

# Exemplar-free Continual Representation Learning via Learnable Drift Compensation

Alex Gomez-Villa<sup>1,2</sup>, Dipam Goswami<sup>1,2</sup>, Kai Wang<sup>1\*</sup>, Andrew D. Bagdanov<sup>4</sup>,  
Bartłomiej Twardowski<sup>1,2,3</sup>, and Joost van de Weijer<sup>1,2</sup>

<sup>1</sup> Computer Vision Center, Barcelona, Spain

<sup>2</sup> Universitat Autònoma de Barcelona, Barcelona, Spain

<sup>3</sup> IDEAS NCBR, Warsaw, Poland

<sup>4</sup> MICC, University of Florence, Florence, Italy

{agomezvi,dgoswami,kwang,btwardowski,joost}@cvc.uab.es,  
andrew.bagdanov@unifi.it

**Abstract.** Exemplar-free class-incremental learning using a backbone trained from scratch and starting from a small first task presents a significant challenge for continual representation learning. Prototype-based approaches, when continually updated, face the critical issue of semantic drift due to which the old class prototypes drift to different positions in the new feature space. Through an analysis of prototype-based continual learning, we show that forgetting is not due to diminished discriminative power of the feature extractor, and can potentially be corrected by drift compensation. To address this, we propose Learnable Drift Compensation (LDC), which can effectively mitigate drift in any moving backbone, whether supervised or unsupervised. LDC is fast and straightforward to integrate on top of existing continual learning approaches. Furthermore, we showcase how LDC can be applied in combination with self-supervised CL methods, resulting in the first exemplar-free semi-supervised continual learning approach. We achieve state-of-the-art performance in both supervised and semi-supervised settings across multiple datasets. Code is available at <https://github.com/alviur/lcd>.

**Keywords:** Continual Learning · Semi-Supervised Learning · Self-Supervised Learning · Exemplar-Free Class Incremental Learning

## 1 Introduction

The standard supervised learning paradigm assumes that data are available in a single training session. This belies, however, the reality of many real-world scenarios in which machine learning methods must deal with data that arrive in a continuous and non-stationary stream of tasks. These shifting distributions lead to catastrophic forgetting if approached naively (for instance, with continuing stochastic gradient descent) [34]. Continual Learning (CL) considers the design of algorithms capable of learning from such streams of non-stationary data.

---

\* Corresponding author

Several techniques have been proposed for continually training models, such as regularization [2, 24, 28, 51], exemplar replay [5, 10, 20, 40, 48], and growing architectures [43]. One group of methods focuses on the more difficult exemplar-free continual learning scenario where the storing of exemplars is prohibited [15, 38, 49, 56, 57]. This setting is relevant in applications where storing previous-task samples is not feasible due to memory limitations, legal constraints arising from new data privacy regulations (like GDPR), or when handling sensitive data such as medical imaging where patient identity must be protected.

One of the main limitations of exemplar-free CL methods is that, as the backbone is updated during the training of new tasks, the network experiences semantic feature drift [49]. This feature drift results in catastrophic forgetting and a significant drop in performance. As a consequence, most exemplar-free methods are evaluated in Warm Start settings (starting from a large first task or from a pretrained network) to start from a strong backbone and thus avoid significant semantic drift of backbone network [1, 15, 20, 35, 48, 49, 55–57]. In fact, a recent paper [38] showed that a simple strategy combined with a frozen backbone after learning the first task can lead to very competitive results in this Warm Start settings. In this paper, we focus on the exemplar-free Cold Start setting [30] in which we train a network from scratch, starting with a small task, and continually update the backbone.<sup>5</sup>

Class Prototype (CP) accumulation is among the most successful approaches to exemplar-free continual learning [15, 22, 35, 37]. CP methods store class prototypes computed using training data and then use them to classify using a distance classifier in feature space. Class Prototype methods are attractive since they can be agnostically applied to any feature extractor, such as multilayer perceptrons, convolutional neural networks, or transformers. CP methods have been successfully applied in Warm Start settings [15, 35, 38] but are expected to struggle in Cold Start settings due to significant semantic drift.

In this article, we aim to apply class prototype methods in Cold Start settings. The main challenge in this setting is the semantic drift of the feature backbone, due to which class prototypes become progressively invalidated since the mean position of each class moves with each new task and cannot be corrected without access to old data in exemplar-free settings. However, in an initial analysis of feature drift (see Fig. 1), we observe that an oracle that perfectly compensates for feature drift leads to a significant recovery of performance. This is an important observation since it shows that semantic drift has not diminished the discriminative power of the feature representations. Consequently, we show that if drift is correctly estimated, feature drift compensation can correct for most of the performance drop resulting from it.

Existing drift estimation methods [46, 49] assume that the drift can locally be approximated with a translation. Instead, we propose a simple Learnable Drift Compensation (LDC) approach to update old class prototypes which has not prior assumptions on the nature of the drift. The learnable compensation maps

---

<sup>5</sup> This setting is predominant in exemplar-based continual learning but hardly explored for exemplar-free methods.

the previous task features to the current feature space via a forward projector network. This compensation can be learned from only the stored old class prototypes and does not require class labels. Owing to the label-agnostic nature of LDC, we propose the first semi-supervised exemplar-free approach, which uses the labeled data only to compute prototypes and all available data to learn a robust forward projector. The main contributions of this work are as follows:

- We show that a large part of forgetting in prototype-based methods can be addressed by prototype compensation. Very little actual forgetting of features crucial for previous tasks is occurring, which motivates the need to develop drift compensation methods for continual representation learning.
- We propose a new exemplar-free continual learning method, which we call Learnable Drift Compensation (LDC), that is based on class prototype accumulation, can compensate the drift of a moving backbone, and is fast and easy to add on top of CP-based methods.
- We experimentally demonstrate that LDC is better than previous drift compensation methods for supervised exemplar-free CL setup. We show that LDC can be applied to any self-supervised CL method to create the first, to the best of our knowledge, exemplar-free semi-supervised CL method.

## 2 Related Work

Continual Learning [24, 34] addresses the issue of *catastrophic forgetting* when training model on a non-i.i.d stream of data [33]. Various methods [33, 54] have been proposed to tackle this problem, including parameter isolation methods [32, 43, 44], regularization techniques [2, 24, 28], and rehearsal strategies [7, 40]. In this paper, we focus on class-incremental learning [47], where the learner has no access to the task oracle during inference, without the need to store any exemplars.

**Exemplar-free Class Incremental Learning.** In this paper, we consider the challenging exemplar-free CIL setting. PASS [56] aligns prototypes of new classes with those of previous tasks via augmentation with noise. SSRE [58] focuses on continually expanding the model’s representations to accommodate new classes. FeTrIL [38] translates old prototype features using the difference between old and new prototypes to create pseudo-features of old data in the new space. In FeCAM [15], the heterogeneous distribution of features in class-incremental learning is exploited using anisotropic Mahalanobis distance instead of Euclidean metric.

**Continual Representation Learning.** CURL [39] performs continual unsupervised representation learning and is based on a variational autoencoder whose latent space can be sampled to replay samples of old classes. CURL is applicable only to small-sized image datasets, requires task boundaries, and is evaluated only on simple classification task. Using self-supervised learning to enhance the learning process of a sequence of supervised tasks was introduced in prior works [50, 56]. Their aim is not to learn directly from unlabeled data, but rather to leverage self-supervised learning to enrich the feature representation

further. Similarly, several studies [4, 6, 12, 29, 36] have used pre-trained models with self-supervision to enhance incremental average classification accuracy through techniques like data augmentation, distillation, and even exemplars. Furthermore, models are trained in self-supervised way in the class-incremental learning (CL) setting without relying on exemplars, as seen in [11, 13].

Continual unsupervised learning can be straightforwardly extended to address the challenge of semi-supervised continual learning, where training phases abstain from utilizing labels entirely, similar to the approach in CaSSLe [11]. However, a limitation of this approach is mitigated by NNCSL [23], which builds upon the PAWS [3] for continual learning. Nonetheless, it is worth noting that these methods typically rely on exemplars to achieve satisfactory performance.

**Drift estimation.** An obstacle to exemplar-free methods is the drift of backbone features, which can result in catastrophic forgetting. Semantic drift compensation (SDC) [49] addresses *drift compensation* by computing the drift of features using new task data and approximating the drift of old class prototypes based on this information. Subsequent work [46] extended SDC by also considering feature drift. Elastic feature consolidation (EFC) [30] regularizes the drift in a direction highly relevant to the previous task using the empirical feature matrix. Recently, Adversarial drift compensation [16] was proposed to generate adversarial samples from new task data which behaves as pseudo-exemplars and are used for estimating drift of old classes. Our method shares a similar objective of estimating feature backbone drift, but we propose a different approach by suggesting the use of a dedicated network for drift compensation.

Memory efficient CIL through feature adaptation (MEA) [21] addresses the problem of feature drift. It stores the features after the backbone and corresponding labels. A learned mapping is applied to translate stored features into the new space to compensate for drift. Unlike our method, MEA requires labels of features, making it unsuitable for unsupervised and semi-supervised settings. Additionally, MEA stores an extensive amount of features for all old classes, leading to higher memory requirements.

### 3 Preliminaries

#### 3.1 The Continual Learning Problem

We consider a class-incremental learning setup where a model learns the dataset  $D$  split into a sequence of tasks  $T = \{t_1, \dots, t_s\}$ . During the training of each task, the model only has access to the data in the split  $t$  (exemplar-free setup). Each task contains a set of  $m$  classes  $C^t = \{c_1^t, c_2^t, \dots, c_m^t\}$  and there is no overlap between classes of different tasks:  $C^t \cap C^s = \emptyset$  for  $t \neq s$ .

Samples inside a task  $t$  can contain different data depending on the learning method. For supervised learning, each sample is a pair  $(x_i, y_i)$  where  $x_i$  is an image of class  $y_i \in C^t$ . In the semi-supervised setting,  $L$  samples have an assigned label  $y_i$  and  $U$  samples are unlabeled images  $x_i$  (usually  $L_t \ll U_t$ ).

### 3.2 Motivation

**Not all forgetting is catastrophic.** Semantic feature drift, as described in prior work [49], manifests when a model learns sequentially. In continual learning, the prototypes computed for task  $t$  may lose relevance for future tasks because the position of class  $c$  after training on task  $t$  differs from that in task  $t - 1$ . This phenomenon is independent of any learning paradigm (whether supervised or unsupervised) and impacts all non-frozen models during incremental training.


To demonstrate the impact of feature drift on Class-Prototype (CP) methods, we conduct experiments using a ResNet-18 backbone trained in a Cold Start setting on 10-task CIFAR-100 and evaluate the performance of classes from the initial 3 tasks after training on all tasks. Fig. 1(a) illustrates the backbone trained with the supervised continual Learning without Forgetting (LwF) strategy [28], while Fig. 1(b) shows the backbone trained with the unsupervised continual learning (CaSSLe) method [11]. We use a Nearest Class Mean (NCM) [40] classifier to predict the labels of the test set for the classes encountered up to each task  $t$ .

In Fig. 1, we show the NCM accuracy using four types of prototypes:

- **Naive prototypes** ( $p_{naive}$ ): These are prototypes without any update or correction. More specifically, prototypes of classes  $C^t$  are computed at the end of  $t$  and used as is afterward.
- **Corrected prototypes** ( $p_{corrected}$ ): At the end of each task, all stored prototype positions are corrected using LDC (see Section 3.3), without the use of exemplars. This correction strategy is the main focus of this paper.
- **Oracle prototypes** ( $p_{oracle}$ ): The true position of all old class prototypes is computed by utilizing all the training data from old tasks at the end of  $t$ .
- **Jointly trained prototypes** ( $p_{joint}$ ): The prototypes are computed using a model which is based on incremental joint training. This serves as the upper bound.

We observe that while a substantial gap in performance exists between the  $p_{naive}$  and  $p_{joint}$  accuracy in the two plots, which is generally referred to as catastrophic forgetting, a major part of this forgetting can be mitigated using  $p_{oracle}$ . This indicates that a significant amount of forgetting can be attributed to the drift in class prototypes which can be reduced by correcting the position of old prototypes. For example, for the case of task 1 and 2 (red lines) we can observe that the features that are discriminative for the task 1&2 are present in the backbone and have not been significantly impaired by the learning of subsequent tasks. However, regularization methods (like LwF and CaSSLe) are not enough for mitigating forgetting. This motivates us to explore drift compensation. We will demonstrate that correcting the prototypes with our method LDC improves the performance of  $p_{naive}$  to  $p_{corrected}$  by big margins (around 38% for tasks 1&2&3 in unsupervised setting).

**Limitations of Semantic Drift Compensation.** The closest methods to LDC rely on SDC. However, this method has an important weakness: it assumes a fixed local transformation for drift (namely locally it can be captured



**Fig. 1:** Last task accuracy of class-prototype accumulation strategies using an NCM classifier in the 10-task CIFAR-100 scenario. The figure row shows regularization-based methods (a) LwF and (b) CaSSLe. Here we show result for two settings: Tasks 1&2 shows how the performance on the first two tasks evolves while incrementally training all tasks. Analogously, Tasks 1&2&3 shows how the performance on the first three tasks evolves over training of all ten tasks.

by translations). If semantic drift follows a different type of transformation, such as scaling, this assumption will lead to errors in the estimated drift.

To illustrate this point, Fig. 2 presents a toy example showcasing translations, rotations, and scaling applied to three sample distributions. In this experiment, we aim to estimate the true mean of  $C_1$  after it has drifted by using the transformation observed in  $D_{t_2}$  as a reference. When only translations are applied (first row), SDC perfectly estimates the distribution’s drift. However, under transformations that include rotations and scaling (second row), SDC incorrectly approximates the true mean of the distributions.

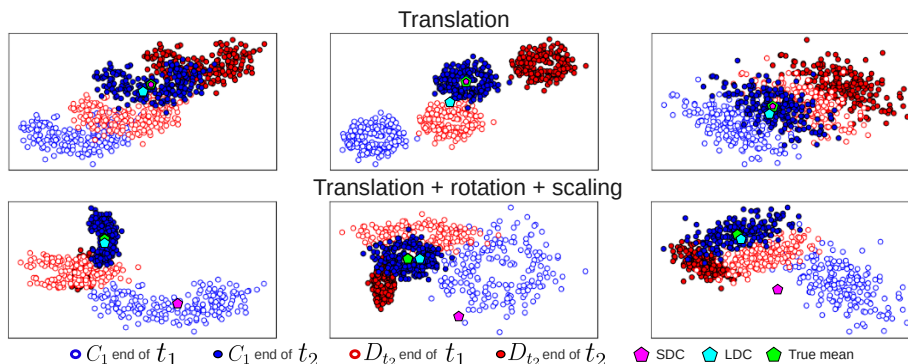
### 3.3 Learnable Drift Compensation

LDC is applied at the end of each training session  $t$ . Let  $f_\theta$  be a feature extractor parameterized by  $\theta$ . To correct the semantic drift, we learn a forward projector  $p_F^t$  to map the features from the old feature space of  $f_\theta^{t-1}$  to the new feature space of  $f_\theta^t$ . In order to learn the drift in the feature space, we use the data of the current task  $D_t$  and minimize the mean squared error (MSE) between the projected features of  $f_\theta^{t-1}$  and  $f_\theta^t$ . Hence, the loss function to learn  $p_F^t$  is:

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N (p_F^t(f_\theta^{t-1}(x_i)) - f_\theta^t(x_i))^2, \quad (1)$$

where  $f_\theta^t$  and  $f_\theta^{t-1}$  are frozen. After this optimization, the old prototypes  $P_{t-1}^c$  for all classes seen till task  $t-1$  are updated using Eq. (2):

$$P_t^c = p_F^t(P_{t-1}^c). \quad (2)$$



**Fig. 2:** Feature drift estimation after applying random translations, rotations and scaling on three sample 2D distributions. We aim to estimate the true mean of  $C_1$  at the end of  $t_2$  using  $D_{t_2}$ . SDC assumes locally that transformations can be captured by translations. LDC can handle rotation and scaling in feature space. Note that LDC (cyan) more accurately approximates the real distribution mean (green).

This process is repeated each time  $f_\theta$  is trained in a new task as shown in Fig. 3. Hence, if we perform class-incremental learning in  $T$  tasks,  $T - 1$  projectors are trained. We do not need to store any old images or features, or use any labels of the current data; we simply store only the old class prototypes and update them using the learned current task projector. The projector  $p_F^t$  can be any trainable mapping function. However, we found that a linear layer performs best in this task (see section 4.4). We summarized the LDC algorithm in the supplementary material.

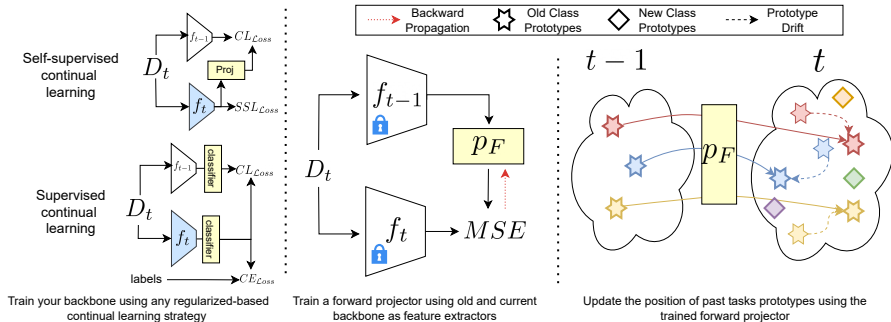
**Nearest Class Mean (NCM) Classifier.** NCM classifier [15, 22, 40, 49] has been found to be very effective for classification in CL settings. We use a NCM classifier on top of the trained feature extractor  $f_\theta^t$  which classifies the test images  $x$  based on their distance to all the class prototypes  $P_t^c$  as follows:

$$y^* = \operatorname{argmin}_{y=1, \dots, Y^t} \|f_\theta^t(x) - P_t^y\|, \quad (3)$$

where  $Y^t$  is the set of classes seen up to task  $t$  and  $P_t^y = p_F^t(P_{t-1}^y) \forall y \in Y^{t-1}$ .

### 3.4 Continual Training Strategies

Our proposed method for drift compensation is general and can be combined with many existing regularization methods commonly used in CL. LDC can be applied at the end of training session  $t$  with the only requirement being access to the old feature extractor  $f_{t-1}$  and the old class prototypes  $P_{t-1}^c$ . A projector  $p_F^t$  must be optimized using Eq. (1) with  $D_t$ , and then old prototypes are updated using Eq. (2) and stored along with current ones. Finally, a new training session begins on  $D_{t+1}$  using any CL strategy (e.g., LwF). It is important to note that our method is independent of the CL strategy used to retain knowledge in the



**Fig. 3: Learnable Drift Compensation.** We train the model on the current task data using regularization-based, self-supervised or supervised continual representation learning methods. After training the new feature extractor  $f_t$ , we learn a forward projector  $p_F$  by minimizing the mean squared error between the projected features from  $f_{t-1}$  and the features from  $f_t$ . We use the learned projector to compensate for the drift of the old class prototypes in the new feature space.

backbone, as LDC is a prototype correction method, not a CL approach. Here, we explain how to extend several existing methods with LDC.

**Supervised CL.** Regularization-based methods retain knowledge while learning new concepts by preventing weights or activations from drifting far from the old model. For instance, LwF [28] aims to match the output of the previous model softmax layer  $h_{\theta}^{t-1}(x)$  with the current model outputs  $h_{\theta}^t(x)$  using knowledge distillation [19] as follows:

$$\mathcal{L} = \mathcal{L}_{ce}(h_t(x), y) + \lambda \mathcal{L}_{ce}(h_{t-1}(x), h_t(x)), \quad (4)$$

where  $\mathcal{L}_{ce}$  refers to the cross-entropy loss and  $\lambda$  denotes the regularization strength. After the training session, the prototypes  $p_t^c$  of  $D_t$  are computed using the labels and added to the prototype pool for NCM classification. This work conducts experiments using LwF as it is a well-established and flexible method in the CL community. However, any exemplar-free supervised CL method can replace the LwF method. Regularization-based methods are particularly well-suited for LDC as they already maintain a copy of the old model at each task.

**Self-supervised CL.** Exemplar-free unsupervised CL methods are based on a projected regularization strategy at the feature level. These approaches perform self-supervised learning (SSL) in  $D_t$  while preserving similarity between old-fixed  $f_{t-1}(D_t)$  and new-projected features  $p(f_t(D_t))$ . PFR [13] project features after the backbone and CaSSLe [11] after a post-backbone layer. POCON [14] uses two networks, one expert network performs SSL in  $D_t$  and the other network (main) distills the knowledge from the expert and the old main network using projected features from the current main network.



The CP accumulation strategy cannot be directly applied in exemplar-free self-supervised CL due to the absence of labels during training. However, if labels become available after training, we can compute prototypes, store them, and even correct them as in the supervised case. After the SSL CL session, we train a projector using the frozen SSL feature extractors and correct the old prototype positions accordingly.

**Semi-supervised CL.** If some labels are accessible during training, we operate within the semi-supervised learning setup. Here, we propose using LDC with exemplar-free self-supervised CL methods, which to the best of our knowledge firstly achieve the exemplar-free continual semi-supervised learning. To implement this, we follow a similar procedure as in the self-supervised CL case but utilize the available labels to compute prototypes. It is important to note that the performance of LDC remains independent of the number of available labels, meaning that the primary source of error would stem from the computation of the prototype, which typically exhibits low variability.

## 4 Experimental Results

### 4.1 Datasets

We use several publicly available datasets for our experiments. CIFAR-100 [26] contains 100 classes of  $32 \times 32$  images with 500 training and 100 testing images per class. Tiny-ImageNet [27] contains 200 classes of  $64 \times 64$  images with a total of 100K training and 10K testing images. ImageNet100 is a subset of one hundred classes from ImageNet [42] containing  $224 \times 224$  images with a total of 130K training and 5K testing images. Stanford Cars [25] consists of 196 finegrained categories of cars and has 8144 images for training and also for testing. We equally split all the datasets into 10 tasks, which is different from the conventional Warm Start settings [15, 38, 49, 56, 57].

### 4.2 Supervised Continual Learning

**Baseline methods.** Since the existing methods were originally proposed for Warm Start settings, we re-implement all methods in Table 1 using PyCIL [54]. We implement a standard distillation baseline method, LwF [28] and also use NCM [40] classification naively. We use SDC [49] with the models trained with LwF. Note that SDC was originally proposed in Warm Start settings with feature distillation. Here, we adapt SDC for Cold Start settings with logits distillation (LwF). We include recent methods like PASS [56], FeTrIL [38], FeCAM [15] and EFC [30] in our comparison. While LwF, PASS and EFC are continual representation learning methods, FeCAM and FeTrIL freeze the model after the first task and continually learn the classifier. We do not compare to the methods [31, 45, 57, 58] proposed in Warm Start settings since they showed poor performance compared to FeTrIL and FeCAM. Following recent practice [17, 35, 52, 53] of using pre-trained ViTs for CIL, we use a ViT-B/16 [8] model pretrained on ImageNet21k [41] and present results with NCM, SDC and LDC in Table 2.

**Table 1:** Results in supervised settings for 10 tasks on CIFAR-100, Tiny-Imagenet, and ImageNet100. The mean and standard deviation using 5 different seeds are reported. Best results are in **bold** and second-best underlined. †: results excerpted from [30].

Method	CIFAR-100		Tiny-ImageNet		ImageNet100	
	$A_{last}$	$A_{inc}$	$A_{last}$	$A_{inc}$	$A_{last}$	$A_{inc}$
LwF+NCM [40]	40.5 ± 2.7	56.2 ± 3.6	28.6 ± 1.1	44.2 ± 1.3	44.0 ± 1.9	64.4 ± 0.4
LwF+SDC [49]	40.6 ± 1.8	56.2 ± 3.0	29.5 ± 0.8	43.8 ± 1.3	42.6 ± 1.9	<u>65.3 ± 0.6</u>
PASS [56]	37.8 ± 0.2	52.3 ± 0.1	31.2 ± 0.4	45.3 ± 0.7	38.7 ± 0.5	55.5 ± 0.5
FeTrIL [38]	37.0 ± 0.6	52.1 ± 0.5	24.4 ± 0.6	38.5 ± 1.1	40.7 ± 0.5	56.3 ± 0.8
FeCAM [15]	33.1 ± 0.9	48.1 ± 1.3	24.9 ± 0.5	38.6 ± 1.3	42.4 ± 0.9	57.9 ± 1.5
EFC [30] †	<u>43.6 ± 0.7</u>	<u>58.6 ± 0.9</u>	<u>34.1 ± 0.8</u>	<b>48.0 ± 0.6</b>	<u>47.3 ± 1.4</u>	59.9 ± 1.4
LwF+LDC	<b>45.4 ± 2.8</b>	<b>59.5 ± 3.9</b>	<b>34.2 ± 0.7</b>	<u>46.8 ± 1.1</u>	<b>51.4 ± 1.2</b>	<b>69.4 ± 0.6</b>

**Training procedure.** We train a ResNet-18 [18] using an SGD optimizer. For all datasets and methods we train the first task for 200 epochs and the incremental tasks for 100 epochs. In the first task for CIFAR-100, we use a learning rate of 0.1, momentum of 0.9, weight decay of 0.0005, and the learning rate is reduced by 10 after 60, 120, and 160 epochs following the default settings of PyCIL. In

the incremental tasks, we use a learning rate of 0.05, which is reduced by 10 after 45 and 90 epochs. For LwF models, we use  $\lambda = 10$ . For LDC, we learn a linear layer using Adam optimizer and a learning rate of 0.001 for 20 epochs using all the data in the current task. We report all details about experimental settings in the supplementary materials. For table 2, we use the same settings and hyperparameters as [52] for finetuning in new tasks and learn a two-layer MLP using Adam optimizer and a learning rate of 0.005 for 100 epochs.

**Results.** We report the last task accuracy as well as the average incremental accuracy with 5 random seeds and class sequences in Table 1. We observe that while the recently proposed methods like PASS [56], FeTrIL [38], and FeCAM [15] performed well in Warm Start settings, LwF with an NCM classifier outperformed these methods in Cold Start settings. While SDC was effective in the originally proposed Warm Start settings, we see that with higher feature drift in our settings, SDC does not improve the performance in most cases. We show that LDC with LwF outperforms all these baselines significantly across all datasets. In CIFAR-100, LDC outperforms SDC by 4.8% and by 3.3% on last task and incremental accuracy respectively. For Tiny-ImageNet, LDC performs similar to the recent EFC method which is also designed for the cold start setting. However, on the more challenging ImageNet100 LDC improves over EFC by a large margin of almost 4% on last task accuracy. Finally, using pre-trained

**Table 2:** Performance on Stanford Cars (10 tasks of 20 classes each) with pre-trained ViT-B/16.

Method	$t1$	$t3$	$t5$	$t7$	$t10$
FT+NCM	85.8	73.3	63.2	58.0	53.4
FT+SDC	85.8	73.5	63.5	57.8	53.3
FT+LDC	85.8	76.0	71.3	67.0	<b>62.9</b>

**Table 3:** Results in semi-supervised settings for 10 tasks on CIFAR-100. The mean and standard deviation using 3 different seeds are reported. The best results for the proposed setting are in **bold**.

Method	Ex Free	Semi-Sup	100%		
LwF+NCM	✓	✗	40.5±2.7		
PASS	✓	✗	37.8±0.2		
			0.8%	5%	25%
PASS	✓	✓	5.91	15.13	15.84
NNCSL (500)	✗	✓	27.4±0.51	31.4±0.40	35.3±0.30
NNCSL (0)	✓	✓	8.7±0.63	9.0±0.42	9.32±0.39
PFR+naive	✓	✓	20.5±0.44	24.8±0.35	25.7±0.12
CaSSLe+naive	✓	✓	19.4±0.57	23.2±0.21	23.6±0.13
POCON+naive	✓	✓	20.1±0.23	21.7±0.27	22.4±0.31
PFR+SDC	✓	✓	16.2±0.66	22.5±0.57	24.8±0.12
CaSSLe+SDC	✓	✓	22.1±0.69	25.8±0.09	26.5±0.04
POCON+SDC	✓	✓	21.8±0.82	25.7±0.23	26.7±0.09
PFR+LDC	✓	✓	18.6±0.20	29.3±0.12	33.4±0.06
CaSSLe+LDC	✓	✓	<b>27.0±0.60</b>	<b>32.8±0.30</b>	<b>35.0±0.21</b>
POCON+LDC	✓	✓	26.1±0.74	31.1±0.62	33.8±0.25

ViTs on Stanford Cars, we observe that LDC improves over SDC by 9.6% after the last task, see Table 2.

### 4.3 Semi-supervised Continual Learning

**Baseline methods.** Since the proposed method is the first exemplar-free semi-supervised CL method, we can only compare it against drift correction methods (SDC). Nevertheless, we present results for the most close approaches. The naive approach achieves the lowest possible accuracy, no prototype correction is performed. SDC [49] is used as a comparison for drift correction. PASS [56] is not a semi-supervised method, but we present our results with all labels (as a reference) and with reduced training data corresponding to the proposed label sets. NNCSL [23] is a semi-supervised method that uses exemplars; we present here results with the same number of exemplars proposed by the authors (as a reference) and a reduced version. Removing exemplars from exemplar-based methods or reducing data to comply with our setup is not fair, and we present these results as a reference to showcase the performance of our approach. We use the code and hyperparameters provided by the authors.

**Training procedure.** We train a ResNet-18 [18] using PFR [13], CaSSLe [11], and POCON [14] as SSL CL strategies. We use the same code, training procedure, and hyperparameters provided by the authors in each method. At the end of each task, we compute the prototypes using the available labels: 0.8%, 5%, 25% for CIFAR-100 and 1%, 5%, 25% for ImageNet100. For LDC, we train a

**Table 4:** Results in semi-supervised settings for 10 tasks on ImageNet100. The mean and standard deviation using 3 different seeds are reported. The best results for the proposed setting are in **bold**.

Method	Ex Free	Semi-Sup	100%		
LwF+NCM	✓	✗	44.0±1.9		
PASS	✓	✗	38.7±0.5		
			1%	5%	25%
PASS	✓	✓	1.76	10.36	23.9
NNCSL (5120)	✗	✓	35.4±0.21	41.6±0.12	47.1±0.08
NNCSL (500)	✗	✓	9.93±0.22	15.1±0.15	16.6±0.1
PFR+naive	✓	✓	19.6±0.16	21.8±0.38	22.0±0.12
CaSSLe+naive	✓	✓	25.37±0.27	26.9±0.25	27.36±0.11
POCON+naive	✓	✓	24.2±0.32	26.1±0.22	25.8±0.20
PFR+SDC	✓	✓	17.2±0.55	19.1±0.16	19.8±0.08
CaSSLe+SDC	✓	✓	22.2±0.03	25.9±0.23	26.6±0.01
POCON+SDC	✓	✓	21.8±0.08	25.1±0.30	25.7±0.03
PFR+LDC	✓	✓	26.5±0.09	32.4±0.35	33.5±0.07
CaSSLe+LDC	✓	✓	<b>29.5±0.58</b>	<b>33.9±0.15</b>	<b>35.8±0.39</b>
POCON+LDC	✓	✓	28.3±0.37	33.2±0.30	34.0±0.28

linear projector using the Adam optimizer (learning rate  $5e - 3$ ) for 100 epochs using all the data in the task. We report the mean last task accuracy in the test set using an NCM classifier.

**Results.** Table 3 presents the results of 10 task split of CIFAR-100. LDC presents the best results for all the pure semi-supervised exemplar-free settings. For 5% and 25% settings, it is reaching the same accuracy as the exemplar-based method NNCSL. Furthermore, LDC is only 3% below the performance of the fully supervised version of PASS. LDC surpasses SDC by at least 7% in all the base SSL CL methods in the pure semi-supervised exemplar-free setting.


Table 4 presents the results of 10 task split of ImageNet100. As in CIFAR-100, LDC performs best for the proposed setting, surpassing SDC by at least 7% in all the SSL CL backbones. However, in this dataset, the replay-based method NNCSL is better with the same number of exemplars proposed by the authors. Nevertheless, a reduction in the number of exemplars in NNCSL leads to a significant performance drop, ending below our proposed exemplar-free method.

#### 4.4 Ablation Studies

In this section we report on a range of ablation studies for LDC. Unless specified, the training and evaluation procedure are the same as in Sections 4.2 and 4.3. We provide more ablation experiments on comparison of LDC with NME and using stored features for learning the projector, in the supplementary materials.

**Table 5:** Ablation study for different complexities of the proposed projector  $p_F$ .


Method	Linear	Linear + bias	Linear + ReLu	MLP
LwF + LDC	45.4±2.8	45.2±3.0	41.2±2.1	43.7±2.3
CaSSLe + LDC	35.0±0.2	34.3±0.2	29.1±0.5	34.9±0.4

**Fig. 4:** Comparison of LDC with NME for varying memory size in the supervised settings on CIFAR-100, Tiny-ImageNet and ImageNet100.

**Architecture of projector.** To ablate the complexity of the proposed projector  $p_F$ , Table 5 presents experiments for LwF+LDC(supervised CL) and CaSSLe+LDC (semi-supervised CL - 25% of labels). We show LDC performance with a linear layer, with and without bias, with an activation layer ReLU after the linear layer and also using a two-layer MLP. We observe that a single linear layer works better in both settings.

**Exemplars versus LDC for prototype correction.** We compare the performance of LDC against Nearest-Mean of Exemplars (NME) which directly uses exemplars for old class prototype compensation and is explored in exemplar-based methods [9, 40]. While it is easy to predict the updated positions of the old prototypes using original samples of old classes by NME, it is very challenging without exemplars in our settings. We compare the last task accuracy of LDC (same setting as in Table 1, with LwF+NCM) with NME which stores a memory of old class samples. We show in Fig. 4 that LDC is on par with using 20 exemplars per class for NME for CIFAR-100 and Tiny-ImageNet. For ImageNet100, LDC is marginally less than NME with 20 exemplars per class. This demonstrates that LDC is able to effectively correct the prototypes without access to exemplars and yet is on par with NME using exemplars.

**Drift Compensation Analysis.** In order to visualize the drift compensation using LDC, Fig. 5 presents a comparison between SDC and LDC. We update all the old prototypes using both SDC and LDC and then compute the cosine distance of the updated prototype to the oracle prototype (computed using all old data for analysis) in the new feature space. We plot the distributions over all old class prototype distances after alternate tasks till the last task. We observe that SDC struggles to estimate the prototypes even after the first increment (task 2) and is centered around distance of 0.1. As we move to the last task, SDC gets even worse with a higher deviation which implies increasing distance of updated prototypes from their oracle positions. On the other hand, LDC is more robust and the distribution is centered between 0 and 0.05 even after



**Fig. 5:** Drift compensation analysis using SDC and LDC in supervised settings on CIFAR-100. We compute the cosine distance between the updated and oracle prototypes. We plot the distributions of the cosine distances for all old classes using both methods after alternate tasks. While SDC fails to estimate the prototypes closer to the oracle (with increasing distance for many classes), LDC predictions are very close to the oracle prototypes (mean of distributions using LDC is close to 0 even after the last task and has low standard deviation).

the last task, with much lower deviation. This analysis suggests that the LDC predicts the old prototypes very close to their oracle positions and thus better tracks the prototypes position when moving from one feature space to another.

## 5 Conclusions

We showed that performance degradation in exemplar-free CP accumulation methods is largely due to feature drift. We proposed a simple yet effective prototype correction method called Learnable Drift Compensation (LDC). LDC learns a projector that maps between the feature spaces of the consecutive tasks using only the data available at each task along with current and previous model. The trained projector is used at the end of each task to correct the stored old prototype positions. We demonstrate the efficacy of LDC in correcting the drift of any moving backbone by using it on top of already established exemplar-free CL strategies. Extensive experiments on CIFAR-100, Tiny-ImageNet, and ImageNet100 show that LDC outperforms previous drift correction methods in supervised and semi-supervised settings.

**Limitations.** While the learned projection  $p_f$  between tasks is effective for correcting drift, it’s not flawless. The gap between  $p_{corrected}$  and  $p_{oracle}$  in Fig. 1 highlights this difference. Without exemplars,  $p_f$  only learns the projection from current task data. Consequently, due to biases in the current data, updates to some old class prototypes may not be perfect. We would like to explore methods to further minimize or eliminate this gap in our future work.

**Acknowledgments.** We acknowledge projects TED2021-132513B-I00 and PID2022-143257NB-I00, financed by MCIN/AEI/10.13039/501100011033 and FSE+ by the European Union NextGenerationEU/PRTR, and by ERDF A Way of Making Europa, and the Generalitat de Catalunya CERCA Program. Bartłomiej Twardowski acknowledges the grant RYC2021-032765-I.

## References

1. Liu, Xialei, Wu, C., Menta, M., Herranz, L., Raducanu, B., Bagdanov, A.D., Jui, S., van de Weijer, J.: Generative feature replay for class-incremental learning. In: Conference on Computer Vision and Pattern Recognition Workshops (2020)
2. Aljundi, R., Babiloni, F., Elhoseiny, M., Rohrbach, M., Tuytelaars, T.: Memory aware synapses: Learning what (not) to forget. In: European Conference on Computer Vision (2018)
3. Assran, M., Caron, M., Misra, I., Bojanowski, P., Joulin, A., Ballas, N., Rabbat, M.: Semi-supervised learning of visual features by non-parametrically predicting view assignments with support samples. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8443–8452 (2021)
4. Caccia, L., Pineau, J.: Special: Self-supervised pretraining for continual learning. In: International Workshop on Continual Semi-Supervised Learning. pp. 91–103. Springer (2022)
5. Castro, F.M., Marín-Jiménez, M.J., Guil, N., Schmid, C., Alahari, K.: End-to-end incremental learning. In: European Conference on Computer Vision (2018)
6. Cha, H., Lee, J., Shin, J.: Co2l: Contrastive continual learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9516–9525 (2021)
7. Chaudhry, A., Ranzato, M., Rohrbach, M., Elhoseiny, M.: Efficient lifelong learning with a-gem. In: International Conference on Learning Representations (2019)
8. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: International Conference on Learning Representations (ICLR) (2021)
9. Douillard, A., Cord, M., Ollion, C., Robert, T., Valle, E.: 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. pp. 86–102. Springer (2020)
10. Douillard, A., Ramé, A., Couairon, G., Cord, M.: Dytox: Transformers for continual learning with dynamic token expansion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9285–9295 (2022)
11. Fini, E., Da Costa, V.G.T., Alameda-Pineda, X., Ricci, E., Alahari, K., Mairal, J.: Self-supervised models are continual learners. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9621–9630 (2022)
12. Gallardo, J., Hayes, T.L., Kanan, C.: Self-supervised training enhances online continual learning. arXiv preprint arXiv:2103.14010 (2021)
13. Gomez-Villa, A., Twardowski, Bartłomiej, Yu, Lu, Bagdanov, A.D., van de Weijer, J.: Continually learning self-supervised representations with projected functional regularization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3867–3877 (2022)

14. Gomez-Villa, A., Twardowski, B., Wang, K., van de Weijer, J.: Plasticity-optimized complementary networks for unsupervised continual learning. In: IEEE/CVF Winter Conference on Applications of Computer Vision (2024)
15. Goswami, D., Liu, Y., Twardowski, B., van de Weijer, J.: Fecam: Exploiting the heterogeneity of class distributions in exemplar-free continual learning. *Advances in Neural Information Processing Systems* **36** (2024)
16. Goswami, D., Soutif-Cormerais, A., Liu, Y., Kamath, S., Twardowski, B., van de Weijer, J.: Resurrecting old classes with new data for exemplar-free continual learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
17. Goswami, D., Twardowski, B., Van De Weijer, J.: Calibrating higher-order statistics for few-shot class-incremental learning with pre-trained vision transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2024)
18. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Conference on Computer Vision and Pattern Recognition (2016)
19. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. In: NIPS Deep Learning Workshop (2014)
20. Hou, S., Pan, X., Loy, C.C., Wang, Z., Lin, D.: Learning a unified classifier incrementally via rebalancing. In: International Conference on Computer Vision (2019)
21. Iscen, A., Zhang, J., Lazebnik, S., Schmid, C.: Memory-efficient incremental learning through feature adaptation. In: European Conference on Computer Vision (2020)
22. Janson, P., Zhang, W., Aljundi, R., Elhoseiny, M.: A simple baseline that questions the use of pretrained-models in continual learning. In: NeurIPS 2022 Workshop on Distribution Shifts: Connecting Methods and Applications (2022)
23. Kang, Z., Fini, E., Nabi, M., Ricci, E., Alahari, K.: A soft nearest-neighbor framework for continual semi-supervised learning. arXiv preprint arXiv:2212.05102 (2022)
24. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A.A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., et al.: Overcoming catastrophic forgetting in neural networks. *National Academy of Sciences* (2017)
25. Krause, J., Stark, M., Deng, J., Fei-Fei, L.: 3d object representations for fine-grained categorization. In: IEEE International Conference on Computer Vision Workshops. pp. 554–561 (2013)
26. Krizhevsky, A.: Learning multiple layers of features from tiny images. Tech. rep., University of Toronto (2009)
27. Le, Y., Yang, X.: Tiny imagenet visual recognition challenge. CS 231N (2015)
28. Li, Z., Hoiem, D.: Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2017)
29. Madaan, D., Yoon, J., Li, Y., Liu, Y., Hwang, S.J.: Representational continuity for unsupervised continual learning. In: ICLR (2022)
30. Magistri, S., Trinci, T., Soutif-Cormerais, A., van de Weijer, J., Bagdanov, A.D.: Elastic feature consolidation for cold start exemplar-free incremental learning. In: International Conference on Learning Representations (2024), <https://openreview.net/forum?id=7D9X2cFnt1>
31. Malepathirana, T., Senanayake, D., Halgamuge, S.: Napa-vq: Neighborhood-aware prototype augmentation with vector quantization for continual learning. In: International Conference on Computer Vision (ICCV) (2023)



32. Mallya, A., Lazebnik, S.: Packnet: Adding multiple tasks to a single network by iterative pruning. In: Conference on Computer Vision and Pattern Recognition (2018)
33. Masana, M., Liu, X., Twardowski, B., Menta, M., Bagdanov, A.D., Van De Weijer, J.: Class-incremental learning: survey and performance evaluation on image classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(5), 5513–5533 (2022)
34. McCloskey, M., Cohen, N.J.: Catastrophic interference in connectionist networks: The sequential learning problem. In: *Psychology of learning and motivation*, vol. 24, pp. 109–165. Elsevier (1989)
35. McDonnell, M.D., Gong, D., Parvaneh, A., Abbasnejad, E., van den Hengel, A.: Ranpac: Random projections and pre-trained models for continual learning. *Advances in Neural Information Processing Systems* **36** (2024)
36. Ostapenko, O., Lesort, T., Rodriguez, P., Arefin, M.R., Douillard, A., Rish, I., Charlin, L.: Foundational models for continual learning: An empirical study of latent replay. *arXiv preprint arXiv:2205.00329* (2022)
37. Panos, A., Kobe, Y., Reino, D.O., Aljundi, R., Turner, R.E.: First session adaptation: A strong replay-free baseline for class-incremental learning. In: *International Conference on Computer Vision (ICCV)* (2023)
38. Petit, G., Popescu, A., Schindler, H., Picard, D., Delezoide, B.: Fetril: Feature translation for exemplar-free class-incremental learning. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 3911–3920 (2023)
39. Rao, D., Visin, F., Rusu, A., Pascanu, R., Teh, Y.W., Hadsell, R.: Continual unsupervised representation learning. In: *Advances in Neural Information Processing Systems* (2019)
40. Rebuffi, S.A., Kolesnikov, A., Sperl, G., Lampert, C.H.: icarl: Incremental classifier and representation learning. In: *Conference on Computer Vision and Pattern Recognition* (2017)
41. Ridnik, T., Ben-Baruch, E., Noy, A., Zelnik-Manor, L.: Imagenet-21k pretraining for the masses. In: *Conference on Neural Information Processing Systems Datasets and Benchmarks Track* (2021)
42. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *International journal of computer vision* (2015)
43. Rusu, A.A., Rabinowitz, N.C., Desjardins, G., Soyer, H., Kirkpatrick, J., Kavukcuoglu, K., Pascanu, R., Hadsell, R.: Progressive neural networks. *arXiv* (2016)
44. Serra, J., Suris, D., Miron, M., Karatzoglou, A.: Overcoming catastrophic forgetting with hard attention to the task. In: *International Conference on Machine Learning* (2018)
45. Shi, W., Ye, M.: Prototype reminiscence and augmented asymmetric knowledge aggregation for non-exemplar class-incremental learning. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 1772–1781 (2023)
46. Toldo, M., Ozay, M.: Bring evanescent representations to life in lifelong class incremental learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 16732–16741 (2022)
47. van de Ven, G.M., Tolias, A.S.: Three scenarios for continual learning. In: *NeurIPS Continual Learning Workshop* (2018)

48. Wu, Y., Chen, Y., Wang, L., Ye, Y., Liu, Z., Guo, Y., Fu, Y.: Large scale incremental learning. In: Conference on Computer Vision and Pattern Recognition (2019)
49. Yu, L., Twardowski, Bartlomiej, Liu, Xialei, Herranz, L., Wang, K., Jui, S., Weijer, J.v.d.: Semantic drift compensation for class-incremental learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6982–6991 (2020)
50. Zbontar, J., Jing, L., Misra, I., LeCun, Y., Deny, S.: Barlow twins: Self-supervised learning via redundancy reduction. arXiv preprint arXiv:2103.03230 (2021)
51. Zenke, F., Poole, B., Ganguli, S.: Continual learning through synaptic intelligence. In: International Conference on Machine Learning (2017)
52. Zhang, G., Wang, L., Kang, G., Chen, L., Wei, Y.: Slca: Slow learner with classifier alignment for continual learning on a pre-trained model. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (2023)
53. Zhou, D.W., Sun, H.L., Ning, J., Ye, H.J., Zhan, D.C.: Continual learning with pre-trained models: A survey. arXiv preprint arXiv:2401.16386 (2024)
54. Zhou, D.W., Wang, F.Y., Ye, H.J., Zhan, D.C.: Pycil: a python toolbox for class-incremental learning. SCIENCE CHINA Information Sciences (2023)
55. Zhu, F., Cheng, Z., Zhang, X.y., Liu, C.l.: Class-incremental learning via dual augmentation. Advances in Neural Information Processing Systems **34**, 14306–14318 (2021)
56. Zhu, F., Zhang, X.Y., Wang, C., Yin, F., Liu, C.L.: Prototype augmentation and self-supervision for incremental learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5871–5880 (2021)
57. Zhu, K., Zhai, W., Cao, Y., Luo, J., Zha, Z.J.: Self-sustaining representation expansion for non-exemplar class-incremental learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9296–9305 (2022)
58. Zhu, K., Zheng, K., Feng, R., Zhao, D., Cao, Y., Zha, Z.J.: Self-organizing pathway expansion for non-exemplar class-incremental learning. In: International Conference on Computer Vision (ICCV) (2023)