

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.

Enhancing Dataset Distillation via Non-Critical Region Refinement

Minh-Tuan Tran¹, Trung Le¹, Xuan-May Le², Thanh-Toan Do¹, Dinh Phung¹ ¹Monash University, ²The University of Melbourne

Abstract

Dataset distillation has become a popular method for compressing large datasets into smaller, more efficient representations while preserving critical information for model training. Data features are broadly categorized into two types: instance-specific features, which capture unique, fine-grained details of individual examples, and classgeneral features, which represent shared, broad patterns across a class. However, previous approaches often struggle to balance these features—some focus solely on classgeneral patterns, neglecting finer instance details, while others prioritize instance-specific features, overlooking the shared characteristics essential for class-level understanding. In this paper, we introduce the Non-Critical Region Refinement Dataset Distillation (NRR-DD) method, which preserves instance-specific details and fine-grained regions in synthetic data while enriching non-critical regions with class-general information. This approach enables models to leverage all pixel information, capturing both feature types and enhancing overall performance. Additionally, we present Distance-Based Representative (DBR) knowledge transfer, which eliminates the need for soft labels in training by relying on the distance between synthetic data predictions and one-hot encoded labels. Experimental results show that NRR-DD achieves state-of-the-art performance on both small- and large-scale datasets. Furthermore, by storing only two distances per instance, our method delivers comparable results across various settings. The code is available at https://github.com/tmtuan1307/ NRR-DD.

1. Introduction

Dataset distillation, also known as dataset condensation [2, 21, 26, 27, 32], has gained considerable attention as an effective method for compressing large datasets into smaller, more efficient representations, while preserving essential information critical for model training. By generating compact, high-quality data representations, dataset distillation reduces both storage requirements and computational costs associated with training on full-sized datasets [5, 8, 17]. This compression is particularly valuable in resource-constrained environments, such as edge devices or memory-limited systems, where training on large datasets is often impractical [1,9,23]. The primary goal of dataset distillation is to achieve high model performance while drastically reducing the amount of data that needs to be stored and processed [26]. During this process, synthetic data instances are generated to capture the key properties of the original data, enabling models to generalize effectively using only a fraction of the dataset [2, 32].

In recent years, several methods have been proposed to address this task, including Gradient Matching [10,32], Distribution Matching [31, 33], Trajectory Matching [2, 6], and more recent approaches for large-scale datasets [3, 21, 27]. However, designing an efficient distillation method remains a challenge, as it must capture both class-general features, which represent shared patterns, and instance-specific features, which highlight unique, fine-grained details. Previous methods often fall short by emphasizing one of these feature types over the other. Approaches focusing on classgeneral features [2,3,27,32,33] risk losing crucial instancespecific information, which hinders fine-grained generalization. Conversely, methods that prioritize instance-specific features [7,21] may neglect broader class patterns, leading to suboptimal class-level representation. To address these challenges, we introduce the Non-Critical Region Refinement Dataset Distillation (NRR-DD) method, which consists of three key stages:

- (i) Critical-based Initial Data Discovery: This stage involves selecting diverse and significant patches from the original dataset, which are then combined to capture instance-specific features.
- (ii) Non-Critical Region Refinement (NRR): In this stage, we apply Class Activation Mapping (CAM) [19, 34] to identify critical and non-critical regions in the images. The model preserves the critical regions, which contain fine-grained, instance-specific features, while refining the non-critical regions with more class-

general information. By balancing these two feature types, NRR enhances the dataset's comprehensiveness, improving both generalization and performance.

(iii) **Knowledge Transfer via Relabeling**: After training, the synthetic images are relabeled and used to transfer knowledge to a student model.

Additionally, recent research [3, 21, 27] has highlighted the importance of soft labels generated by a pretrained model in enhancing dataset distillation performance, particularly for large-scale, high-resolution datasets like ImageNet1k. However, this approach incurs substantial memory overhead; for instance, with 200 images per class (IPC) in ImageNet1k, it can require over 120 GB of storage [27].

To address this challenge, we propose a novel Distance-Based Representative (DBR) knowledge transfer technique that eliminates the need for traditional soft labels. DBR employs a distance-based approach to measure the discrepancy between predictions on synthetic data and one-hot encoded labels, simplifying the training process and reducing label storage requirements. For example, our method requires only 0.2 GB to store ImageNet1k (200 IPC), achieving a 500× reduction in storage while still delivering comparable results. By integrating DBR with NRR, our method enhances dataset distillation by capturing essential features while minimizing training complexity. This results in more compact and efficient datasets, well-suited for a variety of training environments. Figure 1 illustrates the differences between our method and two popular large-scale dataset distillation techniques, RDED [21] and SRe²L. It is evident that RDED focuses on instance-specific features without refinement, while SRe²L updates all pixels to capture classgeneral features, often at the expense of fine-grained details. In contrast, our NRR-DD method effectively preserves finegrained details by updating only non-critical pixels, while still capturing class-general features.

Contributions. Our major contributions are summarized as follows:

- We introduce the Non-Critical Region Refinement Dataset Distillation (NRR-DD) framework, which consists of three key stages: Critical-based Initial Data Discovery (CIDD), Non-Critical Region Refinement (CRR), and Relabeling. This approach generates synthetic data that captures both instance-specific finegrained features and class-general patterns, significantly enhancing performance.
- We propose the Distance-Based Representative (DBR) method for knowledge transfer, eliminating the need for soft labels and drastically reducing memory requirements. Specifically, our method reduces storage requirements by *500-fold* compared to soft labels on ImageNet1k, while recovering up to 80% of the full performance (see Table 4).
- Experimental results demonstrate that our NRR



Figure 1. Comparison of our method with two popular frameworks, SRe²L [27] and RDED [21], for generating synthetic datasets. RDED selects high-confidence, easily classifiable images, while our method focuses on low-confidence, harder-toclassify samples, which helps reduce overfitting and improve model accuracy. Additionally, RDED targets instance-specific features without refinement, and SRe²L updates all pixels to capture class-general features, often at the expense of fine details. In contrast, our NRR-DD method preserves fine-grained details while capturing class-general features by updating only non-critical pixels.

method achieves state-of-the-art performance on both small- and large-scale datasets. Additionally, by storing only two distances per instance, it achieves comparable results across various settings.

2. Related Works

Dataset Distillation. Several dataset distillation methods have been proposed recently [2, 3, 21, 27, 32, 33], which can be classified into two main categories. The first category, Class-General Feature-Based Methods, aims to capture class-wide features. For example, Gradient Matching [10, 32] generates synthetic data by matching gradients across all samples in a class, while Distribution Matching [16, 31, 33] uses distribution prediction. Trajectory Matching [2, 6] aligns training trajectories of original and synthetic data, and SRe²L [27] recovers Batch Normalization statistics to capture class-general features. The second category, Instance-Specific Feature-Based Methods [7, 21], includes approaches such as RDED [21], which extracts high-confidence patches and combines them to create synthetic data, and MDiff [7], which utilizes a diffusion model to generate fine-grained images tailored to the task. However, each approach has its limitations: methods focused on class-general features risk losing crucial instancespecific details, hindering fine-grained generalization, while instance-specific methods may neglect broader class patterns, leading to suboptimal class-level representation.

Large-Scale Dataset Distillation via Soft-Label Knowl-

edge Transfer. Recently, several methods have been proposed for large-scale datasets like ImageNet1k [3, 21, 27]. However, all of these approaches require storing the soft labels of each augmented data point for training new student models, leading to significant memory storage overhead. For example, with ImageNet1k, storing the soft labels for 50 Images-Per-Class (IPC) requires approximately 30 GB, and for 200 IPC, over 120 GB of storage [27]. This high-lights the need for novel knowledge transfer techniques to reduce memory storage requirements.

3. Proposed Method

In this section, we first provide the necessary preliminaries for the Dataset Distillation Method (Section 3.1), followed by an introduction to our Non-Critical Region Refinement Dataset Distillation (NRR-DD) framework. The framework consists of three key stages: (i) Critical-based **Initial Data Discovery** (Section 3.2), which selects diverse, important patches to capture instance-specific features; (ii) Non-Critical Region Refinement (Section 3.3), where Class Activation Mapping (CAM) [19, 34] is employed to identify and refine both critical and non-critical regions, preserving fine-grained details while enriching non-critical areas with class-general information; and (iii) Knowledge Transfer via Relabeling (Section 3.4), in which synthetic images are relabeled for knowledge transfer to a student model. The overall architecture is shown in Figure 2, and the pseudo code can be found in Algorithm 1.

3.1. Preliminaries

Consider a training dataset $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^m$, where each input $x_i \in \mathbb{R}^{c \times h \times w}$ represents a sample, and $y_i \in \{1, \ldots, K\}$ denotes its label. Let $\mathcal{T} = \mathcal{T}_{\theta_{\mathcal{T}}}$ be a pretrained model on \mathcal{D} . The goal of dataset distillation is to generate a synthetic dataset $\tilde{\mathcal{D}} = \{(\tilde{x}_i, \tilde{y}_i)\}_{i=1}^n$ (with $n \ll m$) that retains the essential information from \mathcal{D} , enabling a new model S to achieve performance comparable to that of \mathcal{T} .

3.2. Critical-based Initial Data Discovery

In this section, instead of using Gaussian noise to generate synthetic data for \mathcal{M} , we propose training-free techniques to more effectively select initial data. The motivation behind this approach is that real data inherently contains fine-grained, instance-specific features, which are crucial for training on large-scale datasets [7, 12–14, 21, 25].

Given the training dataset $\mathcal{D} = \{(\boldsymbol{x}_i, \boldsymbol{y}_i)\}_{i=1}^m$, for each data example $(\boldsymbol{x}, \boldsymbol{y}) \in \mathcal{D}$, we first compute the class activation mapping (CAM) [34], producing a matrix $C(\boldsymbol{x}, \boldsymbol{y})$ for the image \boldsymbol{x} and class \boldsymbol{y} :

$$C(\boldsymbol{x}, \boldsymbol{y}) = \sum_{k} \boldsymbol{w}_{k}^{\boldsymbol{y}} \mathcal{T}_{k}(\boldsymbol{x}, \boldsymbol{y}), \qquad (1)$$

Algorithm 1: NRR-DD **Input:** Pre-trained model $\mathcal{T}_{\theta_{\mathcal{T}}}$, training set $\mathcal{D} = \{(\boldsymbol{x}_i, \boldsymbol{y}_i)\}_{i=1}^m$ **Output:** Synthetic dataset $\tilde{\mathcal{D}} = \{(\tilde{x}_i, \tilde{y}_i)\}_{i=1}^n$ 1 Initial $\tilde{\mathcal{D}} = \emptyset$; /* Critical-based Initial Data Discovery */ 2 foreach $(\boldsymbol{x}, \boldsymbol{y}) \sim \mathcal{D}$ do Calculate CAMs matrix $C(\boldsymbol{x})$ for \boldsymbol{x} ; 3 4 Crop \boldsymbol{x} into k patches $\{p_1, \ldots, p_k\}$; Select the top t patches with the highest values in the 5 CAMs matrix; Resize and store them in the patch pool \mathcal{P} ; 6 7 foreach $p \sim \mathcal{D}$ do Compute the score $s(p) = \mathcal{L}_{ce}(p)$; 9 Select the top $q = \beta \times IPC$ patches with the lowest scores; 10 for i = 1 to IPC do 11 Combine β patches into \tilde{x} , store (\tilde{x}, \tilde{y}) in \mathcal{D} /* Non-Critical Region Refinement */ 12 foreach $(\tilde{\boldsymbol{x}}_{org}, \tilde{\boldsymbol{y}}_{org}) \sim \tilde{\mathcal{D}}$ do for I iterations do 13 $(ilde{oldsymbol{x}}_{ ext{aug}}, ilde{oldsymbol{y}}_{ ext{aug}}) \sim ilde{\mathcal{D}};$ 14 $\tilde{\boldsymbol{x}}_{\text{mix}} = A(\tilde{\boldsymbol{x}}_{\text{org}}, \tilde{\boldsymbol{x}}_{\text{aug}});$ 15 Update \tilde{x}_{org} by Eq. 11; 16 $\tilde{\boldsymbol{y}}_{\text{soft}} = \mathcal{T}(\tilde{\boldsymbol{x}}_{\text{mix}});$ 17 $\begin{aligned} & \boldsymbol{y}_{\text{soft}} = \mathcal{L}_{ce}(\boldsymbol{\tilde{y}}_{\text{soft}}, \boldsymbol{\tilde{y}}_{\text{org}}); \\ & \boldsymbol{d}_{\text{aug}}^{T} = \mathcal{L}_{ce}(\boldsymbol{\tilde{y}}_{\text{soft}}, \boldsymbol{\tilde{y}}_{\text{aug}}); \\ & \text{Store} \left(\boldsymbol{\tilde{x}}_{\text{org}}, \left(\boldsymbol{\tilde{y}}_{\text{org}}, \boldsymbol{\tilde{y}}_{\text{aug}}\right), \left(\boldsymbol{d}_{\text{org}}^{T}, \boldsymbol{d}_{\text{aug}}^{T}\right)\right) \text{ in } \boldsymbol{\tilde{\mathcal{D}}}; \end{aligned}$ 18 19 20 /* Distance-Based Representative Knowledge Transfer */ 21 foreach $(\tilde{\boldsymbol{x}}_{org}, (\tilde{\boldsymbol{y}}_{org}, \tilde{\boldsymbol{y}}_{aug}), (d_{org}^T, d_{aug}^T)) \sim \tilde{\mathcal{D}}$ do Update new model S by Eq. 9; 22

where $\boldsymbol{w}_{k}^{\boldsymbol{y}}$ represents the k^{th} weight in the final classification head for class \boldsymbol{y} , and \mathcal{T}_{k} denotes the k^{th} feature map from the final layers of the model.

Next, the images x are randomly cropped to extract multiple patches, and the top t patches, which contain the highest values in the class activation map, are selected. These patches are then resized to the full image size and stored in the patch pool \mathcal{P} . Since the selection is based on the highest values in the class activation map, the chosen patches capture the most important information from the original images.

For each patch p in \mathcal{P} , we feed it into \mathcal{T} to obtain a confidence score $s = \mathcal{T}(p)$, which represents the highest prediction probability. Unlike RDED [21], which selects patches with the highest scores, we choose the top $g = \beta \times IPC$ patches with the lowest scores, where $\beta = 1, 4, 9, \ldots$ denotes the number of patches used to form a single synthetic image, and IPC specifies the number of images per class. Notably, our strategy of selecting the lowest-scoring patches



Figure 2. The architecture of our NRR-DD consists of three key stages: (i) Critical-based Initial Data Discovery (Section 3.2), which selects patches with a high CAM ratio but low confidence level to capture instance-specific features; (ii) Non-Critical Region Refinement (Section 3.3), where CAM [34] is used to identify and refine both critical and non-critical regions, preserving fine-grained details while enriching non-critical areas with class-general information; (iii) Knowledge Transfer, which aims to minimize the distance between $S(\tilde{x}_{mix})$ (student prediction) and $\mathcal{T}(\tilde{x}_{mix})$ (pretrained teacher prediction or soft label) by reducing the distance between d_{org}^{T} and d_{org}^{S} , as well as between d_{aug}^{T} and d_{aug}^{S} . By storing only the two values, d_{org}^{T} and d_{aug}^{T} , the new model can effectively mimic the performance of the pretrained one.

identifies the hardest-to-learn samples, while the method of selecting patches with the highest values in the class activation map ensures that the patches contain important information. This provides more opportunity and flexibility in the next phase to further refine the chosen synthetic images. Figure 1 visualizes the images selected by both the highestand lowest-scoring strategies.

Finally, similar to RDED [21], we construct each synthetic image by combining β patches. The selected patches are resized to $1/\beta$ of their original size and then combined to create *IPC* synthetic images per class, all of which are stored in the dataset \tilde{D} . Unlike RDED [21], our synthetic images in \tilde{D} are refined to include both *fine-grained*, *instance-specific features* and *class-general features*.

3.3. Non-Critical Region Refinement

In contrast to RDED [21], our synthetic images in \mathcal{D} are refined to incorporate both *detailed*, *instance-specific* features and *broader*, *class-general* features. Specifically, for each $(\tilde{x}, \tilde{y}) \in \tilde{\mathcal{D}}$, the image \tilde{x} is refined according to the following loss function:

$$\mathcal{L}_C = \mathcal{L}_{ce}(\mathcal{T}(\tilde{\boldsymbol{x}}), \tilde{\boldsymbol{y}}) + \alpha_{bn} \mathcal{L}_{bn}(\mathcal{T}(\tilde{\boldsymbol{x}})), \qquad (2)$$

where α_{bn} is a parameter.

In this framework, \mathcal{L}_{ce} represents the cross-entropy (CE) loss, which aims to move \tilde{x} with the smallest confidence score into the teacher's high-confidence regions. The batch normalization regularization (\mathcal{L}_{bn}) [22, 24, 27], a standard DFKD loss, aligns the mean and variance at the

BatchNorm layer with its running mean and variance:

$$\mathcal{L}_{bn} = \sum_{l} \left(\|\mu_l(\tilde{\boldsymbol{x}}) - \mu_l\| + \|\sigma_l^2(\tilde{\boldsymbol{x}}) - \sigma_l^2\| \right), \quad (3)$$

where $\mu_l(\tilde{x})$ and $\sigma_l^2(\tilde{x})$ are the mean and variance of the *l*-th BatchNorm layer of \mathcal{T} , and μ_l and σ_l^2 are the running mean and variance of the *l*-th BatchNorm layer in \mathcal{T} .

Non-Critical Region. A naive approach to refining \tilde{x} would involve optimizing all pixels in \tilde{x} using the gradient of \mathcal{L}_C . However, this would lead to significant changes in the images, resulting in the loss of fine-grained, instance-specific features, as shown in Figure 1. Instead, we propose the *Non-Critical Region Refinement Dataset Distillation* (NRR-DD) method, which preserves instance-specific and fine-grained regions in the synthetic data while enriching non-critical regions with more class-general information. This approach enables our models to utilize all pixel information to capture both types of features, thereby enhancing overall performance.

Given a synthetic image \tilde{x} with label \tilde{y} , we use CAM to create a *non-critical mask* M of the same size as \tilde{x} . This mask assigns low or zero values to high-importance pixels in \tilde{x} and higher values to less important pixels. The objective is to control pixel updates in \tilde{x} by limiting updates to important pixels in order to preserve instance-specific features, while allowing less important pixels to update more significantly, thus enhancing the learning of class-general features. The process is detailed as follows.

We first generate the CAM matrix C of (\tilde{x}, \tilde{y}) using Eq. 1. Subsequently, we calculate the pixel-wise non-



Figure 3. Visualization of images from the 'tench' and 'English springer' classes synthesized using various dataset distillation methods, including SRe²L [27], RDED [21], and our NRR-DD. For additional visualizations, please refer to the **Supplementary Material**.

critical mask M matrix using the following formula:

$$M = \max\{0, \epsilon - C\}.$$
 (4)

Here, ϵ serves as the upper threshold for the method. For any value c in C that exceeds ϵ , its corresponding value in M will be set to 0; otherwise, the value will be $\epsilon - c$. Since higher class activation values correspond to lower non-critical scores, the matrix M will be used to weight the gradient update for each pixel in \tilde{x} . With each gradient update, the image \tilde{x} is updated as follows:

$$\tilde{\boldsymbol{x}} = \tilde{\boldsymbol{x}} - \boldsymbol{M} \times \eta \nabla_{\tilde{\boldsymbol{x}}} \mathcal{L}_{\mathcal{C}},\tag{5}$$

where $\nabla_{\tilde{x}} \mathcal{L}_{\mathcal{C}}$ represents the gradient of the loss function $\mathcal{L}_{\mathcal{C}}$ with respect to the image \tilde{x} .

3.4. Knowledge Transfer via Relabeling

Soft-Label Knowledge Transfer. Previous methods [3,21, 27] store soft labels generated by a pretrained teacher model \mathcal{T} to train a student model \mathcal{S} by minimizing the Kullback-Leibler (KL) divergence between the student model's predictions and the teacher's soft labels. These soft labels are generated through augmentation techniques such as Cut-Mix [28] and Mixup [30].

Given a pair of original image $(\tilde{x}_{org}, \tilde{y}_{org})$ and augmented image $(\tilde{x}_{aug}, \tilde{y}_{aug})$ in the synthetic dataset $\tilde{\mathcal{D}}$, we apply CutMix or Mixup to create the mixed image and compute its soft label as:

$$\tilde{\boldsymbol{x}}_{\text{mix}} = A(\tilde{\boldsymbol{x}}_{\text{org}}, \tilde{\boldsymbol{x}}_{\text{aug}})$$
 (6)

$$\tilde{\boldsymbol{y}}_{\text{soft}} = \mathcal{T}(\tilde{\boldsymbol{x}}_{\text{mix}}).$$
 (7)

Here, A represents an augmentation method such as Cut-Mix [28] or Mixup [30], and \tilde{y}_{soft} denotes the soft labels produced by the mixed image \tilde{x}_{mix} .

Therefore, for each pair of original image $(\tilde{x}_{org}, \tilde{y}_{org})$ and augmented image $(\tilde{x}_{aug}, \tilde{y}_{aug})$, we need to store the indices of the two images in the synthetic dataset, the specification of the mixing method A used to construct the mixed image \tilde{x}_{mix} , and the soft label \tilde{y}_{soft} . However, this approach requires substantial memory to store the soft labels, especially when large datasets and numerous augmentations are involved.

Distance-Based Representative Knowledge Transfer. To address the memory limitations of storing soft labels, we propose a more memory-efficient approach. Instead of storing the soft labels \tilde{y}_{soft} , each of which consists of 1,000 real numbers for ImageNet1K, we store only two real numbers representing the cross-entropy (CE) divergences between the soft label and the one-hot vectors of \tilde{y}_{aug} and \tilde{y}_{org} .

Specifically, for each pair of an original image $(\tilde{x}_{org}, \tilde{y}_{org})$ and an augmented image $(\tilde{x}_{aug}, \tilde{y}_{aug})$ in the synthetic dataset $\tilde{\mathcal{D}}$, we apply CutMix or Mixup to create the mixed image \tilde{x}_{mix} and compute its soft label \tilde{y}_{soft} using the teacher model \mathcal{T} . We then calculate and store the crossentropy (CE) divergences between the soft label and the one-hot vectors of \tilde{y}_{aug} and \tilde{y}_{org} as follows:

$$d_{\text{org}}^{T} = \mathcal{L}_{ce}(\tilde{\boldsymbol{y}}_{\text{soft}}, \tilde{\boldsymbol{y}}_{\text{org}}), \\ d_{\text{aug}}^{T} = \mathcal{L}_{ce}(\tilde{\boldsymbol{y}}_{\text{soft}}, \tilde{\boldsymbol{y}}_{\text{aug}}).$$
(8)

Moreover, for each pair of original and augmented images, we store their indices in the synthetic dataset, the details of

		ConvNet					Resnet18			ResNet-101		
	IPC	MTT	IDM	TESLA	RDED	NRR-DD	SRe ² L	RDED	NRR-DD	SRe ² L	RDED	NRR-DD
CIFAR10	1	46.3 ± 0.8	45.6 ± 0.7	48.5 ± 0.8	23.5 ± 0.3	48.4 ± 0.4	16.6 ± 0.9	22.9 ± 0.4	30.3 ± 0.4	13.7 ± 0.2	18.7 ± 0.1	25.7 ± 0.3
	10	65.3 ± 0.7	65.3 ± 0.7	66.4 ± 0.8	50.2 ± 0.3	66.7 ± 0.4	29.3 ± 0.5	37.1 ± 0.3	72.2 ± 0.4	24.3 ± 0.6	33.7 ± 0.3	65.1 ± 0.3
	50	71.6 ± 0.2	67.5 ± 0.1	72.6 ± 0.7	68.4 ± 0.1	73.1 ± 0.1	45.0 ± 0.7	62.1 ± 0.1	84.1 ± 0.1	34.9 ± 0.1	51.6 ± 0.4	78.2 ± 0.4
CIFAR100	1	24.3 ± 0.3	20.1 ± 0.3	24.8 ± 0.5	19.6 ± 0.3	27.3 ± 0.3	6.6 ± 0.2	11.0 ± 0.3	33.3 ± 0.3	6.2 ± 0.0	10.8 ± 0.1	32.9 ± 0.3
	10	40.1 ± 0.4	45.1 ± 0.1	41.7 ± 0.3	48.1 ± 0.3	55.7 ± 0.2	27.0 ± 0.4	42.6 ± 0.2	62.7 ± 0.2	30.7 ± 0.3	41.1 ± 0.2	58.3 ± 0.2
	50	47.7 ± 0.2	50.0 ± 0.2	47.9 ± 0.3	57.0 ± 0.1	61.1 ± 0.1	50.2 ± 0.4	62.6 ± 0.1	67.1 ± 0.1	56.9 ± 0.1	63.4 ± 0.3	65.1 ± 0.3
Tiny-ImageNet	1	8.8 ± 0.3	10.1 ± 0.2	-	12.0 ± 0.1	20.4 ± 0.2	2.62 ± 0.1	9.7 ± 0.4	13.5 ± 0.2	1.9 ± 0.1	3.8 ± 0.1	10.1 ± 0.1
	10	23.2 ± 0.2	21.9 ± 0.3	-	39.6 ± 0.1	44.3 ± 0.2	16.1 ± 0.2	41.9 ± 0.2	45.2 ± 0.2	14.6 ± 1.1	22.9 ± 3.3	26.1 ± 3.3
	50	28.0 ± 0.3	27.7 ± 0.3	-	47.6 ± 0.2	50.2 ± 0.1	41.1 ± 0.4	58.2 ± 0.1	61.2 ± 0.1	42.5 ± 0.2	41.2 ± 0.4	46.2 ± 0.4
ImageNette	1	47.7 ± 0.9	-	-	33.8 ± 0.8	39.3 ± 0.9	19.1 ± 1.1	35.8 ± 1.0	40.1 ± 0.9	15.8 ± 0.6	25.1 ± 2.7	28.1 ± 2.7
	10	63.0 ± 1.3	-	-	63.2 ± 0.7	68.3 ± 0.6	29.4 ± 3.0	61.4 ± 0.4	66.2 ± 0.6	23.4 ± 0.8	54.0 ± 0.4	56.0 ± 0.4
	50	-	-	-	83.8 ± 0.2	86.5 ± 0.3	40.9 ± 0.3	80.4 ± 0.4	85.6 ± 0.3	36.5 ± 0.7	75.0 ± 1.2	78.0 ± 1.2
ImageNet1k	1	-	-	7.7 ± 0.2	6.4 ± 0.1	11.2 ± 0.2	0.1 ± 0.1	6.6 ± 0.2	11.6 ± 0.2	0.6 ± 0.1	5.9 ± 0.4	12.2 ± 0.4
	10	-	-	17.8 ± 1.3	20.4 ± 0.1	25.6 ± 0.2	21.3 ± 0.6	42.0 ± 0.1	46.1 ± 0.2	30.9 ± 0.1	48.3 ± 1.0	51.3 ± 1.0
	50	-	-	27.9 ± 1.2	38.4 ± 0.2	42.1 ± 0.1	46.8 ± 0.2	56.5 ± 0.1	60.2 ± 0.1	60.8 ± 0.5	61.2 ± 0.4	64.3 ± 0.4

Table 1. Comparison with state-of-the-art (SOTA) dataset distillation baselines. Identical neural networks are used for both dataset distillation and evaluation. Following [21, 27], ConvNets used for distillation are Conv-3 for CIFAR10 and CIFAR100, Conv-4 for Tiny-ImageNet and ImageNet-1K, Conv-5 for ImageNette and ImageWoof, and Conv-6 for ImageNet-100. MTT and TESLA use down-sampled images for distillation to 224×224 images. SRe2L and RDED use ResNet-18 for distillation and retrieval, and are evaluated on ResNet-18 and ResNet-101. Entries marked with "-" indicate scalability issues. See **Supplementary Material** for further details.

the data mixing, and the two divergences d_{org}^T and d_{aug}^T (see Figure 2(c) for visualization). Subsequently, we train the student model S by minimizing:

$$\mathcal{L}_{S} = \mathcal{L}_{sce} + \alpha_{dbr} \mathcal{L}_{dbr}, \qquad (9)$$
$$\mathcal{L}_{sce} = max\{0, d_{org}^{S} - r\} + max\{0, d_{aug}^{S} - r\},$$
$$\mathcal{L}_{dbr} = |d_{org}^{S} - d_{org}^{T}| + |d_{aug}^{S} - d_{aug}^{T}|,$$

where r is a threshold and d^S is calculated as follows:

$$d_{\text{org}}^{S} = \mathcal{L}_{ce}(\mathcal{S}(\tilde{\boldsymbol{x}}_{\text{mix}}), \tilde{\boldsymbol{y}}_{\text{org}}) d_{\text{aug}}^{S} = \mathcal{L}_{ce}(\mathcal{S}(\tilde{\boldsymbol{x}}_{\text{mix}}), \tilde{\boldsymbol{y}}_{\text{aug}}).$$
(10)

In Eq. 9, \mathcal{L}_{sce} denotes the soft cross-entropy loss function, which encourages the student model to predict the mixed instance $ilde{x}_{
m mix}$ as a blend of the two labels, $ilde{y}_{
m org}$ and \tilde{y}_{aug} . This helps prevent excessive confidence by imposing a threshold r, which is especially crucial when using techniques such as CutMix or MixUp, where the soft label often blends multiple classes. This mechanism reduces the risk of the model becoming overly confident in any single class. Additionally, \mathcal{L}_{dbr} represents the distance-based representative loss, which allows the student model S to replicate the teacher model's divergences from its predictions on the mixed image $\tilde{x}_{\rm mix}$ to $\tilde{y}_{\rm org}$ and $\tilde{y}_{\rm aug}$. This loss ensures consistency between the teacher and student models, both in terms of learned feature representations and decision boundaries. Memory Reduction. It is important to note that, unlike previous models that require storing augmentation information to generate \tilde{x}_{mix} , along with the indices and labels for both \tilde{x}_{org} and \tilde{x}_{aug} , our approach only necessitates storing two additional distances, $d_{\text{org}}^{\hat{T}}$ and d_{aug}^{T} , for the loss function. This significantly reduces memory requirements.

Label Refinement. To further enhance the benefits of Distance-Based Representative (DBR), we propose refining the images with an additional term. Instead of training \tilde{x}_{org} using the formula from Eq. 2, we train \tilde{x}_{org} with the following formulation:

$$\mathcal{L}_C = \mathcal{L}_{\rm org} + \alpha_{lr} \mathcal{L}_{lr},\tag{11}$$

$$\mathcal{L}_{\text{org}} = \mathcal{L}_{ce}(\mathcal{T}(\tilde{\boldsymbol{x}}_{\text{org}}), \tilde{\boldsymbol{y}}_{\text{org}}) + \alpha_{bn} \mathcal{L}_{bn}(\mathcal{T}(\tilde{\boldsymbol{x}}_{\text{org}}))$$

$$\mathcal{L}_{e} = \max\{0, d^{T} - r\} + \max\{0, d^{T} - r\} \quad (12)$$

$$\mathcal{L}_{lr} = \max\{0, d_{\text{org}}^{I} - r\} + \max\{0, d_{\text{aug}}^{I} - r\}$$
(12)

Minimizing \mathcal{L}_{lr} ensures that the prediction of \mathcal{T} on \tilde{x}_{mix} , which will be used for training the student model, is focused solely on the two designated classes, facilitating alignment with the DBR function.

In Figure 3, we visualize the images generated by our NRR-DD method alongside several state-of-the-art approaches. It is clear that our method captures both instancespecific and class-general features, leading to improved performance.

4. Experiment

This section evaluates the effectiveness of our proposed method compared to state-of-the-art techniques across various datasets and neural architectures, accompanied by comprehensive ablation studies.

4.1. Experimental Setting

For large-scale datasets, we assessed our method using two popular pairs of backbones: ResNet34/ResNet18 [8] and ResNet50/MobileNetV2 [18], applied to three wellknown benchmarks: Tiny-ImageNet [15], which consists of

			Using Soft Label		Using One-Hot	Using Compact Label			
Dataset	Architecture	IPC	RDED (SL)	NRR-DD(SL)	RDED (OH)	RDED (CL)	NRR-DD (DBR)	NRR-DD (DBR+LR)	Recover
		1	33.8 ± 0.8	39.3 ± 0.9	16.3 ± 0.3	23.2 ± 0.5	32.3 ± 0.4	34.5 ± 0.5	79%
	ConvNet	10	63.2 ± 0.7	68.3 ± 0.6	27.3 ± 0.2	37.3 ± 0.4	52.2 ± 0.4	55.1 ± 0.4	68%
ImageNette		50	83.8 ± 0.2	86.5 ± 0.3	41.2 ± 0.3	54.6 ± 0.3	66.9 ± 0.2	69.2 ± 0.2	62%
magenette	Resnet18	1	35.8 ± 1.0	40.1 ± 0.9	16.2 ± 0.4	22.2 ± 0.4	34.6 ± 0.6	36.2 ± 0.4	84%
		10	61.4 ± 0.4	66.2 ± 0.6	25.4 ± 0.3	34.1 ± 0.4	54.7 ± 0.4	57.1 ± 0.4	78%
		50	80.4 ± 0.4	85.6 ± 0.3	41.3 ± 0.3	52.1 ± 0.2	69.4 ± 0.2	72.2 ± 0.3	70%
		1	6.4 ± 0.1	11.2 ± 0.2	2.4 ± 0.5	3.1 ± 0.5	6.2 ± 0.5	8.5 ± 0.5	69%
	ConvNet	10	20.4 ± 0.1	25.6 ± 0.2	8.3 ± 0.4	12.5 ± 0.3	16.1 ± 0.3	19.2 ± 0.5	63%
ImageNet1k		50	38.4 ± 0.2	42.1 ± 0.1	14.1 ± 0.3	22.3 ± 0.3	26.1 ± 0.1	28.1 ± 0.3	50%
mugerterrk		1	6.6 ± 0.2	11.6 ± 0.2	2.3 ± 0.5	3.2 ± 0.5	7.5 ± 0.3	8.9 ± 0.5	71%
	Resnet18	10	42.0 ± 0.1	46.1 ± 0.2	16.3 ± 0.7	22.2 ± 0.4	34.3 ± 0.5	37.2 ± 0.4	70%
		50	56.5 ± 0.1	60.2 ± 0.1	32.4 ± 0.5	39.3 ± 0.2	$\overline{45.1 \pm 0.3}$	49.2 ± 0.5	60%

Table 2. Comparison of various relabeling methods on large-scale datasets, including ImageNette and ImageNet1k. Bold values indicate the highest scores, while underlined values indicate the second-highest scores. 'SL' denotes the use of soft labels, 'OH' represents one-hot vector labels, and 'CL' signifies compact labels (with 2 classes for fair comparison to our DBR). 'DBR' refers to our method using distance-based representation without label refinement, and 'DBR+LR' indicates our method with distance-based representation combined with label refinement. The Recover rate is calculate by using (DBR - One-hot)/(Soft-label - One-hot)

200 object categories, with 500 training images, 50 validation images, and 50 test images per category, all resized to 64×64 pixels; ImageNet1k [4], containing 1,000 object categories and over 1.2 million labeled training images, along with its subset Imagenette, which includes 10 sub-classes. For small-scale datasets, we ran experiments with ResNet [8], VGG [20], and WideResNet (WRN) [29] across CIFAR-10 and CIFAR-100 [11]. Both CIFAR-10 and CIFAR-100 consist of 60,000 images (50,000 for training and 10,000 for testing), with 10 and 100 categories, respectively, and all images have a resolution of 32×32 pixels. Consistent with previous research, we set the IPC to 1, 10, and 50. All experiments were run on a single NVIDIA A100 40 GB GPU. The details of the model architectures, parameters, and additional experimental results are provided in the Supplementary Material.

Compared Baselines. We focus on comparing our method to state-of-the-art dataset distillation methods:

- MTT [2], the first to propose trajectory matching;
- IDM [33], the first to introduce distribution matching;
- TESLA [3], the first method to scale up to full ImageNet1k;
- SRe²L [27], a method that efficiently scales to ImageNet-1k and significantly outperforms existing approaches on large, high-resolution datasets;
- RDED [21], a recent paper that uses instance-specific features with fine-grained details to improve large-scale dataset distillation, which we consider the closest baseline.

4.2. Main Result

Large-scale Dataset. To demonstrate the effectiveness of our methods in real-world applications, we first compared them with various baselines on large-scale datasets

Teacher\S	tudent	Resnet18	MobileNetV2	EfficientNet-B0	
Resnet18	SRe ² L	21.7 ± 0.6	15.4 ± 0.2	11.7 ± 0.2	
	RDED	42.3 ± 0.6	40.4 ± 0.1	31.0 ± 0.1	
	NRR-DD	46.1 ± 0.2	45.0 ± 0.2	34.2 ± 0.1	
MobileNetV2	SRe ² L	19.7 ± 0.1	10.2 ± 2.6	9.8 ± 0.4	
	RDED	34.4 ± 0.2	33.8 ± 0.6	24.1 ± 0.8	
	NRR-DD	36.2 ± 0.2	37.2 ± 0.1	27.3 ± 0.7	
EfficientNet-B0	SRe ² L	25.2 ± 0.2	20.5 ± 0.2	11.4 ± 2.5	
	RDED	42.8 ± 0.5	43.6 ± 0.2	33.3 ± 0.9	
	NRR-DD	47.2 ± 0.2	45.6 ± 0.3	35.1 ± 0.2	

Table 3. Evaluation of ImageNet-1K top-1 accuracy for crossarchitecture generalization. Datasets are distilled using ResNet-18, EfficientNet-B0, and MobileNet-V2, and transferred across different architectures. Note that experiments for SRe^2L could not be conducted when the distillation model lacks batch normalization [27]. All methods are evaluated with IPC = 10.

such as ImageNet1k, Tiny ImageNet, and ImageNette. The results in Table 1 show that our methods outperform all the compared baselines in every scenario. For instance, our method achieves 60.2% accuracy on ImageNet1k using ResNet18, which is 4% higher than RDED and 14% higher than SRe²L. These improvements highlight the superiority of our approach in handling large-scale datasets and its ability to significantly outperform existing methods.

Small-scale Dataset. To validate the robustness of our NRR-DD, we compared it with various baselines on small-scale datasets, including CIFAR-10 and CIFAR-100. As shown in Table 1, our NRR-DD demonstrates significant improvements over the baselines. For example, with 10 IPC, our method achieves 72.2% accuracy on CIFAR-10 using ResNet18, which is 35% higher than the current state-of-the-art, RDED, which achieves 37.1%. Similarly, on CIFAR-100, our method shows a 20% improvement with 10 IPC. These results highlight the effectiveness and generalization ability of our method on small-scale datasets.

		ConvNet	Resnet18			
	CIFAR10	CIFAR100	ImageNet1k	CIFAR10	CIFAR100	ImageNet1k
RDED (baseline)	50.2 ± 0.3	48.1 ± 0.3	20.4 ± 0.1	37.1 ± 0.3	42.6 ± 0.2	42.0 ± 0.1
CIDD	51.3 ± 0.4	50.2 ± 0.3	21.2 ± 0.3	50.4 ± 0.5	51.3 ± 0.2	43.2 ± 0.2
CIDD+NRR	66.7 ± 0.4	55.7 ± 0.2	25.6 ± 0.2	65.1 ± 0.3	58.3 ± 0.2	51.3 ± 1.0

Table 4. Comparison between the baseline RDED [21] and our framework using only the CIDD modules (CIDD) and with Non-Critical Region Refinement (NRR) (CIDD+NRR). All methods are evaluated with IPC = 10.

Different Label Compression Comparison. To demonstrate the benefits of our distance-based representation (DBR), we conducted experiments to compare it with various label compression methods. The results show that, while requiring storage of only two distance values per instance (achieving 5× compression on ImageNette and 500× on ImageNet1k compared to soft-labels), our method provides comparable performance. Notably, in all cases with IPC set to 1, our DBR method outperforms the state-of-the-art RDED with soft-labels. When compared to compact labels (which use only two elements of soft-labels), our method achieves significant improvements across all scenarios. These findings highlight the effectiveness of distance-based representation in this task.

Cross-architecture Generalization. To ensure the generalization capability of our distilled datasets, it is crucial to evaluate their performance across a range of neural architectures not involved in the dataset distillation process. The results in Table 3 demonstrate the robustness and generalization of our method across all cross-comparisons. This can be attributed to the advantage of selecting the critical key patches in CIDD and the effectiveness of non-critical image refinement.

4.3. Ablation Study

Effectiveness of Critical-based Initial Data Discovery. To verify the benefits of our Critical-based Initial Data Discovery (CIDD) in Section 3.2, we conducted experiments using only the data generated by CIDD to train the student model, following the same settings as RDED [21]. The results in Table 4 demonstrate that our CIDD achieves better performance than RDED in all cases, indicating the effectiveness of the module.

Effectiveness of Non-Critical Region Refinement. Table 4 shows the results of our model with NRR (CIDD+NRR) and without NRR (CIDD). The results demonstrate that the NRR method significantly improves model performance. For example, on CIFAR-10 using ResNet18, the addition of NRR leads to an approximate 15% performance increase compared to when NRR is not used. This clearly indicates the effectiveness of incorporating the Non-Critical Region Refinement module.

Effectiveness of Label Refinement (\mathcal{L}_{lr}) . To further validate the benefits of \mathcal{L}_{lr} in Eq. 12, we conducted experiments to determine whether this term enhances model learning without relying on soft-label information. As shown in

Table 2, incorporating Label Refinement NRR-DD (DBR) consistently improves performance over without using Label Refinement NRR-DD (DBR) across all comparisons, demonstrating the effectiveness of this term.

5. Limitation and Future Works

A possible limitation of our work is its reliance on the quality of Class Activation Mapping (CAM) for identifying critical and non-critical regions. While CAM provides useful insights, its performance can be sensitive to the model's initial training, potentially affecting the accuracy of the identified regions. Therefore, a promising direction for future work is exploring how to connect the training of pretrained models to generate better CAMs.

Additionally, while the Distance-Based Representative (DBR) technique reduces memory requirements, it may not fully address the trade-off between memory efficiency and the preservation of fine-grained details for certain complex tasks. This emphasizes the need for further research to develop enhanced solutions for this issue, which we plan to explore in our future work.

6. Conclusion

In this paper, we introduced the Non-Critical Region Refinement Dataset Distillation (NRR-DD) method to address the limitations of current dataset distillation techniques. Our approach comprises three key stages: Criticalbased Initial Data Discovery (CIDD) to capture instancespecific details, Non-Critical Region Refinement (NRR) to balance critical and non-critical regions using Class Activation Mapping, and relabeling to transfer knowledge effectively. Additionally, we proposed the Distance-Based Representative (DBR) technique, which eliminates the need for memory-intensive soft label storage by using a distancebased measure, significantly reducing memory requirements. By combining DBR with NRR, our method generates compact and efficient datasets, achieving state-of-theart performance across small- and large-scale datasets. Experimental results confirm that our approach not only enhances dataset representativeness but also minimizes training complexity, making it highly suitable for various training environments. This work advances dataset distillation by capturing critical features and optimizing storage, paving the way for efficient, scalable data solutions.

Acknowledgements

This work was supported by ARC DP23 grant DP230101176 and by the Air Force Office of Scientific Research under award number FA2386-23-1-4044.

References

- Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient based sample selection for online continual learning. *Advances in neural information processing systems*, 32, 2019. 1
- [2] George Cazenavette, Tongzhou Wang, Antonio Torralba, and Charless Fowlkes. Dataset distillation by matching training trajectories. arXiv preprint arXiv:2203.11938, 2022. 1, 2, 7
- [3] Justin Cui, Ruochen Wang, Si Si, and Cho-Jui Hsieh. Scaling up dataset distillation to imagenet-1k with constant memory. In *International Conference on Machine Learning*, pages 6565–6590. PMLR, 2023. 1, 2, 3, 5, 7
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. 7
- [5] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 1
- [6] 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, 2
- [7] Jianyang Gu, Saeed Vahidian, Vyacheslav Kungurtsev, Haonan Wang, Wei Jiang, Yang You, and Yiran Chen. Efficient dataset distillation via minimax diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15793–15803, 2024. 1, 2, 3
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceed-ings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 1, 6, 7
- [9] Ahmed Imteaj, Urmish Thakker, Shiqiang Wang, Jian Li, and M Hadi Amini. A survey on federated learning for resource-constrained iot devices. *IEEE Internet of Things Journal*, 9(1):1–24, 2021. 1
- [10] Zixuan Jiang, Jiaqi Gu, Mingjie Liu, and David Z Pan. Delving into effective gradient matching for dataset condensation. In 2023 IEEE International Conference on Omni-layer Intelligent Systems (COINS), pages 1–6. IEEE, 2023. 1, 2
- [11] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 7
- [12] Xuan-May Le, Ling Luo, Uwe Aickelin, and Minh-Tuan Tran. Shapeformer: Shapelet transformer for multivariate time series classification. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 1484–1494, 2024. 3
- [13] Xuan-May Le, Minh-Tuan Tran, and Van-Nam Huynh. Learning perceptual position-aware shapelets for time series classification. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 53– 69. Springer, 2022. 3
- [14] Xuan-May Thi Le, Tuan Minh Tran, and Hien T Nguyen. An improvement of sax representation for time series by using complexity invariance. *Intelligent Data Analysis*, 24(3):625– 641, 2020. 3

- [15] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. CS 231N, 7(7):3, 2015. 6
- [16] Hongcheng Li, Yucan Zhou, Xiaoyan Gu, Bo Li, and Weiping Wang. Diversified semantic distribution matching for dataset distillation. In Proceedings of the 32nd ACM International Conference on Multimedia, pages 7542–7550, 2024. 2
- [17] Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. arXiv preprint arXiv:1908.10084, 2019. 1
- [18] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (*CVPR*), pages 4510–4520, 2018. 6
- [19] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017. 1, 3
- [20] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556, 2014. 7
- [21] Peng Sun, Bei Shi, Daiwei Yu, and Tao Lin. On the diversity and realism of distilled dataset: An efficient dataset distillation paradigm. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pages 9390– 9399, 2024. 1, 2, 3, 4, 5, 6, 7, 8
- [22] Minh-Tuan Tran, Trung Le, Xuan-May Le, Jianfei Cai, Mehrtash Harandi, and Dinh Phung. Large-scale data-free knowledge distillation for imagenet via multi-resolution data generation. arXiv preprint arXiv:2411.17046, 2024. 4
- [23] Minh-Tuan Tran, Trung Le, Xuan-May Le, Mehrtash Harandi, and Dinh Phung. Text-enhanced data-free approach for federated class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23870–23880, 2024. 1
- [24] Minh-Tuan Tran, Trung Le, Xuan-May Le, Mehrtash Harandi, Quan Hung Tran, and Dinh Phung. Nayer: Noisy layer data generation for efficient and effective data-free knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23860–23869, 2024. 4
- [25] Minh-Tuan Tran, Xuan-May Le, Van-Nam Huynh, and Sung-Eui Yoon. Pisd: A linear complexity distance beats dynamic time warping on time series classification and clustering. *Engineering Applications of Artificial Intelligence*, 138:109222, 2024. 3
- [26] Tongzhou Wang, Jun-Yan Zhu, Antonio Torralba, and William T Freeman. Dataset distillation. arXiv preprint arXiv:1811.10959, 2018. 1
- [27] Zeyuan Yin, Eric Xing, and Zhiqiang Shen. Squeeze, recover and relabel: Dataset condensation at imagenet scale from a new perspective. Advances in Neural Information Processing Systems, 36, 2024. 1, 2, 3, 4, 5, 6, 7
- [28] Sangdoo Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable

features. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6023–6032, 2019. 5

- [29] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016. 7
- [30] Hongyi Zhang. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 5
- [31] Hansong Zhang, Shikun Li, Pengju Wang, Dan Zeng, and Shiming Ge. M3d: Dataset condensation by minimizing maximum mean discrepancy. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 9314–9322, 2024. 1, 2
- [32] Bo Zhao, Krishna Reddy Mopuri, and Hakan Bilen. Dataset condensation with gradient matching. In *Proceedings of* the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 708–717, 2021. 1, 2
- [33] Ganlong Zhao, Guanbin Li, Yipeng Qin, and Yizhou Yu. Improved distribution matching for dataset condensation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 7856–7865, 2023. 1, 2, 7
- [34] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2921–2929, 2016. 1, 3, 4