

where θ_t^* and θ_{t+M}^* are the expert parameters on epoch t and t + M. $\hat{\theta}_{t+N}$ is the student parameter obtained in the



(a) expert trajectory

(b) training trajectory of traditional DD



Figure 3. Effect of biased expert. (a) An expert trained on an imbalanced dataset leads to increasingly imbalanced weight gradients over classes. (b) Existing dataset distillation methods [4, 13] are ignorant of the distribution gap between S_t and D. This causes the student gradient imbalance to increase in each step. If we match such trajectories, the synthetic dataset will be updated by the increasingly imbalanced gradients over classes, and the model trained on this synthetic dataset is highly biased. (c) With Distribution-agnostic Matching, the increasingly imbalanced gradients over classes will be re-weighted, such that the student model is updated with balanced gradients.

inner optimization updated using the cross-entropy loss.

In Section 2.2, we first discuss why such a formulation in Equation 1 is unreasonable for a long-tailed target dataset. Then we propose Distribution-agnostic Matching to mitigate the negative impact of the long-tailed distribution. To further enhance the distillation performance, an Expert Decoupling strategy is proposed in Section 2.3, aiding our method to achieve improved and even lossless performance.

2.2. Distribution-agnostic Matching

Influence of long-tailed distribution. By extending Equation 1 to the whole distillation process, a single step of the synthetic dataset optimization can be formulated as:

$$S_{t+1} = \underset{S_t}{\operatorname{arg\,min}} ||\theta_{\mathcal{D}} - \underset{\theta_{S}^{t}}{\operatorname{arg\,min}} \mathcal{L}(\theta_{S}^{t}, S_{t})||_{p}, \quad (3)$$

where θ_{S}^{t} is the student model parameter at epoch t and \mathcal{L} is the cross-entropy loss based on the student model parameter and the synthetic dataset S_t . $\theta_{\mathcal{D}}$ is the expert model parameters trained on \mathcal{D} and is determined before distillation. p is the factor for L_p norm and usually takes the value of 2.

Equation 3 is an embedded optimization problem in which the inner loop is performed first. This implies that S_t is updated to train a θ_S^t that can match the fixed θ_D as much as possible. We see that θ_S^t is an important intermediate that connects the synthetic dataset and the target dataset. Under the setting that the target dataset D is imbalanced, the weights in θ_D for different classes follow an imbalanced distribution and are fixed during distillation. Then the outer loop in Equation 3 would force θ_S^t to produce an imbalanced distribution similar to θ_D . As illustrated in Figure 3(a) and (b), the student is forced to mimic the expert during distillation, so that its weight distribution in different classes gradually becomes imbalanced.

However, the outer loop in Equation 3 aims to find a valid S_t rather than a proper θ_S^t . Thus, S_t is optimized to resem-

ble the training behavior of \mathcal{D} . Although the tail class performance on the long-tailed dataset \mathcal{D} is poor, this causes the tail classes of distilled S_t to similarly contain much less useful information than the head classes, even though it is uniformly distributed. As a result, Equation 3 proactively distills less information from tail classes, making the distillation process more biased towards head classes.

Distribution-agnostic Matching. As discussed above, matching the student network with an expert trained on an imbalanced dataset could lead to severe flaws in the synthetic dataset S, where the tail classes are unfortunately distorted. Ideally, we would like S to be a compact data summary such that every distilled image is meaningful and useful. However, such a dataset yields a student model with a balanced distribution (as shown in Figure 4 left), which mismatches the expert. This raises an interesting question: how can we preserve tail class information in the distilled dataset while avoiding the weight mismatching?

A key insight is to weaken the relationship between the data distribution and the weight distribution. In other words, we aim to prevent weight distribution mismatches from impacting dataset optimization. To achieve this, we propose Distribution-agnostic Matching, which seeks to align the gradient distributions of the student and expert without altering the distilled dataset. Practically, we modify the student training procedure to absorb the effect of weight imbalance by deliberately making θ_S^t imbalanced. Inspired by the Balanced Softmax loss [34], we propose a long-tailed distributed loss to aid student training:

$$\mathcal{L}^{c}(\mathbf{L}, \mathbf{y}, \mathbf{s}) = -\frac{1}{n} \sum_{i=1}^{n} g(\mathbf{s}_{\mathbf{y}_{i}}) \cdot \log P(\mathbf{L}_{i, \mathbf{y}_{i}}), \quad (4)$$

where

$$P(\mathbf{L}_{i,\mathbf{y}_i}) = \frac{e^{\mathbf{L}_{i,\mathbf{y}_i} - \lambda \log(\mathbf{s}_{\mathbf{y}_i})}}{\sum_{j=1}^{C} e^{\mathbf{L}_{i,j} - \lambda \log(\mathbf{s}_j)}}.$$
 (5)



Figure 4. Comparison of the internal loop between traditional DD methods and ours. The expert model M_D is trained on a long-tailed dataset, leading to biased weights. Left: For traditional DD methods, directly matching with these weights causes large information loss in the tail classes. The backward propagation updates this imbalance from the expert trajectories to the synthetic dataset. Right: With Distribution-agnostic Matching, the gradients are obtained via \mathcal{L}^c , which revises the student weight for matching in the internal loop. This mitigates the distance between the student weight and the expert trajectory to reduce the influence of imbalance on the synthetic dataset.

In the above equations, **L** is the student model logits, n is the number of samples in the distilled dataset, **y** is the ground-truth label, $\mathbf{s}_{\mathbf{y}_i}$ is the number of samples in the long-tailed dataset for class \mathbf{y}_i , $g(\cdot)$ normalizes $\mathbf{s}_{\mathbf{y}_i}$, and λ is used for smoothing the predictions.

The proposed loss is designed to support student model training during distillation by reducing the dependency of the weight distribution on the data distribution, thereby "mimicking" the training behavior on long-tailed datasets. Classes with more samples contribute more significantly to the loss, simulating a long-tailed training gradient for the student network during matching. This approach helps to minimize the trajectory differences between the student and expert networks caused by distribution mismatches between the target and distilled datasets. The right plot in Figure 4 illustrates the design and effect of our method.

2.3. Improved Guidance with Expert Decoupling

Unlike experts trained on balanced target datasets, the experts trained on long-tailed datasets are less effective for guiding the dataset distillation process. Specifically, when adopting typical dataset distillation methods, the long-tailed experts perform poorly on tail classes, degrading student performance on these classes and leading to unreliable initial soft labels. This degraded student performance restricts the effective learning of synthetic tail-class images, while the unreliable labels introduce misleading information and provide weak supervision.

In the conventional dataset distillation process, soft labels are assigned to initialized synthetic images based on the probability output of the expert model [13]. We found that soft-label predictions from long-tailed experts are distributed differently compared to those from balanced experts. As visualized in Figure 5, while the head class soft labels predicted by the balanced expert have similar confidence values with tail classes (0.88 and 0.89), the long-



Figure 5. **Soft-label initialization on CIFAR-10-LT.** We visualize the average predictive confidence of experts on different classes. For traditional DD, while the soft labels can be initialized well on head classes, they are predicted poorly on the tailed classes, leading to insufficient supervision.

tailed expert predictions between head and tail classes differ drastically. Whereas the head class confidence value can be as large as 0.97, the confidence value of the tail class is surprisingly small (0.38). Such a confidence discrepancy leads to less effective learning in tail classes than in head classes. Moreover, for tail classes with a low number of samples, the number of samples with high confidence is reduced and leads to the diversity decrease of the dataset initialization.

Decoupled training [20] is a widely used technique to alleviate the prediction bias introduced by imbalanced data. It consists of two stages: representation learning and classifier fine-tuning. At the representation learning stage, the entire model is trained on the original dataset \mathcal{D} . At the classifier fine-tuning stage, the backbone of the model is frozen, and only the classifier of the model is trained on a balanced dataset \mathcal{B} , which is created from \mathcal{D} via oversampling or undersampling. However, the integration of this decoupling approach with dataset distillation is not trivial, as it requires careful alignment of the student model with the expert model. To effectively utilize the idea of Expert Decoupling, we propose a two-step strategy to improve teacher

Dataset	CIFAR-10-LT											
Imbalance Factor		10			50			100			200	
IPC	10	20	50	10	20	50	10	20	50	10	20	50
Random	$32.5{\pm}2.2$	$39.6{\pm}0.9$	$51.9{\pm}1.5$	33.2±0.4	$42.0{\pm}1.3$	$51.6{\pm}1.3$	34.4±2.0	$41.4{\pm}0.7$	$52.6{\pm}0.5$	32.5±0.8	$42.2{\pm}1.1$	49.9±1.4
K-Center Greedy [36]	$21.9 {\pm} 0.8$	$24.2 {\pm} 0.8$	$31.7 {\pm} 0.9$	17.8±0.2	$20.8{\pm}0.5$	$26.1 {\pm} 0.2$	$16.2 {\pm} 0.5$	$19.0 {\pm} 1.0$	$24.2 {\pm} 1.2$	16.8±0.3	$17.5 {\pm} 1.4$	$22.6 {\pm} 1.6$
Graph-Cut [17]	$28.7{\pm}0.9$	$34.2{\pm}1.0$	$40.6{\pm}1.0$	24.2 ± 0.7	$28.6{\pm}0.8$	$33.9{\pm}0.4$	22.9±0.9	$26.0{\pm}0.5$	$33.3{\pm}1.0$	22.3±0.9	$25.2{\pm}0.5$	$29.2{\pm}0.4$
DC [46]	$37.9{\pm}0.9$	$38.5{\pm}0.9$	37.4±1.4	37.3±0.9	$38.8{\pm}1.0$	$35.8{\pm}1.2$	36.7±0.8	38.1±1.0	$35.3{\pm}1.4$	35.6±0.8	$35.7{\pm}0.9$	33.3±1.4
MTT [4]	$58.0 {\pm} 0.8$	$59.5 {\pm} 0.4$	$62.0 {\pm} 0.9$	45.8±1.4	$49.9 {\pm} 0.8$	$53.6 {\pm} 0.5$	37.7±0.6	41.6 ± 1.1	$47.8 {\pm} 1.1$	N/A	$22.6{\pm}1.0$	$23.9 {\pm} 0.8$
DREAM [28]	$34.6 {\pm} 0.6$	$42.2 {\pm} 1.5$	$50.5 {\pm} 0.7$	30.8±0.6	$38.4 {\pm} 0.3$	$45.5{\pm}0.9$	30.8±1.7	$34.9{\pm}0.8$	$42.2 {\pm} 0.8$	32.7±1.3	$32.4 {\pm} 0.3$	$38.9 {\pm} 0.4$
IDM [47]	$54.8 {\pm} 0.4$	57.1 ± 0.3	$60.1 {\pm} 0.3$	51.9±0.7	$53.3{\pm}0.6$	$56.1 {\pm} 0.4$	$49.8 {\pm} 0.6$	$50.9 {\pm} 0.5$	$53.1 {\pm} 0.4$	47.0±0.5	$48.1 {\pm} 0.5$	$49.9 {\pm} 0.3$
Minimax [12]	$29.2 {\pm} 0.5$	$28.5 {\pm} 0.6$	$39.9 {\pm} 0.1$	$18.4{\pm}0.3$	$22.5 {\pm} 0.2$	$25.2 {\pm} 0.2$	$19.9 {\pm} 0.4$	$23.3 {\pm} 0.2$	$28.0{\pm}0.6$	19.1±0.4	$20.5 {\pm} 0.2$	22.7 ± 0.3
DATM [13]	57.2 ± 0.4	$60.4 {\pm} 0.2$	$66.7 {\pm} 0.6$	41.6±0.2	$43.4 {\pm} 0.3$	$50.3 {\pm} 0.2$	37.3±0.2	$38.9 {\pm} 0.1$	$44.3 {\pm} 0.1$	N/A	$34.8 {\pm} 0.1$	$40.1 {\pm} 0.2$
Ours	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	52.2±0.6	56.6±0.4	62.3±0.3
Full Dataset		76.7±0.3			69.8±0.3			66.2±0.4			62.1±0.4	

Table 1. Quantitative comparisons with the SOTA methods on CIFAR-10-LT. Our method outperforms all existing approaches, and can even achieve lossless performance under certain settings. Performance is evaluated under various imbalance factors. As the first to explore long-tailed dataset distillation, all experiments are conducted using open-source code. N/A indicates distillation failure.

guidance for student training.

Firstly, we train two types of experts in a decoupled manner. The representation expert with parameter $\theta^{\mathcal{D}}$ is obtained from the representation learning stage, and the classifier expert with parameter $\theta^{\mathcal{B}}$ is obtained from the classifier finetuning stage. Since the classifier experts are not biased and are able to provide more accurate classification, we use the classifier experts $\theta^{\mathcal{B}}$ for the synthetic dataset S initialization. Specifically, we use the logits $L_i = \theta^{\mathcal{B}} \mathbf{x}_i$ to initialize soft labels, with softmax (L_i) as the label of \mathbf{x}_i , and then sample from \mathcal{D} to initialize S.

Secondly, we match the teacher expert's training trajectory with the student's in an adaptive manner. Since the representation expert learns the image features from the original dataset \mathcal{D} but fails to classify with biased classifiers [20], we skip matching the classifier layer when matching the student network trained on the synthetic dataset with the representation expert. On the contrary, the classifier expert $\mathcal{M}_{\mathcal{B}}$ only trains the classifier layer, so we also only train the student classifier layer and only match the classifier layer during the classifier matching. During the whole matching process, we jointly match the representation trajectory and the classifier trajectory using the following loss:

$$\mathcal{L} = \lambda_{\text{rep}} \mathcal{L}_{\text{match}}(\hat{\theta}_{t+N}, \theta_{t+M}^{\mathcal{D}*}, \theta_{t}^{\mathcal{D}*}) + \lambda_{\text{cls}} \mathcal{L}_{\text{match}}(\hat{\theta}_{t+N}, \theta_{t+M}^{\mathcal{B}*}, \theta_{t}^{\mathcal{B}*}),$$
(6)

where λ_{rep} and λ_{cls} are hyper-parameters that control the weighting ratio, $\hat{\theta}_{t+N}$ is the student model at its training epoch t + N, $\theta_t^{\mathcal{D}*}$ and $\theta_t^{\mathcal{B}*}$ are the representation backbone and classifier of the expert at epoch t, respectively. $\hat{\theta}_{t+N}$ is updated by our proposed long-tailed distributed loss \mathcal{L}^c in Equation 4, and $\mathcal{L}_{\text{match}}$ follows the form in Equation 2.

3. Experiments

3.1. Experiment Datasets

We evaluate our methods on four widely used long-tailed datasets of different scales: CIFAR-10-LT, CIFAR-100-LT

[7], TinyImageNet-LT, and ImageNet-LT [29]. The experiments are conducted under various degrees of imbalance and IPC settings. Following [7, 9, 29], these long-tailed datasets are sampled from the original balanced CIFAR [22] and ImageNet [8] datasets. The sampling procedure is done by randomly selecting different amounts of images from each class, where the image number $|\mathcal{D}_c|$ for class *c* is determined by an exponential decay function $|\hat{\mathcal{D}}_c| = |\mathcal{D}_c|\mu^c$, and $\mu^c = \beta^{-(c/C)}$, where *C* is the number of classes and β is the imbalanced factor. The degree of imbalance is measured via the imbalanced factor $\beta = \mathcal{D}_0/\mathcal{D}_C$ [7]. The larger β is, the more imbalanced the dataset is.

3.2. Main Results

CIFAR-10-LT. The experiment results on CIFAR-10-LT are presented in Table 1, showing that our proposed method consistently outperforms all other dataset distillation baselines. Table 1 further highlights the advantages of our approach on long-tailed datasets. When the imbalance factor is low, prior dataset distillation methods like DATM perform reasonably well. For instance, DATM only experiences a 10.0% drop in accuracy when $\beta = 10$ and IPC=50, compared to the accuracy achieved with the full dataset. In these scenarios, our method surpasses the best baseline, DATM, by 0.9% on $\beta = 10$ and 3.8% on $\beta = 50$.

When the imbalance factor is high (e.g., $\beta = 200$), IDM [47] achieves performance comparable to random selection. Other methods, such as MTT [4], DREAM [28], DATM [13], and Minimax [12], struggle with higher imbalance factors, performing worse than random selection. In contrast, our method sustains strong performance on highly imbalanced datasets, surpassing the best existing method IDM by a substantial margin of 10.6% on $\beta = 200$. Remarkably, we achieve lossless performance on $\beta = 200$ using only 50 images per class.

CIFAR-100-LT. The results are presented in Table 2. Our method consistently outperforms all existing dataset distillation methods across all imbalance factors on CIFAR-100-LT. Similar to CIFAR-10-LT, our approach demonstrates

Dataset	CIFAR-100-LT				Tiny-ImageNet-LT							
Imbalance Factor		10			20			10			20	
IPC	10	20	50	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	$30.5{\pm}0.5$	7.4±0.2	13.5±0.4	20.6±0.4	7.6±0.1	13.2±0.4	19.8±0.1
K-Center Greedy [36]	10.7±0.9	$15.9 {\pm} 1.0$	$24.8{\pm}0.2$	10.0 ± 0.5	$15.1 {\pm} 0.6$	$23.8{\pm}0.3$	$10.4{\pm}1.5$	$11.3 {\pm} 1.2$	7.7±1.0	12.6±1.4	$9.9 {\pm} 2.2$	13.2 ± 2.7
Graph-Cut [17]	16.9±0.3	$22.2{\pm}0.4$	$29.9{\pm}0.4$	$16.0 {\pm} 0.5$	$20.7{\pm}0.5$	$28.7{\pm}0.3$	9.8±0.7	$5.6{\pm}0.6$	$10.9{\pm}1.1$	3.4±0.8	$5.0{\pm}1.0$	$4.2{\pm}1.1$
DC [46]	24.0±0.3	27.4±0.3	27.4±0.3	23.2±0.3	26.2±0.3	27.4±0.3	-	-	-	-	-	-
MTT [4]	14.3±0.1	$16.7 {\pm} 0.2$	$13.8 {\pm} 0.2$	12.6±0.3	$15.0 {\pm} 0.2$	$10.6 {\pm} 0.5$	11.1 ± 0.2	$18.1 {\pm} 0.2$	23.1 ± 0.1	7.7±0.1	$14.7 {\pm} 0.2$	$15.6 {\pm} 0.3$
DREAM [28]	10.1±0.4	$12.0{\pm}1.0$	$13.1 {\pm} 0.4$	9.4±0.4	$10.3 {\pm} 0.6$	$12.3 {\pm} 0.3$	$5.4{\pm}0.3$	$6.8 {\pm} 0.1$	$7.8 {\pm} 0.2$	4.8±0.1	$6.0 {\pm} 0.2$	$7.4{\pm}0.1$
DATM [13]	28.2±0.4	$34.1 {\pm} 0.2$	$31.6 {\pm} 0.1$	25.3±0.3	$27.2 {\pm} 0.1$	$27.1 {\pm} 0.1$	21.3 ± 0.1	$14.5 {\pm} 0.1$	$26.8 {\pm} 0.1$	14.0±0.6	$19.0 {\pm} 0.3$	23.1 ± 0.1
Ours	31.5±0.2	37.5±0.4	$40.0{\pm}0.1$	31.4±0.5	35.1±0.4	37.0±0.7	26.0±0.3	$\textbf{27.9}{\pm 0.2}$	30.3±0.2	23.6±0.3	25.5±0.3	$\textbf{28.0{\pm}0.6}$
Full Dataset		40.3±0.4			36.5±0.4			31.2±1.0			28.3±0.3	

Table 2. Quantitative comparisons with the SOTA methods on CIFAR-100-LT and Tiny-ImageNet-LT. When scaling up to larger datasets, our approach still outperforms all the included baselines, and achieves nearly lossless performance under several settings.

Dataset	ImageNet-LT						
Imbalance Factor		5	10				
IPC	10	20	10	20			
Random	3.9±0.1	7.0±0.1	3.9±0.1	$6.8 {\pm} 0.1$			
G-VBSM [39]	-	$1.0{\pm}0.1$	-	$1.0{\pm}0.1$			
TESLA [6]	3.0±0.1	-	2.7±0.1	-			
DATM [13]	$7.4{\pm}0.1$	$8.1 {\pm} 0.2$	7.9 ± 0.1	$8.2{\pm}0.1$			
SRe ² L [44]	6.7 ± 0.1	$10.1 {\pm} 0.1$	7.7±0.3	$10.9 {\pm} 1.0$			
Ours	20.8±0.2	$21.0{\pm}0.1$	20.3±0.1	$20.7{\pm}0.1$			
Full Dataset	27.6±0.3		26.6±0.2				

Table 3. **Quantitative comparisons on ImageNet-LT.** The missing results are due to out-of-memory.

greater improvements as the imbalance factor increases.

Tiny-ImageNet-LT. We conducted experiments on Tiny-ImageNet-LT to evaluate the scalability of our method in comparison with other baselines. The results, shown in Table 2, align with our findings on CIFAR-10-LT and CIFAR-100-LT, where our method surpasses baseline methods by a growing margin as the imbalance factor increases.

ImageNet-LT. We further scale up to ImageNet-LT, with results presented in Table 3, where our model consistently outperforms the baselines. Note that as the dataset size increases, many dataset distillation methods become resource-intensive and are impractical to apply.

Cross-architecture performance. Table 4 presents our cross-architecture experiments, demonstrating the generalizability of our synthetic dataset. We evaluated various methods on the CIFAR-10-LT dataset with an imbalance factor of $\beta = 100$ and IPC=50, comparing against random selection, MTT, and DATM. As shown in the table, our distilled dataset consistently outperforms other methods by a significant margin. Specifically, our method achieves a 12.7% accuracy improvement on ConvNet, a 17.1% improvement on ResNet-18, a 7.7% improvement on VGG-11, and a 6% improvement on AlexNet. The results also indicate that the imbalance distilled by ConvNet in traditional DD methods negatively impacts the performance of other models. Similar to experiments on ConvNets, MTT and DATM perform poorly on unseen models due to the imbalanced dataset and the biased expert trained on it, often yielding worse results than random selection.

Method	ConvNet-3	ResNet-18	VGG-11	AlexNet
Random	52.0	49.2	47.8	47.2
MTT [4]	49.1	42.9	42.3	40.0
DATM [13]	44.4	42.4	42.2	44.9
Ours	64.7	60.0	55.5	53.2

Table 4. Cross-architecture evaluation on CIFAR-10-LT.

3.3. Ablation Studies

Ablation on different components. Figure 6a demonstrates the effectiveness of the two proposed components of our method. These experiments were performed on the CIFAR-10-LT dataset with an imbalance factor of $\beta = 200$ and IPC=50. The "base" method is a trajectory matching method without Distribution-agnostic Matching and Expert Decoupling. "+DAM" refers to our method without Expert Decoupling, while "ours" represents the complete proposed method, with both DAM and ED applied.

The baseline, which suffers from weight inconsistency and poor expert performance as mentioned in Section 1, only achieves 43.8% accuracy. After incorporating Distribution-agnostic Matching, the inconsistency between student and expert weights is mitigated, resulting in a 14.3% performance improvement. Further integrating Expert Decoupling boosts the performance to 62.6%.

Effect of classifier layer matching. As discussed in Section 2.3, during representation learning, the model is influenced by the long-tailed distribution, resulting in a biased classifier. During classifier fine-tuning, the classifier is fine-tuned on balanced data with the backbone frozen. Therefore, it is crucial to explore strategies for matching each expert's different modules during dataset distillation. The experiments are presented in Figure 6b. "B", "C", and "W" refer to matching with the representation backbone, the classifier layer, and the whole model, respectively. "B/C" refers to matching the student backbone with the representation expert's backbone, and the student classifier layer with the classification expert's classifier.

As shown in Figure 6b, the best performance is achieved by matching the backbone with representation experts and the classifier layer with classification experts. This outcome



Figure 6. (a) Ablation study on the proposed components, Distribution-agnostic Matching and Expert Decoupling. (b) Ablation of different matching strategies with expert model layers. (c) & (d) Comparisons of matching strategies for long-tailed dataset distillation under low imbalance factor (c) and high imbalance factor (d). (e) Ablation study on hyperparameters λ_{rep} and λ_{cls} .

aligns with the principles of decoupled learning. Since the classifier layer of representation experts is heavily affected by the long-tailed dataset, including it in the matching process causes a performance drop. In addition, matching the entire classifier expert negatively impacts performance, as only the classifier layer is trained at this stage.

Effect of different matching strategies. When distilling a long-tailed dataset, various intuitive strategies can be employed with decoupling trained experts. We explored several strategies, including sequential matching, classification expert matching, oversampled dataset, and joint matching. Sequential matching first matches representation experts, followed by classification experts. Classification expert matching only matches with the classification expert. The oversampled dataset approach involves oversampling the tail classes before distilling the dataset rather than using the decoupling strategy. Joint matching, used in our Expert Decoupling, involves matching both experts simultaneously in the outer loop of the matching.

As shown in Figures 6c and 6d, joint matching (joint), which we use in Expert Decoupling, achieves the best overall performance. Sequential matching (sequential) and directly matching with classification experts (direct) perform poorly when the imbalance factor is low, meaning the data is "not very imbalanced." Sequential matching suffers due to the trajectory gap between the two types of experts. Classification expert matching is less effective when the imbalance factor is low because the representation layers learn more information in such scenarios, and solely matching the classification layers fails to capture sufficient information. As depicted in Figure 6d, the oversampled dataset strategy (oversample) performs poorly when the imbalance factor is high. This finding aligns with previous work [7], which indicates that oversampling leads to model overfitting.

Effect of hyperparameters. The most important hyperparameters in our method are λ_{rep} and λ_{cls} in Equation 6. Together, they regulate the balance between the matching representation expert and the classification expert. The setting of λ_{rep} and λ_{cls} depends on the imbalance factors. From Figure 6e, we have the following observations. For a "more imbalanced" dataset (e.g., $\beta = 200$), lower λ_{rep} and higher λ_{cls}

Hyperparameters	$\beta = 10$	Hyperparameters	$\beta = 200$
$\begin{array}{l} \lambda_{\rm rep} = 1, \lambda_{\rm cls} = 0.2 \\ \lambda_{\rm rep} = 1, \lambda_{\rm cls} = 0.4 \\ \lambda_{\rm rep} = 1, \lambda_{\rm cls} = 0.6 \end{array}$	70.5 68.9 68.0	$ \begin{aligned} \lambda_{\rm rep} &= 0.2, \lambda_{\rm cls} = 1 \\ \lambda_{\rm rep} &= 0.4, \lambda_{\rm cls} = 1 \\ \lambda_{\rm rep} &= 0.6, \lambda_{\rm cls} = 1 \end{aligned} $	62.4 61.7 61.3

Table 5. Effect of λ_{rep} and λ_{cls} under different imbalance factors.

are preferred. For "less imbalanced" datasets (e.g., $\beta = 10$), higher λ_{rep} and lower λ_{cls} are chosen. This observation is consistent with our experiments in Figure 6c and 6d, where only matching classifier experts perform well when β is high. This is because the representation layers learn more information when β is low, and increasing λ_{rep} leverages information from the representation layers.

We further analyze the sensitivity of the hyperparameters in Table 5. We can observe a minor change in the performance as we change hyperparameters rapidly. When the relative magnitude of the two hyperparameters is preserved, their absolute values do not affect the distillation performance much, suggesting the robustness of our method.

3.4. Long-tailed Dataset Distillation Analysis

Effect of imbalance in long-tailed dataset distillation. In this part, we empirically show that the performance drop in long-tailed dataset distillation is due to the data imbalance instead of the total sample number reduction of the target dataset. We compare the performance of using regular dataset distillation methods for CIFAR-10-LT (with $\beta = 200$ and IPC=50) and a subset of CIFAR-10 that contains the same amount of data but in a uniform distribution. The total number of samples is 11203 under this setting.

As we can observe in Figure 7a, the performance gap between distilling a long-tailed dataset and a balanced subset can be as large as 28.6%, even though the two target datasets have the same total amount of samples. Thus, we conclude that it is the imbalanced distribution in the target dataset that prevents effective distillation. Based on this observation, we delved into it and proposed methods to alleviate the imbalanced distribution's negative effects.

Handling classes with extremely small number of samples. In long-tailed dataset distillation, a scenario exists where the number of samples in tail classes is even less than



Figure 7. (a) Comparison of DD performance between a uniform dataset and a long-tailed dataset with the same number of samples. (b) Comparison of strategies to handle an extremely low number of samples. (c) Comparison of class-wise accuracy between DATM and our method. (d) Comparison between balanced IPC and long-tailed distributed IPC.

the required number of images per class. In this case, we employ and compare three strategies: 1) Inherent strategy, which keeps the number of samples the same in the synthetic dataset; 2) Oversampling, which duplicates the samples of the tail classes before distillation; 3) Noise generation, which randomly generates Gaussian noises to supplement the insufficient samples for tail classes. From figure 7b, we see that the inherent strategy provides a comparatively lower accuracy, while oversampling and noise generation achieve better and similar performance.

Class-wise accuracy for long-tailed dataset. We also provide class-wise accuracy to illustrate the impact of the imbalanced data, and we show that our method can effectively improve the accuracy of tail classes. As shown in Figure 7c, DATM performs poorly on tail classes, and even achieves near-zero performance on the last class. On the contrary, our method can successfully boost the tail class accuracy while reserving performance on head classes.

Choice of distilled dataset distribution. Should we distill a long-tailed dataset into another long-tailed dataset or a balanced one? We compared the performance of distilling the target dataset into either a similarly distributed longtailed synthetic dataset or a balanced synthetic dataset using our method and DREAM [28]. For the long-tailed synthetic dataset, we maintained the same total number of samples as in the synthetic dataset with IPC=50. As shown in Figure 7d, both methods demonstrate that distilling the dataset into a balanced synthetic dataset provides better performance than distilling it into a long-tailed one. This result is reasonable since the model trained on the synthetic long-tailed dataset will be affected by the long-tailed distribution.

4. Related Work

4.1. Dataset Distillation

Dataset distillation [43] was first formulated by Wang et al. The goal of dataset distillation is to synthesize a small dataset, such that models trained on it have good performance on the original dataset. The following works can be divided into kernel ridge regression-based methods [30, 32, 50], gradient matching methods [6, 13, 28, 42, 46], representation matching methods [38, 47], and generative-based methods [12, 40]. Despite the advantages of current methods, they only focus on uniformly distributed datasets, ignoring the prevalent usage of long-tailed datasets in practice. Therefore, we propose to address the long-tailed dataset distillation problem in this paper.

4.2. Long-tailed Recognition

Under real-world settings, data often tends to follow a longtailed distribution [2]. Models trained on such datasets usually exhibit good performance on head classes but low accuracy on tail classes. A straightforward way is to preprocess the data, using oversampling [14, 18, 33], undersampling [3, 14], or data augmentation [5, 11] to make it balancedly distributed. An alternative way is weight balancing [2, 26, 34], which re-weights the layers during training to force the parameters to update unbiasedly. Decoupling methods [20] first train the model to obtain the features of the data and then fine-tune the model on the class-balanced data. In our methods, we use the decoupling method to train the expert networks during the distillation.

5. Conclusion

We introduce long-tailed dataset distillation, a novel yet challenging task. We find that existing DD methods fail when applied to long-tailed datasets and encounter two main problems: weight distribution mismatch between the student and the expert; and suboptimal performance of the expert on the tail classes. To mitigate these issues, we propose two strategies: Distribution-agnostic Matching and Expert Decoupling. Experimental results over four longtailed datasets demonstrate the effectiveness of our proposed method. As the pioneering work in this research area, we are the first to effectively distill long-tailed datasets.

Acknowledgments: This research is supported by NSF IIS-2525840, ECCS-2123521 and Cisco Research unrestricted gift. This article solely reflects the opinions and conclusions of its authors and not the funding agencies.

References

- Shaden Alshammari, Yuxiong Wang, Deva Ramanan, and Shu Kong. Long-tailed recognition via weight balancing. In *CVPR*, 2022. 2
- [2] Shaden Alshammari, Yu-Xiong Wang, Deva Ramanan, and Shu Kong. Long-tailed recognition via weight balancing. In *CVPR*, 2022. 8
- [3] Mateusz Buda, Atsuto Maki, and Maciej A Mazurowski. A systematic study of the class imbalance problem in convolutional neural networks. *Neural networks*, 2018. 8
- [4] George Cazenavette, Tongzhou Wang, Antonio Torralba, Alexei A Efros, and Jun-Yan Zhu. Dataset distillation by matching training trajectories. In CVPR, 2022. 1, 2, 3, 5, 6
- [5] Hsin-Ping Chou, Shih-Chieh Chang, Jia-Yu Pan, Wei Wei, and Da-Cheng Juan. Remix: rebalanced mixup. In *Computer Vision–ECCV 2020 Workshops: Glasgow, UK, August 23–28, 2020, Proceedings, Part VI 16, 2020.* 8
- [6] Justin Cui, Ruochen Wang, Si Si, and Cho-Jui Hsieh. Scaling up dataset distillation to imagenet-1k with constant memory. In *ICML*, 2023. 2, 6, 8, 1
- [7] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *CVPR*, 2019. 5, 7
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 5
- [9] 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 *CVPR*, 2023. 5
- [10] Jiawei Du, Yidi Jiang, Vincent YF Tan, Joey Tianyi Zhou, and Haizhou Li. Minimizing the accumulated trajectory error to improve dataset distillation. In CVPR, 2023. 2
- [11] Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In CVPR, 2018. 8
- [12] Jianyang Gu, Saeed Vahidian, Vyacheslav Kungurtsev, Haonan Wang, Wei Jiang, Yang You, and Yiran Chen. Efficient dataset distillation via minimax diffusion. In *CVPR*, 2024. 1, 5, 8
- [13] Ziyao Guo, Kai Wang, George Cazenavette, Hui Li, Kaipeng Zhang, and Yang You. Towards lossless dataset distillation via difficulty-aligned trajectory matching. In *ICLR*, 2024. 1, 2, 3, 4, 5, 6, 8
- [14] Guo Haixiang, Li Yijing, Jennifer Shang, Gu Mingyun, Huang Yuanyue, and Gong Bing. Learning from classimbalanced data: Review of methods and applications. *Expert systems with applications*, 2017. 8
- [15] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 1
- [16] Geoffrey E Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan R Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. arXiv preprint arXiv:1207.0580, 2012. 1
- [17] Rishabh Iyer, Ninad Khargoankar, Jeff Bilmes, and Himanshu Asanani. Submodular combinatorial information mea-

sures with applications in machine learning. In *Algorithmic Learning Theory*, 2021. 5, 6

- [18] Andrew Janowczyk and Anant Madabhushi. Deep learning for digital pathology image analysis: A comprehensive tutorial with selected use cases. *Journal of pathology informatics*, 2016. 8
- [19] 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 *MICCAI*, 2021. 1
- [20] Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. Decoupling representation and classifier for long-tailed recognition. In *ICLR*, 2020. 2, 4, 5, 8
- [21] Jang-Hyun Kim, Jinuk Kim, Seong Joon Oh, Sangdoo Yun, Hwanjun Song, Joonhyun Jeong, Jung-Woo Ha, and Hyun Oh Song. Dataset condensation via efficient syntheticdata parameterization. In *ICML*, 2022. 2
- [22] Alex Krizhevsky. Learning multiple layers of features from tiny images. University of Toronto, 2012. 1, 5
- [23] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 2017. 2, 1
- [24] Yann Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS* 231N, 2015. 1
- [25] Zekai Li, Xinhao Zhong, Zhiyuan Liang, Yuhao Zhou, Mingjia Shi, Ziqiao Wang, Wangbo Zhao, Xuanlei Zhao, Haonan Wang, Ziheng Qin, Dai Liu, Kaipeng Zhang, Tianyi Zhou, Zheng Zhu, Kun Wang, Guang Li, Junhao Zhang, Jiawei Liu, Yiran Huang, Lingjuan Lyu, Jiancheng Lv, Yaochu Jin, Zeynep Akata, Jindong Gu, Rama Vedantam, Mike Shou, Zhiwei Deng, Yan Yan, Yuzhang Shang, George Cazenavette, Xindi Wu, Justin Cui, Tianlong Chen, Angela Yao, Manolis Kellis, Konstantinos N. Plataniotis, Bo Zhao, Zhangyang Wang, Yang You, and Kai Wang. Dd-ranking: Rethinking the evaluation of dataset distillation. GitHub repository, 2024. 1
- [26] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, 2017. 8
- [27] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. *arXiv preprint arXiv:1806.09055*, 2018. 1
- [28] Yanqing Liu, Jianyang Gu, Kai Wang, Zheng Zhu, Wei Jiang, and Yang You. Dream: Efficient dataset distillation by representative matching. In *ICCV*, 2023. 1, 2, 5, 6, 8
- [29] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. Large-scale long-tailed recognition in an open world. In CVPR, 2019. 5
- [30] Noel Loo, Ramin Hasani, Alexander Amini, and Daniela Rus. Efficient dataset distillation using random feature approximation. In *NeurIPS*, 2022. 1, 8
- [31] Wojciech Masarczyk and Ivona Tautkute. Reducing catastrophic forgetting with learning on synthetic data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020. 1

- [32] Timothy Nguyen, Zhourong Chen, and Jaehoon Lee. Dataset meta-learning from kernel ridge-regression. arXiv preprint arXiv:2011.00050, 2020. 1, 8
- [33] Junran Peng, Xingyuan Bu, Ming Sun, Zhaoxiang Zhang, Tieniu Tan, and Junjie Yan. Large-scale object detection in the wild from imbalanced multi-labels. In CVPR, 2020. 8
- [34] Jiawei Ren, Cunjun Yu, Xiao Ma, Haiyu Zhao, Shuai Yi, et al. Balanced meta-softmax for long-tailed visual recognition. In *NeurIPS*, 2020. 3, 8
- [35] Noveen Sachdeva and Julian McAuley. Data distillation: A survey. arXiv preprint arXiv:2301.04272, 2023. 1
- [36] Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. arXiv preprint arXiv:1708.00489, 2017. 5, 6
- [37] Shai Shalev-Shwartz, Yoram Singer, and Nathan Srebro. Pegasos: Primal estimated sub-gradient solver for svm. In *ICML*, 2007. 1
- [38] Yuzhang Shang, Zhihang Yuan, and Yan Yan. Mim4dd: Mutual information maximization for dataset distillation. In *NeurIPS*, 2024. 1, 8
- [39] Shitong Shao, Zeyuan Yin, Muxin Zhou, Xindong Zhang, and Zhiqiang Shen. Generalized large-scale data condensation via various backbone and statistical matching. In CVPR, 2024. 6
- [40] Duo Su, Junjie Hou, Weizhi Gao, Yingjie Tian, and Bowen Tang. D⁴: Dataset distillation via disentangled diffusion model. In *CVPR*, 2024. 1, 8
- [41] Kai Wang, Bo Zhao, Xiangyu Peng, Zheng Zhu, Shuo Yang, Shuo Wang, Guan Huang, Hakan Bilen, Xinchao Wang, and Yang You. Cafe: Learning to condense dataset by aligning features. In *CVPR*, 2022. 2
- [42] Kai Wang, Zekai Li, Zhi-Qi Cheng, Samir Khaki, Ahmad Sajedi, Ramakrishna Vedantam, Konstantinos N Plataniotis, Alexander Hauptmann, and Yang You. Emphasizing discriminative features for dataset distillation in complex scenarios. arXiv preprint arXiv:2410.17193, 2024. 8
- [43] Tongzhou Wang, Jun-Yan Zhu, Antonio Torralba, and Alexei A Efros. Dataset distillation. arXiv preprint arXiv:1811.10959, 2018. 8
- [44] Zeyuan Yin, Eric Xing, and Zhiqiang Shen. Squeeze, recover and relabel: Dataset condensation at imagenet scale from a new perspective. In *NeurIPS*, 2023. 6
- [45] 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*, 2023. 1
- [46] Bo Zhao, Konda Reddy Mopuri, and Hakan Bilen. Dataset condensation with gradient matching. arXiv preprint arXiv:2006.05929, 2020. 1, 5, 6, 8
- [47] Ganlong Zhao, Guanbin Li, Yipeng Qin, and Yizhou Yu. Improved distribution matching for dataset condensation. In *CVPR*, 2023. 1, 5, 8
- [48] Zhenghao Zhao, Yuzhang Shang, Junyi Wu, and Yan Yan. Dataset quantization with active learning based adaptive sampling. In ECCV, 2024. 1
- [49] Zhisheng Zhong, Jiequan Cui, Shu Liu, and Jiaya Jia. Improving calibration for long-tailed recognition. In CVPR, 2021. 2

[50] Yongchao Zhou, Ehsan Nezhadarya, and Jimmy Ba. Dataset distillation using neural feature regression. In *NeurIPS*, 2022. 1, 8