

ORDisCo: Effective and Efficient Usage of Incremental Unlabeled Data for Semi-supervised Continual Learning

Liyuan Wang¹, Kuo Yang², Chongxuan Li^{1*}, Lanqing Hong², Zhenguo Li², Jun Zhu^{1*}

¹ Dept. of Comp. Sci. & Tech., Institute for AI, BNRist Center,
THBI Lab, Tsinghua University, Beijing, China

² Huawei Noah's Ark Lab

Abstract

Continual learning usually assumes the incoming data are fully labeled, which might not be applicable in real applications. In this work, we consider semi-supervised continual learning (SSCL) that incrementally learns from partially labeled data. Observing that existing continual learning methods lack the ability to continually exploit the unlabeled data, we propose deep Online Replay with Discriminator Consistency (ORDisCo) to interdependently learn a classifier with a conditional generative adversarial network (GAN), which continually passes the learned data distribution to the classifier. In particular, ORDisCo replays data sampled from the conditional generator to the classifier in an online manner, exploiting unlabeled data in a time- and storage-efficient way. Further, to explicitly overcome the catastrophic forgetting of unlabeled data, we selectively stabilize parameters of the discriminator that are important for discriminating the pairs of old unlabeled data and their pseudo-labels predicted by the classifier. We extensively evaluate ORDisCo on various semi-supervised learning benchmark datasets for SSCL, and show that ORDisCo achieves significant performance improvement on SVHN, CIFAR10 and Tiny-ImageNet, compared to strong baselines.

1. Introduction

Current achievements of deep neural networks (DNNs) heavily rely on large amounts of supervised data, which are expensive and difficult to acquire simultaneously. Therefore, the ability of continual learning (CL) on incremental training samples becomes extremely important. Numerous efforts have been devoted to CL, which aim to continually learn new training samples without *catastrophic forgetting* of the learned data distribution [29]. Existing CL methods mainly fall in two categories: weight regularization meth-

ods [29, 11, 1] and replay-based methods [29, 32, 35], and have achieved promising results in purely supervised settings.

In many real-world applications, nevertheless, the incremental data are often partially labeled. For example, in face recognition [33], a device continually obtains user data for unlocking. These increasing data could be used to update the model for better user experience. However, true labels of the incoming data are usually unavailable unless the user provides the password for verification. Since frequently asking for labelling would affect user experience, most of the input data are unlabeled. Similar scenarios occur in fingerprint identification [46] and video recognition [20]. Though such scenarios are common in our daily life, they are seldom studied in the CL literature. Therefore, in this paper, we focus on the challenging and realistic task that continually learns incremental partially labeled data. For simplicity, we refer to it as *semi-supervised continual learning* (SSCL).

Different from supervised CL, SSCL provides insufficient supervision and a large amount of unlabeled data. As is well known, unlabeled data are crucial in semi-supervised scenarios [48] but are massive to exploit. In fact, we conduct preliminary experiments in SSCL and empirically verify that the representative CL methods including the weight regularization methods [29, 11, 1] and the replay-based methods [29, 32, 35] may not effectively exploit the unlabeled data. Specifically, the joint training of a strong semi-supervised classifier significantly outperforms the best of existing CL strategies on the same classifier (see results in Fig. 1 and Appendix A). We identify this as the *catastrophic forgetting of unlabeled data* problem in SSCL.

To this end, we present deep *online replay with discriminator consistency* (ORDisCo), a framework to continually learn a semi-supervised classifier and a conditional generative adversarial networks (GAN) [23] together in SSCL. Specifically, ORDisCo is formulated as a minimax adversarial game [7], where a generator tries to capture the underlying joint distribution of partially labeled data, and help

*Corresponding author: C. Li and J. Zhu.

the classifier make accurate predictions. At the same time, the classifier also predicts pseudo-labels for unlabeled data to improve the training of the conditional generator. In contrast to previous work [35, 43, 28], ORDisCo replays data sampled from the conditional generator to the classifier in an *online* manner, which is time- and storage-efficient to exploit the large amount of unlabeled data in SSCL. Further, to explicitly overcome the catastrophic forgetting of unlabeled data, we selectively stabilize parameters of the discriminator that are important for discriminating the pairs of old unlabeled data and their pseudo-labels predicted by the classifier.

We follow the *New Instance* and *New Class* scenarios of CL [29] to split commonly used SSL benchmark datasets [16, 27], including SVHN, CIFAR10 and Tiny-ImageNet, as the evaluation benchmarks for SSCL. To simulate the practical scenarios, all methods continually receive a batch of data with a few labels during training. Extensive evaluations on such benchmarks show that ORDisCo can significantly outperform strong baselines in SSCL.

In summary, our contributions include: (i) We consider a realistic yet challenging task called semi-supervised continual learning (SSCL). We provide a systematical study of existing CL strategies and show that the catastrophic forgetting of unlabeled data is the key challenge in SSCL. (ii) We present ORDisCo to continually learn a classifier and a conditional GAN in SSCL. The generator replays data in an online manner with a consistency regularization on the discriminator to address the catastrophic forgetting of unlabeled data; (iii) We evaluate ORDisCo on various benchmarks and demonstrate that ORDisCo significantly improves both classification and conditional image generation over strong baselines in SSCL.

2. Related Work

Continual learning (CL) aims to address catastrophic forgetting in DNNs on a dynamic data distribution [29], which requires DNNs to learn from incoming tasks or training samples while retaining previously learned data distribution. Current efforts in CL mainly consider the setting where large amounts of annotated data are available for training. Regularization-based methods selectively penalize changes of the parameters to maintain the learned data distribution, e.g., EWC [11], SI [47] and MAS [1]. While, replay-based methods store a small memory buffer to replay representative training samples [32, 4]. To address the imbalance between the small memory buffer and large amounts of training samples, BiC [44] splits an equal number of samples from both the memory buffer and new training data as a validation set to train an additional linear layer for bias correction, while the Unified Classifier [8] regularizes the cosine similarity of features in an unsupervised fashion to normalize the prediction. [39] maintains a topol-

ogy of learned feature space, but it's difficult to learn such a CL-based topology from semi-supervised data. [19] provides a more effective strategy to select the memory buffer from large amounts of labeled data, which is unavailable in SSCL. To better recover the learned data distribution, generative replay strategies continually learn a generative model to replay generated data [35, 43, 28]. However, conditional generation for CL heavily relies on large amounts of labeled data to assign correct labels. Thus, the extension to large amounts of unlabeled data with only a few labels is highly nontrivial. Also, the generative models are often offline saved in CL, which results in additional time and storage cost.

CL of limited supervised data Few-shot continual or incremental learning (FSIL) [10, 40, 3] incrementally learns new classes from a small amount of labeled data, which generally requires a pretraining step on large amounts of base tasks and training samples for few-shot generalization. Unsupervised continual learning (UCL) [31, 36] continually learns new classes from large amounts of unlabeled data. However, FSIL and UCL only make use of a small amount of labeled data or large amounts of unlabeled data, both of which are easy to acquire in many real-world applications. Also, both FSIL and UCL mainly consider learning new classes from a fixed number of training samples, rather than continually learn new instances of the learned classes. By contrast, our setting of SSCL aims to continually learn new instances of the learned classes and new classes from partially labeled data, which is more realistic in real-world scenarios. [17] continually learns partially labeled new instances of fixed classes through stacking generators and discriminators that significantly expand the model. While, we focus on the setting that the model should maintain a relatively constant size, which is a common assumption for continual learning.

Semi-supervised learning (SSL) provides a powerful framework to leverage unlabeled data from a small amount of labels. Pseudo-Labeling (PL) [14] is an early work to assign the prediction of an unlabeled data that is higher than a threshold as its pseudo-label to augment the labeled data. Many more recent works of SSL on discriminative model define a regularization term to learn the distribution of unlabeled data. Consistency regularization, e.g., PI-model [13], makes use of the stochastic predictions of a network and adds a loss term to regularize the consistency of predictions on different passes of the same data. Mean teacher (MT) [41] enforces the predictions of the classifier closer to its exponential moving average on the same batch of unlabeled data. Virtual Adversarial Training (VAT) [24] adds adversarial perturbations to the unlabeled data and enforces the predictions to be the same as the original one. Graph Laplacian regularization [6] penalizes the variation of labels on the graph of manifold structure. SSL on GAN is

generally built on a two-network architecture, that the discriminator is responsible for both sample quality and label prediction, including CatGAN [38], Improved-GAN [34], CGAN [23] and its variants [30, 26]. Triple-GAN [16] applies a triple-network architecture that includes a classifier taking the role of label prediction. ISL-GAN [42] dynamically assigns more virtual labels to unlabeled data during joint training of all the partially labeled data together. However, existing efforts of SSL on both discriminative model and GAN only consider a static data distribution rather than a dynamic one.

3. Semi-supervised Continual Learning and Its Challenges

In this section, we first introduce the problem formulation of continual learning (CL) of incremental semi-supervised data, i.e. semi-supervised continual learning (SSCL). Then, we provide a systematic study of existing representative continual learning methods in SSCL. They suffer from catastrophic forgetting due to the poor usage of large amounts of unlabeled data.

3.1. Problem Formulation

Compared with supervised CL, SSCL only provides a small amount of labeled data and a large amount of unlabeled data, which is key to achieve good performance. Formally, when training on a task t , SSCL is a special continual learning setting on a partially labeled dataset $D^t = \bigcup_{b=0}^{[B]} D_b^t$ with B batches, where $D_b^t = \{(x_i, y_i)\}_{i \in b_l} \cup \{(x_j)\}_{j \in b_{ul}}$ is a semi-supervised batch consisting of a labeled sub-batch b_l and an unlabeled sub-batch b_{ul} . D_b^t is introduced when training on the current batch b and the performance on the task is evaluated after learning each batch.

The above SSCL setting is similar to continual learning of *New Instance* [29], while the dataset D^t for SSCL only consists of small amounts of labels rather than all the labels. We consider continual learning of *New Class* [29] with partially labeled data as a natural extension, where a collection of T semi-supervised datasets $D = \bigcup_{t=0}^{[T]} D^t$ is sequentially learned. D^t is provided during training on the task t , which includes several classes. After learning each task, all of the classes ever seen are evaluated without access to the task labels, i.e. the single-head evaluation [5].

3.2. A Systematic Study of Existing CL Methods

We adapt the representative continual learning strategies, including weight regularization methods [29] and memory replay ones [29], to SSCL by considering *unlabeled data* and provides a systematical analysis.

In particular, we conduct extensive SSCL experiments on the SVHN dataset, where we equally split the training set

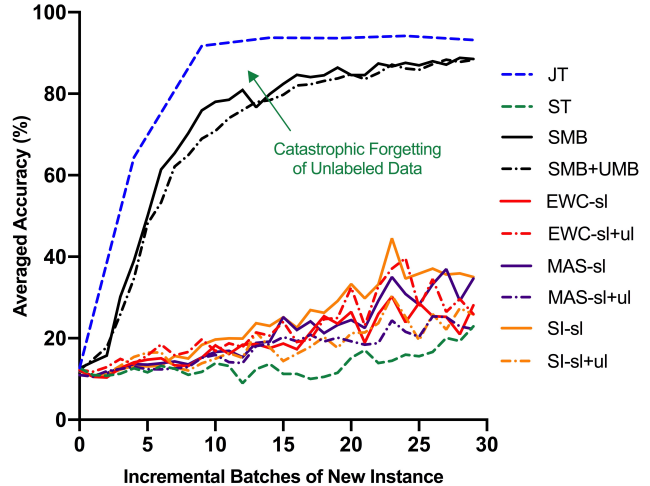


Figure 1. Baselines that combine CL strategies to address SSCL. JT: Joint training of all the training samples ever seen. ST: Sequential training on the incremental data; SMB: ST with replaying the supervised data ever seen. SMB+UMB: SMB with an unsupervised memory buffer of 10000 images to replay unlabeled data. Weight regularization methods (EWC, MAS and SI) are implemented on only the supervised loss (-sl) and both the supervised loss and the unsupervised loss (-sl+ul). The gap between JT and SMB is around 9.70% in average, which is caused by the catastrophic forgetting of unlabeled data.

into 30 batches with 3 labels per class in each batch. Notably, we learn various strong semi-supervised classifiers, including MT [41], VAT [24], PL [14] and PI [13] models in SSCL. We consider four learning strategies for all classifiers to verify the existence of catastrophic forgetting and analyze the underlying reason, as follows:

- **Joint Training (JT):** The classifier is jointly trained on all the partially labeled data ever seen.
- **Sequential Training (ST):** The classifier is sequentially trained on the incremental partially labeled data.
- **Weight Regularization:** Representative approaches of weight regularization, including SI [47], MAS [1] and EWC [11], are implemented to the classifier to selectively stabilize parameters.
- **Memory Replay:** A memory buffer is implemented to replay old training data. Classical methods [32, 4] in CL use mean-of-feature to select samples, while in SSCL we do not have sufficient labels to perform selection. Therefore, we select data through uniform sampling, which is indeed competitive to existing selection methods as analyzed in [5] (See empirical results in Appendix A).

We extensively search the hyperparameters of all baselines and summarize the best results. For details like the

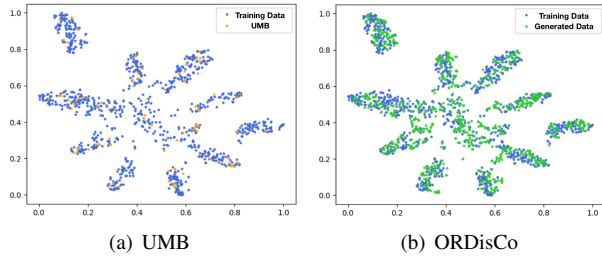


Figure 2. Data coverage of an unsupervised memory buffer, i.e. UMB (a), and the samples generated from ORDisCo (b). ORDisCo has a less storage cost but achieves a better result.

search space, we refer the readers to Appendix B. We compare the performance of different learning strategies based on the MT classifier in Fig. 1. The results of other classifiers are similar and detailed in Appendix A.

In Fig. 1, it can be seen that ST on the incremental partially labeled data significantly underperforms JT, suggesting the existence of catastrophic forgetting. Besides, simply adopting weight regularization methods [29, 11, 1, 47] on a semi-supervised classifier cannot address the issue in SSCL. We hypothesize that the regularization of parameter changes in such methods limits the learning of incremental data, see more experiments and detailed discussion in Appendix B. A potential future work can be designing proper regularization terms in SSCL.

In comparison, using a supervised memory buffer (SMB), whose maximum size is of 900 images, to replay labeled data [29, 5] can effectively reduce the gap from JT but still has a large space to improve. Further, we evaluate a strategy with an additional unsupervised memory buffer (SMB+UMB), whose maximum size is of 10,000 images, to replay unlabeled data. Indeed, the extra buffer is about twice the size of a typical generator used in our method. Given an extra large memory buffer for unlabeled data, however, we observed that the performance of the memory replay approach cannot be further improved, implying that the unlabeled data are not fully exploited. In fact, we visualize the coverage of the UMB on the distribution of all the training data using t-SNE [22] embedding (see details in Appendix A). As shown in Fig. 2 (a), although the UMB has already stored around 10% data of the entire dataset, it still cannot cover the distribution of training data very well.

Therefore, a common issue of existing CL strategies in the semi-supervised scenarios is that they cannot effectively capture and exploit the distribution of unlabeled data. We identify this as the *catastrophic forgetting of unlabeled data* problem, which is challenging because the incremental unlabeled data are massive and lack annotations. This motivates us to propose two new strategies to continually learn a generative model that can successfully capture the data distribution (see Fig. 2, b) given partially labeled data.

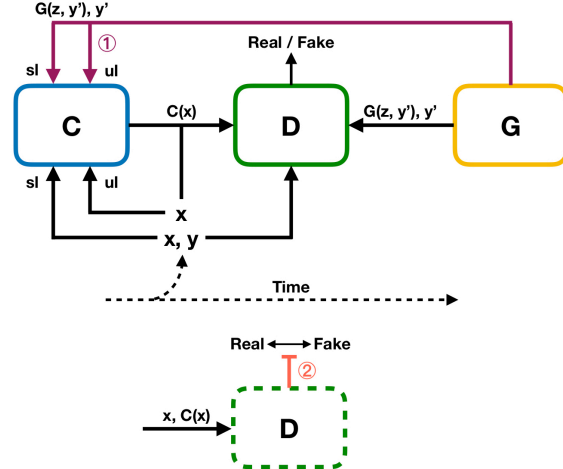


Figure 3. Deep online replay with discriminator consistency (ORDisCo). The partially labeled data are incrementally learned in the interplay of C , G , and D . To mitigate catastrophic forgetting and better exploit the unlabeled data, we implement ① online semi-supervised generative replay in C and ② stabilization of discrimination consistency in D . z and y' are the noise and the labels for conditional sampling, where the generated samples are applied as both labeled and unlabeled data for replay. Here, sl and ul stand for supervised and unsupervised losses, respectively.

4. Method

Based on the empirical analysis above, in this section, we first describe our framework that interdependently trains a classifier, a discriminator and a generator on the incremental semi-supervised data in Sec.4.1. Then, in Sec.4.2, we present two new strategies to mitigate catastrophic forgetting of the unlabeled data.

4.1. Conditional Generation on Incremental Semi-supervised Data

Different from existing generative replay methods [35, 43, 28], we consider a SSCL setting where partially labeled data are given. Inspired by the state-of-the-art semi-supervised learning (SSL) GAN [16, 15], we build up a triple-network structure that continually learns a classifier together with a conditional generative adversarial network (GAN) to better exploit the unlabeled distribution from a few labeled data in SSCL. Both the classifier and the conditional generator produce “fake” data-label pairs and the discriminator focuses on discriminating if a data-label pair is real or fake. Compared to other SSL GANs [38, 34, 23], our framework is theoretically optimal under a nonparametric assumption and ensures that the GAN to capture the data distribution and the classifier to make predictions actually help each other [16, 15].

Formally, the entire model includes a classifier C , a generator G and a discriminator D . Correspondingly, the optimization problem is to learn three groups of parameters as

$\{\theta_C, \theta_G, \theta_D\}$. In such a triple-network architecture, C tries to improve its ability of classification and generate pseudo-labels for the conditional generation in D and G . While, D and G continually learn and recover the data distributions seen in the incremental semi-supervised data. The general framework is shown in Fig. 3 and we will describe the objective function for each component below.

To learn a classifier C from partially labeled data, the loss functions can be generally defined as a supervised loss $L_{sl}(\theta_C)$ and an unsupervised loss $L_{ul}(\theta_C)$. Since our motivation is to continually capture and exploit the unlabeled data, we build up our method with a supervised memory buffer *smb* to replay a few labeled data. We define $L_{sl}(\theta_C)$ and $L_{ul}(\theta_C)$ as a cross entropy term and a consistency regularization term [41]:

$$L_{C,pl}(\theta_C) = L_{sl}(\theta_C) + L_{ul}(\theta_C), \quad (1)$$

$$L_{sl}(\theta_C) = \mathbb{E}_{x, y \sim b_l \cup smb} [-y \log C(x)], \quad (2)$$

$$L_{ul}(\theta_C) = \mathbb{E}_{x \sim b_{ul}, \epsilon, \epsilon'} [\|C(x, \epsilon) - C'(x, \epsilon')\|^2], \quad (3)$$

where C' is the exponential moving average of C [41], $\|\cdot\|^2$ is the square of the l_2 -norm, and ϵ and ϵ' denote two different random noises, e.g. dropout masks. We mention that there are other options of unlabeled losses for SSL such as [24, 13, 37]. We leave a systematical study of this for future work.

Next, we design a semi-supervised discriminator, which aims to distinguish if a data-label pair is from the labeled dataset or is a fake one:

$$\begin{aligned} L_{D,pl}(\theta_D) = & \mathbb{E}_{x, y \sim b_l \cup smb} [\log(D(x, y))] \\ & + \alpha \mathbb{E}_{y' \sim p_{y'}, z \sim p_z} [\log(1 - D(G(z, y'), y'))] \\ & + (1 - \alpha) \mathbb{E}_{x \sim b_{ul}} [\log(1 - D(x, C(x)))], \end{aligned} \quad (4)$$

where the sampling distribution of label y' is uniform $p_{y'} = U\{1, 2, \dots, k\}$ among all the k classes ever seen and the noise z is Gaussian $p_z = N(0, 1)$. The first two terms are the regular loss functions for conditional GAN, computed on the labeled training data and the generated data, respectively. While the third term is computed on the unlabeled sub-batch b_{ul} with the pseudo-labels predicted by C . We assign α and $1 - \alpha$ as the coefficients of the last two terms to keep balance with the first term.

Finally, the generator applies a regular loss function for conditional generation as follows:

$$L_G(\theta_G) = \mathbb{E}_{y' \sim p_{y'}, z \sim p_z} [\log(1 - D(G(z, y'), y'))]. \quad (5)$$

4.2. Improving Continual Learning of Unlabeled Data

In such a framework, after training on a semi-supervised batch b_i , C learns the distribution of $b_{i,ul}$ from $b_{i,l}$ for classification while G and D learns that for conditional generation. When accommodating for the next batch b_{i+1} ,

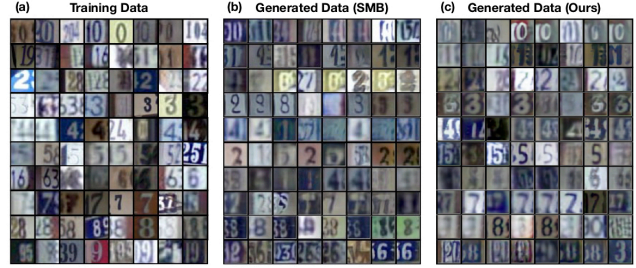


Figure 4. Conditional samples in SSCL on SVHN (after learning 5 batches of partially labeled data). (a) Training data in SVHN. (b) Conditional samples from a conditional GAN baseline with replay of supervised data (SMB). (c) Conditional samples from ORDisCo with the same SMB. Each row has a fixed label. Our method outperforms the baseline in terms of label consistency because of the two strategies proposed.

Table 1. Time and storage complexity of different generative replay strategies. B is the total number of batches to learn.

	Strategy (1) [28]	Strategy (2) [35, 43]	ORDisCo
Algorithm	offline	offline	online
Time	$O(B)$	$O(B^2)$	$O(B)$
Storage	$O(B)$	$O(1)$	$O(1)$

all of the three networks adjust the learned parameters of the previous batch and suffer from *catastrophic forgetting*, especially on unlabeled data. Because unlabeled data are key to solve semi-supervised tasks, without effective strategies to continually learn the unlabeled data from limited labels, the generative model would fail to assign correct labels to the conditional samples during the incremental learning (see Fig. 4 for empirical evidence). To this end, we propose two strategies: online semi-supervised generative replay and stabilization of discrimination consistency.

First, we propose a time- and storage-efficient strategy of online semi-supervised generative replay to improve CL of unlabeled data in C . Existing generative replay methods [35, 43, 28, 45, 18] often offline sample generated data or embedding, or save the old generators, which results in additional time and storage cost. Let's consider two main strategies for offline generative replay: (1) All old generators learned on each task or batch are saved, and replay conditional samples to a classifier [28]; and (2) After learning each task or batch, the generator is saved. In the next task or batch, the generator replays conditional samples with the new training samples to update the generator and the classifier [35, 43, 45, 18]. In contrast, we online generate conditional samples from G and replay them to C rather than offline save a model for replay. As shown in Table 1, our online strategy is more time- and storage-efficient to exploit the large amounts of unlabeled data in SSCL (See detailed calculation in Appendix E). The conditional samples are used as both labeled data and unlabeled data to fully exploit the learned data distribution:

$$L_C(\theta_C) = L_{C,pl}(\theta_C) + L_{G \rightarrow C}(\theta_C), \quad (6)$$

where the first term is defined in Eqn. 1 and the second term exploits the data sampled from the generator in an online manner, as follows:

$$L_{G \rightarrow C}(\theta_C) = \mathbb{E}_{y' \sim p_{y'}, z \sim p_z} [-y' \log C(G(z, y'))] + \mathbb{E}_{y' \sim p_{y'}, z \sim p_z, \epsilon, \epsilon'} [\| C(G(z, y'), \epsilon) - C'(G(z, y'), \epsilon') \|^2], \quad (7)$$

which is in analogy to the losses of labeled and unlabeled training data in Eqn. 2, 3.

Further, to explicitly mitigate catastrophic forgetting of unlabeled data in D , we add a regularization term to selectively penalize changes of the parameters θ_D from the old parameters θ_D^* learned from the previous batches $1 : b$, depending on their contributions to the consistency of discriminating data-label pairs:

$$L_D(\theta_D) = L_{D,pl}(\theta_D) + \lambda \sum_i \xi_{1:b,i} (\theta_{D,i} - \theta_{D,i}^*)^2, \quad (8)$$

where the first term is defined in Eqn. 4 and $\xi_{1:b,i}$ is the strength of penalty on the changes of parameter i . It is computed as

$$\xi_{1:b,i} = ((b-1) \times \xi_{1:b-1,i} + \xi_{b,i})/b, \quad (9)$$

where

$$\xi_{b,i} = \mathbb{E}_{x \sim bu} \left[\left\| \frac{\partial \| (D(x, C(x))) \|^2}{\partial \theta_{D,i}} \right\| \right]. \quad (10)$$

After training on each batch b , we measure $\xi_{b,i}$ through the expected norm of gradients on each parameters to the squared l_2 -norm of discriminating the pairs of unlabeled data and their pseudo-labels predicted by C . Then we update $\xi_{1:b,i}$ by its average. Therefore, the changes of parameters will be penalized if they change the predictions of data-label pairs on the old unlabeled batches and the learned pseudo-labels. Note that, since the regularization is in a re-weighted weight decay form, which is non-negative and is zero in expectation given an optimal discriminator, the regularization won't hurt the consistency and convergence given different initialization and data sequences.

5. Experiment

In this section, we first introduce the benchmark dataset and experiment settings of SSCL. Then we show the results of our method and the baselines on SSCL.

5.1. Experiment Setup

Dataset: SVHN [25] is a dataset of colored digits, including 73,257 training samples and 26,032 testing samples of size 32×32 . CIFAR10 [12] is a dataset of 10-class colored images with 50,000 training samples and 10,000 testing samples of size 32×32 . Tiny-ImageNet is a dataset

of 200-class natural images with 500, 50, 50 samples per class for training, validating and testing, respectively. We randomly choose 10 or 20 classes of images from Tiny-ImageNet, resized to 32×32 or 64×64 . For the experiment of *New Instance*, we randomly split the training images of SVHN, CIFAR10 and Tiny-ImageNet into 30, 30 and 10 batches. Then we allocate a small amount of labels to each batch as the benchmark of SSCL. To simplify the notation, the benchmarks are denoted as “dataset-(number of labels / class / batch)”, including SVHN-1, SVHN-3, CIFAR10-5, CIFAR10-13 and Tiny-ImageNet-5. A model incrementally learns the semi-supervised batch and is evaluated on the test set after training on each batch. For the experiment of *New Class*, we use the same semi-supervised splits as *New Instance*, but construct the sequence as 5 binary classification tasks, where the partially labeled data of two new classes are provided in each task. We follow the single-head evaluation [5] that all of the classes ever seen are evaluated without access to the task labels.

Architecture: In our preliminary experiments (Fig. 1, Appendix A), the classifiers use a Wide ResNet architecture (WRN-28-2), which is widely applied in SSL [27]. To accelerate training, ORDisCo applies a much simplified classifier similar to [15], including 9 convolution layers without residual connections. The simplified classifier achieves a comparable performance on the incremental semi-supervised data as WRN-28-2 (Fig. 1, Appendix A). Our generator and discriminator use a similar architecture as [15], with spectral norm to stabilize training. We keep ORDisCo and all the baselines that directly compare with ORDisCo using the same architecture for fair comparison. Please see Appendix G for hyperparameters of ORDisCo.

Baselines: All the SSL methods in our preliminary experiment (Fig. 1, Appendix A) follow the same implementation as [27], validated in Appendix B. Because mean teacher (MT) [41] can be easily extended to many recent works [9, 2, 21], we choose it as the base model of SSL classifier in ORDisCo. We compare ORDisCo with the strongest baseline in the preliminary experiment, that the classifier is accessible to the labeled data ever seen through a supervised memory buffer (SMB). Since our motivation is to effectively capture and exploit the unlabeled data, we consider two extensions of SMB: (1) We combine SMB with an unsupervised memory buffer (UMB) of a similar size as our generator; and (2) We implement the *Unified Classifier* [8], the state-of-the-art method of incremental learning with memory replay, to better exploit the unsupervised memory buffer (UMB+UC). The Unified Classifier aims to address the imbalance between the small memory buffer and large amounts of training samples through regularizing the cosine distance of feature extractors in an unsupervised fashion to unify the prediction.

Table 2. Averaged accuracy (%) of SSCL on Tiny-ImageNet. We randomly split training images of Tiny-ImageNet into 10 batches with equal number of labels (5 labels / class) in each batch. Here we show the results on the earlier 5 batches.

	Method	Batch 1	Batch 2	Batch 3	Batch 4	Batch 5
10-class 32 × 32	SMB	43.00	46.60	55.40	63.00	59.80
	SMB+UMB	43.80	47.60	55.20	62.20	63.40
	ORDisCo	45.20	54.40	59.60	65.40	69.20
20-class 32 × 32	SMB	24.90	32.50	32.80	38.70	42.30
	SMB+UMB	25.20	33.10	33.60	37.70	42.60
	ORDisCo	25.20	33.30	36.50	42.00	45.10
10-class 64 × 64	SMB	52.00	53.00	58.60	66.00	68.40
	SMB+UMB	51.60	57.00	62.80	68.40	68.80
	ORDisCo	51.20	60.20	65.20	70.40	72.80

Table 3. Ablation study on SVHN-1 and CIFAR10-5. Both online semi-supervised generative replay (Replay) and regularization of discrimination consistency (Reg) substantially improve SSCL.

Method	SVHN-1		CIFAR10-5	
	15 Batch	30 Batch	15 Batch	30 Batch
SMB	38.08	73.32	59.61	67.18
ORDisCo (+Replay, -Reg)	31.14	84.66	64.38	71.69
ORDisCo (+Replay, +Reg)	55.07	85.52	66.58	73.13

5.2. SSCL of New Instance

We report the performance of ORDIsCo and other baselines during SSCL of 30 incremental batches in Fig. 5 (SVHN-1 and CIFAR10-5) and Appendix C (SVHN-3 and CIFAR10-13). ORDIsCo significantly outperforms other baselines, particularly when the numbers of labels are smaller, i.e. SVHN-1 and CIFAR10-5. Compared with SMB, an additional unsupervised buffer (UMB) slightly improves the performance of SSCL on SVHN-3 and CIFAR10-13, but results in overfitting on SVHN-1 and CIFAR10-5 where the labeled data are extremely limited (See empirical results and analysis of the overfitting in Appendix A). The Unified Classifier better exploits the UMB through balancing the predictions on UMB and large amounts of unlabeled data, and thus slightly improves its performance. Also, the performance is only slightly increased by using a much larger UMB (See details in Appendix C). By contrast, ORDIsCo can much better exploit the unlabeled data and substantially improve SSCL. We then evaluate ORDIsCo on Tiny-ImageNet-5 with larger number of classes and larger-scale images in Table 2. Compared with SMB and SMB+UMB, ORDIsCo achieves a much better performance on 10-class and 20-class of size 32×32 and 10-class of size 64×64 , through effective usage of the incremental unlabeled data to improve classification.

Note that the reported performance is the result after training on each semi-supervised batch till convergence. Therefore, the outstanding performance of ORDIsCo on limited partially labeled data, i.e. smaller numbers of labels and semi-supervised batches, indicates that ORDIsCo can *quickly learn a usable model* rather than waiting for more data to be collected.

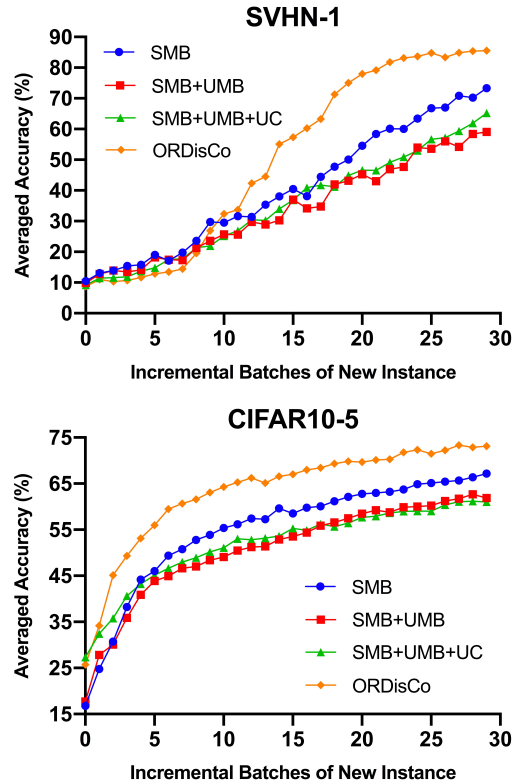


Figure 5. Averaged accuracy (%) of SSCL on SVHN-1 and CIFAR10-5. ORDIsCo (Ours) improves SMB by 12.20% and 5.95% at the final batch on SVHN-1 and CIFAR10-5, respectively. UMB: An unsupervised memory buffer of a similar size as our generator; UC: Using the unified classifier to exploit UMB.

5.3. Ablation Study and Analysis

Next, we analyze why ORDIsCo improves continual learning of unlabeled data. ORDIsCo enables a generative model to continually learn conditional generation from incremental semi-supervised data. Thus, the conditional generator of ORDIsCo can much better cover the training data distribution than an unsupervised memory buffer of a similar size (Fig. 2). We make the ablation study (SVHN-1 and CIFAR10-5 in Table 3, SVHN-3 and CIFAR10-13 in Appendix C) to validate the effects of the proposed strategies in ORDIsCo, that both the online semi-supervised generative replay (refer to as Replay) and the regularizer to stabilize discrimination consistency (refer to as Reg) can significantly improve SSCL. The ablation study shows that the effects of the regularization on discrimination consistency are much more significant when the data sources are *limited*, i.e. fewer labels and semi-supervised batches.

To verify how the discrimination consistency improves SSCL, we use a jointly trained classifier on the semi-supervised dataset, which achieves considerable performance of classification (94.38% for SVHN, 87.37% for CIFAR10), to evaluate conditional samples. We present the

Table 4. Averaged accuracy (%) of a jointly trained classifier to predict conditional samples learned from earlier batches on SVHN-1 and CIFAR10-5, where regularization of discrimination consistency (Reg) significantly improves generative replay.

Method	SVHN-1		CIFAR10-5	
	5 Batch	15 Batch	5 Batch	15 Batch
SMB	27.37	51.45	43.37	70.64
ORDisCo (+Replay, -Reg)	16.95	37.89	49.98	68.67
ORDisCo (+Replay, +Reg)	25.46	64.93	64.15	81.13

Table 5. Averaged accuracy (%) of SSCL of New Class. We follow the same semi-supervised splits as New Instance on SVHN-1, CIFAR10-5, or 10-class Tiny-ImageNet-5 of the size 32×32 , but construct a sequence of 5 binary classification tasks.

Split	Method	Task 1	Task 2	Task 3	Task 4	Task 5
SVHN-1	SMB	81.80	50.73	39.77	40.41	35.46
	SMB+UMB	83.30	50.53	35.92	33.00	37.35
	SMB+UMB+UC	83.04	53.10	36.18	36.83	38.33
	ORDisCo	82.72	54.69	60.68	56.18	53.79
CIFAR10-5	SMB	95.78	74.54	61.63	56.85	55.12
	SMB+UMB	95.50	76.68	63.22	62.34	60.52
	SMB+UMB+UC	95.67	77.94	65.95	63.33	62.43
	ORDisCo	95.60	77.23	68.62	66.49	65.91
TinyImageNet-5	SMB	88.50	75.17	67.17	67.58	65.47
	SMB+UMB	87.83	80.17	68.33	71.33	67.13
	SMB+UMB+UC	88.67	79.08	69.94	71.13	69.87
	ORDisCo	88.50	83.17	70.56	73.92	71.07

accuracy of the classifier to predict conditional samples on the earlier 5 and 15 batches of SVHN-1 and CIFAR10-5, which include only a few labels per batch. As shown in Table 4, the predictions on the conditional samples of ORDisCo are significantly better than the one of a sequentially trained conditional GAN with SMB. Because the quality of conditional samples is generally poor when the partially labeled data are limited, replay of such conditional samples to the classifier in ORDisCo will decrease its performance to predict pseudo-labels of unlabeled data, and further interfere with conditional generation. While, the regularizer on discrimination consistency largely alleviates this issue through stabilizing the learned distribution of unlabeled data in the discriminator, which is critical to SSCL.

5.4. SSCL of New Class

ORDisCo can be naturally extended to incremental learning of New Class in semi-supervised scenarios. We consider an incremental dataset consisting of 5 binary classification tasks, where the partially labeled data of two new classes are introduced in each task. We keep the total number of labels the same as the one used in New Instance. As shown in Table 5, ORDisCo significantly outperforms SMB, SMB+UMB and SMB+UMB+UC through more effective usage of the unlabeled data.

5.5. Conditional Generation in SSCL

We present the conditional samples from ORDisCo and a sequentially trained conditional GAN with SMB during SSCL in Fig. 6 (See complete conditional samples of OR-

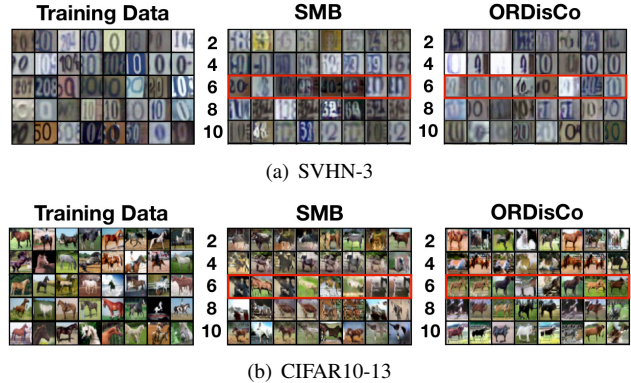


Figure 6. Conditional generation in SSCL. We show the conditional samples of ‘0’ in SVHN-3 and ‘horse’ in CIFAR10-13. The column denotes the number of incremental batches learned by the conditional generator to generate images in each row. Starting from the sixth batch (the red box), ORDisCo can generate high-quality conditional samples, significantly better than the SMB baseline.

DisCo in Appendix D), where conditional generation is improved from incremental learning of partially labeled data. Due to more effective usage of unlabeled data, ORDisCo can generate high-quality conditional samples in SSCL, significantly better than the SMB baseline.

6. Conclusion

In this work, we provide a systematic analysis of semi-supervised continual learning, which is a realistic yet challenging setting without extensive study. Then we propose a novel method to continually learn a conditional GAN with a classifier from partially labeled data. Extensive evaluations on various benchmarks show that our method can effectively capture and exploit the unlabeled data in both classification tasks and conditional generation. Since our method can be a plug-and-play approach for semi-supervised continual learning, a wide range of semi-supervised learning approaches for classification and conditional generation can be flexibly implemented into the framework, which is our further direction. Another further work is to use ORDisCo on Mindspore¹, a new deep learning computing framework.

Acknowledgement

This work was supported by the National Key Research and Development Program of China (Nos. 2017YFA0700904, 2020AAA0104304), NSFC Projects (Nos. 61620106010, 62076145), Beijing NSF Project (No. L172037), Tsinghua-Huawei Joint Research Program. C. Li was supported by the fellowship of China postdoctoral Science Foundation (2020M680572), and the fellowship of China national postdoctoral program for innovative talents (BX20190172) and Shuimu Tsinghua Scholar.

¹<https://www.mindspore.cn>

References

- [1] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 139–154, 2018. 1, 2, 3, 4
- [2] Ben Athiwaratkun, Marc Finzi, Pavel Izmailov, and Andrew Gordon Wilson. There are many consistent explanations of unlabeled data: Why you should average. *arXiv preprint arXiv:1806.05594*, 2018. 6
- [3] Ali Ayub and Alan Wagner. Brain-inspired model for incremental learning using a few examples. *arXiv preprint arXiv:2002.12411*, 2020. 2
- [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 (ECCV)*, pages 233–248, 2018. 2, 3, 11, 14
- [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 Vision (ECCV)*, pages 532–547, 2018. 3, 4, 6, 11, 14
- [6] Chen Gong, Tongliang Liu, Dacheng Tao, Keren Fu, Enmei Tu, and Jie Yang. Deformed graph laplacian for semisupervised learning. *IEEE transactions on neural networks and learning systems*, 26(10):2261–2274, 2015. 2
- [7] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2014. 1
- [8] 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. 2, 6, 13
- [9] Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondrej Chum. Label propagation for deep semi-supervised learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 5070–5079, 2019. 6
- [10] Khurram Javed and Martha White. Meta-learning representations for continual learning. In *Advances in Neural Information Processing Systems*, pages 1818–1828, 2019. 2
- [11] 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, 4
- [12] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 6
- [13] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242*, 2016. 2, 3, 5
- [14] Dong-Hyun Lee. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, 2013. 2, 3
- [15] Chongxuan Li, Kun Xu, Jiashuo Liu, Jun Zhu, and Bo Zhang. Triple generative adversarial networks. *arXiv preprint arXiv:1912.09784*, 2019. 4, 6, 14
- [16] Chongxuan Li, Taufik Xu, Jun Zhu, and Bo Zhang. Triple generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 4088–4098, 2017. 2, 3, 4
- [17] Yanchao Li, Yongli Wang, Qi Liu, Cheng Bi, Xiaohui Jiang, and Shurong Sun. Incremental semi-supervised learning on streaming data. *Pattern Recognition*, 88:383–396, 2019. 2
- [18] Xialei Liu, Chenshen Wu, Mikel Menta, Luis Herranz, Bogdan Raducanu, Andrew D Bagdanov, Shangling Jui, and Joost van de Weijer. Generative feature replay for class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 226–227, 2020. 5
- [19] Yaoyao Liu, Yuting Su, An-An Liu, Bernt Schiele, and Qianru Sun. Mnemonics training: Multi-class incremental learning without forgetting. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 12245–12254, 2020. 2
- [20] Minnan Luo, Xiaojun Chang, Liqiang Nie, Yi Yang, Alexander G Hauptmann, and Qinghua Zheng. An adaptive semisupervised feature analysis for video semantic recognition. *IEEE transactions on cybernetics*, 48(2):648–660, 2017. 1
- [21] Yucen Luo, Jun Zhu, Mengxi Li, Yong Ren, and Bo Zhang. Smooth neighbors on teacher graphs for semi-supervised learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 8896–8905, 2018. 6
- [22] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008. 4
- [23] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014. 1, 3, 4
- [24] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1979–1993, 2018. 2, 3, 5
- [25] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bis-sacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011. 6
- [26] Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional image synthesis with auxiliary classifier gans. In *International Conference on Machine Learning*, pages 2642–2651, 2017. 3
- [27] Avital Oliver, Augustus Odena, Colin A Raffel, Ekin Dogus Cubuk, and Ian Goodfellow. Realistic evaluation of deep semi-supervised learning algorithms. In *Advances in Neural Information Processing Systems*, pages 3235–3246, 2018. 2, 6, 12
- [28] Oleksiy Ostapenko, Mihai Puscas, Tassilo Klein, Patrick Jah-nichen, and Moin Nabi. Learning to remember: A synaptic

- plasticity driven framework for continual learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 11321–11329, 2019. [2](#), [4](#), [5](#), [13](#)
- [29] German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 2019. [1](#), [2](#), [3](#), [4](#)
- [30] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015. [3](#)
- [31] Dushyant Rao, Francesco Visin, Andrei Rusu, Razvan Pascanu, Yee Whye Teh, and Raia Hadsell. Continual unsupervised representation learning. In *Advances in Neural Information Processing Systems*, pages 7647–7657, 2019. [2](#)
- [32] 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](#), [3](#), [11](#), [14](#)
- [33] Fabio Roli and Gian Luca Marcialis. Semi-supervised pca-based face recognition using self-training. In *Joint IAPR International Workshops on Statistical Techniques in Pattern Recognition (SPR) and Structural and Syntactic Pattern Recognition (SSPR)*, pages 560–568. Springer, 2006. [1](#)
- [34] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in Neural Information Processing Systems*, pages 2234–2242, 2016. [3](#), [4](#)
- [35] 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](#), [4](#), [5](#), [13](#)
- [36] James Smith, Seth Baer, Zsolt Kira, and Constantine Dovirolis. Unsupervised continual learning and self-taught associative memory hierarchies. 2019. [2](#)
- [37] Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *arXiv preprint arXiv:2001.07685*, 2020. [5](#)
- [38] Jost Tobias Springenberg. Unsupervised and semi-supervised learning with categorical generative adversarial networks. *arXiv preprint arXiv:1511.06390*, 2015. [3](#), [4](#)
- [39] Xiaoyu Tao, Xinyuan Chang, Xiaopeng Hong, Xing Wei, and Yihong Gong. Topology-preserving class-incremental learning. *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. [2](#)
- [40] Xiaoyu Tao, Xiaopeng Hong, Xinyuan Chang, Songlin Dong, Xing Wei, and Yihong Gong. Few-shot class-incremental learning. *arXiv preprint arXiv:2004.10956*, 2020. [2](#)
- [41] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Advances in Neural Information Processing Systems*, pages 1195–1204, 2017. [2](#), [3](#), [5](#), [6](#)
- [42] Xiaotao Wei, Xiang Wei, Weiwei Xing, Siyang Lu, and Wei Lu. An incremental self-labeling strategy for semi-supervised deep learning based on generative adversarial networks. *IEEE Access*, 8:8913–8921, 2020. [3](#)
- [43] Chenshen Wu, Luis Herranz, Xialei Liu, Joost van de Weijer, Bogdan Raducanu, et al. Memory replay gans: Learning to generate new categories without forgetting. In *Advances in Neural Information Processing Systems*, pages 5962–5972, 2018. [2](#), [4](#), [5](#), [13](#)
- [44] 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. [2](#)
- [45] Ye Xiang, Ying Fu, Pan Ji, and Hua Huang. Incremental learning using conditional adversarial networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6619–6628, 2019. [5](#)
- [46] Chengsheng Yuan, Xianyi Chen, Peipeng Yu, Ruohan Meng, Weijin Cheng, QM Jonathan Wu, and Xingming Sun. Semi-supervised stacked autoencoder-based deep hierarchical semantic feature for real-time fingerprint liveness detection. *Journal of Real-Time Image Processing*, 17(1):55–71, 2020. [1](#)
- [47] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *Proceedings of the 34th International Conference on Machine Learning—Volume 70*, pages 3987–3995. JMLR. org, 2017. [2](#), [3](#), [4](#)
- [48] Xiaojin Zhu and Andrew B Goldberg. Introduction to semi-supervised learning. *Synthesis lectures on artificial intelligence and machine learning*, 3(1):1–130, 2009. [1](#)