

Mimicking the Oracle: An Initial Phase Decorrelation Approach for Class Incremental Learning

Yujun Shi¹ * Kuangqi Zhou¹ Jian Liang³ Zihang Jiang¹
 Jiashi Feng² Philip Torr⁴ Song Bai² Vincent Y. F. Tan¹

¹National University of Singapore ²ByteDance Inc.

³Institute of Automation, Chinese Academy of Sciences (CAS) ⁴University of Oxford

shi.yujun@u.nus.edu vtan@nus.edu.sg

Abstract

Class Incremental Learning (CIL) aims at learning a classifier in a phase-by-phase manner, in which only data of a subset of the classes are provided at each phase. Previous works mainly focus on mitigating forgetting in phases after the initial one. However, we find that improving CIL at its initial phase is also a promising direction. Specifically, we experimentally show that directly encouraging CIL Learner at the initial phase to output similar representations as the model jointly trained on all classes can greatly boost the CIL performance. Motivated by this, we study the difference between a naïvely-trained initial-phase model and the oracle model. Specifically, since one major difference between these two models is the number of training classes, we investigate how such difference affects the model representations. We find that, with fewer training classes, the data representations of each class lie in a long and narrow region; with more training classes, the representations of each class scatter more uniformly. Inspired by this observation, we propose Class-wise Decorrelation (CwD) that effectively regularizes representations of each class to scatter more uniformly, thus mimicking the model jointly trained with all classes (i.e., the oracle model). Our CwD is simple to implement and easy to plug into existing methods. Extensive experiments on various benchmark datasets show that CwD consistently and significantly improves the performance of existing state-of-the-art methods by around 1% to 3%. Code: <https://github.com/Yujun-Shi/CwD>.

1. Introduction

The ability to continually acquire new knowledge is a key to achieve artificial intelligence. To enable this ability for classification models, [8, 12, 19, 20, 26, 29, 33] introduce

*Work done when interning with Song Bai at ByteDance.

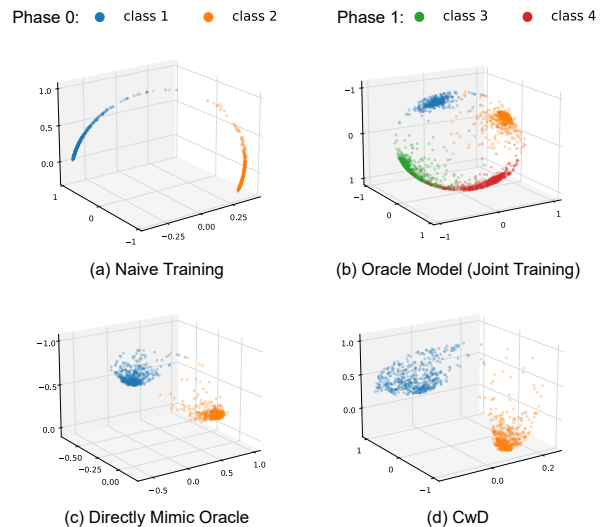


Figure 1. Visualization of representations (normalized to the unit sphere) in a two-phase CIL setting (learning 2 classes for each phase). (a) Naïve training at the initial phase (a.k.a., phase 0). The data representations of each class lie in a long and narrow region. (b) Joint training on all 4 classes (oracle model). The data representations of each class scatter more uniformly. (c) Directly mimicking the oracle model at the initial phase, i.e., training the CIL learner with a regularization term that enforces the learner to output representation that is similar to the oracle model. This makes the representations of each class scatter more uniformly (like (b)). (d) Training at the initial phase with our CwD regularizer, which also yields uniformly-scattered representations (like (b) and (c)). Best viewed in color.

and study *Class Incremental Learning (CIL)*. In CIL, training is conducted phase-by-phase, and only data of a subset of classes are provided at each phase. The goal of CIL is to perform well on classes learned at the current phase as well as all previous phases.

The major challenge of CIL is that the model performance on previously learned classes usually degrades seriously after learning new classes, a.k.a. catastrophic

forgetting [9, 23]. To reduce forgetting, most previous works [8, 12, 18, 26, 29] focus on phases after the initial one, *e.g.* introducing forgetting-reduction regularization terms that enforce the current-phase model and the previous-phase model to produce similar outputs of the same input.

However, the role of the initial phase in CIL (the phase before the CIL learner begins incrementally learning new classes) is largely neglected and much less understood. We argue that the initial phase is of critical importance, since the model trained at this phase implicitly affects model learning in subsequent CIL phases (*e.g.*, through the forgetting-reduction regularization term). In this work, we thus study whether and how we can boost CIL performance by improving the representations of the initial phase.

To start with and to motivate our method, we conduct an exploratory experiment to investigate the potential of improving CIL at its initial phase. Specifically, at the initial phase, we regularize the CIL learner to produce similar representations as the model trained with data of all classes (*i.e.*, the oracle model), since the upper bound of CIL is the oracle model. According to our results, this additional regularization drastically improves CIL performance. In addition, as we experimentally show that, although this term is used in the initial phase, it yields little performance gain in the initial phase. In contrast, it significantly benefits CIL performance in subsequent phases. This demonstrates that the performance improvements are not simply due to a higher accuracy at the initial phase, but because this regularization makes the initial-phase representations more favorable for incrementally learning new classes.

Inspired by this, we consider improving CIL from a novel perspective—encouraging the CIL learner to mimic the oracle model in the initial phase. To achieve this, we first need to understand the difference between representations produced by a naïvely-trained initial-phase model and the oracle model. Specifically, since the oracle model is trained with more classes, we investigate how representations are affected by the number of training classes. To this end, we compute and analyze the eigenvalues of the covariance matrix of representations of each class. Interestingly, we find that when training with fewer classes, the top eigenvalues of the covariance matrix of representations of each class dominate, indicating that the representations of each class lie in a long and narrow region (see Fig. 1 (a) for example). On the other hand, for models trained with more classes (particularly, the oracle model), the top eigenvalues become less dominant, indicating that the representations of each class scatter more uniformly (see Fig. 1 (b)).

We are thus motivated to enforce data representations of each class to be more uniformly scattered at the initial phase, which mimics the representations produced by the oracle model. To this end, we first theoretically show that, a group of embeddings will scatter more uniformly in the

space if its correlation matrix has smaller Frobenius norm. We then propose to minimize the Frobenius norm of the correlation matrix of the data representations for each class. We refer to our regularization term as *Class-wise Decorrelation (CwD)*. We provide a visualization to summarize our motivation and methodology in Fig. 1. Our proposed CwD regularizer can serve as a generic plug-in to other CIL methods and can be easily implemented.

Extensive experiments on various benchmark datasets show that our CwD regularizer works well with state-of-the-art CIL methods, yielding significant and consistent performance gain in different settings. In addition, we also perform detailed ablation studies on how the effectiveness of CwD is influenced by factors such as the number of classes at the initial CIL phase, the number of exemplars for each class and regularization coefficient of the CwD term.

The contributions of this paper are as follows: 1) We empirically discover that encouraging the CIL learner to mimic the oracle model in the initial phase can boost the CIL performance. 2) We find that compared with naïvely-trained initial-phase model, data representations of each class produced by the oracle model scatter more uniformly, and that mimicking such representations at the initial phase can benefit CIL. 3) Based on our findings, we propose a novel *Class-wise Decorrelation (CwD)* regularization technique to enforce representations of each class to scatter more uniformly at the initial CIL phase. 4) Extensive experiments show that our proposed CwD regularization yields consistent improvements over previous state-of-the-art methods.

2. Related Works

Two classic setups of Incremental Learning are Class Incremental Learning (CIL) [1, 12, 13, 19, 20, 22, 26] and Task Incremental Learning (TIL) [2, 16, 21, 24, 28, 30, 31]. CIL and TIL both split all training classes into multiple tasks and learn them sequentially. The difference between these two setups is that TIL allows using task information during inference (*i.e.*, knowing what task does test data belong to) but the CIL does not. In this work, we focus on the setting of CIL. The major challenge of CIL is that model performance on previously learned classes degrades drastically after learning new classes, *i.e.* catastrophic forgetting [9, 23].

Many CIL methods mitigate forgetting through knowledge distillation [8, 12, 18, 26, 29]. In these methods, when learning at a new phase, the model of the previous phase is used as the teacher, and the CIL Learner is regularized to produce similar outputs as the teacher. In this way, the knowledge of previously learned classes can be preserved.

However, distillation-based methods introduce the dilemma of balancing between previously learned classes and current classes. In particular, if the distillation term is too large, then the model’s ability to learn new classes will be limited. In contrast, if the distillation term is too small,

forgetting will be amplified. To mitigate this dilemma, some methods have been proposed to maintain a good balance between old and new classes [4, 19, 32, 33].

The common focus of existing methods is on improving CIL at phases after the initial one. Differently, we study CIL from a less-explored perspective — improving CIL at its initial-phase representations. Previously, [25] also studied the initial-phase representation for CIL. However, their work focuses on the relation between over-fitting at the initial phase and CIL performance. Their main observation is that leveraging some known generalization-improving techniques (e.g., heavy-aug [6], self-distillation [10]) can improve CIL. By contrast, in our work, based on the novel observation that mimicking the oracle model at the initial phase is beneficial, we reveal how number of training classes affects model representations, and further propose our Class-wise Decorrelation (CwD). Unlike techniques such as self-distillation [10] and heavy-aug [6] that brings significant higher accuracy at initial phase, our CwD benefits CIL mainly by making the initial-phase representation more favorable for incrementally learning new classes.

Feature decorrelation has also been explored in some other research fields. For example, [3, 14, 35] rely on feature decorrelation to solve the mode collapse problem in self-supervised learning when negative pairs are missing, and [5, 15, 34] use decorrelation to improve neural networks’ features and hence boost generalization. Differently, our work focuses on CIL, and propose to use class-wise feature decorrelation to mimic the oracle model.

[27] uses eigenvalue analysis on representations of Deep Metric Learning (DML) models. They find that preventing representations to be overly compressed can improve DML generalization. They achieve this by randomly switching negative samples with positive samples in the ranking loss. Differently, our work focuses on CIL, and propose to use class-wise feature decorrelation, which is a more effective way to counter compression of representations.

3. Methodology

In this section, we investigate boosting CIL performance by improving the initial-phase representations. This strategy is different from most previous works.

Firstly, in Sec. 3.1, we investigate the potential of improving CIL through mimicking the oracle model representations at the initial phase.

Motivated by the observation made, in Sec. 3.2, we conduct an eigenvalue analysis on covariance matrix of representations of each class to study how the number of classes used for training affects representations.

We then further develop a novel regularization term, namely the *Class-wise Decorrelation (CwD)* in Sec. 3.3. We show mathematically and experimentally that this regularization term is effective in enforcing data representations

of each class to scatter more uniformly.

3.1. Directly Mimicking the Oracle Model Representation at Initial Phase Improves CIL

In this section, we conduct an exploratory experiment to see whether encouraging the CIL learner to directly mimic the oracle model at the initial phase can improve performance.

Specifically, at the initial CIL phase, we add an additional regularization term to encourage the model to output similar representations as a oracle model, yielding the following objective function:

$$\min_{\theta} L_{ce}(x, y, \theta) + \beta \left(1 - \frac{f_{\theta}(x)^{\top} f_{\theta^*}(x)}{\|f_{\theta}(x)\|_2 \|f_{\theta^*}(x)\|_2} \right), \quad (1)$$

where θ denotes the model parameters, and θ^* denotes parameters of the oracle model (which are fixed). $L_{ce}(x, y, \theta)$ is the standard cross entropy loss, (x, y) is the input training data-label pair, and β is a hyper-parameter controlling the strength of the regularization. $f_{\theta}(x)$ and $f_{\theta^*}(x)$ denote the representations produced by the CIL learner and the oracle model, respectively. The second term in this objective is the regularization that enforces $f_{\theta}(x)$ to be similar to $f_{\theta^*}(x)$.

We experiment on the following two protocols with ImageNet100 and ResNet18 [11]: (1) the CIL learner is initially trained on 50 classes and then incremented with 10 classes per phase for 5 more phases; (2) the CIL learner is initially trained on 10 classes and then incremented 10 classes per phase for 9 more phases. Under these two protocols, we use Eqn. (1) as the optimization objective for the initial phase with the strong baseline of LUCIR [12]. For the following phases, no regularization is added and the original LUCIR [12] is applied. As can be observed in Fig. 2, the regularization term at the initial phase can greatly improves CIL performance. Notably, in the second protocol, although only 10 classes are used at the initial phase, this regularization still brings significant improvements. In addition, in both protocols, although this regularization is only applied at the initial phase, it negligibly improves the accuracy of initial phase, but significantly improves the performance in subsequent phases. This would demonstrate that the improvements are not due simply to an accuracy boost at the initial phase, but because the initial-phase model is more favorable for incrementally learning new classes.

Since the oracle model is not directly available in practice, in the following sections, we explore the characteristics of the oracle model representations and try improving CIL by mimicking this characteristic.

3.2. Class-wise Representations of Oracle Model Scatter More Uniformly

Motivated by the significant improvements yielded by mimicking the oracle model at the initial phase, we inves-




Figure 2. **The effectiveness of directly mimicking the representations of the oracle model at the initial phase.** (a) Initially trained on 50 classes, and then incremented with 10 classes per phase for 5 more phases. (b) Initially trained on 10 classes and then incremented with 10 classes per phase for 9 more phases. The regularization coefficient β is defined in Eqn. (1). We show the accuracy of each CIL phases. Results are averaged over 3 runs.

tigate the difference between naïvely-trained initial-phase model and the oracle.

Specifically, since the oracle model is trained with more classes than naïvely-trained initial-phase model, we conduct an eigenvalue analysis to understand how the number of classes used for training affects representations. Using ImageNet100, we generate four subsets containing 10/25/50/100 classes, where the subset with more classes contains the subset with fewer classes (the 10 classes of the first subset are shared by all 4 subsets). We train four ResNet18 models on each of the subset, and analyze the difference on the representations.

The details of our eigenvalue analysis are elaborated as follows. For a given class c , suppose we have n data points, we denote $Z_i^{(c)} \in \mathbb{R}^d$ as the model output representation on the i -th data point of class c , and the mean vector of all representations of class c is denoted as $\bar{Z}^{(c)} = \frac{1}{n} \sum_{i=1}^n Z_i^{(c)}$. The covariance matrix of class c 's representations is estimated in an unbiased manner as

$$K^{(c)} = \frac{1}{n-1} \sum_{i=1}^n (Z_i^{(c)} - \bar{Z}^{(c)})(Z_i^{(c)} - \bar{Z}^{(c)})^\top. \quad (2)$$

Based on the estimated covariance matrix of class c 's representations, we perform an eigendecomposition $K^{(c)} = U\Sigma^{(c)}U^\top$, where $\Sigma^{(c)}$ is a diagonal matrix with eigenvalues $(\lambda_1^{(c)}, \lambda_2^{(c)}, \dots, \lambda_d^{(c)})$ on the diagonal. Without loss of generality, we assume that the eigenvalues are sorted in descending order. To observe whether the top eigenvalues dominate, we define

$$\alpha_k^{(c)} := \frac{\sum_{i=1}^k \lambda_i^{(c)}}{\sum_{i=1}^d \lambda_i^{(c)}} \in [0, 1], \quad (3)$$




Figure 3. **Visualization on how $\alpha_k^{(c)}$ changes with increasing k for models trained with different number of classes.** $\alpha_k^{(c)}$, which measures the proportion of variance represented by the top k eigenvalues, is defined in Eqn. (3). We plot curve of $\alpha_k^{(c)}$ for ResNet18 models trained with 10/25/50/100 ImageNet classes.

which measures the proportion of variance represented by the top k eigenvalues. If $\alpha_k^{(c)}$ is close to 1 even when k is small, then the top eigenvalues of $K^{(c)}$ dominate.

For one of the 10 shared classes among the four models, we visualize how $\alpha_k^{(c)}$ changes with increasing k . Results on representations of other classes show similar trend, and are in the Appendix. As can be observed in Fig. 3, for the model trained with only 10 classes, $\alpha_k^{(c)}$ increases quickly for $k \in \{1, 2, \dots, 10\}$, and then saturates at a value close to 1 as k keep increasing. This shows that for the 10 class model, the top eigenvalues dominate for covariance matrix of data representations of each class, indicating that data representations lie in a long and narrow region. In addition, for any fixed k , $\alpha_k^{(c)}$ strictly decreases as the model is being trained with more classes. This shows that, as the model is trained with more classes, the top k eigenvalues become less dominant, suggesting that the data representations of each class scatter more uniformly.

Since the oracle model is trained with more classes than the naïvely-trained initial-phase model, class-wise data representations of the oracle model scatter more uniformly.

3.3. Class-wise Decorrelation

The observation in Sec. 3.2 suggests that one way to encourage the CIL learner to mimic the oracle model at the initial phase is to enforce data representations of each class to scatter more uniformly.

This can be achieved by adding the following regularization objective for each class c in the initial phase:

$$L_{\text{shape}}^{(c)} = \frac{1}{d} \sum_{i=1}^d \left(\lambda_i^{(c)} - \frac{1}{d} \sum_{j=1}^d \lambda_j^{(c)} \right)^2, \quad (4)$$

where d is dimension of the representation space. Minimizing this objective will enforce all eigenvalues to be close,

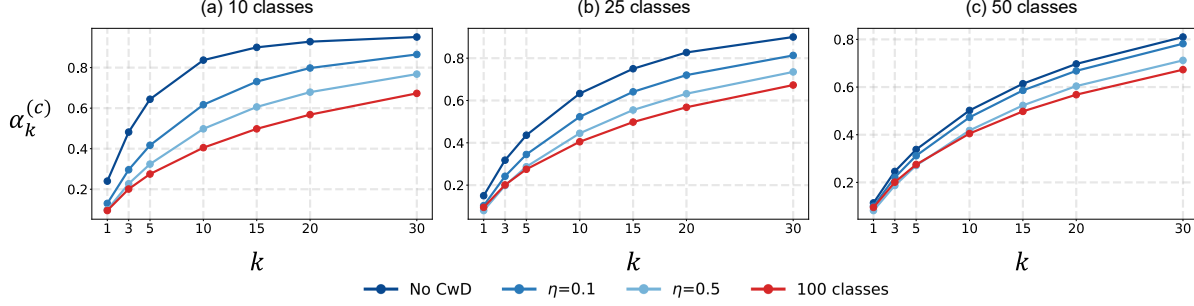


Figure 4. **Effects of class-wise decorrelation on representations of each class.** $\alpha_k^{(c)}$, which measures the proportion of variance represented by the top k eigenvalues, is defined in Eqn. (3). η is the CwD coefficient used in Eqn. (9). We plot curve of $\alpha_k^{(c)}$ with or without CwD objective when training with 10/25/50 classes. We also plot the curve for the model trained on all 100 classes for comparison.

preventing the top eigenvalues to dominate and encouraging the data representations of class c to scatter more uniformly. However, this regularization objective is not practical because computing eigenvalues is expensive.

In order to achieve our desired regularization in an implementation-friendly way, we first normalize all representations of class c by

$$Z_i^{(c)} := \frac{Z_i^{(c)} - \bar{Z}^{(c)}}{\sigma^{(c)}(Z)}, \quad (5)$$

where $\sigma^{(c)}(Z) \in \mathbb{R}^d$ is the vector of standard deviations all the representations, and the division is done element-wise. This normalization results in the covariance matrix $K^{(c)}$ (defined in Eqn. (2)) being equivalent to a *correlation matrix*, which satisfies

$$\sum_{j=1}^d \lambda_j^{(c)} = \text{Tr}(K^{(c)}) = d, \quad (6)$$

where $\text{Tr}(\cdot)$ is the matrix trace operator and d is the dimension of $K^{(c)}$.

Then, by the following proposition, we can relate the Frobenius norm of a correlation matrix and its eigenvalues.

Proposition 1. *For a d -by- d correlation matrix K and its eigenvalues $(\lambda_1, \lambda_2, \dots, \lambda_d)$, we have:*

$$\sum_{i=1}^d \left(\lambda_i - \frac{1}{d} \sum_{j=1}^d \lambda_j \right)^2 = \|K\|_{\text{F}}^2 - d. \quad (7)$$

The proof of this proposition is given in the Appendix. It shows that for any correlation matrix K , minimizing L_{shape} defined in Eqn. (4) is equivalent to minimizing $\|K\|_{\text{F}}^2$.

With this proposition, we convert the impractical regularization in Eqn. (4) into the *Class-wise Decorrelation (CwD)* objective below, which penalizes $\|K^{(c)}\|_{\text{F}}^2$ for every class c :

$$L_{\text{CwD}}(\theta) = \frac{1}{C \cdot d^2} \sum_{c=1}^C \|K^{(c)}\|_{\text{F}}^2, \quad (8)$$

Algorithm 1 PyTorch-style pseudocode for CwD.

```

# N: batch size
# d: representation dimension
# z: a batch of representation, with shape (N, d)
# y: a batch of label corresponding to z
def class_wise_decorrelation_loss(z, y):
    loss_cwd = 0.0 # initialize cwd loss
    unique_y = y.unique() # all classes in the batch
    for c in unique_y:
        # obtain all representation of class c
        z_c = z[y==c, :]
        N_c = z_c.size(0)

        # skip if class c only have 1 sample
        if N_c == 1:
            continue

        # normalize representation as in eq.(5)
        z_c = (z_c - z_c.mean(0)) / z_c.std(0)
        # estimate correlation matrix
        corr_mat = 1/(N_c-1)*torch.matmul(z_c.t(), z_c)
        # calculate CwD loss for class c
        loss_cwd += (corr_mat.pow(2)).mean()
    return loss_cwd

```

where C is the number of classes used when training at the initial phase, $K^{(c)}$ is the correlation matrix of class c estimated over the training data batch. Note that $K^{(c)}$ is a function of the parameters θ through its eigenvalues λ_i .

Therefore, the overall optimization objective at the initial phase is:

$$\min_{\theta} L_{\text{ce}}(x, y, \theta) + \eta \cdot L_{\text{CwD}}(\theta), \quad (9)$$

where η is the hyper-parameter controlling strength of our CwD objective. A Pytorch-style pseudocode for our proposed CwD regularization is given in Algorithm 1.

To empirically verify that our proposed L_{CwD} is indeed effective in encouraging representations of each class to scatter more uniformly, we conduct the same eigenvalue analysis as in Sec. 3.2. We apply L_{CwD} to conduct experiments with the same 10/25/50-class settings in Sec. 3.2. As can be seen in Fig. 4, applying L_{CwD} can effectively decrease $\alpha_k^{(c)}$ for every fixed k and every model, and using larger η will decrease $\alpha_k^{(c)}$ more. These observations show that the data representations of each class scatter more uniformly after applying L_{CwD} .

Method	CIFAR100 ($B=50$)			ImageNet100 ($B=50$)			ImageNet ($B=100$)	
	$S=10$	5	2	10	5	2	100	50
LwF [18]	53.59 \pm 0.51	48.66 \pm 0.58	45.56 \pm 0.28	53.62 \dagger	47.64 \dagger	44.32 \dagger	40.86 \pm 0.13	27.72 \pm 0.12
iCaRL [26]	60.82 \pm 0.03	53.74 \pm 0.25	47.86 \pm 0.41	65.44 \dagger	59.88 \dagger	52.97 \dagger	49.56 \pm 0.09	42.61 \pm 0.15
BiC [33]	51.58 \pm 0.16	48.07 \pm 0.02	43.10 \pm 0.37	70.07 \dagger	64.96 \dagger	57.73 \dagger	43.23 \pm 0.13	38.83 \pm 0.12
LUCIR [12]	66.27 \pm 0.28	60.80 \pm 0.29	52.96 \pm 0.25	70.60 \pm 0.43	67.76 \pm 0.40	62.76 \pm 0.22	56.40 \pm 0.10	52.75 \pm 0.18
+CwD (ours)	67.26 \pm 0.16	62.89 \pm 0.09	56.81 \pm 0.21	71.94 \pm 0.11	69.34 \pm 0.31	65.10 \pm 0.59	57.42 \pm 0.11	53.37 \pm 0.22
PODNet [8]	66.98 \pm 0.13	63.76 \pm 0.48	61.00 \pm 0.18	75.71 \pm 0.37	72.80 \pm 0.35	65.57 \pm 0.41	57.01 \pm 0.12	54.06 \pm 0.09
+CwD (ours)	67.44 \pm 0.35	64.64 \pm 0.38	62.24 \pm 0.32	76.91 \pm 0.10	74.34 \pm 0.02	67.42 \pm 0.07	58.18 \pm 0.20	56.01 \pm 0.14
AANet [19]	69.79 \pm 0.21	67.97 \pm 0.26	64.92 \pm 0.30	71.96 \pm 0.12	70.05 \pm 0.63	67.28 \pm 0.34	51.76* \pm 0.14	46.86* \pm 0.13
+CwD (ours)	70.30 \pm 0.37	68.62 \pm 0.17	66.17 \pm 0.13	72.92 \pm 0.29	71.10 \pm 0.16	68.18 \pm 0.27	52.30* \pm 0.08	47.61* \pm 0.20

Table 1. **Comparison of average incremental accuracy (%) with or without Class-wise Decorrelation (CwD) at the initial phase.** B denotes number of classes learned at initial phase and S denotes number of classes learned per phase after the initial one. Number of exemplars for each class is 20. For AANet, we use its version based on LUCIR [12]. AANet [19] on ImageNet (denoted by *) is running without class-balance finetuning after each phase due to missing implementation in its code. All results are (re)produced by us except the ones denoted by \dagger , which are taken from [19]. Results (re)produced by us are averaged over 3 runs (mean \pm std).

4. Experiments

In this section, we first elaborate on the experimental setups in Sec. 4.1. Next, in Sec. 4.2, we add our proposed Class-wise Decorrelation (CwD) on some State-Of-the-Art (SOTA) methods [8, 12, 19] to validate its effectiveness. Finally, in Sec. 4.3, we provide ablation study on how factors such as number of classes of the initial phase, number of exemplars for each class, and CwD coefficient (η in Eqn. (9)) affects our proposed method. In addition, empirical comparison between our CwD and other decorrelation methods are given in the Appendix.

4.1. Settings

Datasets: We briefly introduce the 3 datasets used in our experiments: CIFAR100 [17] contains 100 classes, with 60000 samples, and the size of each image is 32×32 ; ImageNet [7] contains 1000 classes, with ≈ 1.3 million samples, and the size of each image is 224×224 ; ImageNet100 is a 100-class subset of the full ImageNet [7], which is generated similarly as in [8, 12, 19]. The classes of all datasets are first shuffled using seed 1993 as in [12, 13, 19, 20, 26, 29], and then split into multiple phases.

Implementation Details: For all experiments, we use ResNet18 [11] and the SGD optimizer, with the batch size of 128. The strategy of Herding [8, 12, 19, 26] is used to select exemplars after each phase. For experiments based on CIFAR100, in each CIL phase, all models are trained for 160 epochs and the learning rate is divided by 10 at the 80-th and 120-th epoch. For experiments based on ImageNet100/ImageNet, in each phase, all models are trained for 90 epochs and the learning rate is divided by 10 at the

30-th and 60-th epoch.

When performing CIL, we reserve a constant number of exemplars for each previously learned class as had been done in [8, 12, 19, 20].

Baselines: We apply our proposed CwD to the following three strong SOTA baselines: LUCIR [12], PODNet [8], and AANet [19]. For AANet, we use its version based on LUCIR [12]. In addition, we also report results of some other classical methods, including LwF [18], iCaRL [26] and BiC [33] for comparison.

Evaluation Metric: We use the the average incremental accuracy to evaluate performance of CIL methods as in [8, 12, 13, 19, 20]. Formally, suppose the CIL is conducted for $N + 1$ phases and test accuracy at the i -th phase is A_i , then the average increment accuracy is defined as:

$$\bar{A} = \frac{1}{N+1} \sum_{i=0}^N A_i. \quad (10)$$

4.2. CwD improves previous SOTA methods

In this section, we add our proposed CwD on three previous SOTA methods, namely LUCIR [12], PODNet [8] and AANet [19], to validate the effectiveness of our method. We denote the number of classes learned at initial phase by B and number of new classes learned per phase after the initial one by S . For CIFAR100 and ImageNet100, we evaluate our methods with three protocols: learning with $B = 50$ at the initial phase, and then learning with $S = 10/5/2$ for the rest of the classes. For ImageNet, we evaluate our methods with two protocols: learning with $B = 100$ at the initial phase, and then learning with $S = 100/50$ for the rest of

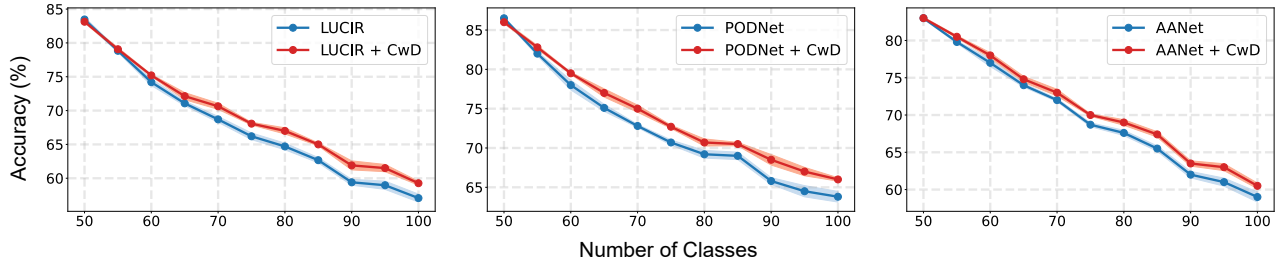


Figure 5. **Accuracy at each phase.** With ImageNet100, learning 50 classes at initial phase and 5 classes per phase for the rest 50 classes. Shading areas denote standard deviation.

the classes. The number of exemplars of each class is 20 for these experiments. The experimental results are shown in Tab. 1.

As shown in Tab. 1, across various CIL benchmarks, CwD can consistently improve the average incremental accuracy of previous SOTAs by around 1% to 3%. In particular, as we increment with fewer classes per phase (*i.e.*, smaller S), the improvements provided by adding CwD at the initial phase is more apparent. For example, with LUCIR [12] on CIFAR100 and ImageNet100 when $S = 2$, CwD can improve the baseline performance by up to 3.85% and 2.34%, respectively. This shows that, in settings that are more challenging (*e.g.*, with longer incremental learning sequences), our CwD can be even more effective.

Furthermore, in the large-scale CIL setting of ImageNet, CwD enables the weaker baseline LUCIR [12] to compete with or even outperform the stronger baseline PODNet [8].

To further understand CwD’s improvement at each phase, we plot how each method’s accuracy at each phase (*i.e.*, A_i in Eqn. (10)) changes as CIL proceeds on ImageNet100 with $S = 5$. From Fig. 5, we observe similar situation as directly mimic oracle model representation (see Fig. 2): our CwD barely affects accuracy at the initial phase, but bring strong improvements when learning at the following phases. This shows that the improvements by CwD on average incremental accuracy is not due to naïve performance improvements on the initial phase, but mainly due to an initial-phase representation that is favorable for incrementally learning new classes.

4.3. Ablation study

In this section, we study how the following factors affect the effectiveness of CwD: (1) number of classes for the initial phase, (2) number of exemplars, and (3) CwD coefficient (η in Eqn. (9)). Experiments in this section are based on LUCIR [12], with the ResNet18 model and the ImageNet100 dataset.

Ablation on the number of classes of the initial phase. We vary the number of classes learned at the initial phase (denoted by B) from 10 to 50 and run LUCIR [12] with or without our CwD. From the results in Tab. 2, we can see that, even when $B = 10$, CwD can bring a notable im-

S	B	LUCIR	+CwD (ours)	\uparrow
10	10	57.01 \pm 0.14	57.90 \pm 0.07	+0.89
	20	61.21 \pm 0.35	62.49 \pm 0.36	+1.28
	30	64.82 \pm 0.38	66.54 \pm 0.35	+1.72
	40	67.68 \pm 0.37	69.70 \pm 0.10	+2.02
	50	70.60 \pm 0.43	71.94 \pm 0.11	+1.33
5	10	50.47 \pm 0.31	51.92 \pm 0.10	+1.45
	20	56.41 \pm 0.37	58.14 \pm 0.13	+1.73
	30	61.00 \pm 0.09	63.18 \pm 0.14	+2.18
	40	63.73 \pm 0.23	66.25 \pm 0.16	+2.52
	50	67.76 \pm 0.40	69.34 \pm 0.31	+1.58

Table 2. **Ablation study on impact of number of classes of the initial phase.** In these experiments, we learn B classes at the initial phase, and then learn the rest of the classes with S classes per phase. Number of exemplars is set to 20 for each class. Average incremental accuracy (%) is reported. All results (mean \pm std) are produced by us, and are averaged over 3 runs.

provements (0.89% when $S = 10$ and 1.45% when $S = 5$). This shows that our CwD can improve CIL even when there is only a few number of classes in the initial phase. Note that this setting is quite challenging, since CwD can only directly affect representations of 10 out of 100 classes.

Additionally, as B increases, the performance of CwD first increases and reaches the peak when $B = 40$ (2.02% and 2.52% when $S = 10$ and 5 respectively), and then decreases. We posit that the reason for this increasing-then-decreasing trend is that, as B increases from 10 to 40, CwD can directly affect representations of more classes (in the initial phase), which enables it to yield greater improvements. However, as B continues increasing from 40 to 50, the gap between CIL and jointly training with all classes shrinks, resulting in less room for improvement of CIL.

Ablation on the number of exemplars. In this ablation, we vary the number of exemplars per class (denoted by R) and run LUCIR [12] with or without CwD.

As can be observed in Tab. 3, with more exemplars per class (*e.g.*, $R = 40$), the gap between CIL and jointly training with all classes is smaller, resulting in a still notable

S	R	LUCIR	+CwD (ours)	\uparrow
10	40	72.41 \pm 0.61	73.29 \pm 0.11	+0.88
	30	71.70 \pm 0.37	72.63 \pm 0.15	+0.93
	20	70.60 \pm 0.43	71.94 \pm 0.11	+1.34
	10	68.73 \pm 0.52	69.77 \pm 0.04	+1.04
	5	66.49 \pm 0.52	67.63 \pm 0.07	+1.14
5	40	70.74 \pm 0.49	72.06 \pm 0.11	+1.32
	30	68.56 \pm 0.42	70.04 \pm 0.12	+1.44
	20	67.76 \pm 0.40	69.34 \pm 0.31	+1.58
	10	64.07 \pm 0.38	66.07 \pm 0.46	+2.00
	5	60.41 \pm 0.77	62.58 \pm 0.53	+2.17

Table 3. **Ablation study on impact of number of exemplars.** In these experiments, we learn 50 classes at the initial phase, and then learn the rest of the classes with S classes per phase. Number of exemplars for each class is R . Average incremental accuracy (%) is reported. All results (mean \pm std) are produced by us, and are averaged over 3 runs.




Figure 6. **Impact of CwD Coefficient (η in Eqn. (9)).** Average incremental accuracy (%) is reported. Results are averaged over 3 runs. Shaded region denotes intervals that CwD is robust to η .

but smaller improvements (0.88% when $S = 10$ and 1.32% when $S = 5$). As R decrease, improvements provided by our CwD generally increase, reaching 1.34% when $S = 10$ and $R = 20$, and 2.17% when $S = 5$ and $R = 5$. This shows that our CwD yields more significant improvements in the more challenging setting with less exemplars.

Ablation on the CwD coefficient. In this ablation, we study how the average incremental accuracy changes when CwD coefficient (*i.e.*, η in Eqn. (9)) varies in the range of

[0, 1.25]. This ablation is based on LUCIR [12] and ImageNet100, with two CIL protocols that initially learn 50 classes and then learn 10/5 classes per phase for the rest.

From Fig. 6, we can see that, with both protocols, as η increases, the performance improvement brought by CwD first increases, then plateaus, and finally decrease.

The stable performance improvement (*i.e.*, the plateau interval) suggests that our CwD is robust to the choice of η . Furthermore, the plateau interval of the second protocol is larger than that of the first protocol ([0.25, 1.0] vs [0.25, 0.75]), indicating that our CwD is even more robust to η when the incremental learning sequence is longer.

When the η is too large (*e.g.*, larger than 0.75 for the first protocol and 1.0 for the second), the performance gain demonstrates a decreasing trend. The reason might be that an excessively large penalization on the Frobenius norm of class-wise correlation matrix would make the representations of each class to spread out drastically, resulting in large overlaps among representations of different classes. This is further studied in the Appendix.

5. Conclusion

In this work, we study Class Incremental Learning (CIL) from a previously underexplored viewpoint — improving CIL by mimicking the oracle model representation at the initial phase. Through extensive experimental analysis, we show the tremendous potential of this viewpoint. We propose a novel CwD regularization term for improving the representation of the initial phase. Our CwD regularizer yields consistent and significant performance improvements over three previous SOTA methods across multiple benchmark datasets with different scales.

Although we have developed our method through observations on the differences between the oracle model and the naïvely-trained initial-phase model, the underlying reason why more uniformly scattered representations for each class benefit CIL is still yet to be explored. Some analysis on why CwD improves CIL are provided in the Appendix, but we leave more rigorous understandings as future works.

Acknowledgement

The authors thank anonymous reviewers for their constructive feedback. Yujun Shi and Vincent Tan are funded by a Singapore National Research Foundation (NRF) Fellowship (R-263-000-D02-281) and a Singapore Ministry of Education AcRF Tier 1 grant (R-263-000-E80-114). Jian Liang is funded by Beijing Nova Program (Z211100002121108). Philip Torr is funded by Turing AI Fellowship EP/W002981/1.

References

- [1] Hongjoon Ahn, Jihwan Kwak, Subin Lim, Hyeonsu Bang, Hyojun Kim, and Taesup Moon. Ss-il: Separated softmax for incremental learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 844–853, 2021.
- [2] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 139–154, 2018.
- [3] Adrien Bardes, Jean Ponce, and Yann LeCun. Vircreg: Variance-invariance-covariance regularization for self-supervised learning. *arXiv preprint arXiv:2105.04906*, 2021.
- [4] Eden Belouadah and Adrian Popescu. I2m: Class incremental learning with dual memory. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 583–592, 2019.
- [5] Michael Cogswell, Faruk Ahmed, Ross Girshick, Larry Zitnick, and Dhruv Batra. Reducing overfitting in deep networks by decorrelating representations. *arXiv preprint arXiv:1511.06068*, 2015.
- [6] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 113–123, 2019.
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 248–255. Ieee, 2009.
- [8] Arthur Douillard, Matthieu Cord, Charles Ollion, Thomas Robert, and Eduardo Valle. Podnet: Pooled outputs distillation for small-tasks incremental learning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 86–102. Springer, 2020.
- [9] Robert M French. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4):128–135, 1999.
- [10] Tommaso Furlanello, Zachary Lipton, Michael Tschannen, Laurent Itti, and Anima Anandkumar. Born again neural networks. In *International Conference on Machine Learning*, pages 1607–1616. PMLR, 2018.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [12] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 831–839, 2019.
- [13] Xinting Hu, Kaihua Tang, Chunyan Miao, Xian-Sheng Hua, and Hanwang Zhang. Distilling causal effect of data in class-incremental learning. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 3957–3966, 2021.
- [14] Tianyu Hua, Wenxiao Wang, Zihui Xue, Sucheng Ren, Yue Wang, and Hang Zhao. On feature decorrelation in self-supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9598–9608, 2021.
- [15] Lei Huang, Dawei Yang, Bo Lang, and Jia Deng. Decorrelated batch normalization. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 791–800, 2018.
- [16] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.
- [17] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [18] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017.
- [19] Yaoyao Liu, Bernt Schiele, and Qianru Sun. Adaptive aggregation networks for class-incremental learning. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 2544–2553, 2021.
- [20] Yaoyao Liu, Yuting Su, An-An Liu, Bernt Schiele, and Qianru Sun. Mnemonics training: Multi-class incremental learning without forgetting. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 12245–12254, 2020.
- [21] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30:6467–6476, 2017.
- [22] Marc Masana, Xialei Liu, Bartłomiej Twardowski, Mikel Menta, Andrew D Bagdanov, and Joost van de Weijer. Class-incremental learning: survey and performance evaluation. *arXiv preprint arXiv:2010.15277*, 2020.
- [23] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier, 1989.
- [24] Seyed Iman Mirzadeh, Mehrdad Farajtabar, Razvan Pascanu, and Hassan Ghasemzadeh. Understanding the role of training regimes in continual learning. *arXiv preprint arXiv:2006.06958*, 2020.
- [25] Sudhanshu Mittal, Silvio Galesso, and Thomas Brox. Essentials for class incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 3513–3522, June 2021.
- [26] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 2001–2010, 2017.
- [27] Karsten Roth, Timo Milbich, Samarth Sinha, Prateek Gupta, Bjorn Ommer, and Joseph Paul Cohen. Revisiting training strategies and generalization performance in deep metric

- learning. In *International Conference on Machine Learning*, pages 8242–8252. PMLR, 2020.
- [28] Yujun Shi, Li Yuan, Yunpeng Chen, and Jiashi Feng. Continual learning via bit-level information preserving. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 16674–16683, 2021.
- [29] Christian Simon, Piotr Koniusz, and Mehrtash Harandi. On learning the geodesic path for incremental learning. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 1591–1600, 2021.
- [30] Shixiang Tang, Dapeng Chen, Jinguo Zhu, Shijie Yu, and Wanli Ouyang. Layerwise optimization by gradient decomposition for continual learning. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 9634–9643, 2021.
- [31] Shipeng Wang, Xiaorong Li, Jian Sun, and Zongben Xu. Training networks in null space of feature covariance for continual learning. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 184–193, 2021.
- [32] Guile Wu, Shaogang Gong, and Pan Li. Striking a balance between stability and plasticity for class-incremental learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1124–1133, 2021.
- [33] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 374–382, 2019.
- [34] Wei Xiong, Bo Du, Lefei Zhang, Ruimin Hu, and Dacheng Tao. Regularizing deep convolutional neural networks with a structured decorrelation constraint. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*, pages 519–528. IEEE, 2016.
- [35] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. *arXiv preprint arXiv:2103.03230*, 2021.