**GyF** 

This WACV 2023 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.

# **Dataset Condensation with Distribution Matching**

Bo Zhao, Hakan Bilen School of Informatics, The University of Edinburgh

{bo.zhao, hbilen}@ed.ac.uk

# Abstract

Computational cost of training state-of-the-art deep models in many learning problems is rapidly increasing due to more sophisticated models and larger datasets. A recent promising direction for reducing training cost is dataset condensation that aims to replace the original large training set with a significantly smaller learned synthetic set while preserving the original information. While training deep models on the small set of condensed images can be extremely fast, their synthesis remains computationally expensive due to the complex bi-level optimization and secondorder derivative computation. In this work, we propose a simple yet effective method that synthesizes condensed images by matching feature distributions of the synthetic and original training images in many sampled embedding spaces. Our method significantly reduces the synthesis cost while achieving comparable or better performance. Thanks to its efficiency, we apply our method to more realistic and larger datasets with sophisticated neural architectures and obtain a significant performance boost<sup>1</sup>. We also show promising practical benefits of our method in continual learning and neural architecture search.

# 1. Introduction

Computational cost for training a single state-of-the-art model in various fields, including computer vision and natural language processing, doubles every 3.4 months in the deep learning era due to larger models and datasets. The pace is significantly faster than the Moore's Law that the hardware performance would roughly double every other year [3]. While training a single model can be expensive, designing new deep learning models or applying them to new tasks certainly require substantially more computations, as they involve to train multiple models on the same dataset for many times to verify the design choices, such as loss functions, architectures and hyper-parameters [5, 14]. For instance, Ying *et al.* [48] spend 100 TPU years of computation time conducting an exhaustive neural architecture search on CIFAR10 dataset [23], while training the bestperforming architectures take only dozens of TPU minutes. Hence, there is a strong demand for techniques that can reduce the computational cost for training multiple models on the same dataset with minimal performance drop. To this end, this paper focuses on lowering the training cost by reducing the training set size.

The traditional solution to reduce the training set size is coreset selection. Typically, coreset selection methods choose samples that are important for training based on heuristic criteria, for example, minimizing distance between coreset and whole-dataset centers [12, 35, 10, 4], maximizing the diversity of selected samples [1], discovering cluster centers [15, 37], counting the mis-classification frequency [41] and choosing samples with the largest negative implicit gradient [7]. Although coreset selection methods can be very computationally efficient, they have two major limitations. First most methods incrementally and greedily select samples, which are short-sighted. Second their efficiency is upper bounded by the information in the selected samples in the original dataset.

An effective way of tackling the information bottleneck is synthesizing informative samples rather than selecting from given samples. A recent approach, dataset condensation (or distillation) [45, 50] aims to learn a small synthetic training set so that a model trained on it can obtain comparable testing accuracy to that trained on the original training set. Wang et al. [45] pose the problem in a learning-to-learn framework by formulating the network parameters as a function of synthetic data and learning them through the network parameters to minimize the training loss over the original data. An important shortcoming of this method is the expensive optimization procedure that involves updating network weights for multiple steps for each outer iteration and unrolling its recursive computation graph. Zhao et al. [50] propose to match the gradients w.r.t. the network weights giving real and synthetic training images that successfully avoids the expensive unrolling of the computational graph. Another efficiency improvement is a closed form optimizer by posing the classification task as a

<sup>&</sup>lt;sup>1</sup>The implementation is available at https://github.com/ VICO-UoE/DatasetCondensation.



Figure 1. Dataset Condensation with Distribution Matching. We randomly sample real and synthetic data, and then embed them with the randomly sampled deep neural networks. We learn the synthetic data by minimizing the distribution discrepancy between real and synthetic data in these sampled embedding spaces.

ridge regression problem to simplify the inner-loop model optimization [6, 31]. In spite of the recent progress, the dataset condensation still requires solving the expensive bilevel optimization which jeopardizes their goal of reducing training time due to the expensive image synthesis process. For instance, the state-of-the-art [49] requires 15 hours of GPU time to learn 500 synthetic images on CIFAR10 which equals to the cost of training 6 deep networks on the same dataset. In addition, these methods also require tuning multiple hyper-parameters, *e.g.* the steps to update synthetic set and network parameters respectively in each iteration, that can be different for different settings such as sizes of synthetic sets.

In this paper, we propose a novel training set synthesis technique that combines the advantages of previous coreset and dataset condensation methods while avoiding their limitations. Unlike the former and like the latter, our method is not limited to individual samples from original dataset and can synthesize training images. Like the former and unlike the latter, our method can very efficiently produce a synthetic set and avoid expensive bi-level optimization. In particular, we pose this task as a distribution matching problem such that the synthetic data are optimized to match the original data distribution in a family of embedding spaces by using the maximum mean discrepancy (MMD) [18] measurement (see Figure 1). Distance between data distributions is commonly used as the criterion for coreset selection [12, 15, 46, 37], however, it has not been used to synthesize training data before. We show that the family of embedding spaces can be efficiently obtained by sampling randomly initialized deep neural networks. Hence, our method is significantly faster (e.g.  $45 \times$  in CIFAR10 when synthesizing 500 images) than the state-of-the-art [49] and involves tuning only one hyperparameter (learning rate for synthetic images), while obtaining comparable or better results. In addition, the training of our method can be independently run for each class in parallel and its computation load can be distributed. Finally, our method provides a different training cost/performance tradeoff for large-scale settings. In contrast to prior works [45, 31] that are limited to learning small synthetic sets on small datasets, our method can be successfully applied in more realistic settings, *i.e.* synthesizing 1250 images per class for CIFAR10 [23], and larger datasets, *i.e.* TinyImageNet [25] and ImageNet-1K [13]. We also validate these benefits in two downstream tasks by producing more data-efficient memory for continual learning and generating more representative proxy dataset for accelerating neural architecture search.

#### 2. Methodology

#### 2.1. Dataset Condensation Problem

The goal of dataset condensation is to condense the large-scale training set  $\mathcal{T} = \{(x_1, y_1), \ldots, (x_{|\mathcal{T}|}, y_{|\mathcal{T}|})\}$  with  $|\mathcal{T}|$  image and label pairs into a small synthetic set  $\mathcal{S} = \{(s_1, y_1), \ldots, (s_{|\mathcal{S}|}, y_{|\mathcal{S}|})\}$  with  $|\mathcal{S}|$  synthetic image and label pairs so that models trained on each  $\mathcal{T}$  and  $\mathcal{S}$  obtain comparable performance on unseen testing data:

$$\mathbb{E}_{\boldsymbol{x} \sim P_{\mathcal{D}}}[\ell(\phi_{\boldsymbol{\theta}}\tau(\boldsymbol{x}), y)] \simeq \mathbb{E}_{\boldsymbol{x} \sim P_{\mathcal{D}}}[\ell(\phi_{\boldsymbol{\theta}}s(\boldsymbol{x}), y)], \quad (1)$$

where  $P_{\mathcal{D}}$  is the real data distribution,  $\ell$  is the loss function (*i.e.* cross-entropy loss),  $\phi$  is a deep neural network parameterized by  $\theta$ , and  $\phi_{\theta^{\mathcal{T}}}$  and  $\phi_{\theta^{\mathcal{S}}}$  are the networks that are trained on  $\mathcal{T}$  and  $\mathcal{S}$  respectively.

**Existing solutions.** Previous works [45, 40, 39, 6, 31, 32] formulate the dataset condensation as a learning-to-learn problem, pose the network parameters  $\theta^{S}$  as a function of synthetic data S and obtain a solution for S by minimizing the training loss  $\mathcal{L}^{\mathcal{T}}$  over the original data  $\mathcal{T}$ :

$$\mathcal{S}^* = \operatorname*{arg\,min}_{\mathcal{S}} \mathcal{L}^{\mathcal{T}}(\boldsymbol{\theta}^{\mathcal{S}}(\mathcal{S}))$$
  
subject to  $\boldsymbol{\theta}^{\mathcal{S}}(\mathcal{S}) = \operatorname*{arg\,min}_{\boldsymbol{\theta}} \mathcal{L}^{\mathcal{S}}(\boldsymbol{\theta}).$  (2)

Recently the authors of [50, 49] show that a similar goal can be achieved by matching gradients of the losses over the synthetic and real training data respectively w.r.t. the network parameters  $\theta$ , while optimizing  $\theta$  and the synthetic data S in an alternating way:

$$S^* = \underset{S}{\operatorname{arg\,min}} \operatorname{E}_{\boldsymbol{\theta}_0 \sim P_{\boldsymbol{\theta}_0}} \Big[ \sum_{t=0}^{T-1} D(\nabla_{\boldsymbol{\theta}} \mathcal{L}^{\mathcal{S}}(\boldsymbol{\theta}_t), \nabla_{\boldsymbol{\theta}} \mathcal{L}^{\mathcal{T}}(\boldsymbol{\theta}_t)) \Big]$$
  
subject to  $\boldsymbol{\theta}_{t+1} \leftarrow \operatorname{opt-alg}_{\boldsymbol{\theta}}(\mathcal{L}^{\mathcal{S}}(\boldsymbol{\theta}_t), \varsigma_{\boldsymbol{\theta}}, \eta_{\boldsymbol{\theta}}),$ 
(3)

where  $P_{\theta_0}$  is the distribution of parameter initialization, T is the outer-loop iteration for updating synthetic data,  $\varsigma_{\theta}$  is the inner-loop iteration for updating network parameters,  $\eta_{\theta}$  is the parameter learning rate and  $D(\cdot, \cdot)$  measures the gradient matching error. Note that all the training algorithms [45, 50, 49] have another loop of sampling  $\theta_0$  over the bilevel optimization.

**Dilemma.** The learning problems in eq. (2) and eq. (3) involve solving an expensive bi-level optimization: first optimizing the model  $\theta^{S}$  in eq. (2) or  $\theta_t$  in eq. (3) at the inner loop, then optimizing the synthetic data S along with additional second-order derivative computation at the outer loop. For example, training 50 images/class synthetic set S by using the method in [50] requires 500K epochs of updating network parameters  $\theta_t$  on S, in addition to the 50K updating of S. Furthermore, Zhao *et al.* [50] need to tune the hyper-parameters of the outer and inner loop optimization (*i.e.* how many steps to update S and  $\theta_t$ ) for different learning settings, which requires cross-validating them and hence multiplies the cost for training synthetic images.

# 2.2. Dataset Condensation with Distribution Matching

Our goal is to synthesize data that can accurately approximate the distribution of the real training data in a similar spirit to coreset techniques (e.g. [47, 37]). However, to this end, we do not limit our method to select a subset of the training samples but to synthesize them as in [45, 50]. As the training images are typically very high dimensional, estimating the real data distribution  $P_D$  can be expensive and inaccurate. Instead, we assume that each training image  $x \in \Re^d$  can be embedded into a lower dimensional space by using a family of parametric functions  $\psi_{\vartheta} : \Re^d \to \Re^{d'}$ where  $d' \ll d$  and  $\vartheta$  is the parameter. In other words, each embedding function  $\psi$  can be seen as providing a partial interpretation of its input, while their combination provides a complete one.

Now we can estimate the distance between the real and synthetic data distribution with commonly used maximum mean discrepancy (MMD) [18]:

$$\sup_{\|\psi_{\vartheta}\|_{\mathcal{H}} \le 1} (\mathbb{E}[\psi_{\vartheta}(\mathcal{T})] - \mathbb{E}[\psi_{\vartheta}(\mathcal{S})]),$$
(4)

where  $\mathcal{H}$  is reproducing kernel Hilbert space. As we do not have access to ground-truth data distributions, we use the empirical estimate of the MMD:

$$\mathbb{E}_{\boldsymbol{\vartheta} \sim P_{\boldsymbol{\vartheta}}} \| \frac{1}{|\mathcal{T}|} \sum_{i=1}^{|\mathcal{T}|} \psi_{\boldsymbol{\vartheta}}(\boldsymbol{x}_i) - \frac{1}{|\mathcal{S}|} \sum_{j=1}^{|\mathcal{S}|} \psi_{\boldsymbol{\vartheta}}(\boldsymbol{s}_j) \|^2, \quad (5)$$

where  $P_{\vartheta}$  is the distribution of network parameters.

Following [49], we also apply the differentiable Siamese augmentation  $\mathcal{A}(\cdot, \omega)$  to real and synthetic data that implements the same randomly sampled augmentation to the real and synthetic minibatch in training, where  $\omega \sim \Omega$ is the augmentation parameter such as the rotation degree. Thus, the learned synthetic data can benefit from semanticpreserving transformations (*e.g.* cropping) and learn prior knowledge about spatial configuration of samples while training deep neural networks with data augmentation. Finally, we solve the following optimization problem:

$$\min_{\mathcal{S}} \mathbb{E}_{\substack{\boldsymbol{\vartheta} \sim P_{\boldsymbol{\vartheta}} \\ \boldsymbol{\omega} \sim \Omega}} \| \frac{1}{|\mathcal{T}|} \sum_{i=1}^{|\mathcal{T}|} \psi_{\boldsymbol{\vartheta}}(\mathcal{A}(\boldsymbol{x}_{i}, \boldsymbol{\omega})) - \frac{1}{|\mathcal{S}|} \sum_{j=1}^{|\mathcal{S}|} \psi_{\boldsymbol{\vartheta}}(\mathcal{A}(\boldsymbol{s}_{j}, \boldsymbol{\omega})) \|^{2}.$$
(6)

We learn the synthetic data S by minimizing the discrepancy between two distributions in various embedding spaces by sampling  $\vartheta$ . Importantly eq. (6) can be efficiently solved, as it requires only optimizing S but no model parameters and thus avoids expensive bi-level optimization. This is in contrast to the existing formulations (see eq. (2) and eq. (3)) that involve bi-level optimizations over network parameters  $\theta$  and the synthetic data S.

Note that, as we target at image classification problems, we minimize the discrepancy between the real and synthetic samples of the same class only. We assume that each real training sample is labelled and we also set a label to each synthetic sample and keep it fixed during training.

#### 2.3. Training Algorithm

We depict the mini-batch based training algorithm in Algorithm 1. We train the synthetic data for K iterations. In each iteration, we randomly sample the model  $\psi_{\vartheta}$  with parameter  $\vartheta \sim P_{\vartheta}$ . Then, we sample a pair of real and synthetic data batches  $(B_c^{\mathcal{T}} \sim \mathcal{T} \text{ and } B_c^{\mathcal{S}} \sim \mathcal{S})$  and augmentation parameter  $\omega_c \sim \Omega$  for every class c. The mean discrepancy between the augmented real and synthetic batches of every class is calculated and then summed as loss  $\mathcal{L}$ . The synthetic data  $\mathcal{S}$  is updated by minimizing  $\mathcal{L}$  with stochastic gradient descent and learning rate  $\eta$ .

#### 2.4. Discussion

Randomly Initialized Networks. The family of embedding functions  $\psi_{\vartheta}$  can be designed in different ways. Here we use a deep neural network with different random initializations rather than sampling its parameters from a set of pre-trained networks which is more computationally expensive to obtain. We experimentally validate that our random initialization strategy produces better or comparable results with the more expensive strategy of using pretrained networks in Section 3.4. However, one may still question why randomly initialized networks provide meaningful embeddings for distribution matching. Here we list two reasons based on the observations from previous work. First, randomly initialized networks are reported to produce powerful representations for multiple computer vision tasks [36, 9, 2]. Second, such random networks are showed to perform a distance-preserving embedding of the data, i.e. smaller distances between samples of same class and larger distances across samples of different classes [16]. In addition, the combination of many weak embeddings provides a complete interpretation of the inputs.

Algorithm 1: Dataset condensation with distribution matching

**Input:** Training set T

**1 Required:** Randomly initialized set of synthetic samples S for C classes, deep neural network  $\psi_{\vartheta}$  parameterized with  $\vartheta$ , probability distribution over parameters  $P_{\vartheta}$ , differentiable augmentation  $\mathcal{A}_{\omega}$  parameterized with  $\omega$ , augmentation parameter distribution  $\Omega$ , training iterations K, learning rate  $\eta$ .

**2** for  $k = 0, \cdots, K - 1$  do

- Sample  $\vartheta \sim P_{\vartheta}$ 3
- Sample mini-batch pairs  $B_c^{\mathcal{T}} \sim \mathcal{T}$  and  $B_c^{\mathcal{S}} \sim \mathcal{S}$  and  $\omega_c \sim \Omega$  for every class c4
- Compute  $\mathcal{L} = \sum_{c=0}^{C-1} \|\frac{1}{|B_c^{\tau}|} \sum_{(\boldsymbol{x},y)\in B_c^{\tau}} \psi_{\boldsymbol{\vartheta}}(\mathcal{A}_{\omega_c}(\boldsymbol{x})) \frac{1}{|B_c^{\mathcal{S}}|} \sum_{(\boldsymbol{s},y)\in B_c^{\mathcal{S}}} \psi_{\boldsymbol{\vartheta}}(\mathcal{A}_{\omega_c}(\boldsymbol{s})) \|^2$ Update  $\mathcal{S} \leftarrow \mathcal{S} \eta \nabla_{\mathcal{S}} \mathcal{L}$ 5
- 6
  - Output: S

Connection to Gradient Matching. While we match the mean features of the real and synthetic image batches, Zhao et al. [50] match the mean gradients of network weights over the two batches. We find that, given a batch of data from the same class, the mean gradient vector w.r.t. each output neuron in the last layer of a network is equivalent to a weighted mean of features where the weights are a function of classification probabilities predicted by the network and proportional to the distance between prediction and groundtruth. In other words, while our method weighs each feature equally, Zhao et al. [50] assign larger weights to samples whose predictions are inaccurate. Note that these weights dynamically vary for different networks and training iterations. We provide the derivation in the supplementary.

Generative Models. The classic image synthesizing techniques, includes AutoEncoders [22] and Generative Adversarial Networks (GANs) [17], aim to synthesize reallooking images, while our goal is to generate data-efficient training samples. Regularizing the images to look real may limit the data-efficiency. Previous work [50] showed that the images synthesized by cGAN [29] are not better than the randomly selected real images for training networks. We further provide the comparison to state-of-the-art VAE and GAN models and GMMN method [28] in the supplementary. Although generative models can be trained to produce data-efficient samples with suitable objectives, e.g. [45, 50] and ours, it is not trivial work to build it and achieve stateof-the-art results [39]. We leave it as the future work.

# **3.** Experiments

# **3.1. Experimental Settings**

**Datasets.** We evaluate the classification performance of deep networks that are trained on the synthetic images generated by our method. We conduct experiments on five datasets including MNIST [26], CIFAR10, CIFAR100 [23], TinyImageNet [25] and ImageNet-1K [13]. MNIST consists of 60K 28×28 gray-scale training images of 10 classes. CIFAR10 and CIFAR100 contain 50k  $32 \times 32$  training images from 10 and 100 object categories respectively. Tiny-ImageNet and ImageNet-1K have 100K training images from 200 categories and 1.3M training images from 1K categories respectively. We resize these ImageNet images with  $64 \times 64$  resolution. These two datasets are significantly more challenging than MNIST and CIFAR10/100 due to more diverse classes and higher image resolution.

Experimental Settings. We first learn 1/10/50 image(s) per class synthetic sets for all datasets by using the same ConvNet architecture in [50]. Then, we use the learned synthetic sets to train randomly initialized ConvNets from scratch and evaluate them on real test data. The default ConvNet includes three repeated convolutional blocks, and each block involves a 128-kernel convolution layer, instance normalization layer [42], ReLU activation function [30] and average pooling. Note that four-block ConvNets are used to adjust to the larger input size  $(64 \times 64)$  of TinyImageNet and ImageNet-1K images. In each experiment, we learn one synthetic set and use it to test 20 randomly initialized networks. We repeat each experiment for 5 times and report the mean testing accuracy of the 100 trained networks. We also do cross-architecture experiments in Section 3.3. where we learn the synthetic set on one network architecture and use them to train networks with different architectures.

Hyper-parameters. Like the standard neural network training, dataset condensation also involves tuning a set of hyperparameters. Our method needs to tune only one hyper-parameter, *i.e.* learning rate for the synthetic images, for learning different sizes of synthetic sets, while existing methods [45, 50, 31, 44, 11] have to tune more hyperparameters such as the steps to update synthetic images and network parameters respectively. We use a fixed learning rate 1 for optimizing synthetic images for all 1/10/50 images/class learning on all datasets. When learning larger synthetic sets such as 100/200/500/1,000 images per class, we use larger learning rate (i.e. 10) due to the relatively smaller distribution matching loss. We train synthetic images for 20,000 iterations on MNIST, CIFAR10/100 and 10,000 iterations on TinyImageNet and ImageNet-1K respectively. The mini-batch size for sampling real data is 256. We initialize the synthetic images using randomly sampled real images with corresponding labels. All synthetic images of a class are used to compute the class mean. We use the same augmentation strategy as [49].

## **3.2.** Comparison to the State-of-the-art

**Competitors.** We compare our method to three standard coreset selection methods, namely, Random Selection, Herding [12, 35, 10, 4] and Forgetting [41]. Herding method greedily adds samples into the coreset so that the mean vector is approaching the whole dataset mean. Toneva et al. [41] count how many times a training sample is learned and then forgotten during network training. The samples that are less forgetful can be dropped. We also compare to four state-of-the-art training set synthesis methods, namely, DD [45], LD [6], DC [50] and DSA [49]. Note that we are aware of concurrent works [27, 21, 11] that largely improves the existing bilevel optimization based dataset condensation solutions. Unlike them, we contribute the first solution that has neither bi-level optimization nor second-order derivative, and provide a different training cost/performance tradeoff. Compared to them, our method is significantly simpler and faster. Thus, it is able to scale to large settings *i.e.* learning 1250 images per class for CI-FAR10 and large datasets i.e. ImageNet-1K, while these concurrent works can't. More detailed comparison and discussion to other methods [39, 31, 32], MMD baseline [18] and generative baselines including DC-VAE [33], BigGAN [8] and GMMN [28] can be found in the supplementary.

**Performance Comparison.** Here we evaluate our method on MNIST, CIFAR10 and CIFAR100 datasets and report the results in Table 1. Among the coreset selection methods, Herding performances the best in most settings. Especially, when small synthetic sets are learned, Herding method performs significantly better. For example, Herding achieves 8.4% testing accuracy when learning 1 image/class synthetic set on CIFAR100, while Random and Forgetting obtains only 4.2% and 4.5% testing accuracies respectively.

Training set synthesis methods have clear superiority over coreset selection methods, as the synthetic training data are not limited to a set of real images. Best results are obtained either by DSA or our method. While DSA produces more data-efficient samples with a small number of synthetic samples (1/10 image(s) per class), our method outperforms DSA at 50 images/class setting in CIFAR10 and CIFAR100. The possible reason is that the inner-loop model optimization in DSA with limited number of steps is more effective to fit the network parameters on smaller synthetic data (see eq. (3)). In case of bigger learned synthetic data, the solution obtained in the inner-loop becomes less accurate as it can use only limited number of steps to keep the algorithm scalable. In contrast, our method is robust to



Figure 2. Visualization of generated 10 images per class synthetic sets of MNIST and CIFAR10 datasets.

increasing synthetic data size, can be efficiently optimized significantly faster than DSA.

TinyImageNet and ImageNet-1K. Due to higher image resolution and more diverse classes, prior bilevel optimization based methods do not scale to TinyImageNet and ImageNet-1K. Our method takes 27 hours with one Tesla V100 GPU to condense TinyImageNet into three condensed sets (1/10/50 images/class synthetic sets), and it takes 28 hours with ten GTX 1080 GPUs to condense ImageNet-1K into these three sets. As shown in Table 1, our method achieves 3.9%, 12.9% and 24.1% testing accuracies when learning 1, 10 and 50 images/class synthetic sets for Tiny-ImageNet, and recovers 60% classification performance of the baseline that is trained on the whole original training set with only 10% of data. Our method significantly outperforms the best coreset selection method - Herding, which obtains 2.8%, 6.3% and 16.7% testing accuracies. On ImageNet-1K dataset, our method achieves 1.3%, 5.7% and 11.4% testing accuracies when learning 1, 10 and 50 images/class synthetic sets, which outperforms random selection (0.52%, 1.94% and 7.54%) by large margins.

**Visualization.** The learned synthetic images of MNIST and CIFAR10 are visualized in Figure 2. We find that the synthetic MNIST images are clear and noise free, while the number images synthesized by previous methods contain obvious noise and some unnatural strokes. The synthetic images of CIFAR10 dataset are also visually recognizable and diverse. It is easy to distinguish the background and foreground object.

Figure 3 depicts the feature distribution of the (50 images/class) synthetic sets learned by DC, DSA and our method (DM). We use a network trained on the whole training set to extract features and visualize the features with T-SNE [43]. We find that the synthetic images learned by DC and DSA cannot cover the real image distribution. In contrast, our synthetic images successfully cover the real image distribution. Furthermore, fewer outlier synthetic samples are produced by our method.

**Learning with Batch Normalization.** Zhao *et al.* [50] showed that instance normalization [42] works better than

	Img/Cls Ratio %		Coreset Selection		tion	Training Set Synthesis				Whole Dataset	
	ing/cis	Katio 70	Random	Herding	Forgetting	$DD^{\dagger}$	$LD^{\dagger}$	DC	DSA	DM	whole Dataset
MNIST	1 10 50	0.017 0.17 0.83	$64.9 \pm 3.5$ $95.1 \pm 0.9$ $97.9 \pm 0.2$	$\substack{89.2 \pm 1.6 \\ 93.7 \pm 0.3 \\ 94.8 \pm 0.2}$	$35.5 \pm 5.6 \\ 68.1 \pm 3.3 \\ 88.2 \pm 1.2$	79.5±8.1 -	$60.9 \pm 3.2 \\ 87.3 \pm 0.7 \\ 93.3 \pm 0.3$	91.7±0.5 97.4±0.2 98.8±0.2	88.7±0.6 97.8±0.1 99.2±0.1	$\substack{89.7 \pm 0.6 \\ 97.5 \pm 0.1 \\ 98.6 \pm 0.1 }$	99.6±0.0
CIFAR10	1 10 50	$\begin{array}{c} 0.02\\ 0.2\\ 1\end{array}$	$^{14.4\pm2.0}_{26.0\pm1.2}_{43.4\pm1.0}$	$^{21.5\pm1.2}_{31.6\pm0.7}_{40.4\pm0.6}$	$^{13.5\pm1.2}_{23.3\pm1.0}_{23.3\pm1.1}$	36.8±1.2	$^{25.7\pm0.7}_{38.3\pm0.4}_{42.5\pm0.4}$	<b>28.3±0.5</b> 44.9±0.5 53.9±0.5	28.8±0.7 52.1±0.5 60.6±0.5	$\substack{ 26.0 \pm 0.8 \\ 48.9 \pm 0.6 \\ \textbf{63.0} \pm \textbf{0.4} }$	84.8±0.1
CIFAR100	1 10 50	0.2 2 10	$\begin{smallmatrix} 4.2 \pm 0.3 \\ 14.6 \pm 0.5 \\ 30.0 \pm 0.4 \end{smallmatrix}$	$8.4 \pm 0.3$ 17.3 \pm 0.3 33.7 \pm 0.5	$4.5 \pm 0.2 \\ 15.1 \pm 0.3 \\ 30.5 \pm 0.3$	- - -	11.5±0.4 -	12.8±0.3 25.2±0.3	<b>13.9±0.3</b> <b>32.3±0.3</b> 42.8±0.4	11.4±0.3 29.7±0.3 <b>43.6±0.4</b>	56.2±0.3
TinyImageNet	1 10 50	0.2 2 10	$1.4 \pm 0.1$ 5.0 $\pm 0.2$ 15.0 $\pm 0.4$	$2.8 \pm 0.2 \\ 6.3 \pm 0.2 \\ 16.7 \pm 0.3$	$1.6 \pm 0.1$ $5.1 \pm 0.2$ $15.0 \pm 0.3$	- -	- -	- -	- -	$\begin{array}{c} 3.9{\pm}0.2\\ 12.9{\pm}0.4\\ 24.1{\pm}0.3\end{array}$	37.6±0.4

Table 1. Comparing to coreset selection and training set synthesis methods. We first learn the synthetic data and then evaluate them by training neural networks from scratch and testing on real testing data. The testing accuracies (%) are reported. Img/Cls: image(s) per class. Ratio (%): the ratio of condensed set size to the whole training set size. Note:  $DD^{\dagger}$  and  $LD^{\dagger}$  use different architectures *i.e.* LeNet for MNIST and AlexNet for CIFAR10. The rest methods all use ConvNet.

	Instanc	eNorm	Batch	Norm
	DSA	DM	DSA	DM
CIFAR10 CIFAR100 TinyImageNet	$\begin{vmatrix} 60.6 \pm 0.5 \\ 42.8 \pm 0.4 \end{vmatrix}$	$^{63.0\pm0.4}_{43.6\pm0.4}_{24.1\pm0.3}$	59.9±0.8 44.6±0.5	$^{65.2\pm0.4}_{48.0\pm0.4}_{28.2\pm0.5}$

Table 2. 50 images/class learning with Batch Normalization.

	C\T	ConvNet	AlexNet	VGG	ResNet
DSA	ConvNet	$59.9{\pm}0.8$	$53.3{\pm}0.7$	$51.0{\pm}1.1$	$47.3 {\pm} 1.0$
DM	ConvNet AlexNet VGG ResNet	$65.2 \pm 0.4$ $60.5 \pm 0.4$ $54.2 \pm 0.6$ $52.2 \pm 1.0$	$61.3 \pm 0.6 \\ 59.8 \pm 0.6 \\ 52.6 \pm 1.0 \\ 50.9 \pm 1.4$	$\begin{array}{c} 59.9{\pm}0.8\\ 58.9{\pm}0.4\\ 52.8{\pm}1.1\\ 49.6{\pm}0.9\end{array}$	$57.0\pm0.9$ $54.6\pm0.7$ $49.1\pm1.0$ $52.2\pm0.4$

Table 3. Cross-architecture testing performance (%) on CIFAR10. The 50 img/cls synthetic set is learned on one architecture (C), and then tested on another architecture (T).

batch normalization (BN) [20] when learning small synthetic sets because the synthetic data number is too small to calculate stable running mean and standard deviation (std). When learning with batch normalization, they first pre-set the BN mean and std using many real training data and then freeze them for synthetic data. Thus, the inaccurate mean and std will make optimization difficult [20]. In contrast, we estimate running mean and std by inputting augmented synthetic data from all classes. Hence, our method benefits from the true mean and std of synthetic data. Table 2 show that using ConvNet with BN can further improve our performance. Specifically, our method with BN achieves 65.2%, 48.0% and 28.2% testing accuracies when learning 50 images/class synthetic sets on CIFAR10, CIFAR100 and Tiny-ImageNet respectively, which means 2.2%, 4.4% and 4.1%improvements over our method with the default instance normalization, and also outperforms DSA with BN by 5.3% and 3.4% on CIFAR10 and CIFAR100 respectively.

**Training Cost Comparison.** Our method is significantly more efficient than those bi-level optimization based methods. Without loss of generality, we compare the training time of ours and DSA in the setting of learning 50 im-

ages/class synthetic data on CIFAR10. Figure 4 shows that our method needs less than 20 minutes to reach the performance of DSA trained for 15 hours, which means less than 2.2% training cost. Note that we run the two methods in the same computation environment with one GTX 1080 GPU.

Learning Larger Synthetic Sets We show that our method can also be used to learn larger synthetic sets, while the bi-level optimization based methods typically requires more training time and elaborate hyper-parameter tuning for larger settings. Figure 5 compares our method to random selection baseline in CIFAR10 in terms of absolute and relative performance w.r.t. whole dataset training performance. Clearly our method outperforms random baseline at all operating points which means that our synthetic set is more data-efficient. The advantage of our method is remarkable in challenging settings, *i.e.* settings with small data budgets. Our method obtains  $67.0\pm0.3\%$ ,  $71.2\pm0.4\%$ ,  $76.1 \pm 0.3\%$ ,  $79.8 \pm 0.3\%$  and  $80.8 \pm 0.3\%$  testing accuracies when learning 100, 200, 500, 1000 and 1250 images/class synthetic sets on CIFAR10 dataset respectively, which means we can recover 79%, 84%, 90%, 94% and 95% relative performance using only 2%, 4%, 10%, 20%and 25% training data compared to whole dataset training. We see that the performance gap between the two methods narrows when we learn larger synthetic set. This is somewhat expected, as randomly selecting more samples will approach the whole dataset training which can be considered as the upper-bound. As we initialize synthetic images from random real images, the initial distribution discrepancy becomes tiny when the synthetic set is large.

#### 3.3. Cross-architecture Generalization

[50, 49] verified the cross-architecture generalization ability of synthetic data in an easy setting - learning 1 image/class for MNIST dataset. In this paper, we implement a more challenging cross-architecture experiment - learning 50 images/class for CIFAR10 dataset. In Table 3, the syn-



Figure 3. Distributions of synthetic images learned by DC, DSA and DM. The red, green and blue points are the real images of first three classes in CIFAR10. The stars are corresponding learned synthetic images.



Figure 4. Training time comparison to DSA when learning 50 img/cls synthetic sets on CIFAR10.

thetic data are learned with one architecture (denoted as C) and then be evaluated on another architecture (denoted as T) by training a model from scratch and testing on real testing data. We test several sophisticated neural architectures namely ConvNet, AlexNet [24], VGG-11 [38] and ResNet-18 [19]. Batch Normalization is used in all architectures.

Table 3 shows that learning and evaluating synthetic set on ConvNet achieves the best performance 65.2%. Comparing with DSA, the synthetic data learned by our method with ConvNet have better generalization performance than that learned by DSA with the ConvNet. Specifically, our method outperforms DSA by 8.0%, 8.9% and 9.7% when testing with AlexNet, VGG and ResNet respectively. These results indicate that the synthetic images learned with distribution matching have better generalization performance on unseen architectures than those learned with gradient matching. The learning of synthetic set can be worse with more sophisticated architecture such as ResNet. It is reasonable that the synthetic data fitted on sophisticated architecture will contain some bias that doesn't exist in other architectures, therefore cause worse cross-architecture generalization performance. We also find that the evaluation of the same synthetic set on more sophisticated architectures will be worse. The reason may be that sophisticated architectures are under-fitted using small synthetic set.

#### 3.4. Ablation Study on Network Distribution

Here we study the effect of using different network distributions while learning 1/10/50 image(s)/class synthetic sets on CIFAR10 with ConvNet architecture. Besides sampling randomly initialized network parameters, we also construct a set of networks that are pre-trained on the original

	Random	10-20	20-30	30-40	40-50	50-60	60-70	$\geq 70$	All
1	26.0	26.2	25.9	26.1	26.7	26.8	27.3	26.5	26.4
10	48.9	48.7	48.1	50.7	51.1	49.9	48.6	48.2	50.7
50	63.0	62.7	62.1	62.8	63.0	61.9	60.6	60.0	62.5

Table 4. The performance of synthetic data learned on CIFAR10 with different network distributions. All standard deviations in this table are < 1. These networks are trained on the whole training set and grouped based on the validation accuracy (%).

training set. In particular, we train 1,000 ConvNets with different random initializations on the whole original training set and also store their intermediate states. We roughly divide these networks into nine groups according to their validation accuracies, sample networks from each group, learn the synthetic data on them and use learned synthetic data to train randomly initialized neural networks. Interestingly we see in Table 4 that our method works well with all nine network distributions and the performance variance is small. The visualization and analysis about synthetic images learned with different network distributions are provided in the supplementary.

# 3.5. Continual Learning

We also use our method to store more efficient training samples in the memory for relieving the catastrophic forgetting problem in continual (incremental) learning [35]. We set up the baseline based on GDumb [34] which stores training samples in memory greedily and keeps class-balance. The model is trained from scratch on the latest memory only. Hence, the continual learning performance completely depends on the quality of the memory construction. We compare our memory construction method *i.e.* training set condensation to the *random* selection that is used in [34], herding [12, 35, 10, 4] and DSA [49]. We implement classincremental learning on CIFAR100 dataset with an increasing memory budget of 20 images/class. We implement 5 and 10 step learning, in which we randomly and evenly split the 100 classes into 5 and 10 learning steps i.e. 20 and 10 classes per step respectively. The default ConvNet is used in this experiment.

As depicted in Figure 6 and Figure 7, we find that our method GDumb + DM outperforms others in both two set-





Figure 6. 5-step class-incremental learn-

ing on CIFAR100.

Figure 5. Learning larger synthetic sets on CIFAR10.

	Random	DSA	DM	Early-stopping	Whole Dataset
Performance (%) Correlation Time cost (min) Storage (imgs)	84.0 -0.04 142.6 <b>500</b>	82.6 0.68 142.6 <b>500</b>	82.8 <b>0.76</b> 142.6 <b>500</b>	$\begin{array}{c} \textbf{84.3} \\ 0.11 \\ 142.6 \\ 5 \times 10^4 \end{array}$	$85.9 \\ 1.00 \\ 3580.2 \\ 5 \times 10^4$

Table 5. We implement neural architecture search on CIFAR10 dataset with the search space of 720 ConvNets.

tings, which means that our method can produce the best condensed set as the memory. The final performances of ours, DSA, herding and random are 34.4%, 31.7%, 28.2% and 24.8% in 5-step learning and 34.6%, 30.5%, 27.4% and 24.8% in 10-step learning. We find that ours and random selection performances are not influenced by how the classes are split namely how many new training classes and images occur in each learning step, because both two methods learn/generate the sets independently for each class. However, DSA and herding methods perform worse when the training classes are densely split into more learning steps. The reason is that DSA and herding needs to learn/generate sets based on the model(s) trained on the current training data, which is influenced by the data split. More details can be found in the supplementary.

#### 3.6. Neural Architecture Search

The synthetic sets can also be used as a proxy set to accelerate model evaluation in Neural Architecture Search (NAS) [14]. Following [50], we implement NAS on CI-FAR10 with the search space of 720 ConvNets varying in network depth, width, activation, normalization and pooling layers. Please refer to [50] for more details. We train all architectures on the learned 50 images/class synthetic set, i.e. 1% size of the whole dataset, from scratch and then rank them based on the accuracy on a small validation set. We compare to random, DSA and early-stopping methods. The same size of real images are selected as the proxy set in random. DSA means that we use the synthetic set learned by DSA in the same setting. In *early-stopping*, we use the whole training set to train the model but with the same training iterations like training on the proxy datasets. Therefore, all these methods have the same training time. We train models on the proxy sets for 200 epochs and whole dataset for 100 epochs. The best model is selected based on validation accuracies obtained by different methods. The Spear-



Figure 7. 10-step class-incremental learning on CIFAR100.

man's rank correlation between performances of proxy-set and whole-dataset training is computed for the top 5% architectures selected by the proxy-set.

The NAS results are provided in Table 5. Although the architecture selected by early-stopping achieves the best performance (84.3%), its performance rank correlation (0.11) is remarkably lower than DSA (0.68) and DM (0.76). In addition, early-stopping needs to use the whole training set, while other proxy-set methods need only 500 training samples. The performance rank correlation of Random (-0.04) is too low to provide a reliable ranking for the architectures. Our method (DM) achieves the highest performance rank correlation (0.76), which means that our method can produce reliable ranking for those candidate architectures while using only around  $\frac{1}{25}$  training time of whole dataset training. Although our method needs 72 min to obtain the condensed set, it is negligible compared to whole-dataset training (3580.2 min). More implementation details and analysis can be found in the supplementary.

#### 4. Conclusion

In this paper, we propose an efficient dataset condensation method based on distribution matching. To our knowledge, it is the first solution that has neither bi-level optimization nor second-order derivative. Thus, the synthetic data of different classes can be learned independently and in parallel. Thanks to its efficiency, we can apply our method to more challenging datasets - TinyImageNet and ImageNet-1K, and learn larger synthetic sets - 1250 images/class on CIFAR10. Our method is 45 times faster than the state-of-the-art for learning 50 images/class synthetic set on CIFAR10. We also empirically prove that our method can produce more informative memory for continual learning and better proxy set for speeding up model evaluation in NAS. Though remarkable progress has been seen in this area since the pioneering work [45] released in 2018, dataset condensation is still in its early stage. We will extend dataset condensation to more complex vision tasks in the future.

Acknowledgment. This work is funded by China Scholarship Council 201806010331 and the EPSRC programme grant Visual AI EP/T028572/1.

# References

- Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient based sample selection for online continual learning. In *Advances in Neural Information Processing Systems*, pages 11816–11825, 2019.
- [2] Ehsan Amid, Rohan Anil, Wojciech Kotłowski, and Manfred K Warmuth. Learning from randomly initialized neural network features. arXiv preprint arXiv:2202.06438, 2022.
- [3] Dario Amodei, Danny Hernandez, Girish Sastry, Jack Clark, Greg Brockman, and Ilya Sutskever. Ai and compute. In *OpenAI Blog*, 2018.
- [4] Eden Belouadah and Adrian Popescu. Scail: Classifier weights scaling for class incremental learning. In *The IEEE Winter Conference on Applications of Computer Vision*, 2020.
- [5] James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. *Journal of machine learning research*, 13(Feb):281–305, 2012.
- [6] Ondrej Bohdal, Yongxin Yang, and Timothy Hospedales. Flexible dataset distillation: Learn labels instead of images. *Neural Information Processing Systems Workshop*, 2020.
- [7] Zalán Borsos, Mojmir Mutny, and Andreas Krause. Coresets via bilevel optimization for continual learning and streaming. Advances in Neural Information Processing Systems, 33:14879–14890, 2020.
- [8] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *ICLR*, 2019.
- [9] Weipeng Cao, Xizhao Wang, Zhong Ming, and Jinzhu Gao. A review on neural networks with random weights. *Neuro-computing*, 275:278–287, 2018.
- [10] Francisco M Castro, Manuel J Marín-Jiménez, Nicolás Guil, Cordelia Schmid, and Karteek Alahari. End-to-end incremental learning. In *Proceedings of the European Conference* on Computer Vision (ECCV), pages 233–248, 2018.
- [11] George Cazenavette, Tongzhou Wang, Antonio Torralba, Alexei A. Efros, and Jun-Yan Zhu. Dataset distillation by matching training trajectories. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- [12] Yutian Chen, Max Welling, and Alex Smola. Super-samples from kernel herding. *The Twenty-Sixth Conference Annual Conference on Uncertainty in Artificial Intelligence*, 2010.
- [13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition*, 2009. CVPR 2009. IEEE Conference on, pages 248–255. Ieee, 2009.
- [14] Thomas Elsken, Jan Hendrik Metzen, Frank Hutter, et al. Neural architecture search: A survey. J. Mach. Learn. Res., 20(55):1–21, 2019.
- [15] Reza Zanjirani Farahani and Masoud Hekmatfar. Facility location: concepts, models, algorithms and case studies. Springer Science & Business Media, 2009.
- [16] Raja Giryes, Guillermo Sapiro, and Alex M Bronstein. Deep neural networks with random gaussian weights: A universal

classification strategy? *IEEE Transactions on Signal Processing*, 64(13):3444–3457, 2016.

- [17] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Advances in neural information processing systems, pages 2672–2680, 2014.
- [18] Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723– 773, 2012.
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 770–778, 2016.
- [20] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. ArXiv, abs/1502.03167, 2015.
- [21] Jang-Hyun Kim, Jinuk Kim, Seong Joon Oh, Sangdoo Yun, Hwanjun Song, Joonhyun Jeong, Jung-Woo Ha, and Hyun Oh Song. Dataset condensation via efficient syntheticdata parameterization. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 11102– 11118, 2022.
- [22] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114, 2013.
- [23] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- [24] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In Advances in neural information processing systems, pages 1097–1105, 2012.
- [25] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015.
- [26] Yann LeCun, Léon Bottou, Yoshua Bengio, Patrick Haffner, et al. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [27] Saehyung Lee, Sanghyuk Chun, Sangwon Jung, Sangdoo Yun, and Sungroh Yoon. Dataset condensation with contrastive signals. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 12352–12364, 2022.
- [28] Yujia Li, Kevin Swersky, and Rich Zemel. Generative moment matching networks. In *International conference on machine learning*, pages 1718–1727. PMLR, 2015.
- [29] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. arXiv preprint arXiv:1411.1784, 2014.
- [30] Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the* 27th international conference on machine learning (ICML-10), pages 807–814, 2010.
- [31] Timothy Nguyen, Zhourong Chen, and Jaehoon Lee. Dataset meta-learning from kernel-ridge regression. In *International Conference on Learning Representations*, 2021.
- [32] Timothy Nguyen, Roman Novak, Lechao Xiao, and Jaehoon Lee. Dataset distillation with infinitely wide convolutional networks. arXiv preprint arXiv:2107.13034, 2021.

- [33] Gaurav Parmar, Dacheng Li, Kwonjoon Lee, and Zhuowen Tu. Dual contradistinctive generative autoencoder. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 823–832, 2021.
- [34] Ameya Prabhu, Philip HS Torr, and Puneet K Dokania. Gdumb: A simple approach that questions our progress in continual learning. In *European Conference on Computer Vision*, pages 524–540. Springer, 2020.
- [35] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2001–2010, 2017.
- [36] Andrew M Saxe, Pang Wei Koh, Zhenghao Chen, Maneesh Bhand, Bipin Suresh, and Andrew Y Ng. On random weights and unsupervised feature learning. In *Icml*, 2011.
- [37] Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. *ICLR*, 2018.
- [38] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556, 2014.
- [39] Felipe Petroski Such, Aditya Rawal, Joel Lehman, Kenneth O Stanley, and Jeff Clune. Generative teaching networks: Accelerating neural architecture search by learning to generate synthetic training data. *International Conference* on Machine Learning (ICML), 2020.
- [40] Ilia Sucholutsky and Matthias Schonlau. Soft-label dataset distillation and text dataset distillation. arXiv preprint arXiv:1910.02551, 2019.
- [41] Mariya Toneva, Alessandro Sordoni, Remi Tachet des Combes, Adam Trischler, Yoshua Bengio, and Geoffrey J Gordon. An empirical study of example forgetting during deep neural network learning. *ICLR*, 2019.
- [42] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. arXiv preprint arXiv:1607.08022, 2016.
- [43] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- [44] Kai Wang, Bo Zhao, Xiangyu Peng, Zheng Zhu, Shuo Yang, Shuo Wang, Guan Huang, Hakan Bilen, Xinchao Wang, and Yang You. Cafe: Learning to condense dataset by aligning features. *CVPR*, 2022.
- [45] Tongzhou Wang, Jun-Yan Zhu, Antonio Torralba, and Alexei A Efros. Dataset distillation. arXiv preprint arXiv:1811.10959, 2018.
- [46] Zheng Wang and Jieping Ye. Querying discriminative and representative samples for batch mode active learning. ACM Transactions on Knowledge Discovery from Data (TKDD), 9(3):1–23, 2015.
- [47] Max Welling. Herding dynamical weights to learn. In Proceedings of the 26th Annual International Conference on Machine Learning, pages 1121–1128. ACM, 2009.
- [48] Chris Ying, Aaron Klein, Eric Christiansen, Esteban Real, Kevin Murphy, and Frank Hutter. Nas-bench-101: Towards reproducible neural architecture search. In *International Conference on Machine Learning*, pages 7105–7114. PMLR, 2019.

- [49] Bo Zhao and Hakan Bilen. Dataset condensation with differentiable siamese augmentation. In *International Conference* on Machine Learning, 2021.
- [50] Bo Zhao, Konda Reddy Mopuri, and Hakan Bilen. Dataset condensation with gradient matching. In *International Conference on Learning Representations*, 2021.