
Dataset Condensation with Contrastive Signals

Saehyung Lee^{1†} Sanghyuk Chun² Sangwon Jung^{1†} Sangdoon Yun² Sungroh Yoon^{1,3}

Abstract

Recent studies have demonstrated that gradient matching-based dataset synthesis, or dataset condensation (DC), methods can achieve state-of-the-art performance when applied to data-efficient learning tasks. However, in this study, we prove that the existing DC methods can perform worse than the random selection method when task-irrelevant information forms a significant part of the training dataset. We attribute this to the lack of participation of the contrastive signals between the classes resulting from the class-wise gradient matching strategy. To address this problem, we propose Dataset Condensation with Contrastive signals (DCC) by modifying the loss function to enable the DC methods to effectively capture the differences between classes. In addition, we analyze the new loss function in terms of training dynamics by tracking the kernel velocity. Furthermore, we introduce a bi-level warm-up strategy to stabilize the optimization. Our experimental results indicate that while the existing methods are ineffective for fine-grained image classification tasks, the proposed method can successfully generate informative synthetic datasets for the same tasks. Moreover, we demonstrate that the proposed method outperforms the baselines even on benchmark datasets such as SVHN, CIFAR-10, and CIFAR-100. Finally, we demonstrate the high applicability of the proposed method by applying it to continual learning tasks.

1. Introduction

Deep neural networks (DNNs) are data hungry; larger datasets make DNNs more generalizable (*e.g.*, by data aug-

[†]Work done at NAVER AI Lab ¹Department of Electric and Computer Engineering, Seoul National University ²NAVER AI Lab ³Interdisciplinary Program in Artificial Intelligence, Seoul National University. Correspondence to: Sungroh Yoon <stryoon@snu.ac.kr>.



Task-irrelevant features: {Wheels, Head lights, Roads, Trees, ...}
Task-relevant features: {Logo, Police sign, Trailers, ...}

Figure 1: Example of a fine-grained Truck classification task where the task-irrelevant common features are dominant and the task-relevant discriminant features are in the minority.

mentation (Zhang et al., 2017; Yun et al., 2019; Lee et al., 2021), or by collecting hyperscale training datasets (Jia et al., 2021)). Unsurprisingly, gigantic datasets (*e.g.*, 410 B language tokens (Brown et al., 2020), 3.5B images (Mahajan et al., 2018), and 1.8B image-text pairs (Jia et al., 2021)) have become central to the training of ground-breaking deep models. However, such large datasets require tremendous computational and infrastructural resources, not only for training deep models but also for collecting and processing data columns. Furthermore, real-world knowledge is increasing exponentially, while machine learning (ML) systems are prone to catastrophic forgetting (Goodfellow et al., 2013; Rebuffi et al., 2017). This necessitates repeated training using massive training samples to ensure that ML applications remain competent and practical. Thus, considering the high computational costs, dataset reduction methods are extremely beneficial in applied ML fields. Zhao et al. (2021) proposed a *dataset condensation* method (DC) to synthesize a small but informative dataset by matching the loss gradients with respect to the training and synthetic datasets. In particular, DC was developed to be suitable for downstream classification tasks by repeating the classifier training during the synthetic data optimization procedure, resulting in reasonable performances with reduced synthetic datasets.

In this study, however, we show that DC primarily focuses on the class-wise gradient while overlooking *contrastive signals*. Thus, DC underperforms even when compared with the random selection baseline when contrastive signals are significant to the task. For example, in fine-grained image classification tasks, such as Truck categorization in Fig. 1, contrastive signals should be considered to encode task-

relevant information (e.g., logo, police sign, trailers) while suppressing task-irrelevant information (e.g., wheels, head lights, roads, trees). In our experiments on the fine-grained Automobile dataset, DC results in a classifier with a test accuracy (11%) lower than that achieved using the random selection method (12.2%). We demonstrate that DC cannot effectively utilize the contrastive signals of interclass samples using a motivating example and qualitative analysis.

To address this issue, we propose the *Dataset Condensation with Contrastive signals (DCC)* method. This introduces a modified gradient matching loss function that enables the optimization of a synthetic dataset to capture the contrastive signals. In contrast to DC, which employs only training data of the same class when synthesizing images for a specific class by the class-wise gradient matching, DCC matches the sum of gradients over all classes with respect to the synthetic and training datasets. Additionally, we analyze our method in terms of training dynamics by tracking the kernel velocity (Fort et al., 2020) and introduce a bi-level warm-up strategy to stabilize the optimization procedure of our method. In our experiments, we demonstrate that the proposed DCC singularly outperforms DC in fine-grained classification tasks and general benchmark datasets, such as SVHN (Netzer et al., 2011), CIFAR-10, and CIFAR-100 (Krizhevsky et al., 2009). Finally, we also demonstrate the superiority of our method compared with baselines on downstream tasks, where small synthetic datasets efficiently reduce the total storage of data (e.g., continual learning). The code of our study is available at: <https://github.com/Saehyung-Lee/DCC>

2. Related Work

Formally, we define *the dataset reduction problem* as follows:

$$\mathcal{S}^* = \arg \max_{\mathcal{S}} I(\mathcal{X}; \mathcal{S} | \tau). \quad (1)$$

Here, $\mathcal{X} = \{X_n\}_{n=1}^N$ and $\mathcal{S} = \{S_k\}_{k=1}^K$ are the training and reduced datasets, respectively, and $K \ll N$. τ is a task-dependent variable, and $I(\mathcal{X}; \mathcal{S} | \tau)$ is the conditional mutual information. Herein, we focus on classification tasks, which is the most widely studied scenario with respect to dataset reduction tasks (Wang et al., 2018; Zhao et al., 2021).

Selection-based methods. Selection-based methods (Mirzasoleiman et al., 2020) find a data subset (coreset) that satisfies the cardinality constraint (i.e., $|\mathcal{S}| = K$) while minimizing the difference between the loss gradient on the training dataset and that on the coreset. Moreover, recent studies (Jiang et al., 2021; Paul et al., 2021) have demonstrated that a large fraction of training dataset can be pruned based on the scores they provide. Jiang et al. (2021) introduced a consistency score (C-score) that represents the expected accuracy for a held-out sample on a training

dataset. By sorting the samples according to their C-scores, we can identify prototypical (high-scoring) samples that can serve as a proxy for the training dataset. Paul et al. (2021) proposed the use of the expected loss gradient norm (GradNd) or the norm of the error vector (EL2N) of each training sample to prune a fraction of the training samples. In our case, we preserve the most typical (low-scoring) samples to obtain a proxy for the training dataset.

The selection-based methods, however, are ineffective, particularly when the task-conditional data information $H(\mathcal{X} | \tau)$ is evenly divided and distributed among the training samples. To be precise, if $H(X_n | \mathcal{X} \setminus \{X_n\}, \tau) = \frac{1}{N}H(\mathcal{X} | \tau)$ for all $n \in \{1, \dots, N\}$, then the mutual information $I(\mathcal{X}; \mathcal{S}^* | \tau)$ found by any selection-based method is always a small value $\frac{K}{N}H(\mathcal{X} | \tau)$. As shown by Zhao et al. (2021), the empirical performance gaps between the existing data selection methods and random selection baselines are of no significance in most realistic evaluation benchmarks.

Synthesis-based methods. Instead of selecting a subset from the training dataset, a small dataset \mathcal{S} that achieves similar performance to \mathcal{X} can be generated. Ideally, assuming that the capacity of a synthetic datum can contain as much information as $\frac{1}{K}H(\mathcal{X} | \tau)$, there exists an \mathcal{S} that achieves the same task performance as that attained through the use of \mathcal{X} . Wang et al. (2018) proposed dataset distillation (DD) to transfer the knowledge from a large dataset to a small dataset. They demonstrated that it is possible to achieve close to original accuracy on MNIST (LeCun, 1998) using merely ten synthetic images. Inspired by DD, Zhao et al. (2021) proposed a DC to synthesize a small set of informative samples for learning downstream tasks. The authors showed that the DC outperformed all the baselines in their experiments. Recently, Nguyen et al. (2020) proposed a meta-learning algorithm called Kernel Inducing Points (KIP) for the dataset reduction problem. Furthermore, they presented state-of-the-art performance by using infinitely wide convolutional neural networks (Nguyen et al., 2021).

3. Method

In this section, we introduce the DC method (Zhao et al., 2021) (Sec. 3.1), and study a motivating example showing the limitation of the class-wise gradient matching strategy employed by DC (Sec. 3.2). To mitigate this issue, we propose a modified gradient matching loss method (Sec. 3.3). Furthermore, we propose a bi-level warm-up strategy to stabilize the optimization of the proposed loss function.

3.1. Preliminary: DC with Gradient Matching

When generating images for class “ c ” (\mathcal{S}^c), DC uses only the training data of class “ c ” (\mathcal{X}^c). In particular, DC first (i) updates a synthetic dataset \mathcal{S} by applying a gradient descent

step toward the minimization of the following loss \mathcal{L} :

$$\mathcal{L} = \sum_{c=0}^{C-1} D(\nabla_{\theta_t} L(\mathcal{X}^c; \theta_t), \nabla_{\theta_t} L(\mathcal{S}^c; \theta_t)), \quad (2)$$

where C , $D(\cdot, \cdot)$ and $L(\cdot; \cdot)$ denote the number of classes, a distance function, and the cross-entropy loss function, respectively, and $\nabla_{\theta_t} L(\mathcal{X}^c)$ is the average loss gradient with respect to a model θ_t ; (ii) before moving on to step $t + 1$, trains the model on \mathcal{S} ; (iii) alternately optimizes the synthetic dataset and the model; and (iv) randomly initializes the model after every pre-defined period T (i.e., $\{\theta_{iT} \mid i \in \mathbb{N}_0\}$ is a set of randomly initialized models). Periodic model initialization plays an important role in ensuring that \mathcal{S} can be used for previously unseen models. In addition, Zhao & Bilen (2021) improved DC by using Differentiable Siamese Augmentation (DSA) to generate more informative synthetic datasets. DSA transforms both \mathcal{X}^c and \mathcal{S}^c with the same random transformation (e.g., color jittering, cropping, cutout, flipping, and scale) at each training step. Except for the transformation part, the DSA and DC methods are identical. That is, DSA also uses class-wise gradient matching loss and periodic model initialization.

3.2. A Motivating Example

In this subsection, we show an example, in which the class-wise gradient matching strategy (employed by DC) is problematic. In particular, we show that the class-wise gradient matching strategy is dominated by task-irrelevant *class-common* features, whereas the *class-discriminative* features are relatively neglected. Fig. 2 presents an overview.




Figure 2: Overview of Sec. 3.2. Red and blue circles denote the data distributions of $y = +1$ and -1 , respectively.

Setup. We define a binary classification dataset $\mathcal{X} = \{(x_n, y_n)\}_{n=1}^N$ sampled from the following distribution:

$$y \stackrel{u.a.r}{\sim} \{-1, +1\}, \quad x \stackrel{i.i.d.}{\sim} \mathcal{N}(y\alpha\phi_1 + \beta\phi_2, 1). \quad (3)$$

Here, $\phi_1 \in \mathbb{R}^2$ and $\phi_2 \in \mathbb{R}^2$ represent class-discriminative and class-common feature basis vectors, respectively, where $\phi_1^\top \phi_2 = 0$ and $\|\phi_1\| = \|\phi_2\| = 1$. α and β denote the strength of the class-discriminative and class-common features, respectively, where $\alpha \geq 1$ and $\beta \geq 0$. We generate a reduced dataset $\mathcal{S} = \{\mathcal{S}^+, \mathcal{S}^-\}$ of \mathcal{X} , where \mathcal{S}^+ and \mathcal{S}^- are $(s_1, +1)$ and $(s_2, -1)$, respectively. We use a linear classifier $f(x) = \text{sign}(w^\top x)$ and the hinge loss function $L(x, y; w) = \max(0, 1 - yw^\top x)$, where $w = \phi_1$. For convenience of description, we define $\mathcal{X}^+ = \{(x_i, y_i) \mid i \in \{1, \dots, N\}, -y_i w^\top x_i < 1, y_i = +1\}$ and $\mathcal{X}^- = \{(x_j, y_j) \mid j \in \{1, \dots, N\}, -y_j w^\top x_j < 1, y_j = -1\}$. We define a ℓ_2 -distance-based gradient matching loss as follows (\mathcal{X} : the training dataset, \mathcal{S} : the synthetic dataset):

$$\begin{aligned} \mathcal{L}(\mathcal{X}, \mathcal{S}; w) &= \frac{\lambda}{|\mathcal{S}|} \sum_{s \in \mathcal{S}} \|s\| \\ &+ \left\| \frac{1}{|\mathcal{X}|} \sum_{(x,y) \in \mathcal{X}} g_w(x, y) - \frac{1}{|\mathcal{S}|} \sum_{(s,t) \in \mathcal{S}} g_w(s, t) \right\|, \end{aligned} \quad (4)$$

where $g_w(\cdot) = \nabla_w L(\cdot; w)$. In our example, $\lambda \in \mathbb{R}^+$ is a control parameter of the capacity of the synthetic dataset \mathcal{S} . Here, we assume that λ is selected by making $\max_{s \in \mathcal{S}} \|s\|$ upper bounded by $\epsilon \leq 1 - \frac{\sqrt{2}}{\sqrt{\pi}}$. Finally, we define a class-discriminative and class-common feature ratio $R(\mathcal{S})$ to evaluate the quality of the generated \mathcal{S} as follows:

$$R(\mathcal{S}) = \frac{1}{|\mathcal{S}|} \sum_{s \in \mathcal{S}} \frac{|s^\top \phi_1|}{|s^\top \phi_1| + |s^\top \phi_2|}, \quad (5)$$

where, $R(\mathcal{S}) = 1$ indicates that \mathcal{S} contains only class-discriminative features, whereas $R(\mathcal{S}) = 0$ indicates that \mathcal{S} holds only class-common features.

Issues with class-wise gradient matching. The optimal solution of Eq. (4) for the class-wise gradient matching strategy, employed by previous DC approaches (Zhao et al., 2021; Zhao & Bilen, 2021) is as the follows:

$$\begin{aligned} \tilde{\mathcal{S}} &= \arg \min_{\mathcal{S}} \mathcal{L}(\mathcal{X}^+, \mathcal{S}^+) + \mathcal{L}(\mathcal{X}^-, \mathcal{S}^-) \\ &= \arg \min_{\mathcal{S}} \|\mu^+ - s_1\| + \|\mu^- - s_2\| + \lambda_{\mathcal{S}} \\ &= \left\{ \left(\frac{\epsilon \mu^+}{\|\mu^+\|}, +1 \right), \left(\frac{\epsilon \mu^-}{\|\mu^-\|}, -1 \right) \right\}, \end{aligned} \quad (6)$$

where $\mu^+ = \frac{1}{|\mathcal{X}^+|} \sum_{x \in \mathcal{X}^+} x$, $\mu^- = \frac{1}{|\mathcal{X}^-|} \sum_{x \in \mathcal{X}^-} x$, and $\lambda_{\mathcal{S}} = \lambda \sum_{s \in \mathcal{S}} \|s\|$. More detailed equations can be found in Appendix B. Equation (6) demonstrates that the class-wise gradient matching method optimizes \mathcal{S} for each class to ensure that it has the same direction as the average of the training samples that generate gradients. Then, $R(\tilde{\mathcal{S}})$ is:

$$R(\tilde{\mathcal{S}}) \leq \frac{\alpha}{\alpha + \beta}. \quad (7)$$

The equality holds when $\beta = 0$, and the inequality is due to $\left| \frac{\phi_1^\top \mu^+}{\|\mu^+\|} \right| = \left| \frac{\phi_1^\top \mu^-}{\|\mu^-\|} \right| < \alpha$. Equation (7) shows that when $\alpha \ll \beta$ (i.e., class-common features are dominant and class-discriminative features are minority), $R(\hat{S}) \rightarrow 0$, that is, the class-wise gradient matching method can result in synthetic datasets that are ineffective for the classification task. For example, as shown in Table 1, the class-wise gradient matching method can fail when applied to fine-grained classification tasks that include shared appearance between classes and can be discriminated only by fine-grained appearances.

Leveraging contrastive signals. The class-wise gradient matching method has a limitation when class-common features are dominant. We need a different approach to capture only class-discriminative features for better downstream task performance. The following simple modification of Eq. (6) can mitigate this issue:

$$\begin{aligned} \hat{S} &= \arg \min_{\mathcal{S}} \mathcal{L}(\mathcal{X}^+ \cup \mathcal{X}^-, \mathcal{S}^+ \cup \mathcal{S}^-) \\ &= \arg \min_{\mathcal{S}} \left\| (\mu^+ - \mu^-) - (s_1 - s_2) \right\| + \lambda_{\mathcal{S}} \quad (8) \\ &= \{(\epsilon\phi_1, +1), (-\epsilon\phi_1, -1)\}. \end{aligned}$$

Equation (8) considers the loss gradients for all classes collectively, whereas Eq. (6) considers the loss gradients for each class separately. Moreover, Eq. (8) reveals that the sum of loss gradients between classes is important because it contains contrastive signals between classes ($(\mu^+ - \mu^-)$ and $(s_1 - s_2)$). Here, $R(\hat{S})$ is calculated as follows:

$$R(\hat{S}) = 1. \quad (9)$$

In other words, \hat{S} contains only class-discriminative features to ensure that it is independent of the proportion of class-common features in the original training dataset \mathcal{X} .

Empirical evidence. Here, we empirically demonstrate that the arguments developed above, based on a simple theoretical model, can also be applied to modern machine learning settings. To be specific, we (i) define a binary classification task (3 vs. 8) using MNIST; (ii) train a convolutional neural network (CNN) model on the binary task using the cross-entropy loss; (iii) generate reduced datasets of the task (3 vs. 8) by applying the class-wise gradient matching method (DC) and the class-collective gradient matching method (Eq. (10)), respectively; and (iv) horizontally flip all training images from the class “3” and repeat (ii) to (iii). Digits “3” and “8” can be easily classified by the difference in shape on the left halves (discriminative features), while the right halves look almost identical (common features).

Figure 3 illustrates images synthesized by the DC and our proposed method. The figure shows that the class-wise gradient matching method generates near-prototype images for

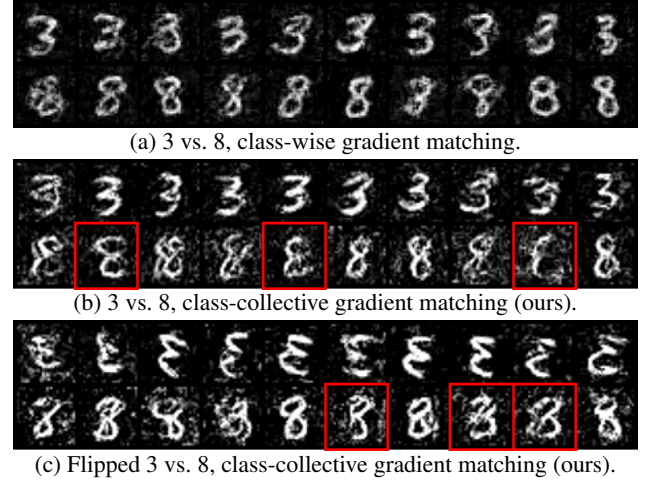


Figure 3: Generated images (10 images per class) for each setting (shown below each subfigure). We mark the images we want to emphasize with red boxes.

each class. In contrast, the class-collective gradient matching method optimizes synthetic images by prioritizing the difference between the two classes. For example, the red boxes in Fig. 3b show that our class-collective gradient matching method synthesizes the images of class “8” with an emphasis on the left half. The same trend can be found in Fig. 3c, indicating that the results are not due to chance or dataset bias, but because the class-collective method leverages contrastive signals. For simple tasks such as MNIST, however, our motivation may not lead to improvements compared to DC, because the number of features in the training dataset is limited to ensure the efficiency of the condensation method. However, for complex tasks that need to capture subtle differences between classes, our approach can result in significant improvements in dataset condensation.

3.3. Dataset Condensation with Contrastive Signals

Based on Sec. 3.2, we propose Dataset Condensation with Contrastive signals (DCC). The DCC optimizes a synthetic dataset by minimizing the following objective function:

$$\begin{aligned} &\mathbb{E}_{\theta_0 \sim P_{\theta_0}} \left[\sum_{t=0}^{T-1} D \left(\sum_{c=0}^{C-1} g_{\theta_t}(\mathcal{X}^c), \sum_{c=0}^{C-1} g_{\theta_t}(\mathcal{S}^c) \right) \right], \quad (10) \\ &\text{subject to } \theta_{t+1} = \theta_t - \frac{\eta}{|\mathcal{S}|} \sum_{(s,t) \in \mathcal{S}} \nabla_{\theta_t} L(s, t; \theta_t). \end{aligned}$$

Here, $g_{\theta_t}(\mathcal{X}^c) = \frac{1}{|\mathcal{X}^c|} \sum_{(x,y) \in \mathcal{X}^c} \nabla_{\theta_t} L(x, y; \theta_t)$, where $\mathcal{X}^c = \{(x, y) \mid (x, y) \in \mathcal{X}, y = c\}$. $D(\cdot, \cdot)$ and $L(\cdot; \cdot)$ denote the distance function and cross-entropy loss function, respectively. We find the solution to Eq. (10) by alternately training the network parameters θ_t and synthetic dataset \mathcal{S} , with the periodic initialization of the classifier as in DC.

We name the loops initializing θ and updating \mathcal{S} “outer-loop” and “inner-loop,” respectively. The primary difference between Eq. (10) and the objective functions of existing methods (Eq. (2)) are the locations of the summation over classes $\sum_{c=0}^{C-1}$. Existing methods first determine the gradient distance for each class and then sum them up, while DCC sums up the gradients over the classes first and then measures the gradient distance between the training and synthetic datasets. Therefore, as implied in Sec. 3.2, DCC can effectively leverage the contrastive signals present in the sum of loss gradients over classes, thereby synthesizing small datasets that are more suitable for classification tasks.

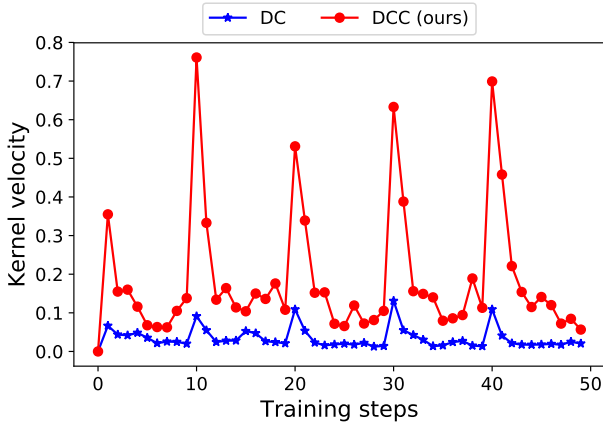


Figure 4: NTK velocity during the synthetic dataset optimization using DC and DCC on CIFAR-10.

A bi-level warm-up strategy. DNNs are known to undergo chaotic transience during the early phase of training (Fort et al., 2020; Liu et al., 2020). In addition, during the dataset condensation process, the classifier is periodically initialized, as described in Sec. 3.1, thereby repeatedly inducing the chaotic training phase of the classifier. We analyze the impact of this periodic transience on the training dynamics of the DCC by measuring the Neural Tangent Kernel (NTK) velocity (Fort et al., 2020) on the synthetic dataset. Fort et al. (2020) introduced the NTK velocity to characterize the loss landscape geometry and training dynamics of DNNs. The NTK velocity is the time evolution of the data-dependent NTK, which, in our case, is the Gram matrix of the Jacobian of the gradient matching loss with respect to the synthetic data samples. The high NTK velocity indicates that the loss landscape is highly nonlinear, and thus, the update direction of the synthetic dataset changes rapidly.

Figure 4 shows the NTK velocity during synthetic dataset optimization using DC and DCC on CIFAR-10. As shown, the NTK velocity periodically repeats the process of peak-

Algorithm 1 Dataset condensation with contrastive signals

Require: Training dataset \mathcal{X} , synthetic dataset \mathcal{S} , outer/inner-loop iterations K_o, K_i , network training iterations T , outer/inner-loop level warm-up iterations γ_o, γ_i , learning rate for synthetic images and network τ, η , number of images per class ζ

- 1: Initialize \mathcal{S} with a subset of \mathcal{X} s.t. $|\mathcal{S}^c| = \zeta, \forall$ class c
- 2: **for** $k_o = 0$ **to** $K_o - 1$ **do** # outer-loop
- 3: Initialize the network parameter θ
- 4: warmup, $k_i \leftarrow \text{True}, 0$
- 5: **while** warmup **do** # inner-loop with warm-up
- 6: **if** $k_o > \gamma_o$ **or** $k_i > \gamma_i$ **then** # bi-level warm-up
- 7: break;
- 8: **end if**
- 9: # class-wise gradient matching loss
- 10: Compute \mathcal{L} by Eq. (2)
- 11: $\mathcal{S} \leftarrow \mathcal{S} - \tau \cdot \nabla_{\mathcal{S}} \mathcal{L}$ # synthetic images update
- 12: Update θ using \mathcal{S} for T iterations
- 13: $k_i \leftarrow k_i + 1$
- 14: **end while**
- 15: $g_{\mathcal{X}}, g_{\mathcal{S}} \leftarrow 0, 0$
- 16: **while** $k_i < K_i$ **do** # inner-loop without warm-up
- 17: **for** $c = 0$ **to** $C - 1$ **do**
- 18: Sample a minibatch pair $\bar{\mathcal{X}}^c \sim \mathcal{X}$ and $\bar{\mathcal{S}}^c \sim \mathcal{S}$
- 19: $g_{\mathcal{X}}, g_{\mathcal{S}} \leftarrow g_{\mathcal{X}} + g_{\theta}(\bar{\mathcal{X}}^c), g_{\mathcal{S}} + g_{\theta}(\bar{\mathcal{S}}^c)$
- 20: **end for**
- 21: # class-collective gradient matching loss
- 22: $\mathcal{L} \leftarrow D(g_{\mathcal{X}}, g_{\mathcal{S}})$
- 23: $\mathcal{S} \leftarrow \mathcal{S} - \tau \cdot \nabla_{\mathcal{S}} \mathcal{L}$ # synthetic images update
- 24: Update θ using \mathcal{S} for T iterations
- 25: $k_i \leftarrow k_i + 1$
- 26: **end while**
- 27: **end for**
- 28: **Output:** a synthetic dataset \mathcal{S}

ing at the classifier initialization and then rapidly stabilizes. Moreover, the peaks of DCC are much higher than those of DC. This difference is reasonable, because when synthesizing images with the class label “ c ,” DC obtains the loss gradient using only the training data of the class “ c ,” while DCC obtains the loss gradient from all the classes. Thus, noisy gradients from the other classes can be excluded in DC, whereas DCC may accumulate noisy gradients from all the classes. Although the higher peaks are not detrimental in terms of optimization (Jastrzebski et al., 2020; Fort et al., 2020), we empirically determine that the peaks during the early phase of dataset condensation can suppress the effectiveness of DCC (see Table 4).

To address this issue, we introduce a *bi-level warm-up strategy* for the DCC. We define the inner-loop level (updating \mathcal{S}) and outer-loop level (initializing θ) warm-up and apply class-wise gradient matching under the two warm-up con-

ditions. The overall procedure for the proposed method is described in Algorithm 1.

4. Experimental Results and Discussion

4.1. Experimental Setup

Datasets. We complement our analysis with experiments conducted on SVHN, CIFAR-10, CIFAR-100, and the fine-grained image classification datasets (Automobile, Terrier, Fish, Truck, Insect, and Lizard) subsampled from ImageNet32x32 (Chrabaszcz et al., 2017) using the WordNet hierarchy (Miller, 1998). A detailed description of the datasets is summarized in Appendix D.

Implementation Details. In our experiments, we compare the proposed method with the baseline methods for the settings of learning 1, 10, and 50 image(s) per class as in Zhao et al. (2021). We use ConvNet (Gidaris & Komodakis, 2018) as a classifier from which the gradients for matching are obtained in the dataset condensation process. We set $K_o = 1000$, $\gamma_o = 250$, $\gamma_i = 10$, and $\tau = 0.1$. For settings of learning 1, 10, and 50 image(s) per class, (K_i, T) is set to (10,5), (10,50), and (50,10), respectively. To reduce the training data based on the selection-based methods, we use the pre-computed scores provided by Jiang et al. (2021) (C-scores) and the average of the scores computed for 10 independently pre-trained models (GraNd and EL2N). To evaluate the selection-based methods, we train 100 classifiers on the coreset from scratch and obtain the mean and standard deviation of their test accuracies. In addition, to evaluate the synthesis-based methods, including our proposed method, we learn five synthetic datasets and train 20 classifiers from scratch on each synthetic dataset to obtain the mean and standard deviation of 100 test accuracies. Please refer to (Zhao et al., 2021; Zhao & Bilen, 2021) for more details. For the implementation of KIP (Nguyen et al., 2020), we use the code¹ provided by the authors. Note that we denote the DCC with differentiable Siamese augmentation (Zhao & Bilen, 2021) as DSAC.

4.2. Dataset Condensation

Results on Fine-Grained Datasets. We first evaluate the improvements over the baselines of the proposed method on fine-grained image classification datasets. Tab. 1 shows that the results of DC are consistent with those of the motivation example described in Sec. 3.2. In particular, DC perform worse than random selection (Random) on the Automobile, Terrier, and Fish datasets. In contrast, DCC always performs better than Random, implying that the proposed

Table 1: Comparison of the proposed method with the baselines (Random, DC and DSA) on fine-grained image classification datasets. Each number is the average over 100 different runs. The blue and red numbers denote worse than Random and the best results, respectively.

Dataset	Img/cls	Random	Baselines		Ours	
			DC	DSA	DCC	DSAC
Automobile	10	12.2	11.0	19.1	18.6	22.1
	50	19.5	16.8	24.1	28.3	29.2
Terrier	10	5.6	4.6	5.1	6.4	6.2
	50	7.8	4.8	7.2	10.7	10.7
Fish	10	14.7	13.5	18.7	20.4	22.3
	50	15.3	17.0	19.4	28.4	23.3
Lizard	10	13.3	23.5	29.1	30.0	34.2
	50	20.9	32.6	32.9	38.8	34.8
Truck	10	21.2	24.8	36.5	39.4	48.1
	50	31.8	43.5	57.9	57.4	60.6
Insect	10	27.6	41.8	47.4	48.7	50.0
	50	42.7	49.6	51.3	55.8	51.9



(a) DC

(b) DCC

Figure 5: Visualization of the generated 10 images per class of Automobile. From the top row, “ambulance”, “beach wagon”, “cab”, “convertible”, “jeep”, “limousine”, “Model T”, “racer”, and “sports car”.

method effectively considers the class-discriminative features. Moreover, DSAC always outperforms DSA, showing that improved methods using diverse image transformations do not effectively detect differences between classes. In addition, although the proposed method is largely orthogonal to differentiable Siamese augmentation, it can be observed that DCC is more effective than DSAC in certain cases. Considering image transformations as a form of regularization (Hernández-García & König, 2018), we hypothesize that such cases indicate that applying additional regularization may hinder the optimization process of the proposed method. Finally, we qualitatively compare the synthetic images for the classes of Automobile generated through each method of DCC and DC. Figure 5 shows that the images learned by the proposed method are sharper than those learned by DC and display the unique patterns of each class more prominently.

¹<https://colab.research.google.com/github/google-research/google-research/blob/master/kip/KIP.ipynb>

Table 2: Comparison of the performance (mean \pm std %) of the proposed method with the selection-based (Random, C-score, GraNd, and EL2N) and synthesis-based (KIP, DC, and DSA) methods on benchmark datasets. Img/cls stands for the number of images per class, and Ω denotes the upper bound of the performance, which can be obtained by learning the original training dataset. The best results within each setting (Dataset, Img/cls) are indicated in bold.

Dataset	Img/cls	Selection-based				Synthesis-based			Ours		Ω
		Random	C-score	GraNd	EL2N	KIP	DC	DSA	DCC	DSAC	
SVHN	1	14.6 \pm 1.6	-	19.6 \pm 0.5	19.1 \pm 0.6	23.3 \pm 2.7	34.6 \pm 2.0	36.0 \pm 2.0	34.3 \pm 1.6	47.5\pm2.6	92.1 \pm 0.2
	10	35.1 \pm 4.1	-	37.5 \pm 1.6	32.5 \pm 1.2	62.4 \pm 0.5	76.2 \pm 0.6	78.9 \pm 0.5	76.2 \pm 0.8	80.5\pm0.6	
	50	70.9 \pm 0.9	-	69.1 \pm 0.7	68.7 \pm 0.7	69.6 \pm 0.5	82.7 \pm 0.3	84.4 \pm 0.4	83.3 \pm 0.2	87.2\pm0.3	
CIFAR-10	1	14.4 \pm 2.0	21.7 \pm 0.6	21.8 \pm 0.5	20.9 \pm 0.6	37.6\pm1.0	28.2 \pm 0.7	28.7 \pm 0.7	32.9 \pm 0.8	34.0 \pm 0.7	81.6 \pm 0.3
	10	26.0 \pm 1.2	31.6 \pm 0.4	32.3 \pm 0.4	32.3 \pm 0.4	47.3 \pm 0.3	44.7 \pm 0.6	52.1 \pm 0.6	49.4 \pm 0.5	54.5\pm0.5	
	50	43.4 \pm 1.0	39.8 \pm 0.4	41.2 \pm 0.3	40.7 \pm 0.3	50.1 \pm 0.2	54.8 \pm 0.5	60.6 \pm 0.4	61.6 \pm 0.4	64.2\pm0.4	
CIFAR-100	1	4.2 \pm 0.3	8.0 \pm 0.3	8.8 \pm 0.3	8.8 \pm 0.3	14.8\pm1.2	12.8 \pm 0.3	13.9 \pm 0.4	13.3 \pm 0.3	14.6 \pm 0.3	52.5 \pm 0.3
	10	14.6 \pm 0.5	18.1 \pm 0.2	17.8 \pm 0.2	17.3 \pm 0.2	13.4 \pm 0.2	26.6 \pm 0.3	32.4 \pm 0.3	30.6 \pm 0.4	33.5\pm0.3	
	50	29.7 \pm 0.4	30.4 \pm 0.3	27.6 \pm 0.2	27.7 \pm 0.2	-	32.1 \pm 0.3	38.6 \pm 0.3	40.0\pm0.3	39.3 \pm 0.4	

For example, the long body and multiple windows of the “limousine” or the distinctive body frame of “Model T” are clearly exhibited in the results of DCC (red box in Fig. 5), while the differences between the classes are ambiguous and difficult to distinguish in the results of DC.




Figure 6: Alignment and uniformity loss for the features of synthetic images generated by DC and DCC on CIFAR-10 (50 images per class). We used 6 pre-trained networks and 10 synthetic datasets (5 DC + 5 DCC), which result in total 60 points in this plot.

Results on Benchmark Datasets. Table 2 presents the comparison results of selection-based methods, synthesis-based methods, and our methods on SVHN, CIFAR-10 and CIFAR-100. First, in contrast to the observation by Zhao et al. (2021), recent selection-based methods achieve better results than Random for the settings of 1 and 10 image(s) per class. However, their performance is still worse than that of the synthesis-based methods with large gaps. Although KIP performs well in some settings, it can be observed that

its effectiveness is unstable depending on the dataset or Img/cls. Unlike fine-grained classification tasks, DC and DSA always show improvements compared to Random and the selection-based methods with large gaps (e.g., +15% in SVHN). Nevertheless, we observe that our method achieves the best performance not only for fine-grained tasks, but also for general vision classification benchmarks.

To understand the significant improvements achieved by our method, we provide an additional analysis of our method and DC from the perspective of representation learning. Wang & Isola (2020) recently demonstrated that two metrics, *i.e.*, uniformity (how features are uniformly distributed on the feature space) and alignment (how two features with the same class are close) are highly correlated to the quality of the learned representations. Following Wang & Isola (2020), we plot the uniformity and alignment losses for the features of DC and DCC in Fig. 6. We extract the features using ResNet-18 (He et al., 2016) and VGG-11 (Simonyan & Zisserman, 2014), which are pre-trained on CIFAR-10. In the figure, DCC shows lower uniformity loss (*i.e.*, samples are more uniformly distributed) and lower alignment loss (*i.e.*, positive samples are closer) than DC. That is, as in our observations, DCC can capture task-relevant features that help classification tasks by using contrastive signals.

Cross-Architecture Generalization. We test the generalizability of the synthetic dataset learned through the the proposed method. In particular, we train various CNN architectures, including ConvNet, LeNet (LeCun et al., 1998), AlexNet (Krizhevsky et al., 2012), VGG-11, and ResNet-18, on small datasets created using Random, DC, DSA, DCC, and DSAC and list their mean test accuracy in Tab. 3. As shown, our proposed method achieves improved results not only for the architecture used for condensation, but also for other CNN architectures tested.

Table 3: Comparison of the cross-architecture generalization performance (the mean test accuracy over 100 runs) of the proposed method with the baseline methods using ConvNet as the source network on CIFAR-10 (50 images per class). The best results for each target network are shown in bold.

Method	Target network				
	ConvNet	LeNet	AlexNet	VGG	ResNet
Random	43.2	30.8	35.9	36.8	26.1
DC	54.8	33.8	40.9	39.3	23.9
DSA	60.4	40.3	46.0	50.7	49.7
DCC (ours)	61.6	38.0	45.2	46.3	27.3
DSAC (ours)	64.1	42.6	48.2	56.0	53.9

Table 4: Improvement in the effectiveness (the mean of test accuracies over 100 runs) of the proposed method on the CIFAR datasets following the application of the bi-level warm-up strategy.

Dataset	Method	Bi-level warm-up	Img/cls		
			1	10	50
CIFAR-10	DCC	✗	28.3	49.2	61.3
		✓	32.9	49.4	61.6
	DSAC	✗	32.3	54.0	63.9
		✓	34.0	54.5	64.2
CIFAR-100	DCC	✗	12.0	28.5	40.5
		✓	13.3	30.6	40.0
	DSAC	✗	12.9	29.3	37.8
		✓	14.6	33.5	39.3

Ablation Study. We demonstrate the importance of the bi-level warm-up strategy when applying the proposed method. We compare the performance of the proposed method with and without the bi-level warm-up strategy and present the results in Tab. 4. In the table, ✗ is equivalent to Algorithm 1 where $\gamma_o = 0$ and $\gamma_i = 0$, while ✓ indicates $\gamma_o = 250$ and $\gamma_i = 10$. From the results, it can be seen that when the capacity of the synthetic dataset is relatively large, the bi-level warm-up has a negligible effect, whereas when the budget is limited, the bi-level warm-up affects significant performance improvements. In particular, without the bi-level warm-up strategy, DCC yields worse results than the baseline methods (DC and DSA) for settings of learning 1 and 10 images per class of CIFAR-100, highlighting the importance of the bi-level warm-up strategy for small budgets. In addition, Tab. 5 shows the ablation study on designing warm-up: None (no warm-up), Simple (warm-up dependent only on inner-loops), Proposed (bi-level). As shown, the NTK velocity peaks during the later phase of dataset condensation do not suppress the effectiveness of DCC.

Table 5: The results of ablation study on designing warm-up

Dataset	Method	Img/cls	None	Simple	Proposed
CIFAR-100	DCC	1	11.98	13.01	13.27
		10	28.46	29.74	30.59
	DSAC	1	12.91	14.21	14.57
		10	29.27	31.97	33.47

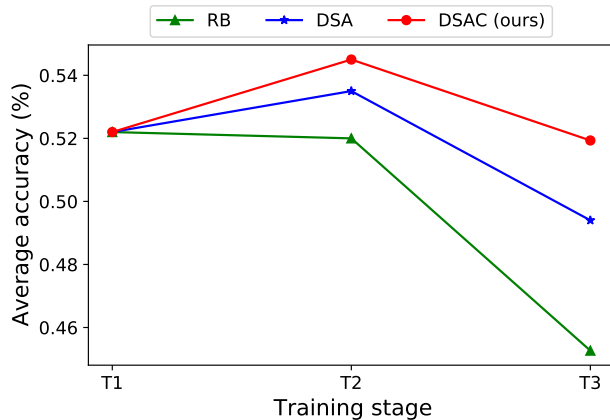


Figure 7: Performance improvements (average accuracy %) on the continual learning task for a sequence of three fine-grained image datasets {Lizard-Truck-Insect} following the application of the proposed method.

4.3. Application: Continual Learning

We apply our method to a continual learning task, where the training datasets are sequentially input with task labels. We build our method on a popular memory-based continual learning baseline, called Experience Replay with Ring Buffer strategy (ER-RB) (Chaudhry et al., 2019). This baseline randomly stores the same amount of data per class of old tasks and replays them to avoid forgetting old tasks while learning a new task. To observe the effectiveness of the applications of the dataset condensation methods on continual learning tasks, we substitute DSA and DSAC for the ring buffer strategy. We train models on a sequence of three fine-grained image datasets (*i.e.*, {Lizard-Truck-Insect}) with the ER-RB, DSA, and DSAC methods (10 images per class) and compare them in terms of the average accuracy of seen tasks in Figure 7. As shown, DSAC outperforms RB and DSA by 6.7% and 2.5% in T_3 , respectively. These results indicate that the dataset generated by DSAC is more informative than those created by other baselines and hence more helpful in preventing memory loss of past tasks. We provide the results of the continual learning task on general vision classification benchmarks in Appendix C.

5. Conclusion and Future Directions

In this study, we demonstrate that the existing dataset condensation methods perform poorly on fine-grained tasks owing to their bias toward reconstructing the prototype of each class. Based on the example providing the motivation for the study, we propose the DCC method, which can effectively capture subtle differences between classes through the application of class-collective gradient matching. In addition, inspired by the training dynamics of the DCC, we introduce a bi-level warm-up strategy that can stabilize the optimization of the proposed loss function. Our experiments demonstrate that the proposed method significantly outperforms the baselines not only for fine-grained tasks, but also for general vision classification benchmarks. However, the proposed method can be further improved by optimizing the strategy used to avoid learning instability. To be precise, the DCC involves contrastive signals from all classes included in a training dataset. Therefore, for a large number of classes, the bi-level warm-up strategy might not be sufficient to control the enormous contrastive signals. Methods such as class subgrouping could remedy this, which will be our focus in future studies. In addition, the proposed method provides potential for further development of dataset condensation by enabling a combination with mixed-class data augmentation methods such as Mixup or CutMix.

Acknowledgements: Most experiments were conducted on NAVER Smart Machine Learning (NSML) platform (Sung et al., 2017; Kim et al., 2018). This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. 2022R1A3B1077720), Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) [NO.2021-0-01343, Artificial Intelligence Graduate School Program (Seoul National University)], and the BK21 FOUR program of the Education and Research Program for Future ICT Pioneers, Seoul National University in 2022.

References

- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- Chaudhry, A., Rohrbach, M., Elhoseiny, M., Ajanthan, T., Dokania, P. K., Torr, P. H., and Ranzato, M. On tiny episodic memories in continual learning. *arXiv preprint arXiv:1902.10486*, 2019.
- Chrabaszcz, P., Loshchilov, I., and Hutter, F. A downsampled variant of imagenet as an alternative to the cifar datasets. *arXiv preprint arXiv:1707.08819*, 2017.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.
- Fort, S., Dziugaite, G. K., Paul, M., Kharaghani, S., Roy, D. M., and Ganguli, S. Deep learning versus kernel learning: an empirical study of loss landscape geometry and the time evolution of the neural tangent kernel. *Advances in Neural Information Processing Systems*, 33, 2020.
- Gidaris, S. and Komodakis, N. Dynamic few-shot visual learning without forgetting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4367–4375, 2018.
- Goodfellow, I. J., Mirza, M., Xiao, D., Courville, A., and Bengio, Y. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*, 2013.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Hernández-García, A. and König, P. Data augmentation instead of explicit regularization. *arXiv preprint arXiv:1806.03852*, 2018.
- Jastrzebski, S., Szymczak, M., Fort, S., Arpit, D., Tabor, J., Cho, K., and Geras, K. The break-even point on optimization trajectories of deep neural networks. *arXiv preprint arXiv:2002.09572*, 2020.
- Jia, C., Yang, Y., Xia, Y., Chen, Y.-T., Parekh, Z., Pham, H., Le, Q. V., Sung, Y., Li, Z., and Duerig, T. Scaling up visual and vision-language representation learning with noisy text supervision. *arXiv preprint arXiv:2102.05918*, 2021.
- Jiang, Z., Zhang, C., Talwar, K., and Mozer, M. C. Characterizing structural regularities of labeled data in overparameterized models. In Meila, M. and Zhang, T. (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 5034–5044. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/jiang21k.html>.
- Kim, H., Kim, M., Seo, D., Kim, J., Park, H., Park, S., Jo, H., Kim, K., Yang, Y., Kim, Y., et al. Nsm1: Meet the mlaas platform with a real-world case study. *arXiv preprint arXiv:1810.09957*, 2018.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.

- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25: 1097–1105, 2012.
- LeCun, Y. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Lee, S., Park, C., Lee, H., Yi, J., Lee, J., and Yoon, S. Removing undesirable feature contributions using out-of-distribution data. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=eIHYL6fpbkA>.
- Liu, J., Jiang, G., Bai, Y., Chen, T., and Wang, H. Understanding why neural networks generalize well through gsnr of parameters. *arXiv preprint arXiv:2001.07384*, 2020.
- Mahajan, D., Girshick, R., Ramanathan, V., He, K., Paluri, M., Li, Y., Bharambe, A., and Van Der Maaten, L. Exploring the limits of weakly supervised pretraining. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 181–196, 2018.
- Miller, G. A. *WordNet: An electronic lexical database*. MIT press, 1998.
- Mirzasoleiman, B., Bilmes, J., and Leskovec, J. Coresets for data-efficient training of machine learning models. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 6950–6960. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/mirzasoleiman20a.html>.
- Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., and Ng, A. Y. Reading digits in natural images with unsupervised feature learning. 2011.
- Nguyen, T., Chen, Z., and Lee, J. Dataset meta-learning from kernel ridge-regression. In *International Conference on Learning Representations*, 2020.
- Nguyen, T., Novak, R., Xiao, L., and Lee, J. Dataset distillation with infinitely wide convolutional networks. In Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=hXWPpJedrVP>.
- Paul, M., Ganguli, S., and Dziugaite, G. K. Deep learning on a data diet: Finding important examples early in training. *Advances in Neural Information Processing Systems*, 34, 2021.
- Rebuffi, S.-A., Kolesnikov, A., Sperl, G., and Lampert, C. H. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp. 2001–2010, 2017.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Stallkamp, J., Schlipsing, M., Salmen, J., and Igel, C. The german traffic sign recognition benchmark: a multi-class classification competition. In *The 2011 international joint conference on neural networks*, pp. 1453–1460. IEEE, 2011.
- Sung, N., Kim, M., Jo, H., Yang, Y., Kim, J., Lausen, L., Kim, Y., Lee, G., Kwak, D., Ha, J.-W., et al. Nsm1: A machine learning platform that enables you to focus on your models. *arXiv preprint arXiv:1712.05902*, 2017.
- Wang, T. and Isola, P. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, pp. 9929–9939. PMLR, 2020.
- Wang, T., Zhu, J.-Y., Torralba, A., and Efros, A. A. Dataset distillation. *arXiv preprint arXiv:1811.10959*, 2018.
- Yun, S., Han, D., Oh, S. J., Chun, S., Choe, J., and Yoo, Y. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 6023–6032, 2019.
- Zhang, H., Cisse, M., Dauphin, Y. N., and Lopez-Paz, D. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.
- Zhao, B. and Bilen, H. Dataset condensation with differentiable siamese augmentation. In *International Conference on Machine Learning*, 2021.
- Zhao, B., Mopuri, K. R., and Bilen, H. Dataset condensation with gradient matching. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=mSAKhLYLSsl>.

A. A Motivating Example

An issue with class-wise gradient matching. The class-wise gradient matching strategy is employed by previous DC approaches (Zhao et al., 2021; Zhao & Bilen, 2021). The optimal solution $\tilde{\mathcal{S}}$ of the class-wise gradient matching strategy for Eq. (4) can be found as follows:

$$\begin{aligned}
 \tilde{\mathcal{S}} &= \arg \min_{\mathcal{S}} \mathcal{L}(\mathcal{X}^+, \mathcal{S}^+) + \mathcal{L}(\mathcal{X}^-, \mathcal{S}^-) = \arg \min_{\mathcal{S}} \left\| \frac{1}{|\mathcal{X}^+|} \sum_{(x,y) \in \mathcal{X}^+} g_w(x,y) - \frac{1}{|\mathcal{S}^+|} \sum_{(s,t) \in \mathcal{S}^+} g_w(s,t) \right\| \\
 &+ \left\| \frac{1}{|\mathcal{X}^-|} \sum_{(x,y) \in \mathcal{X}^-} g_w(x,y) - \frac{1}{|\mathcal{S}^-|} \sum_{(s,t) \in \mathcal{S}^-} g_w(s,t) \right\| + \frac{\lambda}{|\mathcal{S}^+|} \sum_{s \in \mathcal{S}^+} \|s\| + \frac{\lambda}{|\mathcal{S}^-|} \sum_{s \in \mathcal{S}^-} \|s\| \\
 &= \arg \min_{\mathcal{S}} \left\| \sum_{(x,y) \in \mathcal{X}^+} \frac{-x}{|\mathcal{X}^+|} + \sum_{(s,t) \in \mathcal{S}^+} \frac{s}{|\mathcal{S}^+|} \right\| + \left\| \sum_{(x,y) \in \mathcal{X}^-} \frac{x}{|\mathcal{X}^-|} - \sum_{(s,t) \in \mathcal{S}^-} \frac{s}{|\mathcal{S}^-|} \right\| + \lambda_{\mathcal{S}} \\
 &= \arg \min_{\mathcal{S}} \|\mu^+ - s_1\| + \|\mu^- - s_2\| + \lambda_{\mathcal{S}} = \left\{ \left(\frac{\epsilon \mu^+}{\|\mu^+\|}, +1 \right), \left(\frac{\epsilon \mu^-}{\|\mu^-\|}, -1 \right) \right\},
 \end{aligned} \tag{6}$$

where $\mu^+ = \frac{1}{|\mathcal{X}^+|} \sum_{x \in \mathcal{X}^+} x$, $\mu^- = \frac{1}{|\mathcal{X}^-|} \sum_{x \in \mathcal{X}^-} x$, and $\lambda_{\mathcal{S}} = \lambda \sum_{s \in \mathcal{S}} \|s\|$. Eq (6) demonstrates that the class-wise gradient matching method optimizes \mathcal{S} , for each class, to have the same direction as the average of training samples that generate gradients. Then, $R(\tilde{\mathcal{S}})$ is as follows:

$$R(\tilde{\mathcal{S}}) = \frac{1}{2} \left(\frac{\left| \frac{\phi_1^\top \mu^+}{\|\mu^+\|} \right|}{\left| \frac{\phi_1^\top \mu^+}{\|\mu^+\|} \right| + \left| \frac{\phi_2^\top \mu^+}{\|\mu^+\|} \right|} + \frac{\left| \frac{\phi_1^\top \mu^-}{\|\mu^-\|} \right|}{\left| \frac{\phi_1^\top \mu^-}{\|\mu^-\|} \right| + \left| \frac{\phi_2^\top \mu^-}{\|\mu^-\|} \right|} \right) = \frac{\left| \frac{\phi_1^\top \mu^+}{\|\mu^+\|} \right|}{\left| \frac{\phi_1^\top \mu^+}{\|\mu^+\|} \right| + \beta} \leq \frac{\alpha}{\alpha + \beta}. \tag{7}$$

Here, without loss of generality, we assume the cardinality of \mathcal{X} is sufficiently large such that $\left| \frac{\phi_1^\top \mu^+}{\|\mu^+\|} \right| = \left| \frac{\phi_1^\top \mu^-}{\|\mu^-\|} \right|$. The equality holds when $\beta = 0$, and the inequality is due to $\left| \frac{\phi_1^\top \mu^+}{\|\mu^+\|} \right| < \alpha$. Eq. (7) shows that when $\alpha \ll \beta$ (*i.e.*, when class-common features are dominant features and class-discriminative features are minority), $R(\tilde{\mathcal{S}}) \rightarrow 0$, *i.e.*, the class-wise gradient matching method can result in synthetic datasets that are ineffectual for the classification task. For example, as shown in Table 1, the class-wise gradient matching method can fail on fine-grained classification tasks that include shared appearance between classes, and can be discriminative by only fine-grained appearances.

Leveraging contrastive signals. The class-wise gradient matching method has a limitation when the class-common features are dominant. We need a different approach to capture only class-discriminative features for better target task performances. The following simple modification of Eq. (6) can mitigate the issue:

$$\begin{aligned}
 \hat{\mathcal{S}} &= \arg \min_{\mathcal{S}} \mathcal{L}(\mathcal{X}^+ \cup \mathcal{X}^-, \mathcal{S}^+ \cup \mathcal{S}^-) \\
 &= \arg \min_{\mathcal{S}} \left\| \frac{1}{|\mathcal{X}^+| + |\mathcal{X}^-|} \sum_{(x,y) \in \mathcal{X}^+ \cup \mathcal{X}^-} g_w(x,y) - \frac{1}{|\mathcal{S}^+| + |\mathcal{S}^-|} \sum_{(s,t) \in \mathcal{S}^+ \cup \mathcal{S}^-} g_w(s,t) \right\| + \lambda_{\mathcal{S}} \\
 &= \arg \min_{\mathcal{S}} \left\| \left(\sum_{(x,y) \in \mathcal{X}^+} \frac{-x}{2\bar{N}} + \sum_{(x,y) \in \mathcal{X}^-} \frac{x}{2\bar{N}} \right) - \left(\sum_{(s,t) \in \mathcal{S}^+} \frac{-s}{2} + \sum_{(s,t) \in \mathcal{S}^-} \frac{s}{2} \right) \right\| + \lambda_{\mathcal{S}} \\
 &= \arg \min_{\mathcal{S}} \|(\mu^+ - \mu^-) - (s_1 - s_2)\| + \lambda_{\mathcal{S}} = \{(\epsilon \phi_1, +1), (-\epsilon \phi_1, -1)\}.
 \end{aligned} \tag{8}$$

In Eq. (3), we can see that \mathcal{X} is balanced ($y \stackrel{u.a.r.}{\sim} \{-1, +1\}$). Hence, we can set $|\mathcal{X}^+| = |\mathcal{X}^-| = \hat{N}$, without loss of generality. Eq. (8) considers loss gradients for all classes collectively, while Eq. (6) considers loss gradients for each class separately. Moreover, Eq. (8) reveals that the sum of loss gradients between classes is important because it contains contrastive signals between classes ($(\mu^+ - \mu^-)$ and $(s_1 - s_2)$). Here, $R(\hat{\mathcal{S}})$ is as follows:

$$R(\hat{\mathcal{S}}) = \frac{1}{2} \left(\frac{2\epsilon |\phi_1^\top \phi_1|}{\epsilon |\phi_1^\top \phi_1| + \epsilon |\phi_1^\top \phi_2|} \right) = 1, \tag{9}$$

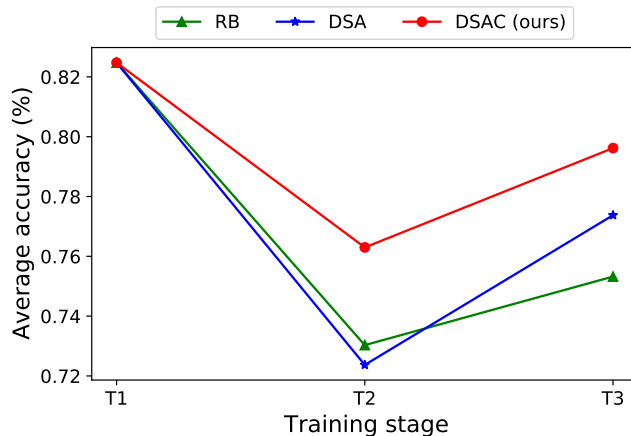


Figure 8: Performance improvements (average accuracy %) on the continual learning task for a sequence of benchmark datasets {CIFAR10-SVHN-TrafficSigns} following the application of the proposed method.

in other words, $\hat{\mathcal{S}}$ contains only class-discriminative features, so that it is independent of the proportion of class-common features in the original training dataset \mathcal{X} .

B. Application: Continual Learning

Figure 8 shows the effectiveness of DSAC, DSA and RB on the continual learning task which is composed of CIFAR10, SVHN and TrafficSigns (Stallkamp et al., 2011). DSAC and DSA utilize the condensed datasets as rehearsal examples as in Figure 7. We note that our DSAC again dominates other baselines for T2 and T3 tasks.

C. Datasets

SVHN (Netzer et al., 2011) consists of 73,257 training images and 26,032 test images in 10 classes. CIFAR-10 (Krizhevsky et al., 2009) consists of 50,000 training images and 10,000 test images in 10 classes. CIFAR-100 (Krizhevsky et al., 2009) consists of 50,000 training images and 10,000 test images in 100 classes. SVHN, CIFAR-10, and CIFAR-100 images have sizes of 32×32 pixels. ImageNet (Deng et al., 2009) consists of 1,281,167 training images and 100,000 test images in 1,000 classes. Chrabaszcz et al. (2017) provided downsampled variants of the ImageNet dataset. The ImageNet32x32 dataset (Chrabaszcz et al., 2017) have the same number of classes and images as ImageNet, but the images are downsampled to sizes of 32×32 pixel. We constructed the fine-grained image classification datasets by subsampling from the ImageNet32x32 dataset using the WordNet (Miller, 1998) hierarchy. The subsampled ImageNet classes are summarized in Tab. 6.

Table 6: The subsampled ImageNet classes for each fine-grained image classification dataset.

Fine-grained dataset	ImageNet classes
Automobile	beach wagon, convertible, sports car, ambulance, jeep, limousine, racer, cab, Model T
Terrier	Lakeland terrier, Scotch terrier, cairn, Airedale, Tibetan terrier, Yorkshire terrier, Norfolk terrier, Staffordshire bullterrier, Sealyham terrier, standard schnauzer, Norwich terrier, Bedlington terrier, Lhasa, Irish terrier, silky terrier, Dandie Dinmont, Boston bull, Border terrier, soft-coated wheaten terrier, Australian terrier, American Staffordshire terrier, West Highland white terrier, giant schnauzer, miniature schnauzer, Kerry blue terrier, wire-haired fox terrier
Fish	tench, stingray, tiger shark, barracouta, coho, gar, electric ray, great white shark, sturgeon puffer, anemone fish, goldfish, eel, rock beauty, lionfish, hammerhead
Lizard	agama, banded gecko, Komodo dragon, frilled lizard, African chameleon, American chameleon, green lizard, whiptail, common iguana, alligator lizard, Gila monster
Truck	pickup, police van, trailer truck, minivan, moving van, tow truck, fire engine, garbage truck, tractor
Insect	cricket, ant, leafhopper, walking stick, grasshopper, dung beetle, tiger beetle, lacewing, rhinoceros beetle, ringlet, long-horned beetle, ladybug, ground beetle, cicada, cabbage butterfly, leaf beetle, lycaenid, bee, monarch, damselfly, admiral, sulphur butterfly, dragonfly, fly, weevil, cockroach, mantis