

TARGET: Federated Class-Continual Learning via Exemplar-Free Distillation

Jie Zhang*
ETH Zurich

zj.jayzhang@gmail.com

Chen Chen & Weiming Zhuang & Lingjuan Lyu†
Sony AI

{ChenA.Chen, weiming.zhuang, Lingjuan.Lyu}@sony.com

Abstract

This paper focuses on an under-explored yet important problem: Federated Class-Continual Learning (FCCL), where new classes are dynamically added in federated learning. Existing FCCL works suffer from various limitations, such as requiring additional datasets or storing the private data from previous tasks. In response, we first demonstrate that non-IID data exacerbates catastrophic forgetting issue in FL. Then we propose a novel method called TARGET (federatEd clAss-continual leaRninG via Exemplar-free disTillation), which alleviates catastrophic forgetting in FCCL while preserving client data privacy. Our proposed method leverages the previously trained global model to transfer knowledge of old tasks to the current task at the model level. Moreover, a generator is trained to produce synthetic data to simulate the global distribution of data on each client at the data level. Compared to previous FCCL methods, TARGET does not require any additional datasets or storing real data from previous tasks, which makes it ideal for data-sensitive scenarios.

1. Introduction

Federated Learning (FL) is a privacy-aware learning paradigm that facilitates collaborations among multiple entities (e.g., edge devices or organizations) [38, 20, 28, 66, 17]. Each entity or client in FL retains data locally and transfers only training updates to the central server for aggregation.

Conventional FL studies assume that the data classes and domains are static, but the new classes could emerge and data domains could change over time in reality [64,

40, 1, 65]. For example, multiple health institutions could use FL to collaborate and train models to identify COVID-19 [56, 9, 2] strains; new COVID-19 strains, however, continue to emerge due to high mutation rate of virus. An intuitive solution to this issue of continuously emerging data classes is training new models from scratch, but this is impractical as it would require significant extra computation cost. Another method is transfer learning from the previously trained model, but this method suffers from catastrophic forgetting [23, 22, 47, 16], degrading performance on the previous classes.

To address these issues, recent research [36, 11, 3, 19] has introduced the concept of Continual Learning (CL) [61, 42, 51, 49, 48] within the FL framework. These methods, collectively referred to as Federated Continual Learning (FCL), aim to mitigate the problems of catastrophic forgetting in FL. In most FCL scenarios, new classes are dynamically added, which we call Federated Class-Continual Learning (FCCL). FCCL allows local clients to continuously collect new data, and new classes can be added at any time.

Unfortunately, existing FCCL works suffer from various limitations. For example, Ma et al. [36] utilize an unlabeled surrogate dataset to address the catastrophic forgetting problem, which may be difficult to obtain in some data sensitive scenarios. Furthermore, the usage of an unlabeled surrogate dataset may not be ideal for certain types of data, as it may not capture the full complexity of the original data. In CL, exemplar-based methods [44, 33, 54, 13] have achieved leading performance. An exemplar refers to a sample or instance of a previously seen data point that is retained in a memory buffer for future use in the learning process. Dong et al. [11] propose a exemplar-based method that stores historical data to address catastrophic forgetting. However, in many privacy-sensitive scenarios (e.g., hospitals and medical research institutions), users are

*Work done during internship at Sony AI.

†Corresponding author.

not permitted to store data from previous tasks due to privacy and policy concerns and data will not be kept for a long time [21, 53, 34, 50]. In summary, the majority of FCCL methods train the global model with additional datasets or previous task data, which could potentially violate data privacy regulations. This dilemma prompts us to consider the following question:

Question: *How to effectively alleviate the catastrophic forgetting problem in the FCCL without storing the local private data of the client or any additional datasets?*

To address this question, we conduct a systematic analysis and observe that the imbalanced distribution of data among clients in FL exacerbates the catastrophic forgetting problem (see Section 3.3). In order to fix this problem: 1) at the model level, we leverage the previously trained global model to transfer knowledge of the old tasks to the current task. 2) at the data level, we train a generator to produce synthetic data that aims to simulate the global distribution of data on each client. Drawing on these insights, we present a method called TARGET (federated class-continual learning via Exemplar-free distillation) that mitigates catastrophic forgetting in FCCL without compromising clients' data privacy.

Our contributions can be concluded as follows:

- We are the first to demonstrate that non-independent and identically distributed (non-IID) data exacerbates catastrophic forgetting issue in FL. Then we propose a novel method called TARGET, which alleviates the catastrophic forgetting in FCCL by leveraging global information.
- Compared to previous FCCL methods, TARGET doesn't require extra datasets or data from previous tasks, it can be applied in data sensitive scenarios.
- Extensive experiments demonstrate the efficacy of our proposed method. For example, when partitioning the CIFAR-100 dataset into five tasks, our method achieves an accuracy of 36.31%, which is about 6% higher than the best baseline method.

2. Related Work

2.1. Federated Continual Learning

In recent years, deep neural networks (DNNs) have been widely used as a fundamental technology for the development of artificial intelligence, both in established and emerging fields [68, 67, 32, 31, 10]. With the advancement of deep learning, an increasing number of researchers have started to focus on the construction of privacy-preserving deep learning frameworks. FL is a paradigm for collaboratively building a model across multiple clients [24, 39], which is gaining momentum in recent years. However, Federated Continual Learning (FCL), which focused on both FL and CL simultaneously is just emerging and still

remains pending further research [52, 43, 58]. Apart from the client-wise catastrophic forgetting, FCL paradigm also poses new challenges such as inter-client interference and communication-efficiency [58]. FedWeIT solves these challenges through decomposing the parameters into three parts, i.e., global parameters, local based parameters, and task-adaptive parameters [58]. Concurrently, Concept-Drift-Aware Federated Averaging (CDA-FedAvg) extends the popular FL algorithm, Federated Averaging (FedAvg), to tackle the CL problem by introducing the concept drift detection and adaptation [14]. FCL with Distillation (CFedD) treats the model trained with the last task as the teacher model to perform a knowledge distillation and proposes a server distillation mechanism to deal with non-i.i.d. issue [37]. Global-Local Forgetting Compensation (GLFC) designs a class-aware gradient compensation loss and a class-semantic relation distillation loss to prevent forgetting, and a proxy server to mitigate the non-independent and identically distributed (non-IID) problem [12]. Additionally, Federated Selective Inter-client Transfer (FedSelIT) applies FCL to NLP through selectively combining model parameters of foreign clients and selecting informative tasks to perform knowledge transfer [5]. These papers illustrate the growing interest in FCL and the need for novel approaches to address its unique challenges. As research in this area continues to advance, we can expect to see more innovative approaches to FCL that further improve its performance, scalability, and privacy preservation.

2.2. Continual Learning

Continual Learning has been studied extensively, several training methods have been proposed to address the catastrophic forgetting challenge it presents. Regularization-based approaches: Elastic Weight Consolidation (EWC) selectively penalizes the network parameters that are important for old tasks [23]. Synaptic intelligence (SI) uses a memory buffer to store important network parameters [61]. Incremental moment matching (IMM) models the posterior distribution after learning multiple tasks as a mixture of Gaussian models [27]. Stable SGD proposes to carefully design the training regimes such as learning rate decay, batch size, dropout, and optimizer to alleviate forgetting. Replay-based approaches: The methods in this family construct a memory to store the past information which will be presented to the model for reviewing in future tasks. Some store the knowledge of the previous tasks, known as experience replay [35, 45, 49, 46]. iCaRL preserves the most representative samples of each class [44]. Averaged Gradient Episodic Memory (A-GEM) builds an episodic memory of model parameter gradients [8]. Architecture-based: DEN expands the model size [60] and RCL utilizes reinforcement learning [55]. APD divides the model parameters into shared and task-specific parameters to restrict the

model complexity [59].

Exemplar-Free Continual Learning Besides, a promising line of work focuses on data-free Continual Learning. DeepDream perturbs current training samples into images that maximize “forgetting” from the previous tasks [41]. DeepInversion proposes a model inversion technique and evaluates its performance in a Continual Learning scenario but found limited success [57]. DFCIL [50] investigates the failure and decomposes the CE-loss into two different losses which guarantee the learning of effective features. These methods enable Continual Learning in scenarios where storing old task data is not feasible or desirable, such as in privacy-sensitive applications.

3. Catastrophic Forgetting in FCCL

This section conducts an in-depth analysis of catastrophic forgetting problem in federated class-continual learning (FCCL). We start by providing a formal definition of the problem. Then, we investigate the forgetting issue in FCCL and discuss potential methods to mitigate it.

3.1. Problem Definition

Federated Class-Continual Learning (FCCL) focuses on the problem of learning models for new classes over time in FL. An FCCL framework consists of a central server and multiple clients. All clients do not share their raw data with any other client or the central server. Each client learns from a sequence of n tasks, where k -th task contains non-overlapping subsets of classes $C_k \in C$, where C is the set of all possible classes. In our privacy-aware scenario, the task stream is presented in an unknown order, and each client can only access its local data from task k during that task’s training period, which is no longer accessible thereafter. Note that the models are trained in a distributed manner, where each party has access to only a subset of the classes C_k (*i.e.* non-IID). We also consider a more challenging and practical setting where the data in each client is heterogeneous. In this paper, we assume the label distribution of data in each client is skewed [29, 63].

Forgetting Issue in FCCL The goal of the global model optimization problem at task k is to minimize the overall classification error on the current set of classes C_k . However, when a new task arises, clients are not able to access data from previous (old) tasks due to privacy concerns and can only update their local model with data from the new task. This often leads to a significant decrease in performance on previous tasks, which is known as catastrophic forgetting [23, 22, 47, 16]. To mitigate catastrophic forgetting in the global model, we aim to minimize the overall classification error on the current set of classes C_k , while

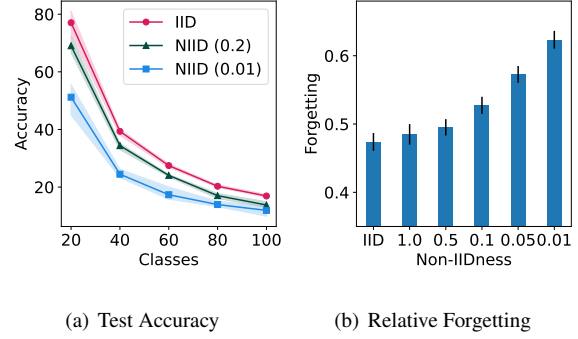


Figure 1: Forgetting under different data partitions. We partitioned the CIFAR-100 dataset into five tasks with 20 classes each, distributed among five clients. “NIID” refers to non-IID, where a lower value represents a more imbalanced or skewed distribution of data. Non-IIDness refers to the Dirichlet parameter.

simultaneously minimizing the changes to the previously learned classes. Formally, the objective function can be written as:

$$\min_{\theta_k} \sum_{c \in C_k} \sum_{i=1}^{m_c} L(f_k(x_{i,c}; \theta_k), c) + \alpha R(\theta_k, \theta_{k-1}) \quad (1)$$

where θ_k is the model parameter at round k , L is a loss function that measures the classification error, R is a regularization term that penalizes changes to the previous model parameters, m_c is the number of data in class c , and α is a hyper-parameter that controls the strength of the regularization. In this formula, $f_k(x_{i,c}; \theta_k)$ represents the classification model that takes as input a data point $x_{i,c}$ associated with class c and outputs a probability distribution over the set of classes in C_k . The regularization term R encourages the new model parameters to be close to the previous model parameters θ_{k-1} , in order to prevent catastrophic forgetting of the previously learned classes.

3.2. Heterogeneous Data Exacerbates Forgetting

We argue that the degree of data heterogeneity has a substantial impact on catastrophic forgetting. To verify this, we conduct an experiment on CIFAR-100 dataset [25] with different degrees of data heterogeneity. Inspired by Backward Transfer (BwT) [6], we derive the following formula to measure the severity of forgetting, a popular forgetting measure in CL [4, 7, 8]:

$$\mathcal{F}_k = \frac{1}{k-1} \sum_{j=1}^{k-1} f_j^k, \quad (2)$$

where \mathcal{F}_k denotes the average forgetting at k -th task and f_j^k quantifies forgetting for the j -th ($j < k$) task after the model

has been continually trained up to task k . Specifically, for a given data distribution, f_j^k can be expressed as follows:

$$f_j^k = \frac{1}{|\mathcal{C}^j|} \sum_{c \in \mathcal{C}^j} \max_{t \in \{1, \dots, N-1\}} \left(\mathcal{A}_c^{(n)} - \mathcal{A}_c^{(N)} \right), \quad (3)$$

where \mathcal{C}^j is a set of classes related to the j -th task, $\mathcal{A}_c^{(n)}$ is the accuracy on class c at round t , and $\mathcal{A}_c^{(N)}$ is the final accuracy on class c after learning all tasks. Note that f_j^k captures the average gap between the peak accuracy and the final accuracy for each class of the j -th task after learning the k -th task.

We further extend the catastrophic forgetting measurement \mathcal{F} in Equation 2 to FL under different data partitions and introduce a relative metric \mathcal{R} to measure forgetting as follows:

$$\mathcal{R}_k = \frac{\sum_{j=1}^{k-1} f_j^k}{\sum_{j=1}^{k-1} \mathcal{A}_{(j,k)}}, \quad (4)$$

where $\mathcal{A}_{(j,k)}$ is the accuracy on task j after learning task k . For different data partitions, an increased \mathcal{R}_k indicates a more serious forgetting of previous tasks.

Figure 1 illustrates the impact of catastrophic forgetting under independent and identically distributed (IID) and different levels of non-independent and identically distributed (non-IID) data partitions. In particular, we employ the Dirichlet distribution, which is widely used in FL [29, 28], to simulate the imbalanced label distribution among different clients. Figure 1(a) shows that the accuracy of the model degrades as training proceeds to new tasks in the IID setting. The performance is even worse in the non-IID settings. These results suggest that FCCL faces significant challenge on extreme non-IID settings. We further analyze the forgetting phenomenon in the Figure 1(b). The higher degree of non-IID exacerbates the forgetting phenomenon in FCCL. These empirical studies motivate us to further investigate the catastrophic forgetting issue in FCCL.

3.3. Alleviating Forgetting via Global Information

To tackle the issue of catastrophic forgetting in the FCCL, we argue utilizing global information to improve performance. The global information can derive from the global model. We explore this approach and empirically demonstrate that the integration of global information can effectively mitigate catastrophic forgetting.

Learn from the Global Model Inspired by LwF [30], we leverage the knowledge from the previously trained global model to the current task via Knowledge Distillation (KD), using only the data from the current task. We conduct experiments on CIFAR100 dataset with 5 continual tasks, where each task contains data of 20 classes. We evaluate the accuracy of the model on each previous task after training on

Table 1: Test accuracy on previous tasks after training on all tasks. In this context, ‘‘FedAvg’’ denotes a naive approach whereby clients learn tasks sequentially.

Partition	Method	Task 1	Task 2	Task 3	Task 4	Task 5	Final
IID	FedAvg	0	0	0	0.05	82.6	16.53
	FedLwF	6.8 \uparrow	11.5 \uparrow	27.1 \uparrow	44.45 \uparrow	63.2 \downarrow	30.61 \uparrow
NIID (1)	FedAvg	0	0	0	0	81.65	16.33
	FedLwF	6.65 \uparrow	13.71 \uparrow	29.60 \uparrow	45.41 \uparrow	59.35 \downarrow	30.94 \uparrow
NIID (0.5)	FedAvg	0	0	0	0.15	77.30	15.49
	FedLwF	1.6 \uparrow	10.65 \uparrow	27.75 \uparrow	41.3 \uparrow	56.65 \downarrow	27.59 \uparrow
NIID (0.02)	FedAvg	0	0	0	0.1	59.00	11.82
	FedLwF	0.1 \uparrow	1.7 \uparrow	3.35 \uparrow	28.2 \uparrow	54.0 \downarrow	17.47 \uparrow

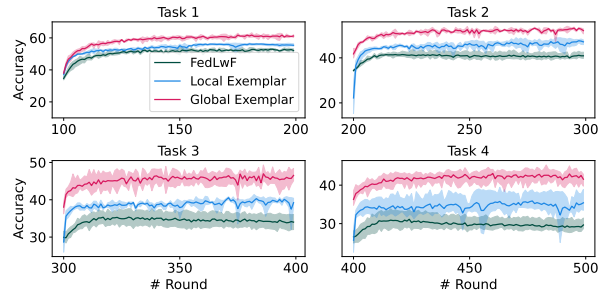


Figure 2: Test accuracy by using current task’s data, local exemplar and global exemplar.

all five tasks and report the final accuracy as the average of these results. Table 1 shows that FCCL without knowledge from the global model (*i.e.*, FedAvg [38]) suffers severely from catastrophic forgetting under both IID and non-IID data partitions. Although FedAvg achieves competitive performance on Task 5, it results in 0% accuracy on previously trained tasks (Task 1, 2, and 3). In contrast, FedLwF achieves better accuracy on Tasks 1, 2, 3, and 4 and final accuracy. The experiments demonstrate that the global model can indeed alleviate the forgetting issue.

Learn From Global Exemplar In continual learning, an exemplar refers to a sample from a previous task that is stored in a memory buffer for future training. Assuming that data stored in clients are compliant to use for future training, we adopt the idea proposed in iCaRL [44] to save a small proportion of prior training data in memory. The selection of this small proportion of data assumes that the data distribution pertaining the old task is known, but this assumption does not hold in FL because clients are unable to know the data distributions of others due to data privacy concerns. Nevertheless, we employ two exemplar selection methods to illustrate the utility of global data: *global exemplar* and *local exemplar*. Global exemplar assumes that the server aggregates a subset of data from clients and distribute these data to clients in training (Note that this method does

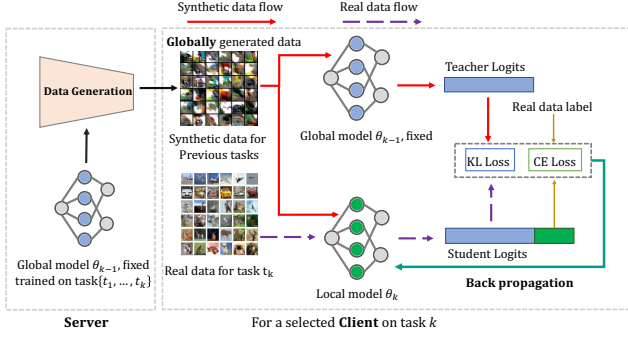


Figure 3: Pipeline of TARGET. We utilize the global model (trained on task $k - 1$) to synthesize the data with a global distribution, subsequently employing this synthesized data for training the k -th task.

not conform to FL and is only used for comparison). Local exemplar means that each client retains a subset of data from local data in the previous tasks for future training.

As shown in Figure 2, we show the performance on the current and all previous tasks after learning the current task by using only the new task data, with local exemplar, and with global exemplar. It can be clearly observed that using exemplars can significantly improve the model performance, especially when using global exemplars, which can achieve much higher accuracy than using local exemplars. However, this raises a critical challenge of how to select such exemplars without violating the data privacy.

4. Our Method: TARGET

4.1. Overview

To utilize the global information without touching on the real exemplars from clients, we present a method called TARGET (federatEd clAss-continual leaRninG via Exemplar-free disTillation), which utilizes global information without storing any real data. A detailed procedure is provided in Algorithm 1. Figure 3 presents an illustration of TARGET, wherein we synthesize data by inverting the global model θ_{k-1} (which was trained on task $k - 1$), followed by combining the synthesized data with real data for local model update on task k .

4.2. Server Side: Synthesizing Data for Old Tasks

As demonstrated in Figure 2, data with global distributional information is more effective in mitigating the problem of catastrophic forgetting. Therefore, we propose a method of synthesizing data that can model the data distribution of the global model, without the need to preserve any client's privacy data. Specifically, given a global (teacher) model θ_{k-1} trained on task $k - 1$, we first initialize a generator G and a student model θ_S . We then repeatedly

Algorithm 1: Procedure of TARGET.

Input: B : local minibatch size, E : local epochs, η : learning rate, synthetic data: $X_{syn} = \emptyset$.

```

1 foreach each task  $\tau = 0, 1, \dots$  do
2   Initialize  $w_0$ 
3   foreach each round  $t = 1, 2, \dots$  do
4      $S_t \leftarrow$  (random set of  $m$  clients)
5     foreach each client  $k \in S_t$  in parallel do
6        $w_{t+1}^k \leftarrow$  ClientUpdate( $k, w_t, \tau, X_{syn}$ )
7      $w_{t+1} \leftarrow \sum_{k=1}^m \frac{n_k}{n} w_{t+1}^k$ 
8    $X_{syn} =$  DataGeneration( $w_{t+1}$ )
9 ClientUpdate( $k, w, \tau, X_{syn}$ ):
10   $\mathcal{B} \leftarrow$  (split  $\mathcal{P}_k \cup X_{syn}$  into batches)
11  set global model  $\mathcal{T} \leftarrow w$ 
12  foreach each local epoch  $i$  from 1 to  $E$  do
13    foreach batch  $(b, b_{syn}) \in \mathcal{B}$  do
14       $\ell(w; b) = L_{ce}(w; b)$ 
15      if  $\tau \neq 0$  then
16         $\ell(w; b) += \alpha L_{ki}(w, \mathcal{T}; b_{syn})$ 
17  return  $w$  to server
18 DataGeneration( $w$ ):
19  Initialize parameter  $\theta_G, \theta_S, X_{syn} = \emptyset$ 
20  foreach round  $i = 1, 2, \dots$  do
21    Sample noises and labels  $\{z_i, y_i\}_{i=1}^b$ 
22    // data generation stage
23    foreach  $j = 1, 2, \dots, b$  do
24      Generate  $\{\hat{x}_i\}_{i=1}^b$  with  $\{z_i\}_{i=1}^b$  and  $G(\cdot)$ 
25      Update  $\theta_G$  using Equation 9
26    Add a batch of data into  $X_{syn}$ 
27    // model distillation stage
28    Update  $\theta_S$  using Equation 10
29  return  $X_{syn}$ 

```

perform the following two training steps (see line 19~28 in Algorithm 1): 1) update the generator by continuously optimizing it to generate data that conforms to the global model distribution; 2) update the student model by distilling knowledge from the teacher model with the synthetic data, hoping that the student model can learn the knowledge of the teacher model sufficiently, which demonstrates the effectiveness of the synthesized data.

Data Generation First, we utilize G to generate synthetic data from noise z , we need to ensure that the synthetic data $\hat{x} = G(z)$ is similar to the training dataset. If the synthetic data is similar to the training dataset, their predictions should also be similar. We minimize the cross-entropy (CE) loss on the output of global model $\theta_{k-1}(\hat{x})$ and random la-

bels \hat{y} ,

$$\mathcal{L}_G^{ce} = CE(\theta_{k-1}(\hat{x}), \hat{y}). \quad (5)$$

It is expected that the synthetic data generated by generator can be classified into a particular class with a high degree of confidence. However, utilizing only the CE loss will cause the generator overfitting to the synthetic data that are far away from the decision boundary (of the global model) [62, 15], thus failing to deliver a good performance. In order to generate samples that are closer to the decision boundary (of the global model) with better transferability, following previous work [62], we introduce a boundary support loss. Additional weight is given to the data on which the global model and the student model diverge in decision making.

$$\mathcal{L}_G^{div} = -\omega KL(\theta_{k-1}(\hat{x}), \theta_S(\hat{x})), \text{ and} \quad (6)$$

$$\omega = \mathbb{1}(\arg \max \theta_{k-1}(\hat{x}) \neq \arg \max \theta_S(\hat{x})), \quad (7)$$

where KL denotes the Kullback-Leibler (KL) divergence loss, $\mathbb{1}(a)$ output 1 if a is true and output 0 if a is false. By maximizing the KL divergence loss, the generator can generate more representative data.

Motivated by [57], in order to further improve the stability of generator training, we introduce Batch Normalization (BN) loss to make synthetic data conform with the batch normalization statistics.

$$\mathcal{L}_G^{bn} = \sum_l (|\mu_l(\hat{x}) - \mu_l| + |\sigma_l^2(\hat{x}) - \sigma_l^2|), \quad (8)$$

where $\mu_l(\hat{x})$ and $\sigma_l^2(\hat{x})$ are the batch-wise mean and variance estimate of the l -th BN layer of the generator, μ_l and σ_l^2 are the mean and variance of the l -th BN layer of $f_S(\cdot)$.

Combining the above losses, we can obtain the loss of the generator as follows,

$$\mathcal{L}_G = \mathcal{L}_G^{ce} + \lambda_1 \mathcal{L}_G^{div} + \lambda_2 \mathcal{L}_G^{bn}, \quad (9)$$

where λ_1 and λ_2 is the weight for different loss functions.

Model Distillation In Equation 7, we introduce a student model to assist in training the generator to produce data with greater diversity. A better student model should lead to a better generator. Therefore, after training the generator for several rounds, we subsequently train the student model using the saved synthesized data and the output of the teacher model, using KL loss for knowledge distillation:

$$\mathcal{L}_S = KL(\theta_{k-1}(\hat{x}), \theta_S(\hat{x})). \quad (10)$$

In this way, we can train a student model with better performance, and then further use it to update the generator. An ideal synthetic dataset should be able to efficiently enable student θ_S to fully learn the knowledge of teacher model.

Note that when the training of the whole process is over (*i.e.* the student model can use the synthesized data to obtain high performance), we only retain the synthetic dataset X_{syn} and transfer it to the clients, without saving the generator and student model.

4.3. Client Side: Update with Global Information

On the client side, we can obtain the data synthesized for the previous task X_{syn} and the real training data of the current task X_{local} , then we train the local model θ_k for task k on the two datasets at the same time. We showed in Section 3.3 that the use of global models and global data can alleviate forgetting. Thus we distill the knowledge of global teacher model and global synthetic data by minimizing the following objective function,

$$\mathcal{L}_{client} = \underbrace{CE(\theta_k(x), y)}_{\text{for current task}} + \alpha \cdot \underbrace{KL(\theta_{k-1}(\hat{x}), \theta_k(\hat{x}))}_{\text{for previous tasks}}, \quad (11)$$

where $(x, y) \in X_{local}$ and $(\hat{x}) \in X_{syn}$. The utilization of the distillation loss facilitates efficient transfer of knowledge from the previous task to the current task model. And α is a hyper-parameter that controls the strength of the regularization for the previous tasks.

5. Experiments

5.1. Experimental Settings

We experiment on two datasets, namely CIFAR-100 [25], and Tiny-ImageNet [26], to evaluate the performance of our proposed approach. To establish the order of the continual tasks, we adopt the widely used protocols [49, 61, 42, 7]. Specifically, we divide all classes of each dataset equally into multiple tasks by default, *i.e.* we evenly divide the classes into 5 and 10 tasks to simulate class continual learning scenarios. We employ ResNet18 [18] as the backbone for the classification model.

To evaluate our approach, we employ the standard continual learning metrics, as used in prior works [23, 22, 47, 16], which include **average accuracy** across all tasks and a **forgetting measure** [6] (see Equation 2). For a fair comparison with the baseline class continual learning methods in the FCCL setting, we implement three types of baselines: 1) Finetune, in which each client simply learns tasks in sequence; 2) FedWeIT [58], a regularization-based method in Federated Continual Learning that maximizes the knowledge transfer between clients; 3) Examples of typical continual learning methods that do not store training data for rehearsal, including EWC [27] and LwF [30]. In addition, we compare our method with methods that store real training data of old tasks, such as iCaRL [44]. We implement these traditional continual learning algorithms in the FCCL scenario and name them as FedEWC, FedLwF, and

Table 2: The Average Accuracy (%) and Forgetting for all learned tasks on CIFAR-100 for various numbers of tasks (5, 10) under both IID and non-IID settings. Results are reported as an average of 3 runs. 'Acc' refers to average accuracy, and ' \mathcal{F} ' represents the forgetting measure utilized in Equation 2. The best results are in bold.

Data partition	IID				NIID (1)				NIID (0.5)			
Tasks	T=5		T=10		T=5		T=10		T=5		T=10	
Method	Acc(\uparrow)	\mathcal{F} (\downarrow)	Acc(\uparrow)	\mathcal{F} (\downarrow)	Acc(\uparrow)	\mathcal{F} (\downarrow)	Acc(\uparrow)	\mathcal{F} (\downarrow)	Acc(\uparrow)	\mathcal{F} (\downarrow)	Acc(\uparrow)	\mathcal{F} (\downarrow)
Finetune	16.12	0.78	7.83	0.75	16.33	0.77	8.45	0.74	15.49	0.74	7.64	0.71
FedEWC	16.51	0.71	8.01	0.65	16.06	0.68	8.84	0.62	16.86	0.66	8.04	0.65
FedWeIT	28.45	0.52	20.39	0.43	28.56	0.49	19.68	0.45	24.57	0.54	15.45	0.48
FedLwF	30.61	0.45	23.27	0.37	30.94	0.42	21.16	0.41	27.59	0.44	17.98	0.45
Ours	36.31	0.22	24.76	0.26	34.89	0.24	22.85	0.26	33.33	0.27	20.71	0.29

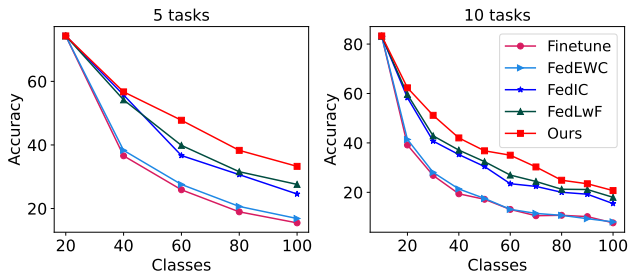


Figure 4: Average accuracy on previous tasks and current task after the model trained on current task.

FedIcaRL. For detailed information on the task configuration, default hyper-parameters and additional experimental results, please refer to the **Appendix**.

5.2. Experiments on CIFAR-100

For CIFAR-100 dataset, we conduct experiments on two sets of tasks consisting of 5 and 10 tasks, respectively. We run the experiment on both IID and non-IID scenarios. For non-IID setting, the Dirichlet parameter is set to 0.5 and 1, *i.e.* NIID(0.5) and NIID(1). Table 2 shows the final average accuracy of the FL model trained on all tasks, along with the corresponding forgetting measure for each experiment. It is important to emphasize that an optimal method is characterized by high average accuracy and low forgetting measure.

Table 2 indicates that Finetune exhibits the poorest performance when attempting to learn the continuously incoming tasks sequentially, thereby experiencing the issue of catastrophic forgetting. Moreover, we observed that methods based on regularization constraints, such as FedEWC and FedWeIT, were often ineffective in preventing the model from forgetting old tasks due to the lack of available data. On the other hand, we found that distillation-based approaches such as FedLwF and our proposed method were

capable of improving the final average accuracy while simultaneously mitigating the issue of catastrophic forgetting. This is due to the transfer of model knowledge learned from the old task to the new task in a distillation manner, enabling the model to prevent catastrophic forgetting. Obviously, we found that applying our proposed method to generate synthetic datasets for the federated models trained on old tasks and subsequently performing model distillation on these synthetic datasets can lead to a substantial improvement in the average accuracy and reduction in the forgetting measure. This observation underscores the ability of our synthetic data to capture the distribution characteristics of historical task data accurately. For example, in IID setting, when partitioning the CIFAR-100 dataset into five tasks, our method achieves an accuracy of 36.31%, which is about 6% higher than the best baseline method FedLwF. It is worth noting that we observed a decrease in the average accuracy of all methods to varying degrees as the number of tasks increased from 5 tasks to 10 tasks due to increased task complexity and the associated forgetting phenomenon. Nonetheless, our proposed method maintains the highest average accuracy and the lowest forgetting measure even under these more demanding conditions.

Figure 4 illustrates the performance of the models trained by various methods on all previously learned tasks after the completion of each task. Specifically, it shows the average accuracy of the model on both the current and previous tasks after the completion of each task (e.g., after learning the second task, the average accuracy of the model on both the first and second tasks is measured). Based on these curves, it is evident that our proposed model outperforms other competing baseline methods in all incremental tasks, regardless of the number of tasks involved. This finding underscores the effectiveness of our approach in facilitating multiple local clients to learn new classes in a streaming manner while mitigating the forgetting problem.

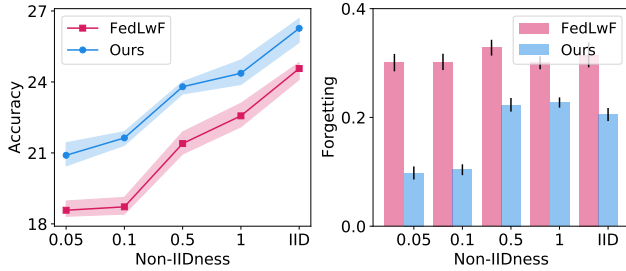


Figure 5: The average accuracy (%) and Forgetting for all learned tasks on Tiny-ImageNet for 5 tasks under both IID and non-IID settings.

5.3. Experiments on Tiny-ImageNet

We also evaluated the performance of our proposed method on the more challenging Tiny-ImageNet dataset and obtained similar results to those observed in the CIFAR-100 experiments. Specifically, in Figure 5, we present the final average accuracy and forgetting measure for all learned tasks in both IID and non-IID settings for the case of 5 tasks. The Dirichlet parameter is set to $\{0.05, 0.1, 0.5, 1\}$. Based on Figure 5, it is evident that our proposed method consistently outperforms FedLwF in terms of average accuracy across all data partitions. Moreover, our method demonstrates a significantly lower forgetting measure than FedLwF under both IID and non-IID settings. Our proposed method achieves an average accuracy that is approximately 3% higher than that of FedLwF even in the most challenging scenario (*i.e.* NIID(0.05)). This result highlights the effectiveness of our method in mitigating catastrophic forgetting in the presence of extreme data distributions.

5.4. Comparison with Exemplar-based Method

In CL, the most successful approaches to alleviate forgetting require extensive replay of previously seen data, which can be problematic when data legality and privacy concerns exist. Among them, iCaRL [44] is a classic but unrealistic algorithm that relies on the stored exemplars in addition to the network parameters, and it is intuitive that using old task’s real data could be beneficial to alleviate forgetting problem. To further understand the performance gap, we compare our proposed TARGET (which uses synthetic data) with iCaRL (which requires storing exemplars from old task’s real data), and study the effects of different exemplar memories on the performance of our method and iCaRL in Figure 6. We set the stored exemplar size to $\{1000, 1500, 2000\}$ for iCaRL, and $\{2000, 3000\}$ for our method on the CIFAR-100 dataset.

The accuracy curve in Figure 6 represents the average accuracy rate measured over all 100 classes learned by the model during the last learning task. While our method

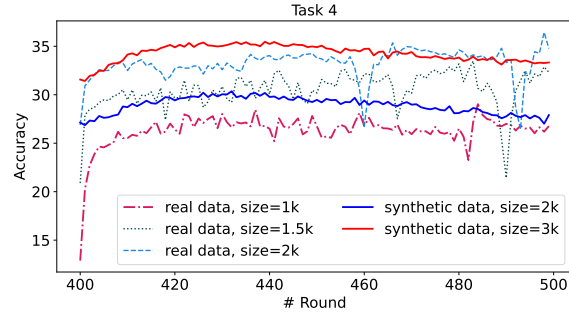


Figure 6: Accuracy comparison between real data-based method (iCaRL) and synthetic data-based method (TARGET). We show the result for the last task.

Table 3: The effect of α on the performance of both new and old tasks in CIFAR-100, 2 tasks.

α	3	5	10	15	25
Old Task (0-49)	28.34	34.91	49.02	58.48	66.02
New Task (50-99)	62.46	60.26	53.74	45.26	29.32
Average (0-99)	45.41	47.58	<u>51.38</u>	51.87	47.76

can achieve similar performance to storing 1k real training data by storing 2k synthetic data, it still cannot outperform storing 2k real training data. However, it is worth noting that our method does not require storing any real training data, which can be a significant advantage in scenarios where storing real data is difficult or not allowed due to privacy or legal concerns. Additionally, our method achieves better performance than storing only 1k real training data, which indicates that our synthetic data is effective in mitigating catastrophic forgetting. We observed that when our method stores 3k synthetic data, it achieves better accuracy than when it stores 2k synthetic data. However, surpassing iCaRL in performance with an equal amount of data remains a challenge for our method. How to effectively use fewer synthetic data with more valuable knowledge from previous tasks will be left as a future research direction.

5.5. Analysis of Our Method

Trade-off between Backward and Forward Transfer.

Continual learning presents a challenge in balancing the trade-off between maintaining high accuracy on old tasks (backwards transfer) and achieving high accuracy on new tasks (forward transfer). In Table 3, we evenly split the CIFAR-100 dataset into two tasks and test the average accuracy of our proposed method under different values of α . We observe that the trade-off between backward and forward transfer is not always balanced. In order to achieve good backward transfer, a large value of α should be used

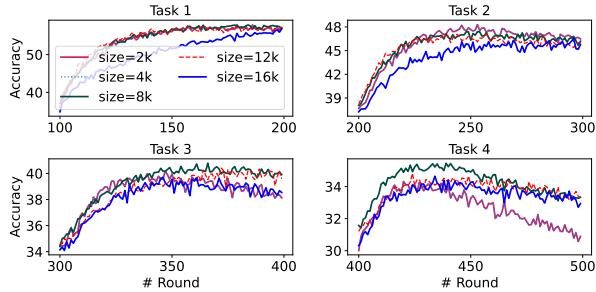


Figure 7: Accuracy for different size of synthetic dataset.

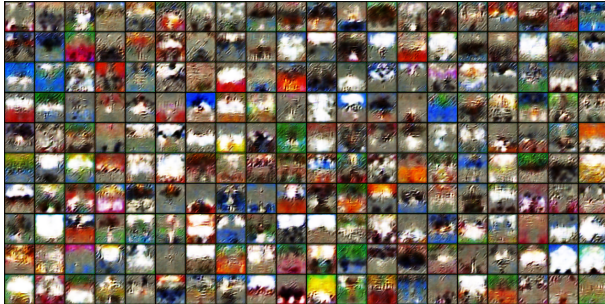


Figure 8: Visualization of randomly synthesised data.

to prevent the model from forgetting previous tasks and to encourage it to focus more on the old synthetic data. Conversely, to achieve good forward transfer, a small value of α should be used, allowing the model to learn quickly and effectively from new tasks while still maintaining some knowledge from the old tasks. The experimental results in Table 3 indicates that when the value of α is between 10-15, the model achieves a good balance between accuracy on new tasks and accuracy on old tasks. We also partition the CIFAR-100 dataset into 5 equal tasks and test our method’s performance on all previous tasks after learning each new task (refer to Appendix).

Memory of Synthetic Data. The optimal amount of synthetic data generated for the old task plays a crucial role in determining the final performance of our method. Insufficient synthetic data may not adequately facilitate knowledge transfer from the old tasks, while excessive data generation can lead to increased memory and communication costs. Therefore, determining an appropriate amount of synthetic data that strikes a balance between knowledge transfer and computational cost is critical for the effectiveness and efficiency of our approach. As shown in Figure 7, we test our method on CIFAR-100, which is divided into 5 tasks with different data sizes ranging from 2k to 16k. It can be observed that when the data volume is relatively small, such as 2k and 4k, the performance of the model is poor, especially when the data volume is 2k, the testing curve in task 4 shows a later decline. This is because when the data volume is too small, the model is unable to effectively learn

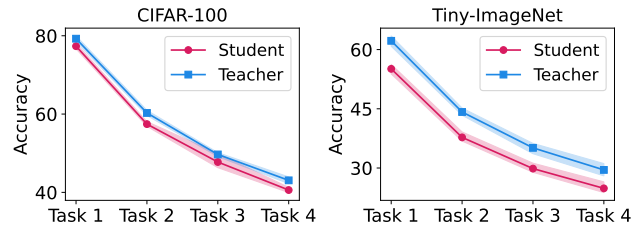


Figure 9: Distillation results for task $\{1, 2, 3, 4\}$ when trained on 5 tasks.

knowledge from old tasks. Increasing the data volume to 8k can effectively alleviate the forgetting phenomenon and achieve good performance. However, continuously increasing the data volume to 12k and 16k do not result in significant improvement in the model’s performance. It is important to note that the size of the data volume alone does not guarantee the effectiveness of synthetic data in improving machine learning models. The quality and relevance of the synthetic data must also be carefully considered to ensure that it accurately represents the underlying distribution of the real-world data.

Visualization on Synthetic Data. To demonstrate the effectiveness of the synthetic data, we present in Figure 8 the visualization results of the synthesized images generated by our method after the model learns the penultimate task on Tiny-ImageNet (CIFAR-100 in Appendix) for 5 tasks. These data have the potential to efficiently enable the student model to approach the performance of the teacher model rapidly. However, it is worth noting that these data exhibit visual dissimilarities from the actual training data.

Analysis on Distillation. Based on the synthetic data, we distill the knowledge of the model trained on old tasks into a student model. The effect of distillation is further demonstrated in Figure 9. It can be observed that even without accessing any private training data from clients, our method can quickly distill the student model on the server to approach the performance of the global model.

6. Conclusion

In conclusion, this paper introduces a novel method, TARGET (federatEd clAss-continual leaRninG via Exemplar-free disTillation), to alleviate the catastrophic forgetting problem in Federated Class-Continual Learning (FCCL). Unlike all the previous methods, our proposed method leverages global knowledge, without requiring any additional datasets or data from previous tasks, making it ideal for privacy-sensitive scenarios. Extensive experimental results demonstrate the effectiveness of our proposed method in comparison to existing FCCL methods.

References

- [1] Eden Belouadah and Adrian Popescu. I2m: Class incremental learning with dual memory. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 583–592, 2019. 1
- [2] Cecilia Bjursell. The covid-19 pandemic as disjuncture: Lifelong learning in a context of fear. *International Review of Education*, 66(5-6):673–689, 2020. 1
- [3] Thang D Bui, Cuong V Nguyen, Siddharth Swaroop, and Richard E Turner. Partitioned variational inference: A unified framework encompassing federated and continual learning. *arXiv preprint arXiv:1811.11206*, 2018. 1
- [4] Sungmin Cha, Hsiang Hsu, Taebaek Hwang, Flavio P Calmon, and Taesup Moon. Cpr: classifier-projection regularization for continual learning. *arXiv preprint arXiv:2006.07326*, 2020. 3
- [5] Yatin Chaudhary, Pranav Rai, Matthias Schubert, Hinrich Schütze, and Pankaj Gupta. Federated continual learning for text classification via selective inter-client transfer. *arXiv preprint arXiv:2210.06101*, 2022. 2
- [6] Arslan Chaudhry, Puneet K Dokania, Thalayasingam Ajanthan, and Philip HS Torr. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *Proceedings of the European conference on computer vision (ECCV)*, pages 532–547, 2018. 3, 6
- [7] Arslan Chaudhry, Naemullah Khan, Puneet Dokania, and Philip Torr. Continual learning in low-rank orthogonal subspaces. *Advances in Neural Information Processing Systems*, 33:9900–9911, 2020. 3, 6
- [8] Arslan Chaudhry, Marc’ Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a gem. *arXiv preprint arXiv:1812.00420*, 2018. 2, 3
- [9] Marco Ciotti, Massimo Ciccozzi, Alessandro Terrinoni, Wen-Can Jiang, Cheng-Bin Wang, and Sergio Bernardini. The covid-19 pandemic. *Critical reviews in clinical laboratory sciences*, 57(6):365–388, 2020. 1
- [10] Jiahua Dong, Yang Cong, Gan Sun, Zhen Fang, and Zhengming Ding. Where and how to transfer: Knowledge aggregation-induced transferability perception for unsupervised domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1(2):1–18, 2021. 2
- [11] Jiahua Dong, Lixu Wang, Zhen Fang, Gan Sun, Shichao Xu, Xiao Wang, and Qi Zhu. Federated class-incremental learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022. 1
- [12] Jiahua Dong, Lixu Wang, Zhen Fang, Gan Sun, Shichao Xu, Xiao Wang, and Qi Zhu. Federated class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10164–10173, 2022. 2
- [13] Arthur Douillard, Matthieu Cord, Charles Ollion, Thomas Robert, and Eduardo Valle. Podnet: Pooled outputs distillation for small-tasks incremental learning. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*, pages 86–102. Springer, 2020. 1
- [14] Fernando Estévez Casado, Dylan Lema Pais, Marcos Fernández Criado, Roberto Iglesias Rodríguez, Carlos Vázquez Regueiro, and Senén Barro Ameneiro. Concept drift detection and adaptation for federated and continual learning. 2021. 2
- [15] Gongfan Fang, Kanya Mo, Xinchao Wang, Jie Song, Shitao Bei, Haofei Zhang, and Mingli Song. Up to 100x faster data-free knowledge distillation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 6597–6604, 2022. 6
- [16] 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, 3, 6
- [17] Andrew Hard, Kanishka Rao, Rajiv Mathews, Swaroop Ramaswamy, Françoise Beaufays, Sean Augenstein, Hubert Eichner, Chloé Kiddon, and Daniel Ramage. Federated learning for mobile keyboard prediction. *arXiv preprint arXiv:1811.03604*, 2018. 1
- [18] 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. 6
- [19] Ziyue Jiang, Yi Ren, Ming Lei, and Zhou Zhao. Fed-speech: Federated text-to-speech with continual learning. *arXiv preprint arXiv:2110.07216*, 2021. 1
- [20] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2):1–210, 2021. 1
- [21] Georgios Kaissis, Alexander Ziller, Jonathan Passerat-Palmbach, Théo Ryffel, Dmitrii Usynin, Andrew Trask, Ionésio Lima Jr, Jason Mancuso, Friederike Jungmann, Marc-Matthias Steinborn, et al. End-to-end privacy preserving deep learning on multi-institutional medical imaging. *Nature Machine Intelligence*, 3(6):473–484, 2021. 2
- [22] Ronald Kemker, Marc McClure, Angelina Abitino, Tyler Hayes, and Christopher Kanan. Measuring catastrophic forgetting in neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018. 1, 3, 6
- [23] 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. 1, 2, 3, 6
- [24] Jakub Konečný, Brendan McMahan, and Daniel Ramage. Federated optimization: Distributed optimization beyond the datacenter. *arXiv preprint arXiv:1511.03575*, 2015. 2
- [25] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 3, 6
- [26] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015. 6
- [27] Sang-Woo Lee, Jin-Hwa Kim, Jaehyun Jun, Jung-Woo Ha, and Byoung-Tak Zhang. Overcoming catastrophic forget-

- ting by incremental moment matching. *Advances in neural information processing systems*, 30, 2017. 2, 6
- [28] Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. Federated learning: Challenges, methods, and future directions. *IEEE signal processing magazine*, 37(3):50–60, 2020. 1, 4
- [29] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *arXiv preprint arXiv:1812.06127*, 2018. 3, 4
- [30] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017. 4, 6
- [31] Zexi Li, Qunwei Li, Yi Zhou, Wenliang Zhong, Guannan Zhang, and Chao Wu. Edge-cloud collaborative learning with federated and centralized features. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '23*, page 1949–1953, New York, NY, USA, 2023. Association for Computing Machinery. 2
- [32] Zexi Li, Tao Lin, Xinyi Shang, and Chao Wu. Revisiting weighted aggregation in federated learning with neural networks. In *International Conference on Machine Learning*. PMLR, 2023. 2
- [33] Yaoyao Liu, Bernt Schiele, and Qianru Sun. Rmm: Reinforced memory management for class-incremental learning. *Advances in Neural Information Processing Systems*, 34:3478–3490, 2021. 1
- [34] Yuang Liu, Wei Zhang, Jun Wang, and Jianyong Wang. Data-free knowledge transfer: A survey. *arXiv preprint arXiv:2112.15278*, 2021. 2
- [35] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30, 2017. 2
- [36] Yuhang Ma, Zhongle Xie, Jue Wang, Ke Chen, and Lidan Shou. Continual federated learning based on knowledge distillation. In Lud De Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 2182–2188. International Joint Conferences on Artificial Intelligence Organization, 7 2022. Main Track. 1
- [37] Yuhang Ma, Zhongle Xie, Jue Wang, Ke Chen, and Lidan Shou. Continual federated learning based on knowledge distillation. In Luc De Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 2182–2188. ijcai.org, 2022. 2
- [38] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-Efficient Learning of Deep Networks from Decentralized Data. In Aarti Singh and Jerry Zhu, editors, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 1273–1282. PMLR, 20–22 Apr 2017. 1, 4
- [39] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017. 2
- [40] 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*, pages 3513–3522, 2021. 1
- [41] Alexander Mordvintsev, Christopher Olah, and Mike Tyka. Inceptionism: Going deeper into neural networks. 2015. 3
- [42] 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, 6
- [43] Tae Jin Park, Kenichi Kumatani, and Dimitrios Dimitriadis. Tackling dynamics in federated incremental learning with variational embedding rehearsal. *arXiv preprint arXiv:2110.09695*, 2021. 2
- [44] 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, 4, 6, 8
- [45] Matthew Riemer, Ignacio Cases, Robert Ajemian, Miao Liu, Irina Rish, Yuhai Tu, and Gerald Tesauero. Learning to learn without forgetting by maximizing transfer and minimizing interference. *arXiv preprint arXiv:1810.11910*, 2018. 2
- [46] Amanda Rios and Laurent Itti. Closed-loop gan for continual learning. *arXiv preprint arXiv:1811.01146*, 2018. 2
- [47] Anthony Robins. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connection Science*, 7(2):123–146, 1995. 1, 3, 6
- [48] Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. In *International Conference on Machine Learning*, pages 4548–4557. PMLR, 2018. 1
- [49] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. *Advances in neural information processing systems*, 30, 2017. 1, 2, 6
- [50] James Smith, Yen-Chang Hsu, Jonathan Balloch, Yilin Shen, Hongxia Jin, and Zsolt Kira. Always be dreaming: A new approach for data-free class-incremental learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9374–9384, 2021. 2, 3
- [51] Sebastian Thrun. Lifelong learning algorithms. *Learning to learn*, 8:181–209, 1998. 1
- [52] Anastasiia Usmanova, François Portet, Philippe Lalanda, and German Vega. A distillation-based approach integrating continual learning and federated learning for pervasive services. *arXiv preprint arXiv:2109.04197*, 2021. 2
- [53] Anamaria Vizitiu, Cosmin Ioan Niță, Andrei Puiu, Constantin Suciuc, and Lucian Mihai Itu. Towards privacy-preserving deep learning based medical imaging applications. In *2019 IEEE international symposium on medical measurements and applications (MeMeA)*, pages 1–6. IEEE, 2019. 2
- [54] Fu-Yun Wang, Da-Wei Zhou, Han-Jia Ye, and De-Chuan Zhan. Foster: Feature boosting and compression for class-incremental learning. In *Computer Vision–ECCV 2022: 17th*

- European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXV*, pages 398–414. Springer, 2022. [1](#)
- [55] Ju Xu and Zhanxing Zhu. Reinforced continual learning. *Advances in Neural Information Processing Systems*, 31, 2018. [2](#)
- [56] Li Yang, Shasha Liu, Jinyan Liu, Zhixin Zhang, Xiaochun Wan, Bo Huang, Youhai Chen, and Yi Zhang. Covid-19: immunopathogenesis and immunotherapeutics. *Signal transduction and targeted therapy*, 5(1):128, 2020. [1](#)
- [57] Hongxu Yin, Pavlo Molchanov, Jose M Alvarez, Zhizhong Li, Arun Mallya, Derek Hoiem, Niraj K Jha, and Jan Kautz. Dreaming to distill: Data-free knowledge transfer via deep-inversion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8715–8724, 2020. [3](#), [6](#)
- [58] Jaehong Yoon, Wonyong Jeong, Giwoong Lee, Eunho Yang, and Sung Ju Hwang. Federated continual learning with weighted inter-client transfer. In *International Conference on Machine Learning*, pages 12073–12086. PMLR, 2021. [2](#), [6](#)
- [59] Jaehong Yoon, Saehoon Kim, Eunho Yang, and Sung Ju Hwang. Scalable and order-robust continual learning with additive parameter decomposition. *arXiv preprint arXiv:1902.09432*, 2019. [3](#)
- [60] Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong learning with dynamically expandable networks. *arXiv preprint arXiv:1708.01547*, 2017. [2](#)
- [61] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International conference on machine learning*, pages 3987–3995. PMLR, 2017. [1](#), [2](#), [6](#)
- [62] Jie Zhang, Chen Chen, Bo Li, Lingjuan Lyu, Shuang Wu, Shouhong Ding, Chunhua Shen, and Chao Wu. Dense: Data-free one-shot federated learning. In *Advances in Neural Information Processing Systems*. [6](#)
- [63] Jie Zhang, Zhiqi Li, Bo Li, Jianghe Xu, Shuang Wu, Shouhong Ding, and Chao Wu. Federated learning with label distribution skew via logits calibration. In *International Conference on Machine Learning*, pages 26311–26329. PMLR, 2022. [3](#)
- [64] Junting Zhang, Jie Zhang, Shalini Ghosh, Dawei Li, Serafettin Tasci, Larry Heck, Heming Zhang, and C-C Jay Kuo. Class-incremental learning via deep model consolidation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1131–1140, 2020. [1](#)
- [65] 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*, pages 13208–13217, 2020. [1](#)
- [66] Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Civin, and Vikas Chandra. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*, 2018. [1](#)
- [67] Didi Zhu, Yinchuan Li, Yunfeng Shao, Jianye Hao, Fei Wu, Kun Kuang, Jun Xiao, and Chao Wu. Generalized universal domain adaptation with generative flow networks. *arXiv preprint arXiv:2305.04466*, 2023. [2](#)
- [68] Didi Zhu, Yinchuan Li, Junkun Yuan, Zexi Li, Yunfeng Shao, Kun Kuang, and Chao Wu. Universal domain adaptation via compressive attention matching. *arXiv preprint arXiv:2304.11862*, 2023. [2](#)