

# Spanning Training Progress: Temporal Dual-Depth Scoring (TDDS) for Enhanced Dataset Pruning

Xin Zhang<sup>1,2\*</sup> Jiawei Du<sup>2,3†</sup> Yunsong Li<sup>1</sup> Weiyang Xie<sup>1‡</sup> Joey Tianyi Zhou<sup>2,3‡</sup>

<sup>1</sup>XiDian University, Xi'an, China

<sup>2</sup>Centre for Frontier AI Research (CFAR), Agency for Science, Technology and Research (A\*STAR), Singapore

<sup>3</sup>Institute of High Performance Computing (IHPC), Agency for Science, Technology and Research (A\*STAR), Singapore

xinzhang01@stu.xidian.edu.cn, {dujw, Joey\_Zhou}@cfar.istar.edu.sg

{ysli, wyxie}@mail.xidian.edu.cn

## Abstract

Dataset pruning aims to construct a coreset capable of achieving performance comparable to the original, full dataset. Most existing dataset pruning methods rely on snapshot-based criteria to identify representative samples, often resulting in poor generalization across various pruning and cross-architecture scenarios. Recent studies have addressed this issue by expanding the scope of training dynamics considered, including factors such as forgetting event and probability change, typically using an averaging approach. However, these works struggle to integrate a broader range of training dynamics without overlooking well-generalized samples, which may not be sufficiently highlighted in an averaging manner. In this study, we propose a novel dataset pruning method termed as **Temporal Dual-Depth Scoring (TDDS)**, to tackle this problem. TDDS utilizes a dual-depth strategy to achieve a balance between incorporating extensive training dynamics and identifying representative samples for dataset pruning. In the first depth, we estimate the series of each sample's individual contributions spanning the training progress, ensuring comprehensive integration of training dynamics. In the second depth, we focus on the variability of the sample-wise contributions identified in the first depth to highlight well-generalized samples. Extensive experiments conducted on CIFAR and ImageNet datasets verify the superiority of TDDS over previous SOTA methods. Specifically on CIFAR-100, our method achieves 54.51% accuracy with only 10% training data, surpassing baselines methods by more than 12.69%. Our codes are available at <https://github.com/zhangxin-xd/Dataset-Pruning-TDDS>.

<sup>†</sup>Equal contribution.

<sup>‡</sup>Corresponding authors.

\*Work completed during internship at A\*STAR




Figure 1. (a) Coresets constructed by snapshot-based dataset pruning tend to fluctuate depending on the epochs selected for data selection. The red and blue points represent the data samples of two coresets constructed with 90% pruning rate at Epoch 10 and 100 from the same training progress, respectively. (b) Snapshot-based coresets fail to generalize to pruning settings, especially pruning rates. The accuracy gap between coresets constructed at Epoch 10 and 100 is increasing as the pruning rate grows. Here, we measure importance using the error vector score, a classic snapshot-based criterion defined in [31]. These scores are obtained by training ResNet-18 on CIFAR-10.

## 1. Introduction

Explosively growing datasets [9, 20, 52, 53] have been crucial for the success of deep neural networks (DNNs) in various applications [7, 21, 35, 50]. However, learning from large-scale datasets is both time-consuming and financially prohibitive [5, 27, 39, 40, 44, 48]. Actually, a substantial portion of data samples are redundant [43, 46, 51], which means they can be excluded from training without compromising performance.

Numerous dataset pruning works [12, 16, 17, 22, 25, 29, 32, 46, 49] have explored various criteria [13, 41] for identifying important data samples, including distribution distance [45, 46], uncertainty [8], loss [31, 43], and gradient [31]. Most of them evaluate sample importance using a snapshot of training progress. For example, Paul et al. [31]

uses the error vector score generated by a few epochs<sup>1</sup> in early training; Xia et al. [46] calculates the distribution distances of features at the end of training.

However, these snapshot-based pruning methods have two potential problems. Firstly, the importance scores of samples fluctuate with epochs during the progress of training. Consequently, the coresets constructed by the importance scores at different snapshots of training progress differ from each other significantly. As verified in Figure 1(a), the coresets constructed at Epoch 10 and 100 obviously contain different samples. Secondly, relying on importance scores from a snapshot can lead to overfitting in the construction of the coreset, as this approach may not generalize well across different scenarios, such as varying pruning rates. As shown in Figure 1(b), the accuracy gap of coresets from different snapshots becomes larger as the pruning rate increases.

To address these two problems, a straightforward approach is to consider a broader range of training dynamics throughout the training process, with a focus on selecting well-generalized data samples. For example, Pleiss et al. [33] involves training dynamics by measuring the probability gap between the target class and the second largest class in each epoch; Toneva et al. [43] monitors the forgetting events throughout the entire training process. These forgetting events occur when samples, classified correctly in a prior epoch, are subsequently predicted incorrectly. Unfortunately, existing studies fall short of balancing the incorporation of training dynamics with the effective identification of well-generalized samples. The use of an averaging down-sampling operation aims to cover a broader range of training dynamics. However, this approach can inadvertently lead to the neglect of well-generalized samples during data selection, ultimately degrading overall performance.

In this paper, we propose **Temporal Dual-Depth Scoring (TDDS)** pruning method, designed to effectively balance and address the aforementioned problems. We employ a dual-depth temporal scoring metric that spans the training process for data selection. *In the inner level (depth one)*, we approximate the actual contribution of each sample in every epoch. This approximation is achieved by projecting the sample-wise gradients onto the accumulated gradient in each epoch, following the approach suggested by Du et al. [11]. We then integrate these contributions over a specified window, typically spanning 10 epochs, to incorporate more training dynamics. *In the outer level (depth two)*, we reduce reliance on the averaging down-sampling operation to prevent overlooking vital samples. Well-generalized samples are identified based on the variability of the projected gradients, calculated in each window from the inner level. This design of the outer level draws inspiration from He et al. [15], who tracked the variance of predicted target probabilities throughout training to enhance pruning effectiveness.

<sup>1</sup>Usually, the first 10 or 20 epochs.

The pipeline of our proposed TDDS is clearly illustrated in Figure 2. Extensive experiments demonstrate the superior generalization capability of the proposed TDDS, enabling it to outperform other comparison methods across different pruning settings on various datasets and networks. For example, on CIFAR-100, our method achieves 54.51% accuracy with only 10% training data, surpassing random selection by 7.83% and other SOTA methods by at least 12.69%.

Our main contributions are summarized as:

- We examine the poor generalization resulting from existing snapshot-based dataset pruning. Subsequently, we analyze the inefficacy of current solutions to this generalization issue, attributing it to the inability to balance the incorporation of training dynamics with the effective identification of well-generalized samples.
- We thereby propose a novel dataset pruning method called **Temporal Dual-Depth Scoring (TDDS)**, resorting to two temporal depths/levels to take respective care of the contribution evaluation and generalization spanning training progress. Experiments on diverse datasets and networks demonstrate the effectiveness of TDDS, achieving SOTA performance.

## 2. Related Works

### 2.1. Dataset Pruning

Dataset Pruning, also known as Coreset Selection aims to shrink the dataset scale by selecting important samples according to some predefined criteria. Herding [6, 45] and Moderate [46] calculate the distance in feature space. Entropy [8] and Cal [24] explore the uncertainty and decision boundary with the predicted outputs. GradND/EL2N [31] quantifies the importance of a sample with its gradient magnitude. Most of these methods are snapshot-based, that is they do not care about the training dynamics that actually distinguish the behaviors of samples in the overall optimization. There are some efficient training methods [10, 28, 42, 47], which focus more on training efficiency and execute online selection spanning training progress. For example, gradient matching-based methods GradMatch [18] and Craig [26] minimize the distance between gradients produced by the full dataset and coreset. ACS [17] selects samples with larger gradient magnitude to accelerate quantization-aware training. These methods can capture training dynamics but often necessitate repeated importance re-evaluating and can not be generalized to different networks. Some methods such as Forgetting [43] and AUM [33] also have proved training dynamics benefits samples importance assessment. However, the downsampling operation, which averages dynamics across training, disregards the variability crucial for evaluating sample generalization, ultimately resulting in suboptimal coresets.

## 2.2. Variance in Deep Learning

Variance, a standard second-order statistic, revealing the variability along a certain dimension, has been explored in numerous works [4, 30, 34]. For example, MVT [30] transfers the observed variance of one class to another to improve the robustness of the unseen examples and few-shot learning tasks. VICReg [4] maintains variance of embeddings to prevent the collapse in which two branches produce identical outputs regardless of inputs. For the gradient, its variance along the parameter dimension is an approximation of Hessian matrix indicating loss landscape. Based on that, Fishr [34] matches domain-level variances of gradients for out-of-distribution generalization. LCMat [37] induces a better adaptation of the pruned dataset on the perturbed parameter region than the exact point matching. Besides, calculating gradient variance along time dimension tracks dynamic optimization change. VoG [1] uses this to estimate sample importance and the visualization shows that high-variance samples are more worthwhile to learn. Dyn-Unc [15] calculates the time-dimension variance of predicted target class probability to evaluate sample importance. Differing from VoG and Dyn-Unc, our TDDS does not need pixel-wise gradient calculation while considering non-target probabilities, which enables a more precise evaluation with lower computational overhead.

## 3. Preliminary

Throughout this paper, we denote the full training dataset as  $\mathbb{U} = \{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N$ , where  $\mathbf{x}_n \in \mathbb{R}^D$  and  $\mathbf{y}_n \in \mathbb{R}^{1 \times C}$  are drawn i.i.d. from a natural distribution  $\mathcal{D}$ . We denote a neural network parameterized with weight matrix  $\theta$  as  $f_\theta$ . Under these definitions, the optimization objective of  $f_\theta$  on  $\mathbb{U}$  is to minimize empirical risk  $\mathcal{L}(\mathbb{U}; \theta) = \frac{1}{N} \sum_{n=1}^N \ell(f_\theta(\mathbf{x}_n), \mathbf{y}_n)$ , where  $f_\theta(\mathbf{x}_n) \in \mathbb{R}^{1 \times C}$  outputs the predicted probabilities of each class. Typically,  $\theta$  is updated as follow<sup>2</sup>,

$$\theta_{t+1} = \theta_t - \eta \sum_{n=1, \mathbf{x}_n \in \mathbb{U}}^N \mathbf{g}_t(\mathbf{x}_n), \quad (1)$$

where  $\eta$  is the learning rate and  $\mathbf{g}_t(\mathbf{x}_n)$  is defined as the gradient calculated over the data pair  $(\mathbf{x}_n, \mathbf{y}_n)$  in epoch  $t$ .

Dataset pruning aims to enhance training efficiency by constructing a coreset  $\mathbb{S} = \{(\mathbf{x}_m, \mathbf{y}_m)\}_{m=1}^M$ ,  $\mathbb{S} \subset \mathbb{U}$  that can achieve comparable performance to  $\mathbb{U}$ . The objective of dataset pruning can be formulated as,

$$\mathbb{E}_{\substack{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D} \\ \theta_0 \sim \mathcal{P}_{\theta_0}}} [\ell(f_{(\mathbb{U}, \theta_0)}(\mathbf{x}), \mathbf{y})] \simeq \mathbb{E}_{\substack{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D} \\ \theta_0 \sim \mathcal{P}_{\theta_0}}} [\ell(f_{(\mathbb{S}, \theta_0)}(\mathbf{x}), \mathbf{y})], \quad (2)$$

<sup>2</sup>Here,  $t$  indicates epoch and the batch-wise updates [2] in one epoch are omitted.

where  $f_{(\mathbb{U}, \theta_0)}$  and  $f_{(\mathbb{S}, \theta_0)}$  represent the networks trained on  $\mathbb{U}$  and  $\mathbb{S}$  with weight  $\theta_0$  initialized from distribution  $\mathcal{P}_{\theta_0}$ .

## 4. Methodology

As discussed earlier, despite that exploring long-range dynamics provides a more comprehensive sample evaluation, the averaging operation hinders the construction of a coreset with strong generalization. In this section, we present the proposed **Temporal Dual-Depth Scoring (TDDS)** strategy, which balances the incorporation of training dynamics and the identification of well-generalized samples. In the outer level, we identify samples with their variability during training which is encouraged theoretically by the optimization space matching between coreset and original dataset (refer to Section 4.1). In the inner level, the gradient in the outer level is replaced with the projected component on the accumulated direction to estimate valid contributions in every epoch (refer to Section 4.2). Finally, to enhance efficiency, we conduct TDDS with a moving window (refer to Section 4.3). The proposed TDDS is summarized in **Algorithm 1**.

### 4.1. Outer Level: Enhance Coreset Generalization by Maximizing Variance

To construct the coreset  $\mathbb{S}$  satisfying **Equation 2**, gradient-based methods [3, 18, 26] propose to minimize the distance between full dataset and coreset gradient,

$$\mathbb{S}^* = \arg \min_{\mathbb{S} \subset \mathbb{U}} \|\mathbf{G}_{t, \mathbb{U}} - \tilde{\mathbf{G}}_{t, \mathbb{S}}\|, \quad \text{where} \\ \mathbf{G}_{t, \mathbb{U}} = \sum_{\substack{n=1, \\ \mathbf{x}_n \in \mathbb{U}}}^{|\mathbb{U}|} \mathbf{g}_t(\mathbf{x}_n), \quad \tilde{\mathbf{G}}_{t, \mathbb{S}} = \sum_{\substack{m=1, \\ \mathbf{x}_m \in \mathbb{S}}}^{|\mathbb{S}|} \mathbf{g}_t(\mathbf{x}_m), \quad (3)$$

$\tilde{\mathbf{G}}_{t, \mathbb{S}}$  can be regarded as the gradient reconstruction in coreset space. To solve this minimization problem, Huang et al. [17] recommends pruning samples with smaller gradient magnitudes to minimize their potential impact on optimization.

However, whenever the distance is minimized, the obtained snapshot-based coreset fails to generalize across training progress. Here are the reasons: (1) the distinctive contributions of samples are temporal statistics. Evaluating them at a specific time point causes significant bias when considering dynamic optimization; (2) samples with larger individual gradients dominate the matching process. Without considering the current optimization context, these samples introduce a substantial error in approximating the accumulated gradient.

To address the first issue, we proposed to extend **Equation 3** to a continuous training duration to capture training dynamics,

$$\mathbb{S}^* = \arg \min_{\mathbb{S} \subset \mathbb{U}} \frac{1}{T} \sum_{t=1}^T \|\mathbf{g}_{t, \mathbb{U}} - \tilde{\mathbf{g}}_{t, \mathbb{S}}\|^2, \quad (4)$$

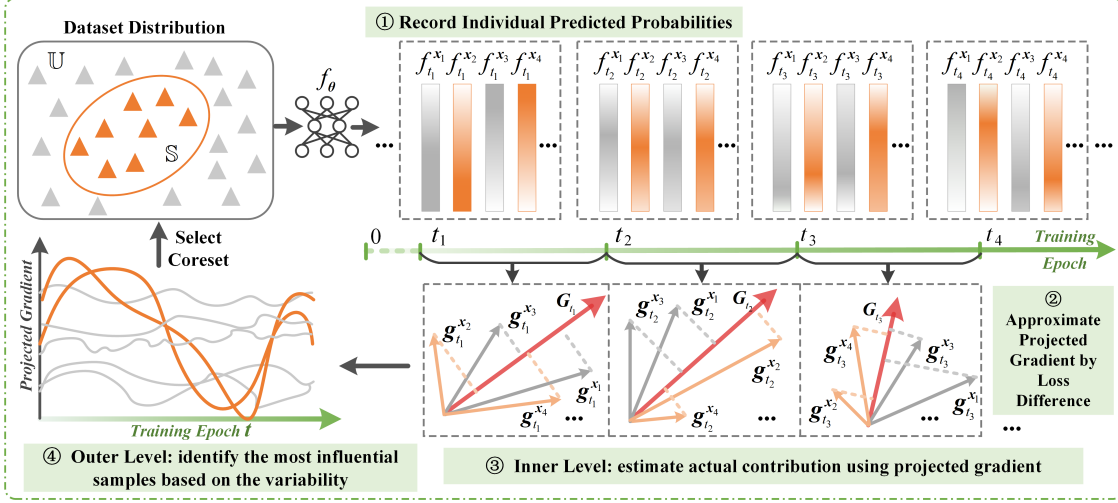


Figure 2. The pipeline of our proposed TDDS. First, we record individual predicted probabilities of samples  $f_{t_i}^{x_n} = f_{\theta_{t_i}}(\mathbf{x}_n)$  during training. Then, we approximate the projected gradient, i.e. the component of gradient  $\mathbf{g}^{x_n} = \mathbf{g}_{t_i}(\mathbf{x}_n)$  on accumulated gradient  $\mathbf{G}_{t_i}$  direction, with KL-based loss difference between two adjacent epochs. Next, the projected gradients are used to estimate individual contributions of samples. Finally, the important samples are identified with larger projected gradient variances over time.

Here, the gradients generated by  $\mathbb{U}$  and  $\mathbb{S}$  at epoch  $t$  are defined as  $\mathcal{G}_{t,\mathbb{U}} = \|\mathbf{g}_t(\mathbf{x}_n)\|_{n=1}^N$  and similarly  $\mathcal{G}_{t,\mathbb{S}}$  is its reconstructed counterpart using  $\mathbb{S}$ . Here we use magnitude instead of gradient itself to balance the training dynamic integration and complexity. Experiments in Section 5.4 also shows that magnitude-based minimization performs better than gradient-based when considering continuous training.

Directly solving Equation 4 is NP-hard. Thus, we turn to its equivalent, temporal gradient variance maximization,

$$\mathbb{S}^* = \arg \max_{\mathbb{S} \subset \mathbb{U}} \mathcal{R}(\mathbb{S}), \quad \mathcal{R}(\mathbb{S}) = \sum_{t=1}^T \|\mathcal{G}_{t,\mathbb{S}} - \bar{\mathcal{G}}_{\mathbb{S}}\|^2, \quad (5)$$

where  $\bar{\mathcal{G}}_{\mathbb{S}} = \frac{1}{T} \sum_{t=1}^T \mathcal{G}_{t,\mathbb{S}}$ ,  $\mathcal{G}_{t,\mathbb{S}} = \|\mathbf{g}_t(\mathbf{x}_m)\|_{m=1}^M$  is the mean gradient over training time. Detailed proof can be founded in Supplementary. Equation 5 measures the variability of samples and enhances the identification of well-generalized samples compared to previous average-based methods [33, 43].

#### 4.2. Inner Level: Estimate Actual Contribution with Projected Gradient

According to Equation 5, sample importance can be evaluated with its gradient variance over time. For any sample  $\mathbf{x}_n$ , we have,

$$\mathcal{R}(\mathbf{x}_n) = \sum_{t=1}^T \|\mathbf{g}_t(\mathbf{x}_n) - \overline{\mathbf{g}(\mathbf{x}_n)}\|^2, \quad (6)$$

where  $\overline{\mathbf{g}(\mathbf{x}_n)} = \frac{1}{T} \sum_{t=1}^T \mathbf{g}_t(\mathbf{x}_n)$ . Equation 6 treats sample-wise gradient independently, which may not pre-




Figure 3. Incorporating sample-wise gradient in the current optimization context. After projecting onto the current accumulated gradient, samples with similar gradients exhibit different behaviors.

cisely describe the valid contribution in the current optimization context. That is the second problem mentioned in Section 4.1. As illustrated in Figure 3(a), the gradient of sample  $\mathbf{x}_n$  remains constant between  $t$  and  $t+1$ . However, when incorporating it into the current optimization context, its valid contribution undergoes a significant shift. Thus, we propose to calculate the gradient variance after projecting it on the accumulated direction, that is,  $|\mathbf{g}_t(\mathbf{x}_n)|$  is reformulated as,

$$|\mathbf{g}_t(\mathbf{x}_n)| = |\mathbf{g}_t(\mathbf{x}_n)| \cdot \sum_{j=1, \mathbf{x}_j \in \mathbb{U}}^{|\mathbb{U}|} \mathbf{g}_t(\mathbf{x}_j). \quad (7)$$

Calculating Equation 6 requires sample-wise gradients over a period of  $T$ , which is not feasible due to time and memory constraints. According to Equation 1, the accumu-

lated gradient is equal to  $\frac{1}{\eta}(\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t+1})$ , thus the projection in Equation 7 is reformed as loss difference,

$$\begin{aligned} |\mathbf{g}_t(\mathbf{x}_n)| &= \frac{1}{\eta} |(\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t+1}) \nabla_{\boldsymbol{\theta}_t} \ell(f_{\boldsymbol{\theta}_t}(\mathbf{x}_n), \mathbf{y}_n)| \\ &\approx \frac{1}{\eta} |\ell(f_{\boldsymbol{\theta}_{t+1}}(\mathbf{x}_n), \mathbf{y}_n) - \ell(f_{\boldsymbol{\theta}_t}(\mathbf{x}_n), \mathbf{y}_n)|, \end{aligned} \quad (8)$$

which can be deduced from the first-order Taylor expansion.

Thus, we can evaluate the per-sample contribution to the accumulated gradient direction with its loss change. In the context of classification,  $\ell$  is typically Cross Entropy (CE) loss,

$$\Delta \ell_t^n = \mathbf{y}_n^\top \cdot \log \frac{f_{\boldsymbol{\theta}_{t+1}}(\mathbf{x}_n)}{f_{\boldsymbol{\theta}_t}(\mathbf{x}_n)}, \quad (9)$$

where  $\Delta \ell_t^n = \ell(f_{\boldsymbol{\theta}_{t+1}}(\mathbf{x}_n), \mathbf{y}_n) - \ell(f_{\boldsymbol{\theta}_t}(\mathbf{x}_n), \mathbf{y}_n)$ . Due to the one-hot  $\mathbf{y}_n$ , non-target probabilities are unfortunately overlooked, obliterating a significant amount of distinctive information. This reminds us of the Kullback-Leibler (KL) divergence loss, which can be achieved by simply replacing  $\mathbf{y}_n$  with  $f_{\boldsymbol{\theta}_{t+1}}(\mathbf{x}_n)$ ,

$$\Delta \ell_t^n = f_{\boldsymbol{\theta}_{t+1}}(\mathbf{x}_n)^\top \cdot \log \frac{f_{\boldsymbol{\theta}_{t+1}}(\mathbf{x}_n)}{f_{\boldsymbol{\theta}_t}(\mathbf{x}_n)}. \quad (10)$$

Considering that  $f_{\boldsymbol{\theta}_t}(\mathbf{x}_n)$  follows a Bernoulli distribution  $\mathcal{B}(p_t)$ , where  $p_t$  represents the probability of the target class at epoch  $t$ . If  $p_t$  equals  $p_{t+1}$ , Equation 9 fails to distinguish, but Equation 10 can still differentiate the contributions of samples with  $q_t$  and  $q_{t+1}$ . Here,  $q_t$  and  $q_{t+1}$  are the non-target probability distributions. This effectiveness has been proved in Section 5.4.

### 4.3. Efficient Implementation of TDDS

When computing  $\mathcal{R}(\mathbf{x}_n)$  with Equation 6, Equation 8, and Equation 10, storing  $T$ -epoch predicted outputs may lead to out-of-memory problems, especially with large datasets like ImageNet (containing 1, 281, 167 samples across 1, 000 classes). As a solution, we suggest an efficient calculation using an exponential moving average (EMA) weighting strategy. For any training point  $t \geq K$ , the variance  $\mathcal{R}_t(\mathbf{x}_n)$ <sup>3</sup> is computed over a  $K$ -epoch window ranging from  $t - K + 1$  to  $t$ ,

$$\mathcal{R}_t(\mathbf{x}_n) = \sum_{t-K+1}^t \|\Delta \ell_t^n - \overline{\Delta \ell_t^n}\|^2, \quad (11)$$

where  $\overline{\Delta \ell_t^n}$  is the average of loss difference in a window. For  $T$ -epoch trajectory,  $T - K + 1$  windows are weighted as,

$$\mathcal{R}(\mathbf{x}_n) = \beta \mathcal{R}_t(\mathbf{x}_n) + (1 - \beta) \mathcal{R}(\mathbf{x}_n), \quad (12)$$

<sup>3</sup>The constant coefficient  $\eta$  has been omitted.

where  $\beta$  is the decay coefficient of EMA whose setting is examined in Section 5.3. With this EMA, we only need to store predicted outputs for one window, significantly alleviating storage burden.

After sorting  $\{\mathcal{R}(\mathbf{x}_n)\}_{n=1}^N$ , the top- $M$  samples are selected to construct coreset denoted as  $\{\mathbb{S}, r\} = \{(\mathbf{x}_m, \mathbf{y}_m), \mathcal{R}(\mathbf{x}_m)\}_{m=1}^M$ . Since, the samples with larger  $\mathcal{R}(\mathbf{x}_m)$  contribute more in the original data training. We remark that this dominance should be maintained in coreset training, i.e.,

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \eta \sum_{m=1}^M \mathbf{g}_t(\mathbf{x}_m) \mathcal{R}(\mathbf{x}_m). \quad (13)$$

This helps to transfer relative sample importance to coreset training. This effectiveness is proved in Section 5.4.

---

#### Algorithm 1 Temporal Dual-Depth Scoring (TDDS)

---

**Input:** Training dataset  $\mathbb{U}$ ; A network  $f_\theta$  with weight  $\theta$ ; Learning rate  $\eta$ ; Epochs  $T$ ; Iterations  $I$  per epoch; Window size  $K$ ; Decay factor  $\beta$ .

- 1: **for**  $t = 1$  to  $T$  **do**
  - 2:   **for**  $i = 1$  to  $I$ , Sample a mini-batch  $\mathbb{B}_i \subset \mathbb{U}$  **do**
  - 3:     Record predicted probabilities  $f_{\boldsymbol{\theta}_t}(\mathbf{x}_n)$ ,  $\mathbf{x}_n \in \mathbb{B}_i$
  - 4:     Calculate  $\Delta \ell_t^n$  for each  $\mathbf{x}_n \in \mathbb{B}_i$   
    $\triangleright$  Defined in Equation 10
  - 5:     Update weight  $\boldsymbol{\theta}_{i+1} = \boldsymbol{\theta}_i - \eta \mathbf{G}_{t, \mathbb{B}_i}$
  - 6:   **end for**
  - 7:   **if**  $K \leq t \leq T$  **then**
  - 8:     Calculate  $\mathcal{R}_t(\mathbf{x}_n)$  in one window,  $\mathbf{x}_n \in \mathbb{U}$   
    $\triangleright$  Defined in Equation 11
  - 9:     Update  $\mathcal{R}(\mathbf{x}_n)$  with EMA,  $\mathbf{x}_n \in \mathbb{U}$   
    $\triangleright$  Defined in Equation 12
  - 10:   **end if**
  - 11: **end for**
  - 12: Select Top- $M$   $\{\mathbb{S}, r\} \leftarrow \{(\mathbf{x}_m, \mathbf{y}_m), \mathcal{R}(\mathbf{x}_m)\}_{m=1}^M$
- Output:** Selected subset  $\mathbb{S} = \{(\mathbf{x}_m, \mathbf{y}_m)\}_{m=1}^M$ ; associated sample-wise importance  $r = \{\mathcal{R}(\mathbf{x}_m)\}_{m=1}^M$
- 

## 5. Experiments

### 5.1. Experiment Settings

**Datasets and Networks.** The effectiveness of the proposed TDDS is evaluated on three popular benchmarks, i.e., CIFAR-10, CIFAR-100 [19], and ImageNet-1K [9] with three networks, i.e., ResNet-18, ResNet-34 [14], and Swin-T [21]. We also verify the cross-architecture performance on four networks including ResNet-50 [14], VGG-16 [38], MobileNet-v2 [36], and ShuffleNet [23].

**Training Hyperparameters.** Theoretically, if pruning rate  $p = 1 - \frac{|\mathbb{S}|}{|\mathbb{U}|}$ , where  $0 < p \leq 1$ , the training process

Table 1. Comparison of different dataset pruning methods on CIFAR-10 and 100 with ResNet-18 under Strategy-E. The model trained with the full dataset achieves 95.23% and 78.21% accuracy.

| $p$             | CIFAR-10                   |                            |                            |                            |                            | CIFAR-100                  |                            |                            |                            |                            |
|-----------------|----------------------------|----------------------------|----------------------------|----------------------------|----------------------------|----------------------------|----------------------------|----------------------------|----------------------------|----------------------------|
|                 | 30%                        | 50%                        | 70%                        | 80%                        | 90%                        | 30%                        | 50%                        | 70%                        | 80%                        | 90%                        |
| Random          | 94.58<br>$\pm 0.04$        | 93.38<br>$\pm 0.17$        | 90.61<br>$\pm 0.44$        | 88.87<br>$\pm 0.47$        | 83.77<br>$\pm 0.26$        | 75.53<br>$\pm 0.04$        | 71.95<br>$\pm 0.16$        | 64.59<br>$\pm 0.32$        | 57.79<br>$\pm 0.24$        | 46.68<br>$\pm 1.07$        |
| Entropy [8]     | 94.45<br>$\pm 0.07$        | 91.90<br>$\pm 0.16$        | 86.24<br>$\pm 0.26$        | 83.49<br>$\pm 0.21$        | 72.06<br>$\pm 0.81$        | 72.39<br>$\pm 0.20$        | 64.44<br>$\pm 0.36$        | 50.73<br>$\pm 0.86$        | 42.86<br>$\pm 0.25$        | 29.56<br>$\pm 0.54$        |
| Forgetting [43] | 95.45<br>$\pm 0.24$        | 95.05<br>$\pm 0.05$        | 89.14<br>$\pm 2.04$        | 76.18<br>$\pm 3.18$        | 45.87<br>$\pm 1.87$        | 77.38<br>$\pm 0.09$        | 70.76<br>$\pm 0.40$        | 49.92<br>$\pm 0.28$        | 38.42<br>$\pm 1.13$        | 25.82<br>$\pm 0.52$        |
| EL2N [31]       | 95.43<br>$\pm 0.10$        | 95.06<br>$\pm 0.04$        | 86.69<br>$\pm 1.71$        | 68.64<br>$\pm 3.70$        | 31.89<br>$\pm 1.51$        | 76.89<br>$\pm 0.31$        | 67.57<br>$\pm 0.15$        | 36.45<br>$\pm 1.36$        | 17.31<br>$\pm 0.33$        | 9.10<br>$\pm 0.69$         |
| AUM [33]        | 95.44<br>$\pm 0.09$        | 95.19<br>$\pm 0.09$        | 91.19<br>$\pm 0.63$        | 69.60<br>$\pm 3.11$        | 34.74<br>$\pm 0.11$        | 77.35<br>$\pm 0.18$        | 68.17<br>$\pm 0.52$        | 31.69<br>$\pm 0.34$        | 18.43<br>$\pm 0.47$        | 9.29<br>$\pm 0.27$         |
| Moderate [46]   | 93.96<br>$\pm 0.06$        | 92.34<br>$\pm 0.09$        | 89.71<br>$\pm 0.14$        | 87.75<br>$\pm 0.27$        | 83.61<br>$\pm 0.24$        | 74.60<br>$\pm 0.10$        | 70.29<br>$\pm 0.31$        | 62.81<br>$\pm 0.08$        | 56.52<br>$\pm 0.37$        | 41.82<br>$\pm 1.12$        |
| Dyn-Unc[15]     | 95.08<br>$\pm 0.02$        | 94.03<br>$\pm 0.14$        | 89.40<br>$\pm 0.13$        | 79.76<br>$\pm 1.09$        | 37.12<br>$\pm 1.12$        | 73.36<br>$\pm 0.10$        | 65.90<br>$\pm 0.25$        | 50.16<br>$\pm 0.47$        | 39.19<br>$\pm 0.27$        | 15.20<br>$\pm 0.41$        |
| <b>TDDS</b>     | <b>95.47</b><br>$\pm 0.06$ | <b>95.21</b><br>$\pm 0.04$ | <b>93.03</b><br>$\pm 0.25$ | <b>91.30</b><br>$\pm 0.25$ | <b>85.46</b><br>$\pm 0.21$ | <b>77.56</b><br>$\pm 0.06$ | <b>74.04</b><br>$\pm 0.34$ | <b>67.78</b><br>$\pm 0.44$ | <b>63.01</b><br>$\pm 0.12$ | <b>54.51</b><br>$\pm 0.22$ |




Figure 4. Parameter analysis of range  $T$  and window size  $K$  on CIFAR-100 with ResNet-18. From left to right, the corresponding pruning rates are 0.3, 0.5, 0.7, 0.8, and 0.9.

Table 2. Comparison of different dataset pruning methods on CIFAR-10 and 100 with ResNet-18 under Strategy-P. The model trained with the full dataset achieves 95.23% and 78.21% accuracy.

| $p$             | CIFAR-10                   |                            |                            | CIFAR-100                  |                            |                            |
|-----------------|----------------------------|----------------------------|----------------------------|----------------------------|----------------------------|----------------------------|
|                 | 30%                        | 50%                        | 70%                        | 30%                        | 50%                        | 70%                        |
| Random          | 94.33<br>$\pm 0.17$        | 93.4<br>$\pm 0.17$         | 90.94<br>$\pm 0.38$        | 74.59<br>$\pm 0.27$        | 71.07<br>$\pm 0.4$         | 65.3<br>$\pm 0.21$         |
| Entropy [8]     | 94.44<br>$\pm 0.2$         | 92.11<br>$\pm 0.47$        | 85.67<br>$\pm 0.71$        | 72.26<br>$\pm 0.08$        | 63.26<br>$\pm 0.29$        | 50.49<br>$\pm 0.88$        |
| Forgetting [43] | 95.36<br>$\pm 0.13$        | 95.29<br>$\pm 0.18$        | 90.56<br>$\pm 1.8$         | 76.91<br>$\pm 0.32$        | 68.6<br>$\pm 1.02$         | 38.06<br>$\pm 1.14$        |
| EL2N [31]       | 95.44<br>$\pm 0.15$        | 94.61<br>$\pm 0.20$        | 87.48<br>$\pm 1.33$        | 76.25<br>$\pm 0.24$        | 65.90<br>$\pm 1.06$        | 34.42<br>$\pm 1.50$        |
| AUM [33]        | 95.07<br>$\pm 0.24$        | 95.26<br>$\pm 0.15$        | 91.36<br>$\pm 1.4$         | 76.93<br>$\pm 0.32$        | 67.42<br>$\pm 0.49$        | 30.64<br>$\pm 0.58$        |
| Moderate [46]   | 93.86<br>$\pm 0.11$        | 92.58<br>$\pm 0.30$        | 90.56<br>$\pm 0.27$        | 74.60<br>$\pm 0.41$        | 71.10<br>$\pm 0.24$        | 65.34<br>$\pm 0.41$        |
| CCS [51]        | 95.40<br>$\pm 0.12$        | 95.04<br>$\pm 0.37$        | 93.00<br>$\pm 0.16$        | 77.14<br>$\pm 0.31$        | 74.45<br>$\pm 0.16$        | 68.92<br>$\pm 0.12$        |
| <b>TDDS</b>     | <b>95.50</b><br>$\pm 0.07$ | <b>95.66</b><br>$\pm 0.08$ | <b>93.92</b><br>$\pm 0.04$ | <b>77.95</b><br>$\pm 0.30$ | <b>74.98</b><br>$\pm 0.16$ | <b>69.46</b><br>$\pm 0.10$ |

speeds up by  $\frac{1}{1-p}$  times. We consider two coreset training strategies: **Strategy-E** prioritizes efficiency with same-epoch training on both full dataset and coreset. **Strategy-P**

Table 3. Comparison of different dataset pruning methods on ImageNet-1K with ResNet-34 under Strategy-E. The model trained with the full dataset achieves 73.54% Top-1 accuracy.

| $p$             | ImageNet-1K  |              |              |
|-----------------|--------------|--------------|--------------|
|                 | 70%          | 80%          | 90%          |
| Random          | 64.19        | 60.76        | 52.63        |
| Entropy[8]      | 62.34        | 56.80        | 43.39        |
| Forgetting [43] | 64.29        | 62.01        | 52.14        |
| EL2N [31]       | 46.92        | 32.68        | 15.90        |
| AUM [33]        | 39.34        | 23.64        | 11.70        |
| Moderate [46]   | 64.04        | 61.35        | 52.45        |
| <b>TDDS</b>     | <b>64.69</b> | <b>62.56</b> | <b>53.91</b> |

prioritizes performance by training coreset for more epochs. If  $\frac{|S|}{|U|} = 1 - p$ , the training process is extended to  $\frac{1}{1-p}$  times of the original one. For CIFAR-10 and -100, we train ResNet-18 for 200 epochs with a batch size of 128. For ImageNet, ResNet-34 and Swin-T are trained for 60 [51] and 300 epochs [15] with batch sizes of 256 and 1024, respec-




Figure 5. Parameter analysis of decay coefficient  $\beta$  on CIFAR-100 with ResNet-18. From left to right, the corresponding pruning rates are 0.7, 0.8, and 0.9.

tively. SGD optimizer with momentum of 0.9 and weight decay of 0.0005 is used for update networks. The learning rate is initialized as 0.1 and decays with the cosine annealing scheduler. During experiments, we found that smaller batch size boosts performance at high pruning rates (shown in Supplementary). Thus, for  $p = 0.8$ , we halved the batch size, and for  $p = 0.9$ , we reduced it to one-fourth. All compared methods use the same settings. With Strategy-P, for example, training a CIFAR coreset with only 50% training samples takes 400 epochs, while keeping the other hyper-parameters unchanged. Considering the time-consuming training, we only conduct Strategy-P experiments at pruning rates of 0.3, 0.5, 0.7.

**Baselines.** Eight baselines are used for comparison, with the latter seven being SOTA methods: 1) **Random**; 2) **Entropy** [8]; 3) **Forgetting** [43]; 4) **EL2N** [31]; 5) **AUM** [33]; 6) **Moderate** [46]; 7) **CCS** [51]; 8) **Dyn-Unc** [15]. Due to the limited page, we provide the technical details of these baselines in Supplementary.

## 5.2. Benchmark Evaluation Results

**ResNet-18 on CIFAR-10 and 100.** Table 1 and Table 2 report the performance of different dataset pruning methods under Strategy-E and P, respectively. Under both strategies, TDDS consistently outperforms other methods. For instance, under Strategy-E, none of the compared methods surpasses Random at aggressive pruning rates, while TDDS achieves 85.46% and 54.51% accuracy after pruning 90% samples of CIFAR-10 and 100 surpassing Random by 1.69% and 7.83%, respectively. By extending the training duration, Strategy-P enhances the performance of all methods to varying degrees. TDDS still maintains its superiority, for example, with only 30% training data, the accuracy on CIFAR-10 and -100 are 93.92% and 69.46% outperforming CCS by 0.92% and 0.54%, respectively. These results highlight TDDS as a promising choice for both efficiency and performance-oriented dataset pruning.

**ResNet on ImageNet-1K.** For ImageNet-1K, we train ResNet-34 models on original dataset for 90 epochs to collect training dynamics and then train the model on coreset for 60 epochs as [51]. Experiment results are given in Table 3. The results show that the proposed TDDS has better performance on ImageNet-1K, especially under aggressive

pruning rates. For example, when pruning 90% samples, TDDS still achieves 53.91% which is 1.77% and 1.46% higher than Forgetting and Moderate, respectively.

**Swin-T on ImageNet-1k.** We also evaluate Swin-T on ImageNet-1K. Like Dyn-Unc [15], Table 4 reports the results under four pruning rates. Notably, TDDS consistently outperforms other competitors, for example, after pruning 50% training data, TDDS achieves 78.33% accuracy, surpassing Forgetting, Moderate, and Dyn-Unc by 4.01%, 3.35% and 0.69%, respectively.

Table 4. Comparison of different dataset pruning methods on ImageNet-1K with Swin-T under Strategy-E. The model trained with the full dataset achieves 79.58% Top-1 accuracy.

| $p$             | ImageNet-1K  |              |              |              |
|-----------------|--------------|--------------|--------------|--------------|
|                 | 25%          | 30%          | 40%          | 50%          |
| Random          | 77.82        | 77.18        | 75.93        | 74.54        |
| Forgetting [43] | 78.70        | 78.27        | 77.55        | 74.32        |
| Moderate [46]   | 77.74        | 77.06        | 75.94        | 74.98        |
| Dyn-Unc [15]    | 79.54        | 79.14        | 78.49        | 77.64        |
| <b>TDDS</b>     | <b>79.76</b> | <b>79.57</b> | <b>79.09</b> | <b>78.33</b> |

## 5.3. Parameter Analysis

**Estimation Range  $T$  and Window Size  $K$ .** Due to the random initialization, models adapt to prominent features (easy samples) during early training. As training progresses, models gradually acquire more abstract and complex features (hard samples) that are challenging to capture in early training. Thus, the importance of samples should be estimated in different training phases with respect to different pruning levels. As shown in Figure 4, we analyze the influence of range  $T$  and window size  $K$  on perfor-

Table 5. Ablation study for Dual-depth strategy. The numbers in square brackets indicate the improvement in accuracy of the method presented in the current column compared to the corresponding method in the previous column.

| $p$   | CIFAR-100 |               |               |
|-------|-----------|---------------|---------------|
|       | EL2N      | EL2N+outer    | TDDS          |
| 30%   | 76.89     | 77.47[+ 0.58] | 77.56[+ 0.09] |
| 50%   | 67.57     | 73.93[+ 6.36] | 74.04[+ 0.11] |
| 70%   | 36.45     | 67.04[+30.59] | 67.78[+ 0.74] |
| 80%   | 17.31     | 58.21[+40.90] | 63.01[+ 1.80] |
| 90%   | 9.10      | 40.36[+31.26] | 54.51[+14.15] |
| outer | ✗         | ✓             | ✓             |
| inner | ✗         | ✗             | ✓             |




Figure 6. Ablation studies for (a) Gradient Magnitude vs. Gradient, (b) KL-based Loss Difference, and (c) Importance-aware Coreset Training. In (a), GM and G represent estimating sample contribution with gradient magnitude and gradient, respectively. In (b), KL and CE represent calculating loss difference with KL and CE loss, respectively. In (c), Imp-aware and Imp-unaware represent training coreset with and without weighting samples with their importance scores, respectively.

mance. Apparently, for a larger pruning rate, we should select samples earlier in training. For CIFAR-100, the optimal  $(p, T, K)$  setting is  $(0.3, 90, 10)$ ,  $(0.5, 70, 10)$ ,  $(0.7, 60, 20)$ ,  $(0.8, 20, 10)$ , and  $(0.9, 10, 5)$ . The parameter settings of CIFAR-10 and ImageNet are reported in Supplementary.

**Decay Coefficient  $\beta$  of EMA.** In contrast to the simple average (SA), EMA employs exponential weighting, giving more weight to the most recent window. We examine four values of  $\beta$ , where  $\beta = 0$  represents SA. As illustrated in Figure 5, EMA outperforms SA regardless of the  $\beta$  value chosen. Among the considered pruning levels,  $\beta = 0.9$  yields the best performance, surpassing  $\beta = 0.95$  and  $\beta = 0.85$ .

## 5.4. Ablation Studies

**Gradient Magnitude vs Gradient.** We use gradient magnitudes in Equation 4. Although gradients themselves appear capable of describing optimization, our experimental findings indicate that the use of gradient magnitudes yields superior performance, as depicted in Figure 6(a). We argue that gradients often prioritize samples with frequently reversing gradients, resulting in unstable training and slower convergence.

**KL-based Loss Difference.** As discussed in Section 4.3, the KL loss provides a more comprehensive assessment of loss difference and is expected to precisely characterize sample contribution. To evaluate its effectiveness, we replace KL loss with CE loss and conduct experiments across all considered pruning rates. As illustrated in Figure 6(b), although CE is slightly higher than KL when the pruning rate is 30%, KL exhibits an absolute advantage at all other pruning levels. For example, when pruning rate is 90%, KL outperforms CE by a substantial margin of 15.97%.

**Importance-aware Coreset Training.** We formulate importance-aware subset training as Equation 13. The maintained relative importance provides optimization guid-

ance for subset training. As shown in Figure 6(c), the advantage becomes increasingly pronounced as the pruning rate rises. When pruning rate is 90%, the accuracy of importance-aware subset training outperforms its unaware counterpart by 6.55%.

**Dual-Depth Strategy.** To investigate the effectiveness of the proposed Dual-Depth strategy, we replace the inner depth with EL2N, and Table 5 indicates that incorporating the outer depth boosts EL2N at all considered pruning rates. Nevertheless, this enhancement still falls short of our dual-depth design, particularly evident in aggressive pruning scenarios.

## 6. Conclusion

In this study, we reveal that the snapshot-based data pruning results in poor generalization across various pruning scenarios. However, existing methods fail to identify well-generalized samples even with evaluation criteria considering the whole training progress. We attribute this to the averaging-based dynamics incorporation and imprecise sample contribution characterization. To address these problems, we propose a dataset pruning method called TDDS, which designs a Temporal Dual-Depth Scoring strategy to balance the incorporation of training dynamics and the identification of well-generalized samples. Extensive experiments on various datasets and networks, demonstrate the SOTA performance of our proposed method.

## Acknowledgements

This research is supported in part by the National Natural Science Foundation of China under Grant 62121001, Grant 62322117, and Grant 62371365. This research is also supported by Joey Tianyi Zhou’s A\*STAR SERC Central Research Fund (Use-inspired Basic Research) and Jiawei Du’s A\*STAR Career Development Fund (CDF) C233312004.



## References

- [1] Chirag Agarwal, Daniel D’souza, and Sara Hooker. Estimating example difficulty using variance of gradients. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 10368–10378, 2022. [3](#)
- [2] Shun-ichi Amari. Backpropagation and stochastic gradient descent method. *Neurocomputing*, 5(4-5):185–196, 1993. [3](#)
- [3] Lukas Balles, Giovanni Zappella, and Cédric Archambeau. Gradient-matching coresets for rehearsal-based continual learning. *arXiv preprint arXiv:2203.14544*, 2022. [3](#)
- [4] Adrien Bardes, Jean Ponce, and Yann LeCun. Vireg: Variance-invariance-covariance regularization for self-supervised learning. In *Int. Conf. Learn. Represent.*, 2021. [3](#)
- [5] George Cazenavette, Tongzhou Wang, Antonio Torralba, Alexei A Efros, and Jun-Yan Zhu. Dataset distillation by matching training trajectories. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4750–4759, 2022. [1](#)
- [6] Yutian Chen, Max Welling, and Alex Smola. Super-samples from kernel herding. *arXiv preprint arXiv:1203.3472*, 2012. [2](#)
- [7] Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1290–1299, 2022. [1](#)
- [8] Cody Coleman, Christopher Yeh, Stephen Mussmann, Baharan Mirzasoleiman, Peter Bailis, Percy Liang, Jure Leskovec, and Matei Zaharia. Selection via proxy: Efficient data selection for deep learning. In *Int. Conf. Learn. Represent.*, 2019. [1](#), [2](#), [6](#), [7](#)
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 248–255, 2009. [1](#), [5](#)
- [10] Hadi M. Dolatabadi, Sarah M. Erfani, and Christopher Leckie. L-robustness and beyond: Unleashing efficient adversarial training. In *Eur. Conf. Comput. Vis.*, 2022. [2](#)
- [11] Jiawei Du, Hanshu Yan, Jiashi Feng, Joey Tianyi Zhou, Liangli Zhen, Rick Siow Mong Goh, and Vincent Tan. Efficient sharpness-aware minimization for improved training of neural networks. In *International Conference on Learning Representations*, 2021. [2](#)
- [12] Jiawei Du, Yidi Jiang, Vincent YF Tan, Joey Tianyi Zhou, and Haizhou Li. Minimizing the accumulated trajectory error to improve dataset distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3749–3758, 2023. [1](#)
- [13] Chengcheng Guo, Bo Zhao, and Yanbing Bai. Deepcore: A comprehensive library for coreset selection in deep learning. In *Int. Conf. Database. and Expert Syst. Appl.*, pages 181–195, 2022. [1](#)
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 770–778, 2016. [5](#)
- [15] MUYANG HE, Shuo Yang, Tiejun Huang, and Bo Zhao. Large-scale dataset pruning with dynamic uncertainty. *arXiv preprint arXiv:2306.05175*, 2023. [2](#), [3](#), [6](#), [7](#)
- [16] Yang He, Lingao Xiao, Joey Tianyi Zhou, and Ivor Tsang. Multisize dataset condensation. *arXiv preprint arXiv:2403.06075*, 2024. [1](#)
- [17] Xijie Huang, Zechun Liu, Shih-Yang Liu, and Kwang-Ting Cheng. Efficient quantization-aware training with adaptive coreset selection. *arXiv preprint arXiv:2306.07215*, 2023. [1](#), [2](#), [3](#)
- [18] Krishnateja Killamsetty, Sivasubramanian Durga, Ganesh Ramakrishnan, Abir De, and Rishabh Iyer. Grad-match: Gradient matching based data subset selection for efficient deep model training. In *Proc. Int. Conf. Mach. Learn.*, pages 5464–5474, 2021. [2](#), [3](#)
- [19] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. [5](#)
- [20] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, et al. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *Int. J. Comput. Vis.*, 128(7):1956–1981, 2020. [1](#)
- [21] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Int. Conf. Comput. Vis.*, pages 10012–10022, 2021. [1](#), [5](#)
- [22] Ilya Loshchilov and Frank Hutter. Online batch selection for faster training of neural networks. *arXiv preprint arXiv:1511.06343*, 2015. [1](#)
- [23] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *Eur. Conf. Comput. Vis.*, pages 116–131, 2018. [5](#)
- [24] Katerina Margatina, Giorgos Vernikos, Loïc Barrault, and Nikolaos Aletras. Active learning by acquiring contrastive examples. *arXiv preprint arXiv:2109.03764*, 2021. [2](#)
- [25] Max Marion, Ahmet Üstün, Luiza Pozzobon, Alex Wang, Marzieh Fadaee, and Sara Hooker. When less is more: Investigating data pruning for pretraining llms at scale. *arXiv preprint arXiv:2309.04564*, 2023. [1](#)
- [26] Baharan Mirzasoleiman, Jeff Bilmes, and Jure Leskovec. Coresets for data-efficient training of machine learning models. In *Proc. Int. Conf. Mach. Learn.*, pages 6950–6960, 2020. [2](#), [3](#)
- [27] Timothy Nguyen, Zhourong Chen, and Jaehoon Lee. Dataset meta-learning from kernel ridge-regression. *arXiv preprint arXiv:2011.00050*, 2020. [1](#)
- [28] Patrik Okanovic, Roger Waleffe, Vasilis Mageirakos, Konstantinos E Nikolakakis, Amin Karbasi, Dionysis Kalogerias, Nezihe Merve Gürel, and Theodoros Rekatsinas. Repeated random sampling for minimizing the time-to-accuracy of learning. *arXiv preprint arXiv:2305.18424*, 2023. [2](#)
- [29] Dongmin Park, Seola Choi, Doyoung Kim, Hwanjun Song, and Jae-Gil Lee. Robust data pruning under label noise via maximizing re-labeling accuracy. *arXiv preprint arXiv:2311.01002*, 2023. [1](#)
- [30] Seong-Jin Park, Seungju Han, Ji-Won Baek, Insoo Kim, Juhwan Song, Hae Beom Lee, Jae-Joon Han, and Sung Ju

- Hwang. Meta variance transfer: Learning to augment from the others. In *Int. Conf. Learn. Represent.*, pages 7510–7520, 2020. [3](#)
- [31] Mansheej Paul, Surya Ganguli, and Gintare Karolina Dziugaite. Deep learning on a data diet: Finding important examples early in training. In *Adv. Neural Inform. Process. Syst.*, pages 20596–20607, 2021. [1](#), [2](#), [6](#), [7](#)
- [32] Mansheej Paul, Brett Larsen, Surya Ganguli, Jonathan Frankle, and Gintare Karolina Dziugaite. Lottery tickets on a data diet: Finding initializations with sparse trainable networks. In *Adv. Neural Inform. Process. Syst.*, pages 18916–18928, 2022. [1](#)
- [33] Geoff Pleiss, Tianyi Zhang, Ethan Elenberg, and Kilian Q Weinberger. Identifying mislabeled data using the area under the margin ranking. In *Adv. Neural Inform. Process. Syst.*, pages 17044–17056, 2020. [2](#), [4](#), [6](#), [7](#)
- [34] Alexandre Rame, Corentin Dancette, and Matthieu Cord. Fishr: Invariant gradient variances for out-of-distribution generalization. In *Int. Conf. Learn. Represent.*, pages 18347–18377, 2022. [3](#)
- [35] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 10684–10695, 2022. [1](#)
- [36] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4510–4520, 2018. [5](#)
- [37] Seungjae Shin, Heesun Bae, Donghyeok Shin, Weonyoung Joo, and Il-Chul Moon. Loss-curvature matching for dataset selection and condensation. In *AISTATS*, pages 8606–8628, 2023. [3](#)
- [38] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *Int. Conf. Learn. Represent.*, 2015. [5](#)
- [39] Ben Sorscher, Robert Geirhos, Shashank Shekhar, Surya Ganguli, and Ari Morcos. Beyond neural scaling laws: beating power law scaling via data pruning. In *Adv. Neural Inform. Process. Syst.*, pages 19523–19536, 2022. [1](#)
- [40] Emma Strubell, Ananya Ganesh, and Andrew McCallum. Energy and policy considerations for deep learning in nlp. *arXiv preprint arXiv:1906.02243*, 2019. [1](#)
- [41] Haoru Tan, Sitong Wu, Fei Du, Yukang Chen, Zhibin Wang, Fan Wang, and Xiaojuan Qi. Data pruning via moving-one-sample-out. In *Adv. Neural Inform. Process. Syst.*, 2023. [1](#)
- [42] Zhenyu Tang, Shaoting Zhang, and Xiaosong Wang. Exploring data redundancy in real-world image classification through data selection. *arXiv preprint arXiv:2306.14113*, 2023. [2](#)
- [43] Mariya Toneva, Alessandro Sordani, Remi Tachet des Combes, Adam Trischler, Yoshua Bengio, and Geoffrey J. Gordon. An empirical study of example forgetting during deep neural network learning. In *Int. Conf. Learn. Represent.*, 2019. [1](#), [2](#), [4](#), [6](#), [7](#)
- [44] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 12196–12205, 2022. [1](#)
- [45] Max Welling. Herding dynamical weights to learn. In *Proc. Int. Conf. Mach. Learn.*, pages 1121–1128, 2009. [1](#), [2](#)
- [46] Xiaobo Xia, Jiale Liu, Jun Yu, Xu Shen, Bo Han, and Tongliang Liu. Moderate coreset: A universal method of data selection for real-world data-efficient deep learning. In *Int. Conf. Learn. Represent.*, 2022. [1](#), [2](#), [6](#), [7](#)
- [47] Xilie Xu, Jingfeng Zhang, Feng Liu, Masashi Sugiyama, and Mohan Kankanhalli. Efficient adversarial contrastive learning via robustness-aware coreset selection. In *Adv. Neural Inform. Process. Syst.*, 2023. [2](#)
- [48] Ruonan Yu, Songhua Liu, and Xinchao Wang. Dataset distillation: A comprehensive review. *arXiv preprint arXiv:2301.07014*, 2023. [1](#)
- [49] Baharan Mirzasoleiman Yu Yang, Kang Hao. Towards sustainable learning: Coresets for data-efficient deep learning. In *Proc. Int. Conf. Mach. Learn.*, 2023. [1](#)
- [50] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Int. Conf. Comput. Vis.*, pages 3836–3847, 2023. [1](#)
- [51] Haizhong Zheng, Rui Liu, Fan Lai, and Atul Prakash. Coverage-centric coreset selection for high pruning rates. In *Int. Conf. Learn. Represent.*, 2022. [1](#), [6](#), [7](#)
- [52] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 633–641, 2017. [1](#)
- [53] Daquan Zhou, Kai Wang, Jianyang Gu, Xiangyu Peng, Dongze Lian, Yifan Zhang, Yang You, and Jiashi Feng. Dataset quantization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17205–17216, 2023. [1](#)