This CVPR 2021 workshop paper is the Open Access version, provided by the Computer Vision Foundation.
 Except for this watermark, it is identical to the accepted version;
 the final published version of the proceedings is available on IEEE Xplore.

Dual-Teacher Class-Incremental Learning With Data-Free Generative Replay

Yoojin Choi, Mostafa El-Khamy SoC R&D, Samsung Semiconductor Inc. San Diego, CA 92121, USA {yoojin.c,mostafa.e}@samsung.com

Abstract

This paper proposes two novel knowledge transfer techniques for class-incremental learning (CIL). First, we propose data-free generative replay (DF-GR) to mitigate catastrophic forgetting in CIL by using synthetic samples from a generative model. In the conventional generative replay, the generative model is pre-trained for old data and shared in extra memory for later incremental learning. In our proposed DF-GR, we train a generative model from scratch without using any training data, based on the pre-trained classification model from the past, so we curtail the cost of sharing pre-trained generative models. Second, we introduce dual-teacher information distillation (DT-ID) for knowledge distillation from two teachers to one student. In CIL. we use DT-ID to learn new classes incrementally based on the pre-trained model for old classes and another model (pre-)trained on the new data for new classes. We implemented the proposed schemes on top of one of the state-ofthe-art CIL methods and showed the performance improvement on CIFAR-100 and ImageNet datasets.

1. Introduction

The natural learning process of humans is incremental, implying that we incrementally learn new knowledge, as we explore the world and observe new data over time. Many of real-world vision applications desire a similar incremental learning property, since their target tasks usually grow progressively by increasing demands of users, and the training data are also often collected incrementally. However, most of conventional supervised learning methods do not adapt well to this situation, since they are developed under the assumption that all the training data for learning are provided and used at once.

Incremental learning is a learning paradigm that enables a model to acquire new knowledge from new data continually, instead of training it once for all [24]. A baseline approach to incremental learning is to fine-tune a pre-trained Jungwon Lee System LSI, Samsung Electronics South Korea

jungwon2.lee@samsung.com



Figure 1: Class-incremental learning with dual-teacher information distillation (DT-ID) and data-free generative replay (DF-GR). DT-ID is used to transfer knowledge from two teachers, which are trained for old and new classes, respectively, to one student. We also propose DF-GR to train a generative model for old classes from scratch without using any training data, given the pre-trained classification model for old classes. The synthetic samples from DF-GR are used for data-free KD to mitigate catastrophic forgetting in CIL.

model on new data when they become available, but naive fine-tuning suffers from severe performance degradation on the old tasks that the model already learned before, which is called *catastrophic forgetting* [10]. Catastrophic forgetting is caused by over-fitting to the new data when the past data are not available and cannot be used during the incremental training stages.

There have been substantial efforts devoted to mitigate catastrophic forgetting in incremental learning.

Replay methods. Reserving some of the original training data of past tasks for future learning is a naive but effective approach to reduce catastrophic forgetting, and it was used in many of the previous works [4, 13, 20, 26, 31]. In [31], a subset of the reserved exemplars was used as a validation set to correct the bias towards a new task due to data imbalance. On the other hand, instead of storing raw-format samples, *generative replay (GR)* proposes utilizing pre-trained generative models to reproduce synthetic samples for past tasks and use them for pseudo-rehearsal [14, 19, 29, 32].



Figure 2: Timeline of our proposed dual-teacher CIL (DT-CIL). We curtail the extra burden of storing any past data or pretrained generative models by using DF-GR for pseudo-rehearsal of the past data. DF-GR can also be used to augment data for new classes, if not enough data are accessible for them.

Parameter-based regularization methods. The methods in this category mainly focus on estimating the importance of each weight in a pre-trained model and add more penalty to the distortion of the significant weights when (re-)training the model for new data [2, 5, 16, 23, 33]. The difference among these methods can be found in the way to compute the importance of weights. For example, in [16], the Fisher information matrix was used to measure the importance. In [23], variational continual learning introduced a prior distribution on the weights to learn the importance of the weights via a variational learning framework.

Distillation-based regularization methods. The methods in this category utilize knowledge distillation (KD) [12] to consolidate the previous knowledge in a pre-trained model, while the model is (re-)trained for new data. A distillationbased method was first introduced to incremental learning in LwF [18] to transfer the previous knowledge from a pretrained model to a new model by matching their (softmax) outputs on new data. Then, iCaRL [26] improved the LwF method for CIL with a small number of exemplars reserved for old classes. Moreover, iCaRL proposed nearest-meanof-exemplars classification and prioritized construction of exemplars. The distillation-based method was improved in [4] by sophisticated data augmentation with reserved samples and class-balanced fine-tuning. Recently, it was further refined in LUCIR [13] with cosine-similarity-based classification, less-forget constraint, and inter-class separation. In [9], attention maps were utilized for distillation.

The main contribution of this paper is two-fold. First, we propose *data-free generative replay (DF-GR)* to augment limited (old) training data with synthetic samples from a generative model. In particular, given a pre-trained model for the past tasks, the generative model is trained from scratch without using any training data, so we call it a data-free method. This is different from conventional generative replay methods [19, 29, 32], where auxiliary generative models are trained along with the main objective models by using the original training data, and the pre-trained generative models are shared for later incremental learning. Training good generative models is difficult and time-consuming. Moreover, it is redundant work in the view of current model

trainers who may not care much of later incremental learning. In DF-GR, the generative model is trained by the current model trainer who wants to adapt the model for a new task, without accessing any previous data. Second, we introduce *dual-teacher information distillation (DT-ID)* for knowledge distillation from two teachers to one student. In incremental learning, we use the model (pre-)trained on the new data as the second teacher for a new task, in addition to the previous one-teacher distillation method that uses the past model as the (first) teacher for old tasks. We adopt the information distillation (ID) method [1] and alter the original one-teacher method to a dual-teacher method.

We investigate the usefulness of the proposed DT-ID and DF-GR methods in class-incremental learning (CIL). The goal of CIL is to learn a unified classifier that can recognize all the classes learned so far, when the training data for unseen classes arrive incrementally. Our proposed dualteacher CIL (DT-CIL) employs both DT-ID and DF-GR, as shown in Figure 1. In Figure 2, we illustrate the timeline of DT-CIL. We implemented the proposed schemes on top of one of the state-of-the-art CIL methods, LUCIR [13]. In our experimental results on CIFAR-100 and ImageNet datasets. we show that each proposed component provides performance improvement over the baseline LUCIR method. We also show the potential of the proposed scheme in the datalimited CIL scenario, where the available training data are limited not only for old classes, but also for new classes, due to the cost of data sharing or to preserve data privacy.

2. Data-free generative replay (DF-GR)

Let us assume that we are given a pre-trained classification model t (called teacher) that estimates the probability distribution of class y for input x. Let g be the conditional generator that we train to produce the synthetic data similar to the training data used to train the teacher t. The conditional generator takes a random noise vector z and a label (condition) y to produce a labeled sample. Let p(z) be the random noise distribution, and let p(y) be the label distribution over classes C. We employ the following two losses to train the conditional generator without any training data.



Figure 3: Zero-shot learning of a conditional generator for DF-GR. We match the statistics stored in the batch normalization (BN) layers of the teacher (which were computed for the original data) and the statistics computed for the generator output at the same BN layers of the teacher.

Cross-entropy (CE) loss. We use the teacher t as a fixed discriminator to criticize the labels of the synthetic samples from the conditional generator. We define the cross-entropy loss between the label fed to the generator and the softmax output from the teacher for the generated sample as below:

$$L_{\rm CE}(g|t) = \mathbb{E}_{p(z)p(y)}[\mathcal{H}(y, t(g(z, y)))], \tag{1}$$

where \mathcal{H} denotes the cross-entropy, and the label y is one-hot encoded in \mathcal{H} .

Batch-normalization statistics (BNS) loss [6]. Each batch normalization (BN) layer in the pre-trained teacher tstores the mean and variance of the layer input, which we can utilize as a proxy to verify that the generator output is similar to the original training data. We use the Kullback-Leibler (KL) divergence of two Gaussian distributions to match the statistics (mean and variance) stored in the BN layers of the teacher (which were obtained when trained with the original data) and the statistics computed for the generator output at the same BN layers of the teacher. Let $\mu_{l,c}$ and $\sigma_{l,c}^2$ be the mean and variance stored in batch normalization layer l of the teacher for channel c. Let $\hat{\mu}_{l,c}(g)$ and $\hat{\sigma}_{l,c}^2(g)$ be the corresponding mean and variance computed based on the synthetic samples from generator g. The batch-normalization statistics (BNS) loss is given by

$$L_{\text{BNS}}(g|t) = \sum_{l,c} \mathcal{D}_{\mathcal{N}}((\hat{\mu}_{l,c}(g), \hat{\sigma}_{l,c}^2(g)), (\mu_{l,c}, \sigma_{l,c}^2)), \quad (2)$$

where $\mathcal{D}_{\mathcal{N}}$ is the KL divergence of two Gaussians, i.e.,

$$\mathcal{D}_{\mathcal{N}}((\hat{\mu}, \hat{\sigma}^2), (\mu, \sigma^2)) = \frac{(\hat{\mu} - \mu)^2 + \hat{\sigma}^2}{2\sigma^2} - \log\frac{\hat{\sigma}}{\sigma} - \frac{1}{2}$$

From (1) and (2), we propose the following objective for zero-shot learning of a conditional generator (see Figure 3):

$$\min_{g} \{ L_{\text{CE}}(g|t) + L_{\text{BNS}}(g|t) \}.$$
(3)

In [6], the BNS loss was introduced to train non-conditional generators for data-free KD, in particular, for data-free network quantization. In this paper, we utilize the BNS loss together with the CE loss to train a conditional generator

Table 1: Summary of notations for dual-teacher CIL.

Time	Classes		Da	ita	Teachers (pr	Student	
	Old	New	Old*	New	for C_0^{i-1}	for C_i	for C_0^i
i	C_0^{i-1}	C_i	R_0^{i-1}	D_i	$t_0 \equiv f_{i-1}$	$t_1 \equiv h_i$	$s \equiv f_i$
					*]	Reserved ex	emplars.

without data. The conditional generator is then used in CIL for DF-GR, which will be elaborated in Section 3.2.

In Figure 4, we present some synthetic images from DF-GR, given a pre-trained ResNet-18 model on the ImageNet dataset as the teacher. The ResNet-18 model is pre-trained on the ImageNet data of 500 classes that we randomly select as the base task at time 0 for CIL (see Section 4). In the case of using the CE loss only, the generator fails to synthesize natural images of given classes. If we use the BNS loss only, then the conditional generator gets no feedback from the teacher (discriminator) for the label of its generated sample, so it turns into producing synthetic samples without labels. Once we utilize both CE and BNS losses, DF-GR produces labeled images that we can distinguish their classes.

3. Dual-teacher CIL (DT-CIL)

In CIL, we are given a sequence of classification tasks, denoted by T_i for $i \in \mathbb{Z}^*$, where \mathbb{Z}^* is the set of the nonnegative integers. Task T_i at time i is a classification task for a set of classes C_i , such that $C_i \cap C_j = \emptyset$ for all $i \neq j$, where \emptyset denotes an empty set. At time i = 0, we train a network for (base) task T_0 with base training data D_0 . At each time $i \ge 1$, we are given a new task T_i with a set of new training data D_i that belong to C_i , and we aim to learn the new task without forgetting the past tasks $T_0, T_1, \ldots, T_{i-1}$ that we already learned. We cannot revisit the old training data for the past tasks, unless we reserve a small number of samples, called exemplars. The reserved exemplars for task T_i are denoted by R_i .

Let f_i be the neural network that we use at time $i \in \mathbb{Z}^*$ for CIL. Each network f_i consists of a feature extractor ϕ_i and a classifier based on the extracted feature. For simplicity, we assume one-layer classifier, and let $W_0^i = \{w_c\}_{c \in C_0^i}$ be the set of classification weights used in f_i for classification among all the learned classes $C_0^i = \bigcup_{j=0}^i C_j$, where $W_i = \{w_c\}_{c \in C_i}$ is the classification weights newly introduced at time *i* for classes C_i .

3.1. Dual-teacher information distillation (DT-ID)

The information distillation (ID) scheme [1] is one of the state-of-the-art KD variants, which transfers knowledge from a teacher to a student by maximizing the mutual information at their intermediate layers. We adopt this method for *dual-teacher information distillation (DT-ID)* in CIL.

At time *i*, the first teacher for DT-ID is the model f_{i-1} from time i-1, which was pre-trained for old classes C_0^{i-1} .



(c) DF-GR trained with BNS loss only

Figure 4: Examples of the synthetic images from DF-GR, given a pre-trained ResNet-18 for 500 base classes from ImageNet.

We let h_i be the second teacher at time *i*, which is the model (pre-)trained on the new data at time *i* for new classes C_i . The student is the current model f_i that we train at time *i* for both old and new classes in C_0^i . For notational simplicity, we use $t_0 \equiv f_{i-1}, t_1 \equiv h_i$ and $s \equiv f_i$ to denote two teachers and the student, respectively (see Table 1).

For information distillation at some of intermediate layers, we select the same K intermediate layers of the teachers and the student, respectively—In practice, we select the intermediate layers with different resolutions (e.g., the layers just before down-sampling layers). Let $t_{0,k}$, $t_{1,k}$, and s_k for $1 \le k \le K$ be the feature maps from the k-th layer selected for DT-ID in the first teacher, the second teacher, and the student, respectively. Then, DT-ID aims to minimize

$$L_{\text{DT-ID}} = -\sum_{k=1}^{K} (\mathcal{I}(t_{0,k}, s_k) + \mathcal{I}(t_{1,k}, s_k)),$$

where \mathcal{I} denotes the mutual information (e.g., see [7, Section 2.3]). As proposed in [1], we take the variational lower bound of the mutual information and perform variational information maximization [3] with a Gaussian prior. Then, for each $n \in \{0, 1\}$, it can be shown that

$$-\mathcal{I}(t_{n,k}, s_k) \le \mathbb{E}[\mathcal{J}_{n,k}(t_{n,k}, s_k)] + X, \qquad (4)$$

where

$$\mathcal{J}_{n,k}(t,s) = \sum_{c,h,w} \frac{(t_{c,h,w} - \mu_{k,c,h,w}^n(s))^2}{2\sigma_{k,c}^2} + \log \sigma_{k,c},$$
(5)

for some constant X; $t_{c,h,w}$ is the scalar element of tensor t at channel c, height h, width w, and $\mu_{k,c,h,w}^n(s)$ is the output of the neural network μ_k^n at channel c, height h, width w,

when it takes *s* as its input— μ_k^n is the convolutional network used to transform the student feature maps into the teacher domain at each intermediate layer *k* selected for ID. We propose using a common variance $\sigma_{k,c}^2$ for both $n \in \{0, 1\}$ at each layer *k* and channel *c*, so we transfer information from two teachers without inclining towards either of them. We note that the detailed steps to derive (4) and (5) from variational information maximization can be found in [1, 3] and are omitted here due to the page limit.

Dual-teacher information distillation (DT-ID) loss. In practice, we take empirical expectation in (4) by using the feature maps obtained from all available data at time *i* and define the DT-ID loss as below, for $A_i = D_i \cup R_0^{i-1}$,

$$\bar{L}_{\text{DT-ID}}^{i} = \sum_{k=1}^{K} \mathbb{E}_{(x,y)\in A_{i}} [\mathcal{J}_{0,k}(f_{i-1,k}(x), f_{i,k}(x)) + \mathcal{J}_{1,k}(h_{i,k}(x), f_{i,k}(x))].$$
(6)

For the baseline CIL scheme, we adopt three key components of LUCIR [13], i.e., cosine-similarity-based classifier, cross-entropy loss, and less-forget constraint.

Cosine-similarity-based classifier. In [13], it was proposed for the classifier to use the cosine similarity instead of the conventional dot-product, to resolve the data imbalance problem, when only a small number of reserved exemplars are given for old classes. The cosine-similarity-based classifier yields, before applying the softmax function,

$$l_{i,c}(x) = \langle w_c, \phi_i(x) \rangle = \frac{w_c^T \phi_i(x)}{|w_c||\phi_i(x)|}, \quad c \in C_0^i, \quad (7)$$

for input x, where a^T is the transpose of a, |a| is the l^2 -norm

of a, and $\langle a, b \rangle$ is the cosine similarity between a and b. The network f_i computes the categorical probability distribution of input x over classes C_0^i by encoding the cosine-similarity scores with the softmax function.

Cross-entropy (CE) loss. The categorical cross-entropy loss is computed on the labeled training data for new classes and the reserved exemplars, if present, for old classes:

$$L_{CE}^{i} = \mathbb{E}_{(x,y)\in A_{i}}[\mathcal{H}(y, f_{i}(x))], \quad A_{i} = D_{i} \cup R_{0}^{i-1}, \quad (8)$$

where (x, y) is the pair of a training sample and its ground-truth label.

Less-forget (LF) constraint [13]. To maintain the previous knowledge learned in a pre-trained model, it was proposed to inherit and fix the weights for old classes from the pre-trained model and minimize the following LF loss:

$$L_{\rm LF}^i = -\mathbb{E}_{(x,y)\in A_i}[\langle \phi_{i-1}(x), \phi_i(x) \rangle]. \tag{9}$$

The LF loss constrains the change of the feature extractor, when (re-)training the model on new data so that it does not forget the knowledge for old classes.

By combining the DT-ID, CE and LF losses in (6), (8), and (9), the final objective of our DT-CIL is given by

$$\min_{\phi_i, W_i, \mu' \mathbf{s}, \sigma^2, \mathbf{s}} \{ \bar{L}^i_{\text{DT-ID}} + L^i_{\text{CE}} + \alpha_i L^i_{\text{LF}} \}, \qquad (10)$$

where $\alpha_i \ge 0$ is the factor to control less-forgetting.

3.2. DT-CIL with DF-GR

For DF-GR at time *i* of CIL, the pre-trained model f_{i-1} for old classes from time i - 1 is set to be the teacher *t* in (3), and a conditional generator $g = g_i$ is trained to reproduce the synthetic samples for old classes C_0^{i-1} by (3). We emphasize that at each time *i*, we train a new generator g_i , from scratch without using any training data, based on the previous model f_{i-1} , and no pre-trained generators are transferred to the future (see Figure 2). The synthetic samples from g_i are used in CIL as follows:

- First, we include the synthetic samples when computing the DT-ID and LF losses in (6) and (9), i.e., we add the synthetic samples to the available data A_i in (6) and (9).
- Second, we add the data-free KD loss (see (12)) for the synthetic samples to our DT-CIL objective in (10), which is elaborated below.
- We do not compute the CE loss in (8) with the synthetic samples, since the ground-truth labels for the synthetic samples are unknown although they are machine-labeled in conditional generation.

Data-free KD with weight imprinting. Note that the past model f_{i-1} yields the softmax output for old classes only, while the current model f_i that we train at time *i* produces the softmax output for both old and new classes. For KD from f_{i-1} to f_i , we need to match the number of output

classes. In [18, 31], it was proposed to match the softmax output for old classes only in KD. In this paper, we propose extending the classification layer of f_{i-1} to cover the new classes by *weight imprinting* [25] and matching the softmax output for both old and new classes. We will show the gain of weight imprinting in our experiments (see Figure 6).

For weight imprinting, we collect the output of the feature extractor ϕ_{i-1} for every training sample of each new class and use their average as the classification weight of that class. If we use a cosine-similarity-based classifier (see (7)), we normalize the features before taking their average:

$$w_c = \mathbb{E}_{(x,y)\in D_i(c)} \left[\frac{\phi_{i-1}(x)}{|\phi_{i-1}(x)|} \right], \quad c \in C_i,$$
(11)

where $D_i(c)$ is the samples of class $c \in C_i$ in D_i . In (11), weight imprinting basically finds the weight that maximizes the average cosine similarity to the features extracted from the available training samples for each class.

Let f_{i-1} be the past model with an extended classifier after weight imprinting for new classes. Then, the data-free KD (DF-KD) loss using DF-GR is given by

$$L^{i}_{\text{DF-KD}} = \mathbb{E}_{p(z)p_{i}(y)}[\mathcal{H}(\tilde{f}_{i-1}(g_{i}(z,y)), f_{i}(g_{i}(z,y)))], \quad (12)$$

where $p_i(y)$ is the label distribution, for which we use the uniform distribution over all the past classes C_0^{i-1} at time *i*. The DF-KD loss is finally added to (10) for DT-CIL.

4. Experiments

We perform experiments on our proposed DT-CIL with DT-ID and DF-GR for CIFAR-100 [17] and ImageNet [28] 2012 datasets. For ImageNet, we consider two cases. First, we evaluate CIL on 100 classes randomly selected from the full ImageNet data, which is called ImageNet-Subset. Second, we evaluate CIL on all 1000 classes of the ImageNet dataset, which is called ImageNet-Full. For CIL on CIFAR-100 and ImageNet-Subset, we split the classes into 6 tasks, where the base task at time 0 has 50 classes, and each of the following 5 tasks for time $i \in \{1, 2, 3, 4, 5\}$ has 10 classes. For ImageNet-Full, we split the classes into the base task of 500 classes and the following 5 tasks of 100 classes in each.

4.1. Implementation details

We use 32-layer and 18-layer ResNets [11] (ResNet-32 and ResNet-18) for CIFAR-100 and ImageNet, respectively. We adopt the cosine-similarity-based classifier in (7), while we remove the last rectified linear unit (ReLU) from the original ResNet feature extractor, so the features can be any real values. We mostly follow the training hyper-parameters suggested in LUCIR [13], on top of which we additional implement our DT-ID and DF-GR methods. Due to the page limit, we defer more details to the supplementary materials.

Table 2: Notations used in experimental results.

N_D	# training data (or reserved exemplars) for each new class						
N_R	# reserved exemplars for each old class						
N_t	# teachers,	 The baseline one-teacher method The proposed dual-teacher method using DT-ID 					
Ng	# generators,	0: No DF-GR 1: Single DF-GR for old classes 2: Dual DF-GR for both old and new classes					

For the conditional generator of DF-GR, we adopt the U-Net [27] architecture used in [30]. We replace batch normalization with conditional batch normalization [8,21] for conditional generation. To train the conditional generator for (3), we employ Adam optimizer [15] with momentum 0.5 and learning rate 10^{-3} . We use 80 and 45 epochs of batch size 256 and 100 for CIFAR-100 and ImageNet, respectively, where each epoch has 50, 128, and 1280 batches for CIFAR-100, ImageNet-Subset, and ImageNet-Full, respectively. The trained generator is then used to compute the DF-KD loss in (12). In each batch to compute the DF-KD loss, we generate new synthetic samples, while we continue updating the generator for (3).

For DT-ID, we choose the layers at the end of each group of residual blocks, except the first group. Each network μ_k^n in (5) consists three convolutional layers of kernel size 1×1 , where we double the channel size in the first two convolutional layers and reduce it back to the original size at the last convolution layer to compare its output to the feature map from the teacher. Batch normalization and ReLU activation are included in the first two convolutional layers. To ensure $\sigma_{k,c}^2$ to be always positive, we use a proxy $\omega_{k,c} = \log \sigma_{k,c}^2$ for gradient descent. The initial value of $\omega_{k,c}$ is set to be 0.

4.2. Evaluation metrics

Average accuracy. Let $A_i(c)$ for $c \in C_0^i$ be the accuracy on the test data of class c at time i. The average accuracy at time i is the average of $A_i(c)$ for all seen classes $c \in C_0^i$.

Average forgetting [5]. The forgotten knowledge for an old class $c \in C_0^{i-1}$ at time $i \ge 1$ is measured by

$$F_i(c) = \max_{j \le k \le i-1} \{A_k(c) - A_i(c)\}, \ c \in C_j, \ 0 \le j \le i-1,$$

and the average forgetting at time *i* is the average of $F_i(c)$ for all old classes $c \in C_0^{i-1}$.

4.3. Experimental results

First, Table 2 summarizes the notations used in our experimental results. In Figure 5, we show the results of CIL on CIFAR-100, ImageNet-Sub, and ImageNet-Full, in the conventional CIL scenario, where the training data for new classes in the datasets are all available. In particular, we



Figure 5: Average accuracy and average forgetting in CIL on CIFAR-100, ImageNet-Subset, and ImageNet-Full. For ours, we used both DT-ID and DF-GR (N_t , $N_g = 2, 1$).

evaluate two cases of $N_R = 0$ and $N_R = 20$. For the results of ours in Figure 5, we use both DT-ID and DF-GR, i.e., $N_t, N_q = 2, 1$. We compare ours with the baseline LUCIR [13] and iCaRL [26]. For iCaRL, we followed the implementation in [13], which is slightly different but better than the original one. We show the results for iCaRL in two cases when we use the nearest-mean-of-exemplars classifier (iCaRL-NME) and when we use the softmax scores for classification (iCaRL-CNN). In Figure 5, observe that our proposed DT-CIL method with DT-ID and DF-GR provides gains in accuracy over the baseline LUCIR [13]. In particular, the gain is significant in case that we have no reserved old exemplars ($N_R = 0$). From the plots for average forgetting, one can see that the proposed methods help to reduce catastrophic forgetting when reserved old exemplars are limited.

Table 3 provides the results of an ablation study for each component, DT-ID and DF-GR, of our proposed method for CIL on ImageNet-Full. First, $N_t, N_g = 1, 0$ implies the baseline LUCIR [13]. Second, using DT-ID on top of LU-CIR corresponds to the case of $N_t, N_g = 2, 0$. Third, using DF-GR with LUCIR is denoted by $N_t, N_g = 1, 1$. Finally, in the case of $N_t, N_g = 2, 1$, we use both DT-ID and DF-GR

Time				Average accuracy (%) for all seen classes					Average forgetting (%) for old classes					
(# classes)			0 (500)	1 (+100)	2 (+100)	3 (+100)	4 (+100)	5 (+100)	0 (500)	1 (+100)	2 (+100)	3 (+100)	4 (+100)	5 (+100)
LUCIR [13]*	1.3k, 0	1,0	75.04	62.02	53.91	48.36	44.70	41.68	0.00	16.86	25.82	31.35	34.51	37.29
Ours	1.3k, 0 1.3k, 0 1.3k, 0	2, 0 1, 1 2, 1	75.04 75.04 75.04	62.62 63.41 64.23	54.63 55.71 56.77	49.05 50.13 51.40	45.11 46.12 47.40	41.83 43.01 44.46	0.00 0.00 0.00	16.02 14.70 13.78	24.83 22.70 21.50	30.29 27.83 26.42	33.69 31.03 29.55	36.73 33.66 31.86
LUCIR [13]*	1.3k, 20	1,0	75.04	70.64	65.81	61.42	58.10	56.05	0.00	5.25	10.48	14.38	17.04	18.83
Ours	1.3k, 20 1.3k, 20 1.3k, 20	2, 0 1, 1 2, 1	75.04 75.04 75.04	70.71 70.92 71.05	66.41 66.24 66.97	62.20 62.39 62.93	59.18 59.33 59.90	57.04 57.07 57.89	0.00 0.00 0.00	4.72 4.06 3.64	8.85 8.53 7.38	12.43 11.67 10.58	14.74 14.02 12.60	16.58 15.76 14.16
LUCIR [13]* Ours	1.3k, 20 1.3k, 20 1.3k, 20 1.3k, 20	1, 0 2, 0 1, 1 2, 1	75.04 75.04 75.04 75.04	70.64 70.71 70.92 71.05	65.81 66.41 66.24 66.97	61.42 62.20 62.39 62.93	58.10 59.18 59.33 59.90	56.05 57.04 57.07 57.89	0.00 0.00 0.00 0.00	5 4 4 3	.25 .72 .06 .64	.25 10.48 .72 8.85 .06 8.53 .64 7.38	.25 10.48 14.38 .72 8.85 12.43 .06 8.53 11.67 .64 7.38 10.58	.25 10.48 14.38 17.04 .72 8.85 12.43 14.74 .06 8.53 11.67 14.02 .64 7.38 10.58 12.60 * Reproduce

Table 3: Average accuracy and average forgetting at each time of CIL on ImageNet-Full in the conventional CIL scenario.

DF-KD with weight imprinting (proposed) Match the softmax output for old classes only (conventional) • Use the target label of conditional generation CIFAR-100 CIFAR-100 80 80 (%) (%) accuracy 70 accuracy 60 60 Average Average 50 50 40∟ 50 40 L 50 80 90 100 60 80 90 100 60 70 70 # classes # classes (a) $N_D, N_R = 500, 0$ (b) $N_D, N_R = 500, 20$

Figure 6: Comparison of three different methods of utilizing the synthetic samples from DF-GR in CIL on CIFAR-100. We used both DT-ID and DF-GR ($N_t, N_g = 2, 1$).

(see Table 2 for the notations). Table 3 shows that each proposed component provides performance improvement over the baseline in both cases when we have a small number of reserved exemplars (20 per each old class) and when we have no reserved exemplars for old classes.

In Figure 6, we compare three different methods of utilizing the synthetic samples from DF-GR in CIL. The first method is our proposed DF-KD with weight imprinting to match the softmax output for both old and new classes. The second method is to match the softmax output for old classes only, which was used in [18,31]. The third method is to use the target label of conditional generation directly to compute the cross-entropy $\mathbb{E}_{p(z)p_i(y)}[\mathcal{H}(y, f_i(g_i(z, y)))]$, instead of (12). We observe that the proposed DF-KD with weight imprinting outperforms the other methods.

In Figure 7, we investigate the impact of the number of reserved exemplars for old classes. In particular, we show the average accuracy for 80 and 100 seen classes at time 3 and 5 of CIL on CIFAR-100, respectively. Recall that the major reason of proposing DF-GR is to curtail the burden of storing any past data (exemplars) or pre-trained generative models. We observe that the gain of using DF-GR is promi-

DF-GR (proposed) GR using GAN trained for reserved exemplars



Figure 7: Average accuracy for 80 and 100 seen classes at time 3 and 5 of CIL on CIFAR-100, respectively, when using different numbers of reserved exemplars per old class. We used DT-ID in all cases ($N_t = 2$).

nent when the number of reserved exemplars is small, and it shows the usefulness of the proposed DF-GR. We also compare our DF-GR to GR using the generative adversarial network (GAN) trained for reserved old exemplars only. The results show that our proposed DF-GR yields better CIL performance than GR using the GAN, when the training data is limited to a small number of reserved exemplars.

In Figure 8(a,b), we show the average accuracy of CIL on ImageNet-Subset, when we change the number of new classes added at each time in CIL to 25 and 5, respectively. In Figure 8(c), we also show the average accuracy when CIL starts with 10 base classes instead of 50 base classes. Observe that our proposed method provides gains consistently over the baseline [13] in all cases.

4.4. Experimental results for data-limited CIL

Data sharing becomes more and more difficult due to the large size and is often restricted to preserve data privacy. To adapt to this emerging challenge, we explore the datalimited CIL scenario, where data are limited not only for old classes but also for new classes. In the data-limited CIL scenario, instead of sharing a large number of data, we as-



(a) 50 base classes + 25 new classes per task (b) 50 base classes + 5 new classes per task (c) 10 base classes + 10 new classes per task

Figure 8: Average accuracy in CIL on ImageNet-Subset for different numbers of new classes added at each time of CIL with 50 base classes (see (a) and (b)). We also show the average accuracy when CIL starts from 10 base classes instead of 50 base classes in (c). For ours, we used both DT-ID and DF-GR (N_t , $N_q = 2, 1$).

Table 4: Average accuracy for all seen classes at each time of CIL on ImageNet-Subset in the data-limited CIL scenario, where we are not able to access the training data not only for old classes but also for new classes. We are given two pre-trained models for old and new classes with a small number of exemplars, respectively.

	N_D,N_R	N_t, N_g	Avera	age acc	uracy (%) for	seen c	lasses
Time (# classes)			0 (50)	1 (+10)	2 (+10)	3 (+10)	4 (+10)	5 (+10)
[13]*	20, 0	1,0	87.52	79.40	70.91	61.35	52.09	43.00
Ours	20, 0 20, 0	2, 1 2, 2	87.52 87.52	80.17 79.67	74.54 74.09	68.50 68.55	62.53 63.00	56.80 57.42
[13]*	20, 20	1,0	87.52	80.07	73.69	67.60	63.78	59.86
Ours	20, 20 20, 20	2, 1 2, 2	87.52 87.52	80.77 81.17	75.20 75.14	69.75 69.88	65.42 65.73	61.84 62.40





Figure 9: Average accuracy for different numbers of available data for old and new classes in DT-CIL on CIFAR-100 and ImageNet-Subset.

sume that two pre-trained models are provided for CIL. The first one is the model pre-trained for old classes, as in the conventional CIL. The second one is the pre-trained model on the data for new classes. We employ our proposed DF-GR to reproduce not only the data for old classes but also the data for new classes. That is, we train two generators based on two pre-trained models provided for data-limited CIL, respectively. Then, we add the synthetics samples from two generators when computing the DT-ID and LF losses in (6) and (9), similar to the single DF-GR case (see Section 3.2). In additional to that, we have two DF-KD losses from dual DF-GR for old and new classes, respectively. For the synthetic samples from the generator for old classes, the DF-KD loss is computed after weight imprinting, as described in Section 3.2. For the synthetic samples from the generator for new classes, we only match the softmax output for new classes, without weight imprinting, in DF-KD, since weight imprinting for a large number of old classes was not effective in our trials.

For the data-limited CIL scenario, we perform experiments on two cases of N_D , $N_R = 20, 20$ and N_D , $N_R = 20, 0$ (see Table 2). Table 4 shows the results for these two cases on ImageNet-Subset. Observe that our proposed DT-CIL with DT-ID and DF-GR outperforms the baseline LU-CIR [13]. In Figure 9, we compare the average accuracy for different numbers of available data for old and new classes. Observe that the average accuracy of the data-limited CIL with a small number of exemplars approaches to the conventional CIL performance.

5. Conclusion

In this paper, we proposed a novel dual-teacher knowledge transfer method with data-free generative replay for class-incremental learning. First, DF-GR was used to produce the synthetic samples for old and even for new classes. No original training data and no pre-trained generative models are used in DF-GR. Second, we introduced DT-ID to transfer knowledge from two pre-trained models for old and new classes to one fused model. Our proposed dual-teacher CIL deploys both DF-GR and DT-ID. In our experiments, we implemented the proposed schemes on top of one of the state-of-the-art CIL methods and showed their gains. Last but not least, we demonstrated the potential of DF-GR for data-limited CIL, where available training data are limited due to the cost of data sharing and to preserve data privacy.

References

- [1] Sungsoo Ahn, Shell Xu Hu, Andreas Damianou, Neil D Lawrence, and Zhenwen Dai. Variational information distillation for knowledge transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9163–9171, 2019. 2, 3, 4, 11
- [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*, pages 139– 154, 2018. 2
- [3] David Barber and Felix V Agakov. The IM algorithm: A variational approach to information maximization. In Advances in Neural Information Processing Systems, page None, 2003. 4
- [4] 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, pages 233–248, 2018. 1, 2
- [5] Arslan Chaudhry, Puneet K Dokania, Thalaiyasingam Ajanthan, and Philip HS Torr. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *Proceedings of the European Conference on Computer Vi*sion, pages 532–547, 2018. 2, 6
- [6] Yoojin Choi, Jihwan Choi, Mostafa El-Khamy, and Jungwon Lee. Data-free network quantization with adversarial knowledge distillation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 710–711, 2020. 3
- [7] Thomas M Cover and Joy A Thomas. Elements of Information Theory. John Wiley & Sons, 2012. 4
- [8] Harm De Vries, Florian Strub, Jérémie Mary, Hugo Larochelle, Olivier Pietquin, and Aaron C Courville. Modulating early visual processing by language. In *Advances in Neural Information Processing Systems*, pages 6594–6604, 2017. 6
- [9] Prithviraj Dhar, Rajat Vikram Singh, Kuan-Chuan Peng, Ziyan Wu, and Rama Chellappa. Learning without memorizing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5138–5146, 2019. 2
- [10] Ian J Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. arXiv preprint arXiv:1312.6211, 2013. 1
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 770–778, 2016. 5
- [12] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531, 2015. 2
- [13] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 831–839, 2019. 1, 2, 4, 5, 6, 7, 8, 11, 12

- [14] Ronald Kemker and Christopher Kanan. FearNet: Braininspired model for incremental learning. In *International Conference on Learning Representations*, 2018. 1
- [15] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015. 6
- [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. 2
- [17] Alex Krizhevsky. Learning multiple layers of features from tiny images. *Technical report, Univ. of Toronto*, 2009. 5
- [18] Zhizhong Li and Derek Hoiem. Learning without forgetting. IEEE Transactions on Pattern Analysis and Machine Intelligence, 40(12):2935–2947, 2017. 2, 5, 7
- [19] Xialei Liu, Chenshen Wu, Mikel Menta, Luis Herranz, Bogdan Raducanu, Andrew D Bagdanov, Shangling Jui, and Joost van de Weijer. Generative feature replay for classincremental learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 226–227, 2020. 1, 2
- [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 Conference on Computer Vision and Pattern Recognition*, pages 12245–12254, 2020. 1
- [21] Takeru Miyato and Masanori Koyama. cGANs with projection discriminator. In *International Conference on Learning Representations*, 2018. 6
- [22] Yurii Nesterov. A method for unconstrained convex minimization problem with the rate of convergence $O(1/k^2)$. In *Doklady AN USSR*, volume 269, pages 543–547, 1983. 11
- [23] Cuong V Nguyen, Yingzhen Li, Thang D Bui, and Richard E Turner. Variational continual learning. In *International Conference on Learning Representations*, 2018. 2
- [24] 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. 1
- [25] Hang Qi, Matthew Brown, and David G Lowe. Low-shot learning with imprinted weights. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5822–5830, 2018. 5
- [26] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. iCaRL: Incremental classifier and representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2001–2010, 2017. 1, 2, 6
- [27] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 234–241. Springer, 2015. 6
- [28] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy,

Aditya Khosla, Michael Bernstein, et al. ImageNet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015. **5**

- [29] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In Advances in Neural Information Processing Systems, pages 2990–2999, 2017. 1, 2
- [30] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition, pages 9446– 9454, 2018. 6
- [31] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 374–382, 2019. 1, 5, 7
- [32] Ye Xiang, Ying Fu, Pan Ji, and Hua Huang. Incremental learning using conditional adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6619–6628, 2019. 1, 2
- [33] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, pages 3987–3995, 2017.
 2