

Accelerating Dataset Distillation via Model Augmentation

Lei Zhang^{1*} Jie Zhang^{1*} Bowen Lei² Subhabrata Mukherjee³
Xiang Pan⁴ Bo Zhao⁵ Caiwen Ding⁶ Yao Li⁷ Dongkuan Xu^{8†}

¹Zhejiang University ²Texas A&M University ³Microsoft Research

⁴New York University ⁵Beijing Academy of Artificial Intelligence ⁶University of Connecticut

⁷University of North Carolina, Chapel Hill ⁸North Carolina State University

{zl_leizhang, zj_zhangjie}@zju.edu.cn dxu27@ncsu.edu

Abstract

*Dataset Distillation (DD), a newly emerging field, aims at generating much smaller but efficient synthetic training datasets from large ones. Existing DD methods based on gradient matching achieve leading performance; however, they are extremely computationally intensive as they require continuously optimizing a dataset among thousands of randomly initialized models. In this paper, we assume that training the synthetic data with diverse models leads to better generalization performance. Thus we propose two **model augmentation** techniques, i.e. using **early-stage models** and **parameter perturbation** to learn an informative synthetic set with significantly reduced training cost. Extensive experiments demonstrate that our method achieves up to 20× speedup and comparable performance on par with state-of-the-art methods.*

1. Introduction

Dataset Distillation (DD) [3, 48] or Dataset Condensation [55, 56], aims to reduce the training cost by generating a small but informative synthetic set of training examples; such that the performance of a model trained on the small synthetic set is similar to that trained on the original, large-scale dataset. Recently, DD has become an increasingly more popular research topic, and has been explored in a variety of contexts, including federated learning [17, 42], continual learning [33, 40], neural architecture search [43, 57], medical computing [25, 26] and graph neural networks [21, 30].

DD has been typically cast as a meta-learning problem [16] involving bilevel optimization. For instance, Wang *et al.* [48] formulate the network parameters as a function of the learnable synthetic set in the inner-loop

*Equal contribution.

†Corresponding author: Dongkuan Xu.




Figure 1. Performances of condensed datasets for training ConvNet-3 v.s. GPU hours to learn the 10 images per class condensed CIFAR-10 datasets with a single RTX-2080 GPU. Ours₅, Ours₁₀, and Ours₂₀ accelerates the training speed of the state-of-the-art method IDC [22] 5×, 10×, and 20× faster.

optimization; then optimize the synthetic set by minimizing classification loss on the real data in the outer-loop. This recursive computation hinders its application to real-world large-scale model training, which involves thousands to millions of gradient descent steps. Several methods have been proposed to improve the DD method by introducing ridge regression loss [2, 36], trajectory matching loss [3], *etc.* To avoid unrolling the recursive computation graph, Zhao *et al.* [57] propose to learn synthetic set by matching gradients generated by real and synthetic data when training deep networks. Based on this surrogate goal, several methods have been proposed to improve the informativeness or compatibility of synthetic datasets from other perspectives, ranging from data augmentation [55], contrastive signaling [24], resolution reduction [22], and bit encoding [41].

Although model training on a small synthetic set is fast, the dataset distillation process is typically expensive. For instance, the state-of-the-art method IDC [22] takes approximately 30 hours to condense 50,000 CIFAR-10 im-

ages into 500 synthetic images with a single RTX-2080 GPU, which is equivalent to the time it takes to train 60 ConvNet-3 models on the original dataset. Furthermore, the distillation time cost will rapidly increase for large-scale datasets *e.g.* ImageNet-1K, which prevents its application in computation-limited environments like end-user devices. Prior work [56] on reducing the distillation cost results in significant regression from the state-of-the-art performance. In this paper, we aim to speed up the dataset distillation process, while preserving even improving the testing performance over state-of-the-art methods.

Prior works are computationally expensive as they focus on generalization ability such that the learned synthetic set is useful to train many different networks as opposed to a targeted network. This requires optimizing the synthetic set over thousands of differently initialized networks. For example, IDC [22] learns the synthetic set over 2000 randomly initialized models, while the trajectory matching method (TM) [3] optimizes the synthetic set for 10000 distillation steps with 200 pre-trained expert models. Dataset distillation, which learns the synthetic data that is generalizable to unseen models, can be considered as an orthogonal approach to model training which learns model parameters that are generalizable to unseen data. Similarly, training the synthetic data with diverse models leads to better generalization performance. This intuitive idea leads to the following research questions:

Question 1. *How to design the candidate pool of models to learn synthetic data, for instance, consisting of randomly initialized, early-stage or well-trained models?*

Prior works [3, 22, 48, 57] use models from all training stages. The underlying assumption is that models from all training stages have similar importance. Zhao *et al.* [56] show that synthetic sets with similar generalization performance can be learned with different model parameter distributions, given an objective function in the form of feature distribution matching between real and synthetic data. In this paper, we take a closer look at this problem and show that learning synthetic data on **early-stage** models is more efficient for gradient/parameter matching based dataset distillation methods.

Question 2. *Can we learn a good synthetic set using only a few models?*

Our goal is to learn a synthetic set with a small number of (pre-trained) models to minimize the computational cost. However, using fewer models leads to poor generalization ability of the synthetic set. Therefore, we propose to apply **parameter perturbation** on selected early-stage models to incorporate model diversity and improve the generalization ability of the learned synthetic set.

In a nutshell, we propose two **model augmentation** techniques to accelerate the training speed of dataset distillation, namely using **early-stage models** and **parameter**

perturbation to learn an informative synthetic set with significantly less training cost. As illustrated in Fig. 1., our method achieves up to 20× speedup and comparable performance on par with state-of-the-art DD methods.

2. Related Work

2.1. Dataset Distillation

Recent advances in deep learning [6, 7, 13, 14, 53, 54] rely on massive amounts of training data that not only consume a lot of computational resources, but it is also time-consuming to train these models on large data. Dataset Distillation (DD) is introduced by Wang *et al.* [48], in which network parameters are modeled as functions of synthetic data, and learned by gradient-based hyperparameter optimization [32]. Subsequently, various works significantly improve the performance by learning on soft labels [2, 44], optimizing via infinite-width kernel limit [36, 37], matching on gradient-space [19, 57], model parameter-space [3], and distribution space [47, 56], amplifying contrastive signals [24], adopting data augmentations [55], and exploring regularity of dataset [22]. DD has been applied to various scenarios including continual learning [33, 38, 40], privacy [8], federated learning [11, 17, 52], graph neural network [20, 21], neural architecture search [43] for images [4], text [29], and medical imaging data [27]. In addition to the efforts made to improve performance and expand applications, few studies have focused on the efficiency of DD. This is a critical and practical problem closely related to the real-world application of DD.

2.2. Efficient Dataset Distillation

In this work, we focus on the efficiency of dataset distillation algorithm, which is under-explored in previous works. Zhao *et al.* [56] make improvements in efficiency via distribution matching in random embedding spaces, which replaces expensive bi-level optimization in common methods [22, 57]. However, the speed-up of DD in their work results in a significant drop in performance, which exhibits a large gap between their method and other SOTA DD methods [22]. Cazenavette *et al.* [4] improve efficiency via parameter matching in pre-trained networks. However, they need to pre-train 100 networks from scratch on real data, which leads to massively increased computational resources. In this work, we seek to significantly reduce training time and lower computational resources, while maintaining comparable performance.

3. Preliminary

The goal of dataset distillation is to generate a synthetic dataset \mathcal{S} from the original training dataset \mathcal{T} such that an arbitrary model trained on \mathcal{S} is similar to the

one trained on \mathcal{T} . Among various dataset distillation approaches [3, 22, 37, 56], gradient-matching methods have achieved state-of-the-art performance. However, they require a large amount of training time and expensive computational resources. In this paper, we propose to use gradient matching to reduce the computational requirement while maintaining similar performance.

Gradient Matching. Gradient-matching dataset distillation approach [57] matches the network gradients on synthetic dataset \mathcal{S} to the gradients on real dataset \mathcal{T} . The overall training object can be formulated as:

$$\begin{aligned} & \underset{\mathcal{S}}{\text{maximize}} \sum_{t=0}^T \text{Cos}(\nabla_{\theta} \ell(\theta_t; \mathcal{S}), \nabla_{\theta} \ell(\theta_t; \mathcal{T})) \\ & \text{w.r.t. } \theta_{t+1} = \theta_t - \eta \nabla_{\theta} \ell(\theta_t; \mathcal{S}) \end{aligned} \quad (1)$$

where θ_t denotes the network weights at the t^{th} training step from the randomly initialized weights θ_0 given \mathcal{S} , $\ell(\theta, \mathcal{S})$ denotes the training loss for weight θ and the dataset \mathcal{S} , ℓ denotes loss function, and $\text{Cos}(\cdot, \cdot)$ denotes the channel-wise cosine similarity.

In addition, recent works have made various efforts to enhance the performance of gradient-matching from the perspective of data diversity. Zhao *et al.* [55] utilize differentiable siamese augmentation to synthesize more informative images. Kim *et al.* [22] explore the regularity of dataset to strengthen the representability of condensed datasets.

Discussion on Efficiency. Current works [22, 55, 57] use a large number of randomly initialized networks (*e.g.*, 2000) to improve the generalization performance of condensed dataset. The huge number of models makes the DD process time-consuming and computation-expensive. For instance, condensing 1 image per class in a synthetic dataset of CIFAR-10 by using state-of-the-art method IDC [22] consumes 200k epochs of updating network, in addition to the 2,000k epochs of updating \mathcal{S} , which requires over 22.2 hours on a single RTX-2080 GPU. While Zhao *et al.* [56] make efforts to solve computation the challenge by using distribution-matching instead of gradient-matching – reducing number of updates from 200k to 20k and training time from 22.2 hours to 0.83 hours – the accuracy of condensed data also degrades dramatically from 50.6% to 26.0%. This potentially results from the redundant learning on randomly initialized networks.

4. Method

4.1. Overview

We illustrate the framework of our proposed efficient dataset distillation method in Fig. 2. Our method consists of three stages: 1) Early-stage Pre-training, 2) Parameter Perturbation, and 3) Distillation via gradient-matching. In stage 1, we utilize pre-trained networks at the early stage

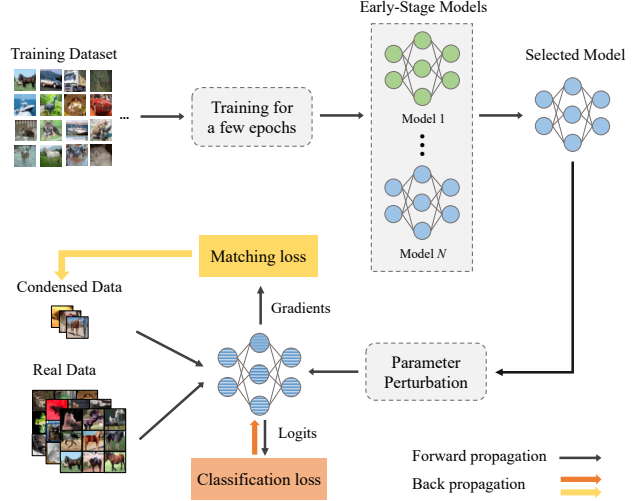


Figure 2. The illustration of our proposed fast dataset distillation method. We perform early-stage pretraining and parameter perturbation on models in dataset distillation.

as an informative parameter space for dataset distillation. In stage 2, we conduct parameter perturbation on models selected from stage 1 to further augment the diversity of model parameter distribution. In stage 3, the synthetic dataset is optimized with gradient-matching strategy on these augmented models from early stages.

4.2. Early-Stage Models: Initializing with Informative Parameter Space

Existing gradient-matching methods [22, 55, 57] train synthetic data on a large number of randomly initialized networks for learning to generalize to unseen initializations. Furthermore, the initialized networks will be updated for many SGD steps in the inner-loop for learning better synthetic data, which requires much computational resources.

Data augmentation is frequently used to prevent overfitting and improve generalization performance when optimizing deep networks [49, 51]. Similarly, we propose to use model augmentation to improve the generalization performance when learning condensed datasets. Inspired by ModelSoups [31, 50], a practical method to improve performance of model ensembles, we pre-train a set of networks with different hyper-parameters, including learning rate, random seed, and data augmentation, so that we construct a parameter space with rich diversity. Instead of leveraging randomly initialized networks in each outer loop in traditional methods, we sample those early-stage networks as the initialization, which are more informative for implementing gradient matching.

Comparing with well-trained networks, using early-stage networks have two benefits. First, early-stage networks require less training cost. Second, the early-stage

networks have rich diversity [1, 12, 39] and provide large gradients [10], which leads to better gradient matching. More discussion can be found in the supplementary.

4.3. Parameter Perturbation: Diversifying Parameter Space

Motivated by the data perturbation which is widely used to diversify the training data for better knowledge distillation [34, 35], we propose to conduct the model perturbation in dataset distillation for further diversifying the parameter space. We implement perturbation after sampling the network (parameters) from the early-stage parameter space in each outer loop.

We formulate our fast dataset distillation as the gradient-matching on parameter-perturbed early-stage models between real data and synthetic data:

$$\min_{\mathcal{S}} D(\nabla_{\theta} \ell(\hat{\theta}; \mathcal{S}), \nabla_{\theta} \ell(\hat{\theta}; \mathcal{T})) \quad (2)$$

$$\text{w.r.t. } \hat{\theta} \leftarrow \theta^{\mathcal{T}} + \alpha \cdot \mathbf{d},$$

where $\theta^{\mathcal{T}}$ represents network weights trained on real data \mathcal{T} , D denotes a distance-based matching objective, and α is the magnitude of parameter perturbation. \mathbf{d} is sampled from a Gaussian distribution $\mathcal{N}(0, \mathbf{I})$ with dimensions compatible with network parameter θ and filter normalized by

$$\mathbf{d}_{l,j} \leftarrow \frac{\mathbf{d}_{l,j}}{\|\mathbf{d}_{l,j}\|_F + \epsilon} \|\theta_{l,j}\|_F \quad (3)$$

to eliminate the scaling invariance of neural networks [28], where $\mathbf{d}_{l,j}$ is the j -th filter at the l -th layer of \mathbf{d} and $\|\cdot\|_F$ denotes the Frobenius norm. ϵ is a small positive constant.

4.4. Training Algorithm

We depict our method in Algorithm 1. We build our training algorithm on the state-of-the-art method IDC [22]. Before dataset distillation, we pre-trained N models on real data for only a few epochs. This is significantly cheaper than existing methods that well-train many networks till convergence. We train the condensed dataset \mathcal{S} for T outer loops and M inner loops. At each outer loop, we randomly select a model from N early-stage models as initialization and employ parameter perturbation on it. At each inner loop, we optimize the synthetic samples \mathcal{S} by minimizing the gradient matching loss with regard to the sampled real batch \mathcal{T}_c and real synthetic batch \mathcal{S}_c of the same class c , respectively. The network θ_m is then updated on real data. Please refer to [22] for more details. The numbers of pre-train epochs P and outer loop K are relatively small. In experiments, we set $P = 2$ compared with 300 for a well-trained network and $K = 400$ compared with 2000 in SOTA DD method IDC [22]. Note that our method can also be easily applied to other dataset distillation methods for reducing training time, and we explore it in Sec. 5.3.

Algorithm 1: Efficient Dataset Distillation

Input: Training data \mathcal{T} , loss function l , number of classes C , number of model N , magnitude α , augmentation function \mathcal{A} , multi-information function f , deep neural network ψ_{θ} parameterized with θ

Output: Condensed dataset \mathcal{S}

Definition: $D(B, B'; \theta) = \|\nabla_{\theta} \ell(\theta; B) - \nabla_{\theta} \ell(\theta; B')\|$

```

/* Early-Stage Pre-train */
1 Randomly initialize  $N$  networks  $\{\tau_1 \dots \tau_N\}$ ;
2 for  $n \leftarrow 1$  to  $N$  do
3   Update network  $\tau_n$  on real data  $\mathcal{T}$ :
4   for  $p \leftarrow 1$  to  $P$  do
5      $\tau_{n,p+1} \leftarrow \tau_{n,p} - \eta \nabla_{\tau_{n,p}} \ell(\tau_{n,p}; \mathcal{A}(\mathcal{T}))$ 
6   end
7 end
8 Initialize condensed dataset  $\mathcal{S}$ 
9 for  $t \leftarrow 0$  to  $T$  do
10  Randomly load one checkpoint from  $\{\tau_1 \dots \tau_N\}$ 
    to initialize  $\psi_{\theta}$ ;
    /* Parameter Perturbation */
11  Sample vector  $\mathbf{d}$  from Gaussian distribution
    Parameter perturbation on  $\psi_{\theta}$ :  $\theta \leftarrow \theta + \alpha \cdot \mathbf{d}$ 
12  for  $m \leftarrow 0$  to  $M$  do
13    for  $c \leftarrow 0$  to  $C$  do
14      Sample an intra-class mini-batch
15       $\mathcal{T}_c \sim \mathcal{T}, \mathcal{S}_c \sim \mathcal{S}$ 
16      Update synthetic data  $\mathcal{S}_c$ :
17       $\mathcal{S}_c \leftarrow \mathcal{S}_c - \lambda \nabla_{\mathcal{S}_c} D(\mathcal{A}(f(\mathcal{S}_c)), \mathcal{A}(\mathcal{T}_c))$ 
18    end
19    Sample a mini-batch  $T \sim \mathcal{T}$ 
20    Update network  $\psi_{\theta}$  w.r.t classification loss:
21     $\theta_{m+1} \leftarrow \theta_m - \eta \nabla_{\theta} \ell(\theta_m; \mathcal{A}(T))$ 
22  end
23 end

```

5. Experiments

In this section, we first evaluate our method on various datasets against state-of-the-art baselines. Next, we examine the proposed method in depth with ablation analysis.

5.1. Experimental Setups

Datasets. We evaluate performance of neural networks trained on condensed datasets generated by several methods as baselines. Following previous works [4, 22, 57], we conduct experiments on both low- and high-resolution datasets including CIFAR-10, CIFAR-100 [23], and ImageNet [5].

Network Architectures. Following previous works [22, 56], we use a depth-3 ConvNet [39] on CIFAR-10 and CIFAR-100. For ImageNet subsets, we follow IDC [22] and adopt ResNetAP-10 for dataset distillation, a modified

Dataset	Method	Img/Cls			Speed Up	Acc. Gain
		1	10	50		
CIFAR-10	Full Dataset	88.1	88.1	88.1	-	-
	IDC [22]	50.6 (21.7h)	67.5 (22.2h)	74.5 (29.4h)	1.00×	1.00×
	CAFE [47]	30.3	46.3	55.5	-	0.54×
	DSA [55]	28.2 (0.09h)	52.1 (1.94h)	60.6 (11.1h)	85.0×	0.71×
	DM [56]	26.0 (0.25h)	48.9 (0.26h)	63.0 (0.31h)	89.0×	0.69×
	TM [3]	46.3 (6.35h)	65.3 (6.69h)	71.6 (7.39h)	3.57×	0.94×
	Ours ₅	49.2 (4.44h)	67.1 (4.45h)	73.8 (6.11h)	4.90×	0.99×
	Ours ₁₀	48.5 (2.22h)	66.5 (2.23h)	73.1 (3.05h)	9.77×	0.97×
CIFAR-100	Full Dataset	56.2	56.2	56.2	-	-
	IDC [22]	25.1 (125h)	45.1 (127h)	-	1.00×	1.00×
	CAFE [47]	12.9	27.8	37.9	-	0.56×
	DSA [55]	13.9 (0.83h)	32.3 (17.5h)	42.8 (221.1h)	78.9×	0.63×
	DM [56]	11.4 (1.67h)	29.7 (2.64h)	43.6 (2.78h)	61.4×	0.55×
	TM [3]	24.3 (7.74h)	40.1 (9.47h)	47.7 (-)	14.7×	0.92×
	Ours ₅	29.8 (25.1h)	45.6 (25.6h)	52.6 (42.00h)	4.97×	1.10×
	Ours ₁₀	29.4 (12.5h)	45.2 (12.8h)	52.2 (21.00h)	9.96×	1.09×
Ours ₂₀	29.1 (6.27h)	44.1 (6.40h)	52.1 (10.50h)	19.9×	1.07×	

Table 1. Comparing efficiency and performance of dataset distillation methods on CIFAR-10 and CIFAR-100. Speed up represents the average acceleration amount of training time on a single RTX-2080 GPU with the same batch size 64. Acc. Gain represents the average improvement in test accuracy of network trained on the condensed dataset over IDC [22]. Training time is not reported for CAFE [47] that does not provide official implementation and IDC [22] that requires more than one GPU on CIFAR-100 for Img/Cls=50.

Dataset	Method	Img/Cls		Speed Up	Acc. Gain
		10	20		
ImageNet-10	Full Dataset	90.8	90.8	-	-
	IDC [22]	72.8 (70.14h)	76.6 (92.78h)	1.00×	1.00×
	DSA [55]	52.7 (26.95h)	57.4 (51.39h)	2.20×	0.73×
	DM [56]	52.3 (1.39h)	59.3 (3.61h)	38.1×	0.74×
	Ours ₅	74.6 (15.52h)	76.3 (20.05h)	4.57×	1.01×
ImageNet-100	Full Dataset	82.0	82.0	-	-
	IDC [22]	46.7 (141h)	53.7 (185h)	1.00×	1.00×
	DSA [55]	21.8 (9.72h)	30.7 (23.9h)	14.1×	0.51×
	DM [56]	22.3 (2.78h)	30.4 (2.81h)	58.2×	0.52×
	Ours ₅	48.4 (29.8h)	56.0 (38.6h)	4.76×	1.04×

Table 2. Comparing efficiency and performance of dataset distillation methods on ImageNet-10 and ImageNet-100. We measure the training time on a single RTX-A6000 GPU with the same training hyperparameters. For ImageNet-100, we follow IDC [22] to split the whole dataset into five tasks with 20 classes each for faster optimization. The training time reported in ImageNet-100 is for one task.

ResNet-10 [15] by replacing strided convolution as average pooling for downsampling.

Evaluation Metrics. We study several methods in terms of performance and efficiency. The performance is measured by the testing accuracy of networks trained on condensed datasets. The efficiency is measured by GPU hours required by the dataset distillation process [9]. For a fair comparison, all GPU hours are measured on a single GPU. The training

time of condensing CIFAR-10, CIFAR-100 and ImageNet subsets is evaluated on RTX-2080 GPU and RTX-A6000 GPU, respectively. We adopt FLOPs as a metric of computational efficiency.

Baselines. We compare our method with several prominent dataset condensation methods like (1) gradient-matching method including DSA [55] and IDC [22] (2) distribution-matching including DM [56] and CAFE [47] (3) parameter-




Figure 3. Performance comparison across a varying number of training steps.




Figure 4. Performance comparison across varying training time and FLOPs.

matching including TM [3]. We use the state-of-the-art dataset distillation method IDC as the strongest baseline to calculate the gap between other methods on performance and efficiency.

Training Details. We adopt IDC as the backbone of our method, which is the state-of-the-art gradient-matching dataset distillation method. The outer loops and learning rate of condensed data are 400/100 and 0.01/0.1 for CIFAR-10/100 and ImageNet-Subsets. We employ 5/10 pre-trained models for CIFAR-10/100 and ImageNet. The number of pre-train epochs is 2/5/10 for CIFAR-10/100, ImageNet-10, and ImageNet-100. The setting of other hyperparameters follows IDC [22] including the number of inner loops, batch size, and augmentation strategy.

5.2. Condensed Data Evaluation

CIFAR-10 & CIFAR-100. Our method achieves a better trade-off in task performance vs. the amount of training time and computation compared to other state-of-the-art baselines on CIFAR-10 and CIFAR-100. For instance, as shown in Tab. 1, our method is comparable to IDC while achieving $5\times$ and $10\times$ **speed ups** on CIFAR-10. Our method shows 10%, 9%, and 7% performance improvements over IDC on CIFAR-100 while achieving $5\times$, $10\times$, and $20\times$ **acceleration**, respectively.

To further demonstrate the advantages of our method, we report the evaluation results across a varying amount

of computational resources in the form of the number of training steps in Fig. 3, training time, and FLOPs in Fig. 4. We observe that our method consistently outperforms all the baselines across different training steps, training times, and FLOPs. This demonstrates the effectiveness of our distillation method in capturing informative features from early-stage training; and enhanced diversity of the models for better generalizability. Interestingly, our method obtains better performance and efficiency over state-of-the-art baselines on CIFAR-100 as compared to CIFAR-10. This demonstrates the effectiveness and scalability of our method on large-scale datasets which makes it more appealing for all practical purposes.

ImageNet. Apart from CIFAR-10/100, we further investigate the performance and efficiency of our method on the high-resolution dataset ImageNet. Following previous baselines [22, 46], we evaluate our method on ImageNet-subset consisting of 10 and 100 classes.

We observe that the dataset distillation methods on ImageNet suffer from severe efficiency challenges. As shown in Tab. 2, dataset distillation method IDC [22] achieves high performance while requiring almost 4 days on ImageNet-10; while DSA [57] and DM [55] are more efficient in training time with significantly poor performance. The accuracy of networks trained on condensed data generated by our method outperforms all existing state-of-the-art baselines with the least training time. For instance, our method

requires **less than 1 day** to condense ImageNet-10, which leads to $5\times$ **speedup** over SOTA methods.

As shown in Fig. 3 and Fig. 4, we conduct extensive experiments with various training budgets. The results demonstrate that our method requires significantly fewer training steps, time, and computation resources to reach the same performance as the SOTA method IDC and achieves higher performance with the same training budgets. This indicates that utilizing early-stage models as initialization guides dataset distillation to focus on distinguishing features at the beginning of distillation. The exploration of diversity expands the parameter space and reduces the amount of time on learning repeated and redundant features.

Dataset	Method	Evaluation model		
		ConvNet-3	ResNet-10	DenseNet-121
CIFAR-100	IDC [22]	45.1	38.9	39.5
	Ours ₅	46.5	38.4	39.6

(a) The performance of condensed CIFAR-100 dataset (10 images per class) trained on ConvNet-3 on different network architectures.

Dataset	Method	Evaluation model		
		ResNetAP-10	ResNet-18	EfficientNet-B0
ImageNet-10	IDC [22]	74.0	73.1	74.3
	Ours ₅	74.6	74.5	75.4

(b) The performance of condensed ImageNet-10 dataset (10 images per class) trained on ResNetAP-10 on different network architectures.

Table 3. Performance of synthetic data learned on CIFAR100 and ImageNet-10 datasets with different architectures. The networks are trained on condensed dataset and validated on test dataset.

Cross-Architecture Generalization. We also evaluate the performance of our condensed data on architectures different from the one used to distill it on the CIFAR-100 (1 and 10 images per class) and ImageNet-10 (10 images per class). In Tab. 3, we show the performance of our baselines ConvNet-3 and ResNetAP-10 evaluated on ResNet-18 [15], DenseNet-121 [18], and EfficientNet-B0 [45].


For IDC [22], we use condensed data provided by the official implementation for evaluation of their method. Our method obtains the best performance on all the transfer models except for ResNet-10 on CIFAR-100 (10 images per class) where we lie within one standard deviation of IDC – demonstrating the robustness of our method to changes in network architecture.

5.3. Analysis

We perform ablation studies on our efficient dataset distillation method described in Sec. 4. Specifically, we measure the impact of (1) the number of epochs of pre-training on real data, (2) the magnitude of parameter perturbation, (3) the number of early-stage models, and (4) the acceleration of training.

Epochs of Pre-training. We study the effect of pre-training

epochs on networks used in our method in terms of test accuracy on CIFAR-10 (10 images per class) and demonstrate results in Fig. 5a. We observe that early-stage networks pre-trained with 2 epochs perform significantly better than randomly initialized networks and well-trained networks with 300 epochs. The results demonstrate that early-stage networks contain a more informative parameter space than randomly initialized networks, thereby helping the condensed datasets to capture features more efficiently. While it is generally known that well-trained networks perform better, well-trained networks tend to get stuck in local optima and lack diversity among parameter spaces. On the other hand, early-stage models provide flexible and informative guidance for dataset distillation.



(a) Effect of pre-train epochs (b) Effect of perturbation magnitude

Figure 5. Condensation performance from networks pre-trained for different epochs and varying magnitudes of parameter perturbation. The networks are trained with same hyper-parameters except for training epochs and perturbation magnitudes, respectively. Evaluation is performed on CIFAR-10 (10 images per class).

Magnitude of Parameter Perturbation. We study the effect of the magnitude α of parameter perturbation in terms of test accuracy on CIFAR-10 (10 images per class) and report results in Fig. 5b. We observe that condensed dataset achieves better performance on both accuracy and efficiency when magnitude α is carefully set as shown in Fig. 5b. When the magnitude is large, e.g., 10, the perturbed networks diverge from the original space; the perturbed parameter space contains less relevant and inconsistent information, thereby impacting performance and efficiency. When the magnitude is small, such as not employing parameter perturbation, the parameter space lacks diversity compared to well-designed perturbed parameter space. Experimental results show that $\alpha = 1$ is optimal for CIFAR in our setting which works consistently better across all training steps. Well-designed magnitude makes perturbed networks concentrated around the original network, thereby augmenting the parameter space with diversified and relevant information.

Number of Early-Stage Models. We study the effect of the number of early-stage models in our experiment and show the results in Fig. 6. It is observed that the number of

early-stage models N has less impact on the test accuracy of the condensed dataset. We argue that parameter perturbation in our method plays an important role in exploring the diversity of early-stage models; such that the description of parameter space depends on the representation of models rather than the number of models. In our method, a few models, *e.g.* 5, can achieve comparable performance to SOTA [22], with two significant advantages. The first is to shorten training time as the number of outer loops in DD is closely related to the number of models N . The second is to reduce computation resources in network pre-training. TM [4] also utilizes network pre-training in DD, however, the number of models in their method is relatively large, *e.g.* 50, which is $10\times$ more than ours. Parameter perturbation in our method augments the diversity of models and improves efficiency with only a small number of models.




Figure 6. Condensation performance from a varying number of early-stage models. Performances with a varying number of models are similar, which demonstrates that our method is not sensitive to the number of models to achieve high performance.




Figure 7. Performance of our method applied to different dataset distillation methods on CIFAR-10 dataset (10 images per class). Our results are reported with $5\times$ **training acceleration**.

Acceleration of Training. We study the effect of acceleration of training on existing DD methods [22, 55, 57] and our method. We observe our method to retain similar performance with minor regression to increased training acceleration / speed-ups – while the performance of existing methods drops dramatically in Tab. 4. Our method achieves

Speed up	DC [57]	DSA [55]	IDC [22]	Ours
$1\times$	44.9	52.1	67.5	-
$5\times$	41.6 (-3.3)	47.0 (-5.1)	66.2 (-1.3)	67.1
$10\times$	39.2 (-5.7)	46.2 (-5.9)	65.0 (-2.5)	66.5 (-0.6)
$20\times$	37.8 (-7.1)	44.8 (-7.3)	63.7 (-3.8)	65.2 (-1.9)

(a) CIFAR-10 (Img/Cls=10)

Speed up	DC [57]	DSA [55]	IDC [22]	Ours
$1\times$	53.9	60.6	74.5	-
$5\times$	50.3 (-3.6)	56.5 (-4.1)	73.3 (-1.2)	73.8
$10\times$	47.3 (-6.6)	55.7 (-4.9)	72.0 (-2.5)	73.1 (-0.7)
$20\times$	42.0 (-11.9)	54.1 (-6.5)	71.1 (-3.4)	71.7 (-2.1)

(b) CIFAR-10 (Img/Cls=50)

Speed up	DC [57]	DSA [55]	IDC [22]	Ours
$1\times$	29.5	32.3	45.1	-
$5\times$	23.1 (-6.4)	29.3 (-3.0)	43.4 (-1.9)	46.2
$10\times$	21.1 (-8.4)	28.7 (-3.6)	41.6 (-3.5)	45.6 (-0.6)
$20\times$	18.6 (-10.9)	27.9 (-4.4)	40.5 (-4.6)	45.0 (-1.2)

(c) CIFAR-100 (Img/Cls=10)

Table 4. Condensation performance with different acceleration / speed ups over state-of-the-art dataset distillation approaches. We show performance drop between increased speed up in brackets. Our method achieves higher performance over baseline methods at all levels of speed up. With increased speed up, our method shows minor regression in performance.

better performance than baselines at all levels of speed-up. This demonstrates the informativeness of our parameter space in terms of diversity and reduced redundancy; such that the condensed dataset does not learn similar information repeatedly and captures sufficient features efficiently. It is worth noting that our method performs better with less regression at higher levels of speed up on the more complex dataset, *e.g.*, CIFAR-100. We also demonstrate our method can be orthogonally applied to other dataset distillation methods in Fig. 7. We apply parameter perturbation on other DD methods to accelerate the training $5\times$ faster. This indicates better scalability and improved efficiency of our method in condensing large-scale datasets.

6. Conclusion

In this work, we introduce a novel method for improving the efficiency of gradient-matching based dataset distillation approaches. We leverage model augmentation strategies with early-stage training and parameter perturbation to increase the diversity of the parameter space as well as massively reduce the computation resource for dataset distillation. Our method is able to achieve $10\times$ acceleration on CIFAR and $5\times$ acceleration on ImageNet. As the first attempt to improve the efficiency of gradient-matching based dataset distillation, the proposed method successfully crafts a condensed dataset of ImageNet in 18 hours, making dataset distillation more applicable in real-world settings.

References

- [1] Alessandro Achille, Matteo Rovere, and Stefano Soatto. Critical learning periods in deep neural networks. *CoRR*, abs/1711.08856, 2017. 4
- [2] Ondrej Bohdal, Yongxin Yang, and Timothy M. Hospedales. Flexible dataset distillation: Learn labels instead of images. *CoRR*, abs/2006.08572, 2020. 1, 2
- [3] George Cazenavette, Tongzhou Wang, Antonio Torralba, Alexei A. Efros, and Jun-Yan Zhu. Dataset distillation by matching training trajectories. In *CVPR*, pages 10708–10717, 2022. 1, 2, 3, 5, 6
- [4] George Cazenavette, Tongzhou Wang, Antonio Torralba, Alexei A. Efros, and Jun-Yan Zhu. Wearable imagenet: Synthesizing tileable textures via dataset distillation. In *CVPR Workshops*, pages 2277–2281, 2022. 2, 4, 8
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. 4
- [6] Jiahua Dong, Yang Cong, Gan Sun, Bineng Zhong, and Xiaowei Xu. What can be transferred: Unsupervised domain adaptation for endoscopic lesions segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4022–4031, June 2020. 2
- [7] Jiahua Dong, Lixu Wang, Zhen Fang, Gan Sun, Shichao Xu, Xiao Wang, and Qi Zhu. Federated class-incremental learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022. 2
- [8] Tian Dong, Bo Zhao, and Lingjuan Lyu. Privacy for free: How does dataset condensation help privacy? In *ICML*, volume 162, pages 5378–5396, 2022. 2
- [9] Gongfan Fang, Kanya Mo, Xinchao Wang, Jie Song, Shitao Bei, Haofei Zhang, and Mingli Song. Up to 100x faster data-free knowledge distillation. In *AAAI*, pages 6597–6604, 2022. 5
- [10] Jonathan Frankle, David J. Schwab, and Ari S. Morcos. The early phase of neural network training. In *ICLR*, 2020. 4
- [11] Jack Goetz and Ambuj Tewari. Federated learning via synthetic data. *CoRR*, abs/2008.04489, 2020. 2
- [12] Guy Gur-Ari, Daniel A. Roberts, and Ethan Dyer. Gradient descent happens in a tiny subspace. *CoRR*, abs/1812.04754, 2018. 4
- [13] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. 2
- [14] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. 2
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 5, 7
- [16] Timothy Hospedales, Antreas Antoniou, Paul Micaelli, and Amos Storkey. Meta-learning in neural networks: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 44(9):5149–5169, 2021. 1
- [17] Shengyuan Hu, Jack Goetz, Kshitiz Malik, Hongyuan Zhan, Zhe Liu, and Yue Liu. FedSynth: Gradient compression via synthetic data in federated learning. *CoRR*, abs/2204.01273, 2022. 1, 2
- [18] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *CVPR*, pages 2261–2269, 2017. 7
- [19] Zixuan Jiang, Jiaqi Gu, Mingjie Liu, and David Z. Pan. Delving into effective gradient matching for dataset condensation. *CoRR*, abs/2208.00311, 2022. 2
- [20] Wei Jin, Xianfeng Tang, Haoming Jiang, Zheng Li, Danqing Zhang, Jiliang Tang, and Bing Yin. Condensing graphs via one-step gradient matching. In *KDD*, pages 720–730, 2022. 2
- [21] Wei Jin, Lingxiao Zhao, Shichang Zhang, Yozen Liu, Jiliang Tang, and Neil Shah. Graph condensation for graph neural networks. In *ICLR*, 2022. 1, 2
- [22] Jang-Hyun Kim, Jinuk Kim, Seong Joon Oh, Sangdoon Yun, Hwanjun Song, Joonhyun Jeong, Jung-Woo Ha, and Hyun Oh Song. Dataset condensation via efficient synthetic-data parameterization. In *ICML*, volume 162, pages 11102–11118, 2022. 1, 2, 3, 4, 5, 6, 7, 8
- [23] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 4
- [24] Saehyung Lee, Sanghyuk Chun, Sangwon Jung, Sangdoon Yun, and Sungroh Yoon. Dataset condensation with contrastive signals. In *ICML*, volume 162, pages 12352–12364, 2022. 1, 2
- [25] Guang Li, Ren Togo, Takahiro Ogawa, and Miki Haseyama. Soft-label anonymous gastric x-ray image distillation. In *ICIP*, pages 305–309, 2020. 1
- [26] Guang Li, Ren Togo, Takahiro Ogawa, and Miki Haseyama. Compressed gastric image generation based on soft-label dataset distillation for medical data sharing. *Computer Methods and Programs in Biomedicine*, page 107189, 2022. 1
- [27] Guang Li, Ren Togo, Takahiro Ogawa, and Miki Haseyama. Dataset distillation for medical dataset sharing. *CoRR*, abs/2209.14603, 2022. 2
- [28] Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape of neural nets. In *NIPS*, pages 6391–6401, 2018. 4
- [29] Yongqi Li and Wenjie Li. Data distillation for text classification. *CoRR*, abs/2104.08448, 2021. 2
- [30] Mengyang Liu, Shanchuan Li, Xinshi Chen, and Le Song. Graph condensation via receptive field distribution matching. *CoRR*, abs/2206.13697, 2022. 1
- [31] Raphael Gontijo Lopes, Yann Dauphin, and Ekin Dogus Cubuk. No one representation to rule them all: Overlapping features of training methods. In *ICLR*, 2022. 3
- [32] Dougal Maclaurin, David Duvenaud, and Ryan P. Adams. Gradient-based hyperparameter optimization through reversible learning. In *ICML*, volume 37, pages 2113–2122, 2015. 2
- [33] Wojciech Masarczyk and Ivona Tautkute. Reducing catastrophic forgetting with learning on synthetic data. In *CVPR Workshops*, pages 1019–1024, 2020. 1, 2

- [34] Giung Nam, Hyungi Lee, Byeongho Heo, and Juho Lee. Improving ensemble distillation with weight averaging and diversifying perturbation. In *ICML*, volume 162, pages 16353–16367, 2022. 4
- [35] Giung Nam, Jongmin Yoon, Yoonho Lee, and Juho Lee. Diversity matters when learning from ensembles. In *NIPS*, pages 8367–8377, 2021. 4
- [36] Timothy Nguyen, Zhourong Chen, and Jaehoon Lee. Dataset meta-learning from kernel ridge-regression. In *ICLR*, 2021. 1, 2
- [37] Timothy Nguyen, Roman Novak, Lechao Xiao, and Jaehoon Lee. Dataset distillation with infinitely wide convolutional networks. In *NIPS*, pages 5186–5198, 2021. 2, 3
- [38] Andrea Rosasco, Antonio Carta, Andrea Cossu, Vincenzo Lomonaco, and Davide Bacciu. Distilled replay: Overcoming forgetting through synthetic samples. *CoRR*, abs/2103.15851, 2021. 2
- [39] Levent Sagun, Utku Evci, V. Ugur Güney, Yann N. Dauphin, and Léon Bottou. Empirical analysis of the hessian of over-parametrized neural networks. In *ICLR Workshop*, 2018. 4
- [40] Mattia Sangermano, Antonio Carta, Andrea Cossu, and Davide Bacciu. Sample condensation in online continual learning. In *IJCNN*, pages 1–8, 2022. 1, 2
- [41] Robin Tibor Schirrmeyer, Rosanne Liu, Sara Hooker, and Tonio Ball. When less is more: Simplifying inputs aids neural network understanding. *arXiv preprint arXiv:2201.05610*, 2022. 1
- [42] Rui Song, Dai Liu, Dave Zhenyu Chen, Andreas Festag, Carsten Trinitis, Martin Schulz, and Alois C. Knoll. Federated learning via decentralized dataset distillation in resource-constrained edge environments. *CoRR*, abs/2208.11311, 2022. 1
- [43] Felipe Petroski Such, Aditya Rawal, Joel Lehman, Kenneth O. Stanley, and Jeffrey Clune. Generative teaching networks: Accelerating neural architecture search by learning to generate synthetic training data. In *ICML*, volume 119, pages 9206–9216, 2020. 1, 2
- [44] Iliia Sucholutsky and Matthias Schonlau. Soft-label dataset distillation and text dataset distillation. In *IJCNN*, pages 1–8, 2021. 2
- [45] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, volume 97, pages 6105–6114, 2019. 7
- [46] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *ECCV*, volume 12356, pages 776–794, 2020. 6
- [47] Kai Wang, Bo Zhao, Xiangyu Peng, Zheng Zhu, Shuo Yang, Shuo Wang, Guan Huang, Hakan Bilen, Xinchao Wang, and Yang You. CAFE: learning to condense dataset by aligning features. In *CVPR*, pages 12186–12195, 2022. 2, 5
- [48] Tongzhou Wang, Jun-Yan Zhu, Antonio Torralba, and Alexei A. Efros. Dataset distillation. *CoRR*, abs/1811.10959, 2018. 1, 2
- [49] Qingsong Wen, Liang Sun, Fan Yang, Xiaomin Song, Jingkun Gao, Xue Wang, and Huan Xu. Time series data augmentation for deep learning: A survey. In *IJCAI*, pages 4653–4660, 2021. 3
- [50] Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo Lopes, Ari S. Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, and Ludwig Schmidt. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *ICML*, volume 162, pages 23965–23998, 2022. 3
- [51] Sen Wu, Hongyang R. Zhang, Gregory Valiant, and Christopher Ré. On the generalization effects of linear transformations in data augmentation. In *ICML*, volume 119, pages 10410–10420, 2020. 3
- [52] Yuanhao Xiong, Ruochen Wang, Minhao Cheng, Felix Yu, and Cho-Jui Hsieh. Feddm: Iterative distribution matching for communication-efficient federated learning. *CoRR*, abs/2207.09653, 2022. 2
- [53] Jie Zhang, Bo Li, Chen Chen, Lingjuan Lyu, Shuang Wu, Shouhong Ding, and Chao Wu. Delving into the adversarial robustness of federated learning. *arXiv preprint arXiv:2302.09479*, 2023. 2
- [54] Jie Zhang, Bo Li, Jianghe Xu, Shuang Wu, Shouhong Ding, Lei Zhang, and Chao Wu. Towards efficient data free black-box adversarial attack. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15115–15125, 2022. 2
- [55] Bo Zhao and Hakan Bilen. Dataset condensation with differentiable siamese augmentation. In *ICML*, volume 139, pages 12674–12685, 2021. 1, 2, 3, 5, 6, 8
- [56] Bo Zhao and Hakan Bilen. Dataset condensation with distribution matching. In *WACV*, pages 6503–6512, 2023. 1, 2, 3, 4, 5
- [57] Bo Zhao, Konda Reddy Mopuri, and Hakan Bilen. Dataset condensation with gradient matching. In *ICLR*, 2021. 1, 2, 3, 4, 6, 8