

Rebalancing Batch Normalization for Exemplar-based Class-Incremental Learning

Sungmin Cha¹, Sungjun Cho², Dasol Hwang², Sunwon Hong¹, Moontae Lee^{2,3}, and Taesup Moon^{1,4,5*}

¹Department of ECE, Seoul National University ²LG AI Research ³University of Illinois Chicago

⁴ASRI / INMC / IPAI / AIIS, Seoul National University ⁵SNU-LG AI Research Center

sungmin.cha@snu.ac.kr, {sungjun.cho, dasol.hwang}@lgresearch.ai,
zghdtnsz96@snu.ac.kr, moontae.lee@lgresearch.ai, tsmoon@snu.ac.kr

Abstract

Batch Normalization (BN) and its variants has been extensively studied for neural nets in various computer vision tasks, but relatively little work has been dedicated to studying the effect of BN in continual learning. To that end, we develop a new update patch for BN, particularly tailored for the exemplar-based class-incremental learning (CIL). The main issue of BN in CIL is the imbalance of training data between current and past tasks in a mini-batch, which makes the empirical mean and variance as well as the learnable affine transformation parameters of BN heavily biased toward the current task — contributing to the forgetting of past tasks. While one of the recent BN variants has been developed for “online” CIL, in which the training is done with a single epoch, we show that their method does not necessarily bring gains for “offline” CIL, in which a model is trained with multiple epochs on the imbalanced training data. The main reason for the ineffectiveness of their method lies in not fully addressing the data imbalance issue, especially in computing the gradients for learning the affine transformation parameters of BN. Accordingly, our new hyperparameter-free variant, dubbed as Task-Balanced BN (TBBN), is proposed to more correctly resolve the imbalance issue by making a horizontally-concatenated task-balanced batch using both reshape and repeat operations during training. Based on our experiments on class incremental learning of CIFAR-100, ImageNet-100, and five dissimilar task datasets, we demonstrate that our TBBN, which works exactly the same as the vanilla BN in the inference time, is easily applicable to most existing exemplar-based offline CIL algorithms and consistently outperforms other BN variants.

1. Introduction

In recent years, continual learning (CL) has been actively studied to efficiently learn a neural network on sequentially

arriving datasets while eliminating the process of re-training from scratch at each arrival of a new dataset [9]. However, since the model is typically trained on a dataset that is heavily skewed toward the current task at each step, the resulting neural network often suffers from suboptimal trade-off between stability and plasticity [25] during CL. To overcome this issue, various studies have been focused on addressing the so-called *catastrophic forgetting* phenomenon [9, 28].

Among different CL settings, the class-incremental learning (CIL) setting where the classifier needs to learn previously unseen classes at each incremental step has recently drawn attention due to its practicality [1, 3, 6, 12, 15, 26, 39, 42, 44]. Most state-of-the-art CIL algorithms maintain a small *exemplar memory* to store a subset of previously used training data and combine it with the current task dataset to mitigate the forgetting of past knowledge. A key issue of exemplar-based CIL is that the model prediction becomes heavily biased towards more recently learned classes, due to the imbalance between the training data from the current task (that are abundantly available) and past tasks (with limited access through exemplar memory). In response, recently proposed solutions to biased predictions include bias correction [39], unified classifier [15], and separated softmax [1], which greatly improved the overall accuracy of CIL methods across all classes learned so far.

Despite such progress, relatively less focus has been made on backbone architectures under CIL on computer vision tasks. Specifically, the popular CNN-based models (*e.g.*, ResNet [13]) are mainly used as feature extractors, and those models are equipped with Batch Normalization (BN) [17] by default. However, since BN is designed for single-task training on CNNs, applying BN directly to exemplar-based CIL results in statistics biased toward the current task due to the imbalance between current and past tasks’ data in a mini-batch. Recently, [29] pointed out this issue in CIL, dubbed as *cross-task normalization effect*, and proposed a new normalization scheme called Continual Normalization

*Corresponding author (E-mail: tsmoon@snu.ac.kr)

(CN), which applies Group Normalization (GN) [40] across the channel dimension before running BN across the batch dimension. Therefore, the difference in feature distributions among tasks is essentially removed by GN, and the following BN computes task-balanced mean and variance statistics which is shown to outperform vanilla BN on *online* CIL settings.

In this paper, we argue that CN only partially resolves the bias issue in exemplar-based CIL. Specifically, we find that the gradients on the affine transformation parameters remain biased towards the current task when using CN. As a result, it leads to inconsistent performance gains in the *offline* CIL setting, which is considered more practical than *online* CIL. To this end, we propose a simple yet novel hyperparameter-free normalization layer, dubbed as Task-Balanced Batch Normalization (TBBN), which effectively resolves the bias issue. Our method employs adaptive *reshape* and *repeat* operations on the mini-batch feature map during training in order to compute task-balanced normalization statistics and gradients for learning the affine transformation parameters. Our method does not require any hyperparameter as the size-inputs for reshape and repeat operations are determined adaptively. Furthermore, the application of TBBN during testing is identical to vanilla BN, requiring *no* change on the backbone architecture. Through extensive offline CIL experiments on CIFAR-100, ImageNet-100, and five dissimilar task datasets, we show that a simple replacement of BN layers in the backbone CNN model with TBBN benefits most state-of-the-art exemplar-based CIL algorithms towards an additional boost in performance. Our analysis shows that the gain of TBBN is consistent across various backbone architectures and datasets, suggesting its potential to become a correct choice for exemplar-based offline CIL algorithms.

2. Related Work

Exemplar-based CIL. Among various work in offline CIL, iCaRL [30] is the first to propose an exemplar-based method via utilization of nearest-mean-of-exemplars classification and representation learning. EEIL [6] next devised a distillation-based CIL method with both balanced fine-tuning and representative memory updating. BiC [39] focused on the biased prediction aspect due to data imbalance in CIL and proposed a novel yet simple idea to attach an additional layer that corrects biased predictions. Inspired by LWF [21], LUCIR [15] developed a method that overcomes catastrophic forgetting via a more sophisticated cosine normalization-based loss function. PODNet [12] proposed another distillation-based method that focuses on spatial-based loss to reduce representation forgetting. Most recently, two state-of-the-art knowledge distillation-based CIL methods have been proposed. First, SS-IL [1] demonstrated that a pure softmax mainly causes biased predictions in CIL, leading to the Separated Softmax layer that achieved significant

performance gains. Second, AFC [18] devised a novel regularization for knowledge distillation to minimize the changes of important features when learning new tasks.

On the other hand, online CIL also has attracted attention [2,33], but their performance is significantly lower in the benchmark (*e.g.*, CIFAR-100) than offline CIL counterparts, as reported in the survey paper [23].

Normalization layer. One initial motivation for devising Batch Normalization (BN) [17] was to address internal covariate shift of a neural net. However, the belief in internal covariate shift was broken by follow-up studies, and the benefits from BN in terms of training perspective were analyzed in various directions [4, 32]. Thereafter, several normalization layers devised for various computer vision tasks have been proposed, each with their respective advantages [14, 16, 22, 37, 40], and for more diverse tasks, such as meta-learning [5], domain adaptation [38], task-incremental learning [20], and online CL [29].

3. Preliminaries

Notation. We assume the setting of offline CIL as proposed in [1]; namely, a dataset \mathcal{D}_t for each task t arrives incrementally, and we use T to denote the total number of tasks. Each \mathcal{D}_t consists of pairs of an input image \mathbf{x}_t and its target label y_t . We assume that each incrementally added task introduces m new classes that the model has never observed previously, and hence the total number of classes observed until task t equals $C_t = m \cdot t$. Therefore, the target label in each task is labeled as $y_t \in \{C_{t-1} + 1, \dots, C_t\} \triangleq \mathcal{C}_t$. To save exemplars from previous tasks, we use exemplar memory denoted by \mathcal{M} . While \mathcal{M} is updated via a sampling algorithm after training on each task, the memory size is set to not exceed $|\mathcal{M}|$ over the entire course of tasks.

In particular, exemplar memory \mathcal{M}_{t-1} maintains $\lfloor \frac{|\mathcal{M}|}{C_{t-1}} \rfloor$ images per each past class and is used for training task t . To merge instances from current as well as previous tasks, we consider a training mini-batch size of $B = B_c + B_p$, where B_c and B_p denote the number of data samples from \mathcal{D}_t and \mathcal{M}_{t-1} , respectively. Without loss of generality, we assume that each mini-batch $(\mathbf{x}^B, y^B) \sim (\mathcal{D}_t \cup \mathcal{M}_{t-1})$ is a concatenation of B_c samples from \mathcal{D}_t with B_p samples from \mathcal{M}_{t-1} : $\mathbf{x}^B = [\mathbf{x}_1; \dots; \mathbf{x}_{B_c}; \mathbf{x}_{B_c+1}; \dots; \mathbf{x}_B]$ and $y^B = [y_1; \dots; y_{B_c}; y_{B_c+1}; \dots; y_B]$ with $\{(\mathbf{x}_i, y_i)\}_{i=1}^{B_c} \subset \mathcal{D}_t$ and $\{(\mathbf{x}_i, y_i)\}_{i=B_c+1}^B \subset \mathcal{M}_{t-1}$. Note that B_c is greater than B_p in typical CIL settings, and it is this imbalance that is generally known to cause predictions biased towards the current task [1, 3, 39]. Therefore, we set the sampling ratio between B_c and B_p as $B_c : B_p = 3 : 1$ in our experiments. Lastly, we use $\mathbf{h} \in \mathbb{R}^{B \times C \times D}$ to denote the input feature map for an intermediate BN layer given \mathbf{x}^B as input. Here, C and D denote the number of channels and feature dimension, respectively.

Batch Normalization. Due to its powerful practicality and efficiency, Batch Normalization (BN) [17] has been the go-to normalization layer across various state-of-the-art neural network architectures [13, 34, 36]. During training, BN calculates the empirical mean $\hat{\mu} \in \mathbb{R}^C$ and variance $\hat{\sigma}^2 \in \mathbb{R}^C$ of the given mini-batch as

$$\hat{\mu}_c = \frac{1}{BD} \sum_{b=1}^B \sum_{d=1}^D \mathbf{h}_{b,c,d}, \quad \hat{\sigma}_c^2 = \frac{1}{BD} \sum_{b=1}^B \sum_{d=1}^D (\mathbf{h}_{b,c,d} - \hat{\mu}_c)^2. \quad (1)$$

The feature map \mathbf{h} is then normalized to have zero-mean unit-variance using $\hat{\mu}$ and $\hat{\sigma}^2$ to obtain $\hat{\mathbf{h}} \in \mathbb{R}^{B \times C \times D}$. The normalized feature $\hat{\mathbf{h}}$ is lastly affine-transformed with a trainable scale and shift parameters $\gamma, \beta \in \mathbb{R}^C$ to \mathbf{y} :

$$\hat{\mathbf{h}}_b = \frac{\mathbf{h}_b - \hat{\mu} \cdot \mathbf{1}^\top}{\sqrt{\hat{\sigma}^2 \cdot \mathbf{1}^\top + \epsilon \cdot \mathbf{1}^\top}} \quad (2)$$

$$\mathbf{y}_b = \gamma \cdot \mathbf{1}^\top \odot \hat{\mathbf{h}}_b + \beta \cdot \mathbf{1}^\top, \quad (3)$$

in which $\mathbf{1} \in \mathbb{R}^D$ is the all-1 vector, \odot denotes the element-wise multiplication, and b is the data index within the mini-batch. Note the division and $\sqrt{\cdot}$ in (2) are also done in an element-wise sense. The ϵ term is a small value added for numerical stability, which we set as $\epsilon = 10^{-5}$ in the experiments. It is worth noting that γ and β are trained by computing the gradients

$$\frac{\partial \mathcal{L}}{\partial \gamma_c} = \sum_{b,d=1}^{B,D} \frac{\partial \mathcal{L}}{\partial \mathbf{y}_{b,c,d}} \odot \hat{\mathbf{h}}_{b,c,d}, \quad \frac{\partial \mathcal{L}}{\partial \beta_c} = \sum_{b,d=1}^{B,D} \frac{\partial \mathcal{L}}{\partial \mathbf{y}_{b,c,d}} \quad (4)$$

under a loss function \mathcal{L} . For the normalization in the test phase, running μ and σ^2 are updated incrementally at each i -th training iteration via exponential moving average:

$$\begin{aligned} \mu_{(i)} &= \alpha \cdot \mathbf{1}^\top \odot \hat{\mu}_{(i)} + (1 - \alpha) \cdot \mathbf{1}^\top \odot \mu_{(i-1)} \\ \sigma_{(i)}^2 &= \alpha \cdot \mathbf{1}^\top \odot \hat{\sigma}_{(i)}^2 + \left((1 - \alpha) \frac{V - 1}{V} \right) \cdot \mathbf{1}^\top \odot \sigma_{(i-1)}^2 \end{aligned} \quad (5)$$

in which $V = B \cdot D$ represents the vessel’s correction, and α is the momentum hyperparameter generally set as 0.1.

4. Main Method

4.1. Motivation

While exemplar memory is essential for achieving state-of-the-art performance in offline CIL [24], most methods use well-known backbone classification models (*e.g.*, ResNet-18 [13]) with vanilla BN layers without much scrutiny over the issue of data-imbalance between current and previous tasks. One notable exception is Continual Normalization (CN) [29], but CN has only shown promise in online CIL settings where the data-imbalance causes relatively less of an impact than in offline CIL due to the *single* epoch learning. It is not clear whether CN is also effective in offline CIL, where

the imbalanced mini-batches are used for *multiple* epochs. We conjecture that while CN performs task-balancing by normalizing each channel (or groups of channels) along the spatial dimension, this comes at a cost in discriminative power as it normalizes features from all tasks equally without any separation between tasks. In a later section, we find through experimentation that this limitation is especially detrimental in offline CIL, as CN leads to negligible gain over BN or even worse performance in some cases.

For a concrete motivation, we consider a toy example to reflect upon the necessity of devising a correct BN mechanism for the exemplar memory-based CIL. Consider a synthetic input feature distribution for an intermediate BN layer shown in Figure 1. Figure 1(a) visualizes 2-D data samples generated from a mixture of four different Gaussian distributions, each representing the feature distribution for each given task.¹ In a joint training setting, the mini-batch samples are uniformly sampled across four tasks, and thus the empirical mean and variance obtained from BN would approach the true population mean and variance. Such mean and variance will then successfully standardize the samples in a mini-batch during both training and testing phases via the procedure given in (2) and (3).

In an exemplar memory-based CIL setting, however, there is a severe imbalance among samples in the training mini-batch, since the majority of the samples comes from the current task (Task 4) as shown in Figure 1(b). In this case, the mean and variance calculated by the BN layer is clearly biased towards those of the current task, failing to match the true statistics at test time. In other words, using the biased mean and variance obtained during training to normalize test samples leads the samples to deviate from a standardized distribution. It is evident that such mismatch between the train and test sample distributions can cause performance deterioration since the trainable parameters in BN are obtained from normalized *training* samples. An obvious remedy would be to obtain the *task-balanced* empirical mean and variance during training, and use them in the BN layer to normalize the test samples. Figure 1(c) and 1(d) highlight the difference between the test normalization done by CN and our TBBN, respectively, which computes the empirical mean and variance from the mini-batch in Figure 1(b). Note that while both successfully normalize test samples in a task-balanced way, CN tends to lose the discriminative structure among the tasks due to the intermediate GN step. In contrast, our TBBN, as will be described in details below, normalizes the test data features while maintaining the discriminative structure.

In addition to the mismatch during test time, a naïve application of BN on the biased mini-batch samples can also

¹In fact, we generated 20 dimensional Gaussian vectors, which is generated by concatenating i.i.d. 2-D Gaussians 10 times, and only the first two dimensions are shown in the figures. The normalization by CN and TBBN shown in Figure 1(c)(d) are done by obtaining the statistics from the entire 20 dimensions (which exemplifies the channel dimension in a feature map).

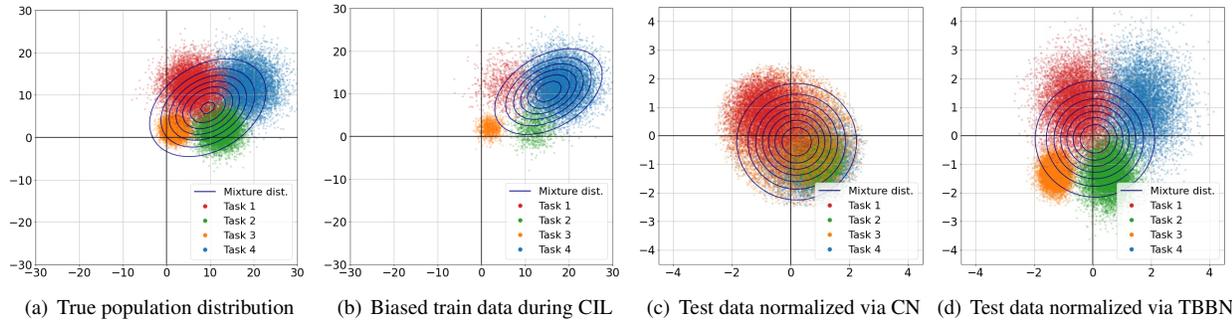


Figure 1. (a) 2-D visualizations on a mixture of four Gaussian distributions, each representing a feature map distribution of each task. (b) Imbalanced training data in exemplar-based CIL while learning Task 4, deviating from true test distribution. (c) Test set normalization result of CN that computes normalization statistics from data in (b). (d) Test set normalization result of our proposed TBBN that computes normalization statistics from data in (b).

cause an issue during training. That is, when training for a current task (*e.g.*, Task 4 in Figure 1) is finished, the empirical mean and variance obtained by a BN layer would be heavily biased towards the most recent task. When a new task arrives (say, Task 5), since only a small number of samples from Task 4 is saved in the exemplar memory, the new empirical mean and variance will now become biased towards Task 5, drifting away from the previous statistics updated via the exponential moving average in (5). This mismatch during training can also significantly alter the learned representations of the past task samples in the exemplar memory after normalization. Therefore, this necessitates re-training of subsequent layers, unintentionally causing forgetting of previously learned representations.

With regards to the two aforementioned issues, we argue that it is crucial to calculate the *task-balanced* mean and variance for BN layers in exemplar memory-based CIL. While a similar observation has also been made in the recent work [29], we go one step further and note that the affine transformation parameters of BN (*i.e.* γ and β in (3)) should also be learned in a task-balanced manner. For a demonstration purpose, we train a ResNet-18 model with vanilla BN and exemplar memory of size $|\mathcal{M}| = 2000$ on ImageNet-100 split to 10 tasks with 10 classes per task. We test four different CIL procedures including two oracle-based approaches that can access the entire training set during test time to exhibit the necessity of task-balanced statistics and affine-transformations. In Figure 2, JOINT denotes the algorithm that jointly trains the model on all tasks at once, providing an upper bound in average accuracy. On the other hand, FT is the fine-tuning baseline that simply fine-tunes the model on samples from the current task and exemplar memory for each task. Note that FT performs poorly on previous tasks except for the final task, clearly reflecting the issue of biased predictions. UPDATE (μ, σ^2) is an oracle scheme that *freezes* all trainable parameters of the trained FT model, and only recomputes the mean and variance of all BN layers using the entire dataset such that they match the task-balanced

population statistics. Finally, UPDATE (μ, σ^2) AND (γ, β) denotes another oracle scheme that also *re-trains* the (γ, β) parameters for all BN layers on the entire dataset, in addition to recomputing the (μ, σ^2) statistics.

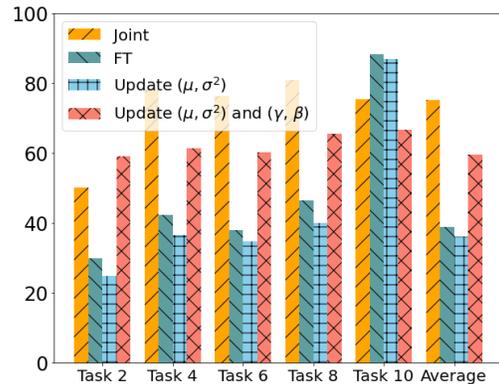


Figure 2. Individual task and average accuracies on ImageNet-100 split into 10 tasks for CIL. y -axis shows the accuracy for each task and the average accuracy across all tasks after full training. Recomputing the affine-transformation parameters (γ, β) as well as empirical statistics (μ, σ^2) based on the full dataset leads to significant boost in performance.

Figure 2 displays the classification accuracy of four methods on every 2nd task and the average accuracy across all tasks, all measured after training up to the final task. The slight drop in performance of UPDATE (μ, σ^2) compared to FT shows that only obtaining the task-balanced (μ, σ^2) is not sufficient due to mismatch between the updated (μ, σ^2) and the frozen (γ, β). In contrast, when affine-transformation parameters are also recomputed with task-balanced data, we observe a significant boost in average accuracy. This is particularly interesting as the affine-transformation parameters account for less than 1% of the entire model parameters, and simply making them task-balanced along with the empirical (μ, σ^2) statistics in BN layers can bring a significant performance boost under a simple FT baseline. Inspired by this example, we develop our Task-Balanced Batch Normalization (TBBN) layer in the next section.

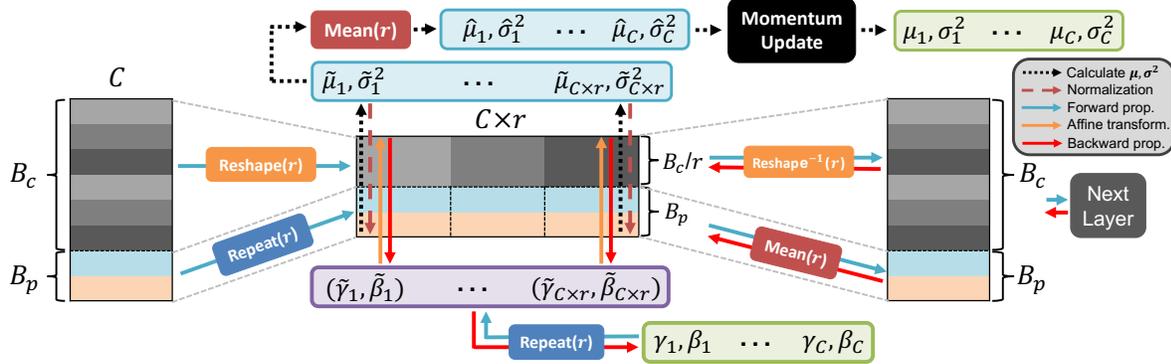


Figure 3. Illustration of the forward (blue arrows) and backward (red arrows) propagations of TBBN. By applying both reshape and repeat operations to a given input of BN layer, TBBN induces that task-balanced calculation and training of (μ, σ^2) and (γ, β) , respectively.

4.2. Task-Balanced Batch Normalization (TBBN)

To solve the problem of BN in the exemplar-based CIL, we propose Task-Balanced Batch Normalization (TBBN) consisting of two components: 1) Task-balanced μ and σ^2 calculation and 2) Task-balanced training of γ and β .

Task-balanced μ and σ^2 calculation. As discussed in the previous section, the class imbalance between the current and previous data biases the mean and variance toward the current task. Instead, we propose a task-balanced algorithm that better computes the empirical mean and variance in the current batch. As shown in Figure 3, we design the two different tensor operations: *reshape* and *repeat*, applying them to the current and previous parts, respectively. We determine the number of repeat/reshape operations r by the following (6). It determines the value of r , which is used to balance the number of batches sampled from the exemplar memory ($\frac{B_p}{t-1}$) and the number of batches of the current task ($\frac{B_c}{r}$) after reshaping with r , so that the ratio between them becomes 1:1. Note that r is a hyperparameter that need not be externally tuned, but that our method provides a task-adaptive guidance for automatically selecting r :

$$\frac{B_c}{r} = \frac{B_p}{t-1}, \text{ hence, } r = \frac{B_c}{B_p} \cdot (t-1), \quad (6)$$

in which $t = 2, \dots, T$, and r is set to 1 when $t = 1$.

Let \mathbf{h}_{B_c} be the input feature map of a BN layer for the current task t . After applying the reshape operation $\mathbf{H}_{B_c} = F_S(\mathbf{h}_{B_c}; r) : \mathbb{R}^{B_c \times C \times D} \rightarrow \mathbb{R}^{B_c/r \times C \cdot r \times D}$ with r , we divide \mathbf{h}_{B_c} into r splits, where each split has $\frac{B_c}{r}$ numbers of data from the current task. Similarly, let \mathbf{h}_{B_p} denote the data from the previous tasks. As \mathbf{h}_{B_p} consists of the data uniformly sampled from task 1 to $t-1$, $\frac{B_p}{t-1}$ numbers of data belongs to \mathbf{h}_{B_p} for each previous task in expectation. The key insight is that we can make every split balanced by repeating \mathbf{h}_{B_p} by r times if the task ratios $\frac{B_c}{r}$ and $\frac{B_p}{t-1}$ become equivalent. So that, we repeat \mathbf{h}_{B_p} for r times by applying a tensor repeat function $\mathbf{H}_{B_p} = F_P(\mathbf{h}_{B_p}; r) : \mathbb{R}^{B_p \times C \times D} \rightarrow \mathbb{R}^{B_p \times C \cdot r \times D}$. Then, we concatenate both tensors along batch

axis, $\mathbf{H} = (\mathbf{H}_{B_c}; \mathbf{H}_{B_p}) \in \mathbb{R}^{(B_c/r+B_p) \times C \cdot r \times D}$, and calculate a task-balanced empirical mean $\tilde{\mu} \in \mathbb{R}^{C \cdot r}$ and variance $\tilde{\sigma}^2 \in \mathbb{R}^{C \cdot r}$ from a horizontally-concatenated task-balanced batch \mathbf{H} . To update running mean $\mu \in \mathbb{R}^C$ and variance $\sigma^2 \in \mathbb{R}^C$ for test phase, we average $\tilde{\mu}$ and $\tilde{\sigma}^2$ over r splits, later applying these to the exponential moving average (5) to update μ and σ^2 .

Task-balanced training of γ and β . Additionally, we propose a way to train learnable parameters of BN (γ and β) in a task-balanced way to the current task. After calculating the task-balanced $\tilde{\mu}$ and $\tilde{\sigma}^2$ from \mathbf{H} , we apply these to normalize the horizontally-concatenated task-balanced batch \mathbf{H} via (2). Then, the normalized feature map $\tilde{\mathbf{H}} \in \mathbb{R}^{(B_c/r+B_p) \times C \cdot r \times D}$ is affine-transformed by (3) with $\tilde{\gamma} = F_{RP}(\gamma) \in \mathbb{R}^{C \cdot r}$ and $\tilde{\beta} = F_{RP}(\beta) \in \mathbb{R}^{C \cdot r}$, resulting in the affine-transformed feature $\mathbf{Y} = (\mathbf{Y}_{B_c}; \mathbf{Y}_{B_p}) \in \mathbb{R}^{(B_c/r+B_p) \times C \cdot r \times D}$. However, due to the different size of batch axis with the original batch size ($B_c/r + B_p \neq B$), \mathbf{Y} cannot be directly forward-propagated toward a next layer. To solve this problem, first, we apply a reverse operation of reshape to \mathbf{Y}_{B_c} to obtain $\mathbf{y}_{B_c} \in \mathbb{R}^{B_c \times C \times D}$, where $F_S^{-1}(\mathbf{Y}_{B_c}; r) : \mathbb{R}^{B_c/r \times C \cdot r \times D} \rightarrow \mathbb{R}^{B_c \times C \times D}$. Next, for the case of \mathbf{Y}_{B_p} , we get $\mathbf{y}_{B_p} \in \mathbb{R}^{B_p \times C \times D}$ by averaging \mathbf{Y}_{B_p} over r repeats along the axis of channel, such as $\mathbf{y}_{B_p} = F_A(\mathbf{Y}_{B_p}; r) : \mathbb{R}^{B_p \times C \cdot r \times D} \rightarrow \mathbb{R}^{B_p \times C \times D}$. Finally, we forward-propagate the concatenated feature $\mathbf{y} = (\mathbf{y}_{B_c}; \mathbf{y}_{B_p}) \in \mathbb{R}^{B \times C \times D}$ toward the next layer.

4.2.1 Why does TBBN work?

Calculating task-balanced μ and σ^2 . When performing the standard BN while fixing the number of data instances from previous tasks as B_p , the individual statistics (μ^{BN}, σ^{BN}) becomes dominated by the current task. As the exponential moving average iteratively averages over biased empirical statistics, BN fails to asymptotically recover population mean and variance. This is especially clear when expanding the BN mean μ_t^{BN} during training on task t in terms of means μ_i of the i -th task. Assuming that each task contributes uniformly in the global data distribution, μ_t^{BN}

deviates from the true global mean μ_t^* by

$$\mu_t^* - \mu_t^{BN} = \frac{B - B_c t}{t(t-1)B} \sum_{i=1}^t \mu_i + \frac{B_c t - B}{(t-1)B} \mu_t. \quad (7)$$

In a nutshell, BN fails to recover the true population mean unless the mini-batch is balanced such that $B_c = B/t$. Our proposed TBBN creates r task-balanced bags of bootstrapped samples such that this balancing is satisfied, and approximates the population statistics (μ^*, σ^{2*}) in an unbiased manner.

Training γ and β with task-balanced gradients. Note that the gradient for training β and γ of the original BN, shown in (4), is biased toward the current task due to the imbalance between B_c and B_p in the mini-batch. On the other hand, the above task-balanced process makes both β and γ to be trained in a task-balanced way during backpropagation. First, a gradient of \mathbf{y} is back-propagated toward the BN layer and it can be denoted as: $\frac{\partial \mathcal{L}}{\partial \mathbf{y}} = \left(\frac{\partial \mathcal{L}}{\partial \mathbf{Y}_{B_c}}; \frac{\partial \mathcal{L}}{\partial \mathbf{Y}_{B_p}} \right) \in \mathbb{R}^{B \times C \times D}$. Then, the gradients of each current and task’s batch pass the backward operation of F_S^{-1} and F_A , respectively. As a result, the gradient of \mathbf{Y} becomes:

$$\frac{\partial \mathcal{L}}{\partial \mathbf{Y}} = \left(F_{S,g}^{-1} \left(\frac{\partial \mathcal{L}}{\partial \mathbf{Y}_{B_c}}; r \right); F_{A,g} \left(\frac{\partial \mathcal{L}}{\partial \mathbf{Y}_{B_p}}; r \right) \right), \quad (8)$$

in which $\frac{\partial \mathcal{L}}{\partial \mathbf{Y}} \in \mathbb{R}^{(B_c/r+B_p) \times C \cdot r \times D}$. $F_{S,g}^{-1}$ and $F_{A,g}$ stands for the backward operation of F_S^{-1} and F_A , respectively. The backward operation of F_S^{-1} is the reverse-direction-reshape of a given gradient: $F_{S,g}^{-1} \left(\frac{\partial \mathcal{L}}{\partial \mathbf{Y}_{B_c}}; r \right) = F_S \left(\frac{\partial \mathcal{L}}{\partial \mathbf{Y}_{B_c}}; r \right) : \mathbb{R}^{B_c \times C \times D} \rightarrow \mathbb{R}^{B_c/r \times C \cdot r \times D}$. In the case of F_A , the backward operation becomes $F_{A,g} \left(\frac{\partial \mathcal{L}}{\partial \mathbf{Y}_{B_p}}; r \right) = \left(\left(\frac{\partial \mathcal{L}}{\partial \mathbf{Y}_{B_p}} \right)_1 / r; \dots; \left(\frac{\partial \mathcal{L}}{\partial \mathbf{Y}_{B_p}} \right)_r / r \right) : \mathbb{R}^{B_p \times C \times D} \rightarrow \mathbb{R}^{B_p \times C \cdot r \times D}$, where distributes and expands the input gradient divided by r , for r times. As a result, a gradient of β and γ becomes:

$$\frac{\partial \mathcal{L}}{\partial \gamma} = F_{P,g} \left(\sum_{b,d=1}^{(B_c/r+B_p),D} \frac{\partial \mathcal{L}}{\partial \mathbf{Y}_{b,d}} \odot \tilde{\mathbf{H}}_{b,d}; r \right) \quad (9)$$

$$\frac{\partial \mathcal{L}}{\partial \beta} = F_{P,g} \left(\sum_{b,d=1}^{(B_c/r+B_p),D} \frac{\partial \mathcal{L}}{\partial \mathbf{Y}_{b,d}}; r \right), \quad (10)$$

in which $F_{P,g}$ denotes the backward operation of F_P , which is the sum of gradients over r : $F_{P,g}(\mathbf{g}; r) = \sum_{i=1}^r ((\mathbf{g}_i); \dots; (\mathbf{g}_r)) : \mathbb{R}^{C \cdot r} \rightarrow \mathbb{R}^C$, and $\frac{\partial \mathcal{L}}{\partial \mathbf{Y}_{b,d}} \in \mathbb{R}^C$ and $\tilde{\mathbf{H}}_{b,d} \in \mathbb{R}^C$. Note that the core differences with the gradients of the original BN shown in (4): First, the summed up gradients are configured to be task-balanced. Second, in the case of the gradient of γ , the task-balanced batch $\tilde{\mathbf{H}}$, normalized by the task-balanced mean and variance, is used to get the gradient. As a result, γ and β can be trained in a more task-balanced way, so that they can do their part of BN at test time where the entire task’s test data is given jointly.

Note that both the pseudo code and implementation details of TBBN are proposed in supplementary materials.

5. Experimental Results

Datasets and Evaluation. We mainly evaluate our method on the offline CIL settings leveraging CIFAR-100 [19] and ImageNet-100 [10] datasets: for each dataset, we divide 100 classes into 10 disjoint sets, and consider each set to constitute a single task (*i.e.* $T = 10$ and $m = 10$). Following previous work [7, 9, 24], we evaluate all normalization methods based on four different metrics, such as final accuracy (A_f), average accuracy (A_a), forgetting measure (F), and learning accuracy (A_l). The detailed explanation for the metrics can be found in the supplementary materials.

Baselines. For backbone architectures, we follow the offline CIL setup proposed in [24] using ResNet-32 [13] for CIFAR-100 [19] and ResNet-18 for ImageNet-100 [10]. In addition, we use exemplar memory with size $|\mathcal{M}| = 2000$ and incrementally update \mathcal{M} via a class-balanced random sampling strategy after each task. To test our method under various offline CIL algorithms, we consider EEIL [6] and LUCIR [15], for which implementations are available in the benchmark environment provided by [24]. We also run SS-IL [1], POD-Net [12] and AFC [18], which are considered state-of-the-art in offline CIL, using their official code. Finetuning (FT) stands for a naive baseline which only utilizes the exemplar memory without using any explicit methods. For comparison across different normalization approaches, we simply replace all BN layers in the backbone model with one of the following: Instance Normalization (IN) [37], Group Normalization (GN) [40], Switch Normalization (SN) [22], Batch Renormalization (BRN) [16], Continual Normalization (CN) [29], and our proposed TBBN. We use official implementations for each baseline method published by the respective authors. Note that CN and GN require a hyperparameter G that sets the group size during feature-level normalization. For our experiments, we use $G = 16$ when training ResNet-32 on CIFAR-100 and $G = 32$ when training ResNet-18 on ImageNet-100. Further details on experimental settings are in the supplementary material.

5.1. Quantitative Results

Effect of normalization layers in offline CIL. Table 1 shows experimental results for FT on 10-task offline CIL setups on CIFAR-100 and ImageNet-100. First, we find that normalization layers previously developed for single task supervised learning (e.g., IN [37], GN [40], SN [22] and BRN [16]) are not effective in offline CIL scenarios compared to BN. Both IN and GN especially show notable performance degradation with respect to all four metrics. Furthermore, CN [29] does not lead to significant improvement, which is in contrast to the performance gain previously demonstrated in online CIL by the original paper. This shows that the benefit of adding the channel-level GN to BN is obscured when naïve BN is exposed to imbalanced mini-batches for multiple epochs. On the other hand, we confirm

that simply replacing BN with TBBN increases both A_f and A_a in both datasets without any additional hyperparameters. We find that the performance gain mainly comes from the increased stability as well as plasticity as we see improvements in both F and A_l . TBBN especially shows significantly less forgetting, supporting our insight that preventing sample statistics from becoming biased to the current task via task-balanced normalization effectively retains information from past tasks leading to a better accuracy.

Table 1. Experimental results for FT with varying normalization layers on CIFAR-100 and ImageNet-100. Bold indicates the best performance in each metric.

Method	CIFAR-100 w/ ResNet-32				ImageNet-100 w/ ResNet-18			
	$A_f(\uparrow)$	$A_a(\uparrow)$	$F(\downarrow)$	$A_l(\uparrow)$	$A_f(\uparrow)$	$A_a(\uparrow)$	$F(\downarrow)$	$A_l(\uparrow)$
BN	35.41	53.88	43.48	78.79	39.40	59.60	48.02	87.42
IN	31.72	46.84	47.72	79.44	33.45	53.59	50.69	84.66
GN	31.26	44.53	44.01	75.27	28.83	47.79	49.79	79.19
SN	36.29	53.64	42.91	79.20	39.45	59.55	48.04	87.79
BRN	36.08	52.58	44.34	80.42	37.49	57.77	48.13	86.57
CN	35.06	54.43	43.82	80.64	41.96	60.02	45.32	87.28
CN*	36.05	54.18	44.81	80.85	40.46	59.03	46.74	87.20
TBBN	38.46	56.17	41.90	80.36	43.20	61.69	43.62	87.82

Applying TBBN to other CIL baselines. We also confirm that the advantage of TBBN is extended to other CIL methods. We consider the same 10-task offline CIL setup from CIFAR-100 and ImageNet-100, and run a total of six different CIL methods while simply replacing the BN layers in the model with either CN or TBBN.

Table 2 shows the corresponding results. First, the most recently proposed CIL methods, SSIL and AFC, achieve the best performance in terms of final accuracy A_f . Especially in ImageNet-100, they surpass all other baselines in both A_f and A_a , as previously reported in the original papers. Second, replacing BN with CN does not show consistent improvement across datasets and CIL methods. For instance, while EEIL+CN produces a considerable performance gain in ImageNet-100, CN fails to outperform BN when combined with more recent CIL algorithms such as PODNet and AFC. On the other hand, our TBBN consistently improves performance in A_f and A_a except when applied on PODNet in CIFAR-100. Following the observation in the previous experiment, we again confirm that the performance gain of TBBN is mostly due to reduction in F (*i.e.* enhanced stability) while maintaining or increasing A_l (*i.e.* enhanced plasticity). In all, applying TBBN to the SOTA baselines achieves new SOTA performance in both datasets. We plotted a graph of these experiments in the supplementary materials, and we could confirm that TBBN enhances the average accuracy at every step of training compared to BN, while CN shows performance loss for some PODNet and AFC.

Applying TBBN to other architectures. To check the applicability of TBBN to other architectures, we selected four CNN-based architectures (*e.g.*, ResNet-34 [13], ShuffleNet-V2 [43], MobileNet-V2 [31], and MnasNet (x0.5) [35])

Table 2. Experimental results for various representative offline CIL methods. Bold indicates the best performance in each metric.

Method		CIFAR-100 w/ ResNet-32				ImageNet-100 w/ ResNet-18			
		$A_f(\uparrow)$	$A_a(\uparrow)$	$F(\downarrow)$	$A_l(\uparrow)$	$A_f(\uparrow)$	$A_a(\uparrow)$	$F(\downarrow)$	$A_l(\uparrow)$
FT	+BN	35.41	53.88	43.38	78.79	39.40	59.60	48.02	87.42
	+CN	35.06	54.43	43.82	80.64	41.96	60.02	45.32	87.28
	+TBBN	38.46	56.17	41.90	80.36	43.20	61.69	43.62	87.82
EEIL	+BN	39.82	55.25	39.40	79.22	40.06	61.15	47.78	87.84
	+CN	39.98	55.09	39.31	79.29	42.48	61.43	44.44	86.92
	+TBBN	41.93	57.53	37.80	79.93	45.18	63.48	42.66	87.84
LUCIR	+BN	38.06	54.20	32.35	70.41	42.26	63.82	41.68	83.94
	+CN	38.07	55.60	33.78	71.85	40.44	61.44	42.04	83.48
	+TBBN	41.45	56.13	29.23	70.68	43.72	64.36	40.18	83.90
PODNet	+BN	38.10	52.95	14.70	52.58	49.05	65.41	22.40	69.99
	+CN	34.80	50.26	15.69	50.52	46.20	62.91	23.66	68.50
	+TBBN	37.90	52.98	13.90	51.78	49.30	65.70	21.85	69.76
SSIL	+BN	41.34	53.00	15.64	56.02	49.56	65.79	21.20	69.94
	+CN	40.74	52.38	14.60	54.44	50.58	64.81	18.56	65.04
	+TBBN	43.80	54.28	15.12	59.37	51.30	66.51	19.58	70.64
AFC	+BN	39.90	53.93	33.17	73.10	52.50	67.53	19.70	72.22
	+CN	37.50	51.16	33.40	70.94	48.00	65.21	20.10	70.68
	+TBBN	41.30	57.31	32.89	73.57	54.00	68.68	19.00	73.22

which contain the BN layer as a default, and conducted experiments by replacing BN with TBBN in the 10 tasks scenario using the ImageNet-100 dataset. Table 3 demonstrates that our TBBN can be successfully applied to various types of architecture, increasing both A_f and A_a concurrently.

Table 3. Experimental results with varying backbone architectures on ImageNet-100. Bold indicates the best performance.

Method		ResNet-34		ShuffleNet-V2		MobileNet-v2		MnasNet (x0.5)	
		$A_f(\uparrow)$	$A_a(\uparrow)$	$A_f(\uparrow)$	$A_a(\uparrow)$	$A_f(\uparrow)$	$A_a(\uparrow)$	$A_f(\uparrow)$	$A_a(\uparrow)$
FT	+BN	41.38	62.00	34.74	55.37	38.86	57.43	36.24	54.96
	+TBBN	46.72	64.76	35.82	58.17	42.66	60.41	40.94	57.09
SSIL	+BN	51.18	65.27	43.86	59.69	47.72	62.25	45.10	59.81
	+TBBN	52.30	66.36	44.72	60.07	48.62	63.37	47.06	60.24

Results on dissimilar tasks. Several benchmark datasets (*e.g.*, CIFAR-10/-100 and ImageNet dataset) have been used to evaluate a CIL method. The CIL scenario made from random class orderings of these datasets has *small domain shifts* due to high similarity between tasks [24]. In this regard, we believe that the CIL scenario using CIFAR-100 and ImageNet-100 datasets is not the best setting to check whether the proposed normalization layer is suitable for the CIL. Therefore, we design the CIL scenario, which has *large domain shifts*, with five distinct datasets such as CIFAR-10 [19], SVHN [27], STL-10 [8], MNIST [11] and Fashion-MNIST [41]. We divide 10 classes of each dataset into 2 tasks and construct the CIL scenario consisting of 10 tasks (2 tasks \times 5 datasets). Note that task ordering is randomly shuffled by random seed and more details for experimental settings are proposed in the supplementary materials.

In Table 4, we observe that applying CN is sometimes harmful and only achieves minor improvement for FT and LUCIR. On the other hand, we again confirm that TBBN constantly makes enhancement of A_f and A_a by diminishing F but well maintaining A_l , for all methods. This result underscores the robustness of our TBBN even for the difficult CIL scenario consisting of dissimilar tasks.

Table 4. Experimental results with dissimilar tasks using ResNet-18. Bold indicates the best performance in each metric.

Method		$A_f(\uparrow)$	$A_a(\uparrow)$	$F(\downarrow)$	$A_t(\uparrow)$
FT	+BN	59.60	65.85	31.48	91.08
	+CN	61.41	71.00	29.90	91.40
	+TBBN	63.42	71.32	28.05	91.47
EEIL	+BN	62.04	72.28	29.46	91.50
	+CN	61.73	72.40	29.27	91.00
	+TBBN	64.88	72.80	26.67	91.55
LUCIR	+BN	62.42	72.91	28.09	90.35
	+CN	62.99	72.81	26.70	89.49
	+TBBN	63.64	73.72	26.73	90.23
SSIL	+BN	66.15	70.62	10.13	72.59
	+CN	65.64	70.56	9.16	69.66
	+TBBN	67.00	71.77	9.74	72.86

5.2. Qualitative Analysis

Reduction in biased predictions. For a more detailed analysis for the effect of TBBN, we investigated a type of misclassification, by following the experiment proposed in [3], and the analysis results for finetuning with the ImageNet-100 dataset (10-tasks) are reported in Figure 4. This figure presents the number of four types of misclassification. For example, $C \rightarrow P$ stands for the number of misclassification current task’s data to another class in previous tasks ($t = 1, \dots, 9$). Among four types, $P \rightarrow C$, which is the number of misclassification of the previous task’s data toward the current task’s class, is called as *biased prediction*, knowing as the major cause of degrading performance in CIL [1, 39]. As already reported in previous works [3, 39], FT + BN seriously suffers from the biased prediction, resulting in the largest number of total misclassification. In the case of both CN variants and TBBN, we confirm that the gain of reducing the total number of misclassifications comes from diminishing the number of biased predictions ($P \rightarrow C$), Especially, our TBBN significantly alleviates the biased prediction than CN variants. As a result, FT + TBBN achieves the lowest number of total misclassifications despite the increase of $P \rightarrow P$ caused by the reduced biased prediction.

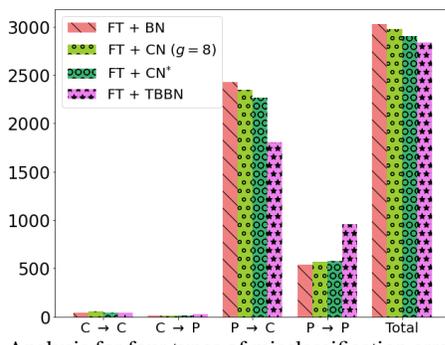


Figure 4. Analysis for four types of misclassification errors (y -axis) after training the model until the final task ($t = 10$).

Ablation studies. Table 5 presents ablation study for TBBN using the CIFAR-100 dataset (10 tasks). First, both Case 1 and Case 2 represent the ablation study for task-balanced mean and variance. We confirm that only calculating task-

balanced running mean μ and variance σ^2 in (5) for test phase (Case 1) or using a task-balanced running mean $\tilde{\sigma}$ and variance $\tilde{\sigma}^2$ for training phase (Case 2) affect final performance negatively, due to mismatch of values (μ, σ^2) for normalization in training and test phases. Case 3 shows the experimental result where calculating task-balanced mean and variance is only applied. However, as we already confirmed in Figure 2 of the Section 4.1, only considering task-balanced mean and variance cannot make a difference in performance with the original BN. The opposite case (Case 4) where γ and β are trained in a task-balanced manner also shows a similar tendency, demonstrating the performance enhancement by TBBN can be obtained only when the proposed components are used together.

Additionally, we compare GPU memory cost, training time and inference time of each case. Because of both reshape and repeat operations consisting of TBBN, TBBN requires maximum $\times 1.6$ of GPU memory and takes more $\times 1.3$ of training time than the original BN. However, note that there is no additional cost at test time because TBBN works to be identical to the original BN.

Table 5. Ablation study using CIFAR-100 with ResNet-32.

Method	Task-balanced μ and σ^2 ?		Task-balanced γ and β ?	$A_f(\uparrow)$	$A_a(\uparrow)$	Peak Memory	Train Time (total)	Test Time (per batch)
	Train	Test						
TBBN	✓	✓	✓	38.46	56.17	2525MiB ($\times 1.6$)	6.0h ($\times 1.3$)	0.6s ($\times 1.0$)
Case 1	✗	✓	✓	29.78	46.99			
Case 2	✓	✗	✓	0.93	9.69			
Case 3	✓	✓	✗	35.45	53.15			
Case 4	✗	✗	✓	36.21	53.25			
BN	✗	✗	✗	35.71	53.88	1545MiB	4.6h	0.6s

6. Concluding Remarks

We propose a simple but effective method, called Task-Balance Batch Normalization, for exemplar-based CIL. Starting from an analysis of the problem of the original BN, namely, the biased mean and variance calculation toward the current task, we devise a novel method to calculate the task-balanced mean and variance for normalization. Furthermore, we propose a method for the task-balanced training of parameters for affine transformation. From extensive experiments with CIFAR-100, ImageNet-100 and dissimilar tasks, we demonstrate that our TBBN can be easily applied to various existing CIL methods, further improving their performance.

7. Acknowledgement

This work was partly done while Sungmin Cha did a research internship at Advanced ML Lab, LG AI Research. The work was also supported in part by NRF grant [NRF-2021R1A2C2007884], IITP grant [No.2021-0-01343, No.2021-0-02068, No.2022-0-00113, No.2022-0-00959] funded by the Korean government (MSIT), and SNU-LG AI Research Center.

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 (CVPR)*, pages 844–853, 2021.
- [2] Rahaf Aljundi, Eugene Belilovsky, Tinne Tuytelaars, Laurent Charlin, Massimo Caccia, Min Lin, and Lucas Page-Caccia. Online continual learning with maximal interfered retrieval. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 32, 2019.
- [3] Eden Belouadah and Adrian Popescu. Il2m: Class incremental learning with dual memory. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 583–592, 2019.
- [4] Nils Bjorck, Carla P Gomes, Bart Selman, and Kilian Q Weinberger. Understanding batch normalization. *Advances in Neural Information Processing Systems (NeurIPS)*, 31, 2018.
- [5] John Bronskill, Jonathan Gordon, James Requeima, Sebastian Nowozin, and Richard Turner. Tasknorm: Rethinking batch normalization for meta-learning. In *International Conference on Machine Learning (ICML)*, pages 1153–1164, 2020.
- [6] Francisco M Castro, Manuel J Marín-Jiménez, Nicolás Guil, Cordelia Schmid, and Karteek Alahari. End-to-end incremental learning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 233–248, 2018.
- [7] Sungmin Cha, Hsiang Hsu, Taebaek Hwang, Flavio Calmon, and Taesup Moon. CPR: Classifier-projection regularization for continual learning. In *International Conference on Learning Representations (ICLR)*, 2021.
- [8] Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTAT)*, pages 215–223, 2011.
- [9] Matthias Delange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Ales Leonardis, Greg Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(7):3366–3385, 2021.
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255. Ieee, 2009.
- [11] Li Deng. The MNIST database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6):141–142, 2012.
- [12] 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, 2020.
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- [14] Elad Hoffer, Itay Hubara, and Daniel Soudry. Train longer, generalize better: closing the generalization gap in large batch training of neural networks. *Advances in Neural Information Processing Systems (NeurIPS)*, 30, 2017.
- [15] 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 (CVPR)*, pages 831–839, 2019.
- [16] Sergey Ioffe. Batch renormalization: Towards reducing mini-batch dependence in batch-normalized models. *Advances in Neural Information Processing Systems (NeurIPS)*, 30, 2017.
- [17] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning (ICML)*, pages 448–456, 2015.
- [18] Minsoo Kang, Jaeyoo Park, and Bohyung Han. Class-incremental learning by knowledge distillation with adaptive feature consolidation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16071–16080, 2022.
- [19] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [20] Janghyeon Lee, Hyeong Gwon Hong, Donggyu Joo, and Junmo Kim. Continual learning with extended kronecker-factored approximate curvature. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9001–9010, 2020.
- [21] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12):2935–2947, 2017.
- [22] Ping Luo, Ruimao Zhang, Jiamin Ren, Zhanglin Peng, and Jingyu Li. Switchable normalization for learning-to-normalize deep representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(2):712–728, 2019.
- [23] Zheda Mai. *Online Continual Learning in Image Classification*. PhD thesis, University of Toronto (Canada), 2021.
- [24] Marc Masana, Xialei Liu, Bartłomiej Twardowski, Mikel Menta, Andrew D Bagdanov, and Joost van de Weijer. Class-incremental learning: Survey and performance evaluation on image classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [25] Martial Mermillod, Aurélie Bugaiska, and Patrick Bonin. The stability-plasticity dilemma: Investigating the continuum from catastrophic forgetting to age-limited learning effects. *Frontiers in Psychology*, 4:504, 2013.
- [26] 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)*, pages 3513–3522, 2021.
- [27] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011.
- [28] German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71, 2019.
- [29] Quang Pham, Chenghao Liu, and Steven HOI. Continual normalization: Rethinking batch normalization for online con-

- tinual learning. In *International Conference on Learning Representations (ICLR)*, 2022.
- [30] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. iCaRL: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2001–2010, 2017.
- [31] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. MobileNetV2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4510–4520, 2018.
- [32] Shibani Santurkar, Dimitris Tsipras, Andrew Ilyas, and Aleksander Mądry. How does batch normalization help optimization? In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2488–2498, 2018.
- [33] Dongsub Shim, Zheda Mai, Jihwan Jeong, Scott Sanner, Hyunwoo Kim, and Jongseong Jang. Online class-incremental continual learning with adversarial shapley value. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 9630–9638, 2021.
- [34] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826, 2016.
- [35] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V Le. MnasNet: Platform-aware neural architecture search for mobile. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2820–2828, 2019.
- [36] Mingxing Tan and Quoc Le. EfficientNet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning (ICML)*, pages 6105–6114, 2019.
- [37] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016.
- [38] Ximei Wang, Ying Jin, Mingsheng Long, Jianmin Wang, and Michael I Jordan. Transferable normalization: Towards improving transferability of deep neural networks. *Advances in Neural Information Processing Systems (NeurIPS)*, 32, 2019.
- [39] 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 (CVPR)*, pages 374–382, 2019.
- [40] Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–19, 2018.
- [41] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- [42] Shipeng Yan, Jiangwei Xie, and Xuming He. DER: Dynamically expandable representation for class incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3014–3023, 2021.
- [43] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. ShuffleNet: An extremely efficient convolutional neural network for mobile devices. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6848–6856, 2018.
- [44] Bowen Zhao, Xi Xiao, Guojun Gan, Bin Zhang, and Shu-Tao Xia. Maintaining discrimination and fairness in class incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13208–13217, 2020.