

Plasticity-Optimized Complementary Networks for Unsupervised Continual Learning

Alex Gomez-Villa^{1,2}, Bartłomiej Twardowski^{1,2,3}, Kai Wang^{1,2*}, Joost van de Weijer^{1,2}

¹Computer Vision Center, Barcelona, Spain

²Universitat Autònoma de Barcelona, Barcelona, Spain

³IDEAS NCBR, Warsaw, Poland

{agomezvi,btwardowski,kwang,joost}@cvc.uab.es

Abstract

Continuous unsupervised representation learning (CURL) research has greatly benefited from improvements in self-supervised learning (SSL) techniques. As a result, existing CURL methods using SSL can learn high-quality representations without any labels, but with a notable performance drop when learning on a many-tasks data stream. We hypothesize that this is caused by the regularization losses that are imposed to prevent forgetting, leading to a suboptimal plasticity-stability trade-off: they either do not adapt fully to the incoming data (low plasticity), or incur significant forgetting when allowed to fully adapt to a new SSL pretext-task (low stability). In this work, we propose to train an expert network that is relieved of the duty of keeping the previous knowledge and can focus on performing optimally on the new tasks (optimizing plasticity). In the second phase, we combine this new knowledge with the previous network in an adaptation-retrospection phase to avoid forgetting and initialize a new expert with the knowledge of the old network. We perform several experiments showing that our proposed approach outperforms other CURL exemplar-free methods in few- and many-task split settings. Furthermore, we show how to adapt our approach to semi-supervised continual learning (Semi-SCL) and show that we surpass the accuracy of other exemplar-free Semi-SCL methods and reach the results of some others that use exemplars.

1. Introduction

Continual learning (CL) designs algorithms that can learn from shifting distributions (non-IID data), generally this is modeled by learning from a sequence of tasks [14]. The main challenge for these methods is the problem of catastrophic forgetting [41], which is a dramatic drop in performance on previous tasks. Most CL approaches, therefore, need

to address the trade-off between acquiring new knowledge (plasticity) and preventing forgetting of previous knowledge (stability). The vast majority of existing methods in continual learning have focussed on supervised learning, where the incoming data stream is fully labeled. In this paper, we focus on continual learning on unsupervised data.

Only recently, some works have explored continual learning on unsupervised non-IID data-streams [22, 19]. Motivated by the tremendous progress in unsupervised learning, notably of contrastive learning approaches [11, 61], methods aim to extend these methods to the continual setting. An additional motivation is the fact that unsupervised learning tends to lead to more generalizable feature representations since features that are not relevant to the specific discriminative task are not automatically discarded. This can potentially lead to representations that can faster incorporate new tasks without incurring significant amounts of forgetting. PFR [22] uses a projection head after the feature extractor of an SSL framework to predict past representations. Projected representations are motivated to be close to past representations; therefore, the present model is encouraged to remember past knowledge. CaSSLe [19] uses a similar strategy as PFR, but the projection head is used after the projector of the SSL approach. Even though these methods obtain satisfactory results, they struggle to adapt to new tasks without jeopardizing the vast knowledge already accumulated by the network. We hypothesize that the regularization imposed by these methods to avoid forgetting hurts the learning process of CURL in the following ways: 1) the SSL component cannot fully adapt to the incoming data (low plasticity), 2) The model will have a significant drift (forgetting) when the current model is unable to perform the SSL pretext-task. These effects increase as the number of tasks increases, and consequently, the data for training reduces.

Complementary learning systems (CLS) theory [40, 31] proposes a computational framework in which the interplay between a fast (episodic memory/specific experience) and

*Corresponding author

a slow (semantic/general structured) memory system is the core of the mechanism of knowledge consolidation. Several existing CL methods have taken inspiration from CLS as a way to find a good stability-plasticity trade-off (see [44, 45] for a review). The fast learner can quickly adapt to new knowledge, which then is carefully absorbed by the slower learner. DualNet [47] proposes to use a self-supervised method to train the slow, more generic learner, which weights can be quickly adapted for solving a supervised task with exemplars from the replay buffer. Recently, in [2] the authors proposed to use a CLS-based approach and maintain separate plastic and stable models for online CL with experience replay. However, existing methods that exploit CLS for continual learning have in common that they only consider a supervised learning scenario.

This paper aims to apply complementary learning systems theory to improve continual learning from unsupervised data streams. The existing methods [22, 19] can suffer from sub-optimal stability and plasticity on longer sequences since they have difficulty adapting to the new knowledge required to address the latest task, while maintaining the vast knowledge already learned on earlier tasks. Instead, we propose to train an expert network that is relieved of the task of keeping the previous knowledge and can focus on the task of performing optimally on new tasks. In the second phase, we combine this new knowledge with the old network in an adaptation-retrospection phase to avoid forgetting.

In conclusion, the main contributions of this work are :

- A new exemplar-free continual unsupervised representation learning (CURL) method called *Plasticity-Optimized COmplementary Networks* (POCON). Existing CURL methods learn new knowledge while imposing regularization to prevent forgetting. Instead, POCON separates the learning of new knowledge from the knowledge integration part. Analysis confirms that this leads to a better stability-plasticity trade-off.
- Extensive experiments confirm that POCON outperforms state-of-the-art CURL on various settings (e.g., a 5-9 % performance gain over CaSSLe on ImageNet100 for a 20-100 task-split). Unlike previous CURL methods, POCON can thrive in low-data regimes (such as small-task incremental learning) and setups without task boundaries. We also demonstrate the application of POCON to semi-supervised continual learning.
- We propose and evaluate a *heterogeneous* version of POCON, where the main network can have a different network architecture than the expert. This opens up the possibility for interesting applications where a slow/big network can be deployed in a cloud environment, while a fast/slow learner can be utilized on an edge device, such as a mobile phone.

2. Related work

Continual Learning and Class Incremental Learning.

Existing continual learning methods can be broadly categorized into three types: replay-based, architecture-based, and regularization-based methods [14, 38]. Replay-based methods either save a small amount of data from previously seen tasks [4, 10] or generate synthetic data with a generative model [56, 64]. The replay data can be used during training together with the current data, such as in iCaRL [49] and LUCIR [25], or to constrain the gradient direction while training, such as in AGEM [9]. Architecture-based methods activate different subsets of network parameters for different tasks by allowing model parameters to grow with the number of tasks. Previous works following this strategy include HAT [52], Piggyback [36], PackNet [37], DER [59] and Ternary Masks [39]. Regularization-based methods add a regularization term derived from knowledge of previous tasks to the training loss. This can be done by either regularizing the weight space, which constrains important parameters [53, 55], or the functional space, which constrains predictions or intermediate features [18, 13, 26]. EWC [29], MAS [1], REWC [34], SI [62], and RWalk [8] constrain the importance of network parameters to prevent forgetting. Methods such as LwF [33], LwM [16], and BiC [57] leverage knowledge distillation to regularize features or predictions. DMC [65] work is more related to POCON as the authors proposed to train the expert network without any regularization for a classification task. However, after that, in a distillation phase an additional auxiliary dataset is used to integrate old and new knowledge.

Self-supervised representation learning. In recent years, unsupervised methods based on self-supervision have become dominant in learning representation for computer vision systems. The aim of self-supervised learning is to acquire high-quality image representations without explicit labeling. Initially, these methods addressed some well-defined pretext tasks, such as predicting rotation [21], determining patch position [17], or solving jigsaw puzzles in images [43], and labels for these discriminative pretext tasks can be automatically computed to enable learning of meaningful feature representations of images. Recently, researchers have adapted contrastive methods for use with unlabeled data and have placed more emphasis on instance-level data augmentation to find similar or contrasting samples [5, 7, 11, 23, 61]. These methods heavily rely on stochastic data augmentation [58, 67] to generate sufficient similar examples for learning representations, and negative examples are either randomly sampled or excluded entirely [12].

Self-supervised learning has also been used to improve the learning of a sequence of supervised tasks [61, 66]. Their objective is not to learn from unlabeled data, but rather to use self-supervised learning to further enrich the

Method	Labels	Exemplars	Regularization
DMC [65]	✓	✓*	✗
DualNet [47]	✓	✓	✓
CLS-ER [2]	✓	✓	✓
LUMP [35]	✗	✓	✓
PFR [22]	✗	✗	✓
CaSSLe [19]	✗	✗	✓
POCON (ours)	✗	✗	✗

Table 1: Comparison of SSL-based CURL. Only POCON does not use any regularization during the training, what results in higher plasticity for the current task. *DMC uses an auxiliary dataset for distillation, instead of exemplars.

feature representation. Similarly, pre-trained models with self-supervision have been used to improve incremental average classification metrics [20] with data augmentation, distillation, and even exemplars. Training self-supervised models directly on class-IL setting without exemplars was also proposed in [22, 19], where both present results that self-supervised learning mitigates the problem of forgetting.

Complementary learning systems. There are CLS-based methods that use several networks in addition to a rehearsal memory or pseudo-sample generator. FearNet [28] uses a hippocampal network for recalling new examples, a PFC network for long-term memories, and a third network to select between the PFC or hippocampal networks for a particular instance. In [46], they propose a G-EM network that performs unsupervised learning from spatiotemporal representations and a G-SM network that uses signals from G-EM and class labels to learn from videos incrementally. Closer to our work are DualNet [47] and CLS-ER [2] (previously explained); however, both models are supervised and use exemplars. Table 1 presents a summary of similar CL methods.

3. Method

In this section, we describe our approach for continual learning of self-supervised representations, referred to as Plasticity-Optimized Complementary Networks (POCON), which eliminates the need for memory or replay. Our method is based on the complementary learning system (CLS) framework, and involves an interplay between a fast-expert learner and a slow-main network. Our work is motivated by the recognition that fast adaptation to new data is crucial in constructing more robust representations during continual unsupervised representation learning. Rather than attempting to maintain network stability in its old state through strict or relaxed distillation methods, as suggested in recent works [19, 22], POCON allows the expert network to learn a new task freely without any restrictions. After acquiring

the new knowledge, we integrate it into the main network. In turn, the main network serves as a good starting point for the new expert. Before presenting the details of the POCON method, we first provide a brief introduction to the problem of continual self-supervised representation learning.

3.1. Self-supervised representation learning

In recent research on self-supervised learning, the objective is to train a network $f_\theta : \mathcal{X} \rightarrow \mathcal{F}$, which maps input data from \mathcal{X} to output feature representations in \mathcal{F} . The network is trained using unlabeled input data x sampled from a distribution \mathcal{D} . The learned feature representation is subsequently used to facilitate various downstream tasks. In this paper, we employ the BarlowTwins approach [61] for self-supervised learning of the representation network g_θ . This approach serves as a common baseline for previous works [19, 22]. However, the proposed method is versatile and can be extended to other self-supervised techniques.

In Fig. 1 (left) the BarlowTwins architecture is presented. Both branches use a projector network z and both share the same parameters, but while computing the empirical cross-correlation loss operates on different views of the same sample x created by data augmentation techniques. For simplicity of notation, we omit the explicit mention of z network parameters, since it is not utilized by downstream tasks. The BarlowTwins method eliminates the need for explicit negative samples and achieves comparable performance while maintaining computational efficiency. It assumes that negatives are accessible in each mini-batch to estimate correlations among all samples in it. The network is trained by minimizing an invariance and a redundancy reduction term in the loss function [61]. Here, different augmented views X_A and X_B of the same data samples X are taken from the set of data augmentations \mathcal{D}^* . This leads to the loss defined as:

$$\mathcal{L}_c = \mathbb{E}_{X_A, X_B \sim \mathcal{D}^*} \left[\sum_i (1 - \mathcal{C}_{ii})^2 + \lambda \sum_i \sum_{j \neq i} \mathcal{C}_{ij}^2 \right], \quad (1)$$

where λ is a positive constant trade-off parameter between both terms, and where \mathcal{C} is the cross-correlation matrix computed between the representations z of all samples X_A and X_B in a mini-batch indexed by b : $\mathcal{C}_{ij} = \sum_b z_{b,i}^A z_{b,j}^B / (\sqrt{\sum_b (z_{b,i}^A)^2} \sqrt{\sum_b (z_{b,j}^B)^2})$. The cross-correlation matrix \mathcal{C} contains values ranging from -1.0 (worst) to 1.0 (best) for the correlation between the projector’s outputs: $Z_A = z(g_\phi(X_A))$ and $Z_B = z(g_\phi(X_B))$. The invariance term of the loss function encourages the diagonal elements to have a value of 1. This ensures that the learned embedding is invariant to the applied data augmentations. Meanwhile, the second term (redundancy reduction) maintains the off-diagonal elements close to zero and decorrelates the outputs of non-related images.

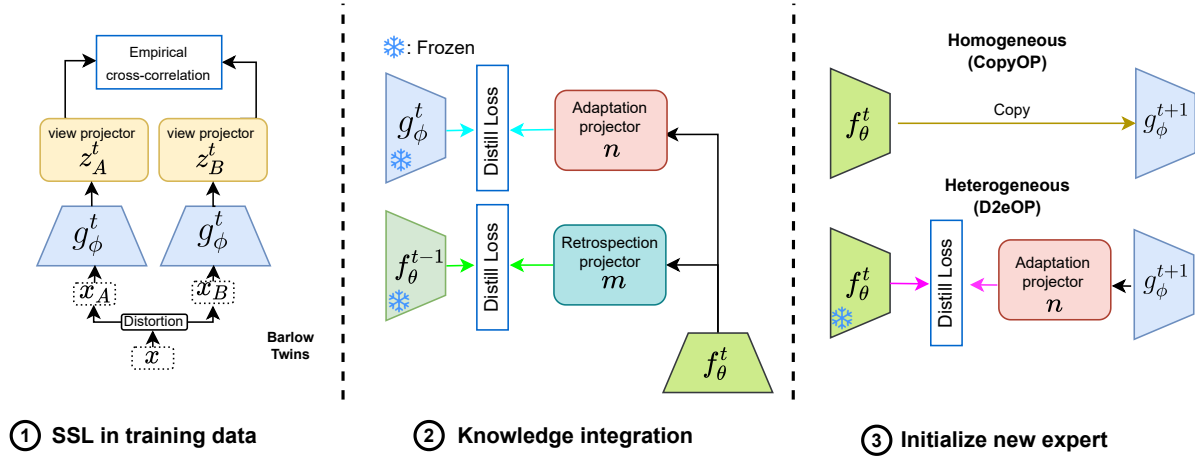


Figure 1: Overview of the POCON method that uses multi-stage training of complementary learning system for unsupervised continual representation learning. Three stages of training allow POCON to maintain fast and efficient adaptation of the fast expert network while maintaining stable representation in a slow learner – the main network. To avoid forgetting, knowledge from the new task expert is integrated (adaptation) with the previous state of the slow network (retrospection). The last stage of the training is dedicated for the best preparation of the new expert for the new task.

3.2. Continual SSL Problem Definition

In this work, we consider a CL scenario in which the feature extractor f_θ must learn from a set of task $\{1, \dots, T\}$ from different distributions, where each task t from that set follows the distributions \mathcal{D}_t . We would like to find the parameters θ of the feature extractor f_θ that minimizes the summed loss over all tasks T :

$$\arg \min_{\theta} \sum_{t=1}^T \mathcal{L}_c^t, \quad (2)$$

where $\mathcal{L}_c^t = \mathbb{E}_{X_A, X_B \sim \mathcal{D}_t^*}[\mathcal{L}_c]$ and \mathcal{L}_c is defined as in Eq. 1. However, finding the right θ poses the main problem in continual learning, as previous data D_1, \dots, D_{t-1} is not available at time t and Eq. 2 cannot be minimized directly.

3.3. Plasticity-Optimized Complementary Networks

We propose Plasticity-Optimized Complementary Networks (POCON) based on CLS framework. POCON is composed of three stages training (see Fig. 1 for details): ① learn expert knowledge from current data task, ② integrate new expert knowledge to the main network, and ③ initialize the new expert from the updated main network. Each stage is explained in the details in the next sections.

Stage ①: SSL of expert for the current task. In this step, we are interested to fully-adapt to the input training data of the current task. Hence, a feature extractor g_ϕ^t is used to learn in a self-supervised way (following Eq. 1) on the data D_t . Note that unlike previous methods (PFR and CaSSLe), we do not constraint in any way our expert during the training (like

imposing regularization to prevent forgetting). We allow the expert network to be fully plastic and optimal for learning representation for the current task.

Stage ②: Knowledge integration. Once we absorb the knowledge of the current task in the expert network at Stage ①, it needs to be transferred and accumulated to the main feature extractor f_θ without forgetting previous tasks (see Fig. 1 ②). To do so, we employ an additional *adaptation projector* $n: \mathcal{Z} \rightarrow \mathcal{W}$, which maps the embedding space from the main network f_θ to the embedding space learned in the expert network g_ϕ^t . Then, to avoid forgetting, we use a *retrospection projector* $m: \mathcal{Z} \rightarrow \mathcal{Z}$ that maps the embedding space learned on the current task back to the embedding space learned on the previous ones. The final loss function for the knowledge integration stage consists of an adaptation and retrospection component:

$$\mathcal{L}_{INT}^t = \mathbb{E}_{X_A \sim \mathcal{D}_t^*} \left[\sum_{i=0}^{|X_A|} \| n(g_{\phi^t}(x_a)) - f_{\theta^t}(x_a) \| + \| m(f_{\theta^t}(x_a)) - f_{\theta^{t-1}}(x_a) \| \right] \quad (3)$$

where both sources from which we integrate our knowledge: previous main network $f_{\theta^{t-1}}$ and the current expert g_ϕ^t are frozen, and only f_{θ^t} and adaptor networks n and m are being updated by this loss function. The goal is to integrate knowledge using distillation and current task data.

Stage ③: New expert initialization. In order to begin the training on the next task, the expert g_ϕ^{t+1} must be prepared to solve the new task with the best prior knowledge for training

new task representation efficiently. In this stage, we need to initialize a new expert (see Fig. 1 ③) in the best way. Improper initialization can influence the training epochs in Stage ① and make the problem of adaptation for projector n more difficult in Stage ②. We looked into two potential initialization setups based on the similarities between the expert and main backbones. The homogeneous setup, where both the main network and the expert network share the same architecture. This is the default setup in our experiments. In addition, we will also consider the heterogeneous setup, where the expert has another architecture than the main network. This allows, for example, to apply smaller networks when per task data is limited, or computation should be performed on edge devices.

Homogeneous setup (CopyOP): In order to begin the training of g_ϕ^{t+1} with all the accumulated knowledge till t we can just copy the weight of the main network f_θ^t into g_ϕ^{t+1} . This operation avoids the recency bias of (g_ϕ^t) and provides an excellent initialization point for g_ϕ^{t+1} to continue learning. Furthermore, CopyOP makes the problem of the adaptation projector n easier since g_ϕ^{t+1} will have a similar representation to f_θ^t . The main drawback of CopyOP is that it constrains of POCON to use the same architecture for the main f_θ^t and the expert g_ϕ^{t+1} networks.

Heterogeneous setup (D2eOP): This distillation-based initialization method allows the use of heterogeneous network architectures for f_θ^t and g_ϕ^{t+1} in POCON. In order to transfer the knowledge, we propose a projected distillation as in Stage ② using a fresh adaptation projector n . Despite being more computationally demanding, this way of transferring knowledge offers one big advantage – different architecture for g_ϕ^{t+1} allows the use of a smaller backbone network or even very different ones, e.g. ViT. Using a smaller backbone is useful for low-data regimes (as we will show later) or for devices low in computational power at the time of learning the expert (robotics, edge devices). The loss function for D2eOP in Stage ③ is given as:

$$\mathcal{L}_{D2eOP}^t = \mathbb{E}_{X_A \sim \mathcal{D}_t^*} \left[\sum_{i=1}^{|X_A|} \| n(g_\phi^{t+1}(x_a)) - f_{\theta^t}(x_a) \| \right] \quad (4)$$

where n is a projector that adapts the embedding space of new expert g_ϕ^{t+1} to the previous one from the main network with current task data D^t . The difference of D2eOP and CopyOP is presented in the right column of Fig. 1.

Other initialization options exists. We discuss them and provide more results in the Appendix.

3.4. Plasticity-stability trade-off in CURL with SSL

Our work is motivated by two key observations:

- Regularized-based CURL models (PFR and CaSSLe) have an implicit handicapped plasticity. During the

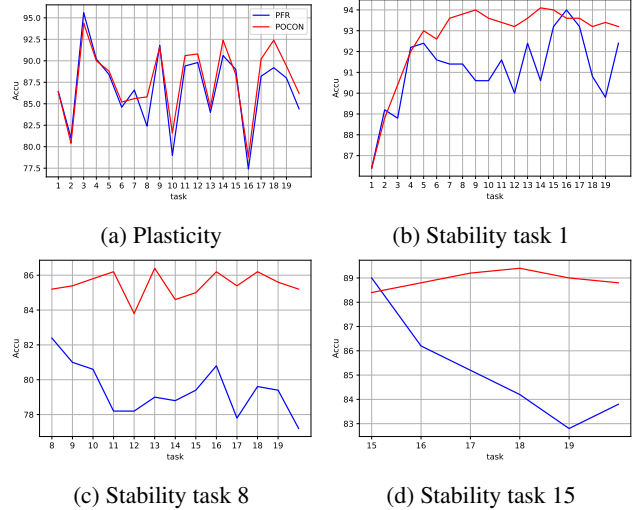


Figure 2: Plasticity-stability trade-off in POCON and PFR: (a) accuracy of a current task during the continual learning session, (b-d) accuracy of task 1, 8, and 15 in the following incremental steps of learning the representation with POCON and PFR. POCON presents superior stability to PFR at the same time still maintaining better plasticity, with its non-restrictive training of the expert network.

training, they learn new tasks data using SSL while maintaining the backbone representations close to the previous tasks to avoid forgetting. Therefore, the backbone network cannot fully adapt to current data due to the regularization imposed by the CL method.

- As the number of tasks increases, current CURL models lose stability. In this case, the data distribution changes more abruptly and the backbone is pushed to follow an imperfect estimation of the new distribution restricted by the regularization of the old model (see first observation). Hence, these models are not as plastic as finetuning and do not present stable behavior.

In Fig. 2 we present the stability-plasticity tradeoff for twenty task partition of CIFAR100 in PFR and our method (POCON). For the plasticity plot, we evaluate the accuracy at the end of each task t using task t data for training and evaluation. The stability plots were made by training and evaluating in a fixed task partition (1, 8, and 15 for Fig. 2) after the training of each task t .

As Fig. 2 a) shows POCON has a higher accuracy (more plastic) in most of the task during the whole training. As training an expert in stage ① is not constrained by any regularization. Moreover, Fig. 2 b), c), and d) display how POCON is able to retain previous representations over the whole incremental learning session, where our double distillation (adaptation and retrospection) is able to retain the learned representation correctly.

4. Experimental Results

4.1. Experimental setup

Datasets We use the following datasets: *CIFAR-100* [30], consists of 100 object classes in 45,000 images for training, 5,000 for validation, and 10,000 for test with 100 classes. All images are 32×32 pixels; *TinyImageNet* a rescaled subset of 200 ImageNet [15] classes used in [54] and containing 64×64 pixel images. Each class has 500 training images, 50 validation images, and 50 test images; *ImageNet100* a subset of one hundred classes from ImageNet [15] that contains 130k images of size 224×224 pixels.

Training procedure In all experiments, we train ResNet-18 [24] (or ResNet-9 for the heterogeneity experiment) for expert and main network using SGD with an initial learning rate of 0.01 and a weight decay of 0.0001 for 250 epochs (200 for ImageNet100) in *Stage ①*. For *Stage ②* same optimization procedures as *Stage ①* is followed but for 500 epochs (400 for ImageNet100). The data augmentation used in *Stage ①* is the same as in BarlowTwins [11]). Based on the self-supervised distillation ideas of [51, 42], we use a four-layer MLP projector as adaptation and retrospection projectors following the architecture of [42].

Downstream classifiers are by default linear and trained with a CE-loss and Adam optimizer with a learning rate $5e-2$ on CIFAR-100, and 3 on TinyImageNet. We use validation data to implement a patience scheme that lowers the learning rate by a factor of 0.3 and 0.06 up to three times while training a downstream task classifier. For ImageNet100 we use the same training and evaluation procedure as [19].

Baseline methods We only compare to exemplar-free methods and exclude methods that require replay from our comparison¹. *Fine-tuning (FT)*: The network is trained sequentially on each task without access to previous data and with no mitigation of catastrophic forgetting. *Joint*: We perform joint training with fine-tuning on all data which provides an upper bound. Equivalent to having a single-task scenario. *PFR* [22] and *CaSSLe* [19] with Barlow Twins: We use the code and hyperparameters provided by the authors, in PFR we used $\lambda = 25$ for all the experiments. In continual semi-supervised learning (section 4.4), we consider the following methods. Regularization-based methods: Learning without Forgetting (LwF) [33], online Elastic Weight Consolidation (oEWC) [29], Replay-based methods: Experience Replay (ER) [50], iCaRL [49] and GDumbciteprabhu2020gumb, and Continual semi-supervised learning methods: CCIC [6], PAWS [3] and NNCSL [27].

4.2. Continual representation learning

In this experiment, we evaluate all methods in the continual representation learning setting, where each task consist

CIFAR-100 (32x32)					
Method	4 tasks	10 tasks	20 tasks	50 tasks	100 tasks
FT	54.8	50.94	44.95	38.0	27.0
CaSSLe	59.80	52.5	49.6	45.3	42.10
PFR	59.70	54.33	44.80	46.5	43.30
POCON	63.7	60.5	56.8	48.9	48.94
Joint	65.4				
TinyImagenet (64x64)					
Method	4 tasks	10 tasks	20 tasks	50 tasks	100 tasks
FT	41.95	36.55	32.29	22.34	2.80
CaSSLe	46.37	41.53	38.18	28.08	25.38
PFR	42.23	39.20	31.22	25.87	21.20
POCON	40.97	41.06	41.14	37.20	30.24
Joint	50.18				
ImageNet100 (224x224)					
Method	5 tasks	10 tasks	20 tasks	50 tasks	100 tasks
FT	56.10	48.13	42.73	39.64	21.03
CaSSLe	67.56	59.78	53.92	46.64	36.44
PFR	66.12	60.46	54.84	42.18	38.34
POCON	66.30	61.36	59.32	53.50	45.40
Joint	71.06				

Table 2: Accuracy of a linear evaluation on equal split various datasets and different number of tasks with ResNet-18. POCON present better result than other regularization-based methods and maintain high accuracy even with increasing number of tasks.

of a distinct set of classes from a single dataset. Splits are prepared similarly to the class incremental learning setting, but without access to labels. Specifically, we split datasets into four, ten, twenty, fifty, and one hundred equal tasks as done in [49]. In each task, we perform SSL (*Stage ①*), knowledge integration (*Stage ②*), and a new expert initialization (*Stage ③*). In the evaluation phase, we train a linear classifier using the learned representation of the main network encoder. Please note that the expert network is never used for evaluation (unless specified). We use all available test data to obtain the overall task-agnostic performance evaluation². In all our tables we report the accuracy in the last task.

Homogeneous setup. Table 2 presents the results for all the methods on three commonly used dataset - CIFAR100, TinyImageNet, and ImageNet100. For CIFAR100 the upper bound of Joint training for a single task is 65.4%. Note, that the gap between FT and *Joint* is getting bigger with more tasks, from 10.6% for four task up to 38.4% for extreme case of one hundred tasks. POCON is significantly better than CaSSLe and PFR for different number of tasks. In

¹Code available at https://github.com/alviur/pocon_wacv2024

²We use *task-agnostic* evaluation in this paper to refer to the class-incremental learning evaluation [38] as in CaSSLe and PFR methods

Method	Arch	4 tasks	10 tasks	20 tasks	50 tasks	100 tasks
POCON	R18-R18	63.7	60.5	56.8	48.9	48.94
POCON	R18-R9	62.05	60.5	57.48	49.7	47.94
POCON	R9-R9	58.34	58.32	56.07	53.3	51.22

Table 3: Accuracy of linear evaluation with heterogeneous POCON architecture using ResNet-18 (R18) and ResNet-9 (R9) on split CIFAR-100. With increasing number of tasks and lower data POCON benefits from a smaller ResNet-9 network, first, only for the expert (20 tasks) and later (50 tasks) for both networks, expert and main.

four tasks setting, POCON is only 1.7% points lower than Joint, while second next, CaSSLe reaches 3.3% points lower accuracy. With the increasing number of tasks, the CaSSLe performance drops faster than PFR, while POCON maintains superior results against others and presents the lowest decrease in performance.

The results for TinyImagenet and ImageNet100 follow the same procedure as for CIFAR-100, but here we have larger images (64x64 and 224x224) and more classes for TinyImagenet (200). In this case, POCON outperforms other methods whenever the number of tasks is higher, 20 for TinyImageNet and 10 for ImageNet100. For a few tasks scenario, POCON will be close to other CL methods due to high data availability. The accuracy gap between POCON and other methods increases with the number of tasks.

Heterogeneous setup. An expert in POCON can use a different network architecture than the main network. That opens the possibility of using a smaller network for the expert whenever this can be beneficial, e.g., tasks are small with not enough data to train a large ResNet-18 network, or the device where the expert network is being trained is not powerful enough (robot, edge). We investigated heterogeneous architecture use in POCON with a smaller network, ResNet-9 which has 6.5M number of parameters instead of 11M in ResNet-18. The results are presented in Table 3. The different combinations of POCON are presented for using smaller network in expert only, or for both, expert and main networks. With increasing number of tasks it is more beneficial to use smaller expert (20 tasks). And having less data per task (50 and 100 tasks) we see improved results when as well the main network is smaller, we gain 4.6% changing from ResNet-18 to ResNet-9 for one hundred tasks.

4.3. Task-free setting

To show the ability of POCON to handle varying data streams, we test it in the task-free setting. In this setting, there is no explicit boundary between tasks, the data from one task changes smoothly to the data from the next [63, 32]. This prevents methods from having a fixed point where the network can change and prepare for the new task; the adaptation is ongoing. For instance, when we receive data

Method	exemplar	CIFAR-100		
		0.8%	5.0%	25.0%
<i>Exemplar-based methods</i>				
ER [50]	✓	8.2 ± 0.1	13.7 ± 0.6	17.1 ± 0.7
iCaRL [49]	✓	3.6 ± 0.1	11.3 ± 0.3	27.6 ± 0.4
GDumb [48]	✓	8.6 ± 0.1	9.9 ± 0.4	10.1 ± 0.4
CCIC [6] (500)	✓	11.5 ± 0.7	19.5 ± 0.2	20.3 ± 0.3
PAWS [3] (500)	✓	16.1 ± 0.4	21.2 ± 0.4	19.2 ± 0.4
NNCSL [27] (500)	✓	27.4 ± 0.5	31.4 ± 0.4	35.3 ± 0.3
<i>Exemplar-free methods</i>				
Fine-tuning	✗	1.8 ± 0.2	5.0 ± 0.3	7.8 ± 0.1
LwF [33]	✗	1.6 ± 0.1	4.5 ± 0.1	8.0 ± 0.1
oEWC [29]	✗	1.4 ± 0.1	4.7 ± 0.1	7.8 ± 0.4
Prototypes	✗	19.2 ± 0.9	23.5 ± 0.4	24.1 ± 0.1
Prototypes+SDC	✗	22.7 ± 0.6	27.6 ± 0.4	28.5 ± 0.1

Table 4: Semi-supervised continual learning comparison on CIFAR100 dataset. The number between brackets indicates the size of the memory buffer. We highlight the best method in each group with **bold** fonts.

D_t there is a mix with the data D_{t-1} . At some point, we only get data from D_t , but later on, we will get a mix with D_{t+1} . For this experiment, we employ the data partition of [63] for the CIFAR100 dataset with beta equal to 4 (please see the Appendix for more details).

Since there is no clear boundary, we cannot perform distillation without losing data from the stream of mini-batches. In this setting, POCON performs *Stage 2* in parallel with *Stage 1*. In order to do so, a frozen copy of the expert g_ϕ is used for the *Stage 2* while a new expert is learning on the current data. After s steps, a copy of the expert g_ϕ is passed to perform a distillation for ds steps. Note that we do not store any data; distillation for *Stage 2* and SSL in *Stage 1* is always performed using the current mini-batch data. *Stage 3* is omitted, as the expert network is not changed and initialized, as there is no task switch.

To compare to other method, we adapted PFR to work in the task-free setting similarly. In this case, the feature extractor of past data is updated after s steps (copy of the current feature extractor), and the regularization is performed during the whole training as in normal PFR. We also present results with a simple fine-tuning (FT).

Fig. 3 presents the results of linear evaluation of the learned representation in continual learning for the ten, twenty, and fifty data partitions settings. Only at the first steps at the beginning, for 10 and 20 tasks, PFR and FT accuracy is above POCON. However, that changes significantly in the following steps in favor of POCON, which continues improving learned representations. Unlike PFR and FT, POCON has a more stable learning curve for all task splits. The improvements over PFR and FT are better for more tasks, since the regularization-based method hurts plasticity while learning new tasks with PFR.




Figure 3: Task-task setting for CIFAR-100 with different length of incremental learning sequence. Plots present linear evaluation of learned representation for the blurred 10, 20, and 30 data partitions respectively (like in [63, 32]). POCON presents significantly better results with stable learning accuracy curves.

4.4. Continual semi-supervised learning

We propose a simple extension of POCON to the semi-supervised setting, where a small percentage of the data is labeled. We train the method in the same fashion as the unsupervised case, and we initially ignore available labels. After updating f_{θ}^t until convergence with POCON, we initialize the prototypes of each class as the average of the labeled samples. Then we assign all unlabeled data to the nearest prototype center, and we again compute the average of all samples assigned to each prototype. We perform the same procedure after each of the tasks. We call this method *Prototypes*. We also show for *Prototypes+SDC* where we use Semantic Drift Compensation [60] to prevent forgetting - we estimate and compensate for learned class prototypes drift.

In Table 4 we present the results for continual semi-supervised learning under three levels of supervision, namely with only 0.8%, 5.0% and 25.0% of samples labeled (following the settings of NNCSL [27] and CCIC [6]). Simply applying prototypes with POCON achieves much better results than the other exemplar-free methods (LwF [33] and oEWC [29]). The method outperforms the continual learning methods which only exploit the labeled data (such as ER, iCaRL, and GDumb). Furthermore, it also outperforms the semi-supervised methods CCIC and PAWS. Note that our method, without storing any exemplars, is only outperformed by the recent NNCSL method which requires exemplars and is dedicated to the semi-supervised learning setting.

5. Conclusions and Future directions

We proposed a method for exemplar-free continual unsupervised representation learning called *Plasticity-Optimized COmplementary Networks*. POCON trains an expert network that performs optimally on the current data without any regularization (optimizing plasticity). The knowledge is transferred to the main network through distillation while retaining similarity with the old main network to avoid for-

getting. Experiments of CIFAR100, TinyImagenet, and ImageNet100 show that our approach outperforms other methods for SSL exemplar-free CL learning (PFR and CaSSL), and it is especially good for many tasks scenarios.

The POCON method presents several opportunities for improvement. One promising direction is to extend the method to the scenario where multiple experts can be trained in parallel on different tasks, similar to the clients in federated learning. Secondly, heterogeneity of the fast and slow learner can be better investigated – how we can benefit from having different architectures (even mixed ones with transformer-based network).

Limitations. Although we show how to squeeze (and retain) the knowledge learned by the SSL method, when the number of samples per tasks is too low, the knowledge transfer (stages ② and ③) and the expert training (stage ①) degrades. Note that in this extreme scenario, we are going towards online continual learning.

Impact Statement. Continual learning systems do not require data to be stored. As such, they can contribute to data privacy, and reduce vulnerabilities related to data storage. As with all machine learning algorithms, special care should be taken to address biases present in the data (and the data collection process). Continual learning could exacerbate biases in the data because of the task-recency bias, which refers to the problem that continual learning algorithms tend to be biased towards the last data fed to the algorithm.

Acknowledgments. We acknowledge the support from the Spanish Government funding for projects PID2022-143257NB-I00, TED2021-132513B-I00 funded by MCIN/AEI/10.13039/501100011033 and by FSE+ and the European Union NextGenerationEU/PRTR, and the CERCA Programme of Generalitat de Catalunya. Bartłomiej Twardowski acknowledges the grant RYC2021-032765-I.

References

- [1] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *European Conference on Computer Vision*, 2018. 2
- [2] Elahe Arani, Fahad Sarfraz, and Bahram Zonooz. Learning fast, learning slow: A general continual learning method based on complementary learning system. In *International Conference on Learning Representations*, 2022. 2, 3
- [3] Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Armand Joulin, Nicolas Ballas, and Michael Rabbat. 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*, pages 8443–8452, 2021. 6, 7
- [4] Jihwan Bang, Heesu Kim, YoungJoon Yoo, Jung-Woo Ha, and Jonghyun Choi. Rainbow memory: Continual learning with a memory of diverse samples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8218–8227, 2021. 2
- [5] Adrien Bardes, Jean Ponce, and Yann LeCun. VICReg: Variance-invariance-covariance regularization for self-supervised learning. In *International Conference on Learning Representations*, 2022. 2
- [6] Matteo Boschini, Pietro Buzzega, Lorenzo Bonicelli, Angelo Porrello, and Simone Calderara. Continual semi-supervised learning through contrastive interpolation consistency. *Pattern Recognition Letters*, 162:9–14, 2022. 6, 7, 8
- [7] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *arXiv preprint arXiv:2006.09882*, 2020. 2
- [8] Arslan Chaudhry, Puneet K Dokania, Thalaiyasingam Ajanthan, and Philip HS Torr. Riemannian walk for incremental learning: understanding forgetting and intransigence. In *European Conference on Computer Vision*, 2018. 2
- [9] Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a-gem. In *International Conference on Learning Representations*, 2019. 2
- [10] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and Marc’Aurelio Ranzato. On tiny episodic memories in continual learning. *arXiv preprint arXiv:1902.10486*, 2019. 2
- [11] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 1, 2, 6
- [12] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15750–15758, 2021. 2
- [13] Ali Cheraghian, Shafin Rahman, Pengfei Fang, Soumava Kumar Roy, Lars Petersson, and Mehrtash Harandi. Semantic-aware knowledge distillation for few-shot class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2534–2543, 2021. 2
- [14] Matthias Delange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Ales Leonardis, Greg Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 1, 2
- [15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Conference on Computer Vision and Pattern Recognition*, 2009. 6
- [16] Prithviraj Dhar, Rajat Vikram Singh, Kuan-Chuan Peng, Ziyang Wu, and Rama Chellappa. Learning without memorizing. In *Conference on Computer Vision and Pattern Recognition*, 2019. 2
- [17] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE international conference on computer vision*, pages 1422–1430, 2015. 2
- [18] Arthur Douillard, Matthieu Cord, Charles Ollion, Thomas Robert, and Eduardo Valle. Podnet: Pooled outputs distillation for small-tasks incremental learning. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*, pages 86–102. Springer, 2020. 2
- [19] Enrico Fini, Victor G Turrissi Da Costa, Xavier Alameda-Pineda, Elisa Ricci, Karteek Alahari, and Julien Mairal. Self-supervised models are continual learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9621–9630, 2022. 1, 2, 3, 6
- [20] Jhair Gallardo, Tyler L Hayes, and Christopher Kanan. Self-supervised training enhances online continual learning. *arXiv preprint arXiv:2103.14010*, 2021. 3
- [21] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*, 2018. 2
- [22] Alex Gomez-Villa, Twardowski, Bartłomiej, Yu, Lu, Andrew D Bagdanov, and Joost van de Weijer. Continually learning self-supervised representations with projected functional regularization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3867–3877, 2022. 1, 2, 3, 6
- [23] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, koray kavukcuoglu, Remi Munos, and Michal Valko. Bootstrap your own latent - a new approach to self-supervised learning. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 21271–21284. Curran Associates, Inc., 2020. 2

- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Conference on Computer Vision and Pattern Recognition*, 2016. 6
- [25] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *International Conference on Computer Vision*, 2019. 2
- [26] Xinting Hu, Kaihua Tang, Chunyan Miao, Xian-Sheng Hua, and Hanwang Zhang. Distilling causal effect of data in class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3957–3966, 2021. 2
- [27] Zhiqi Kang, Enrico Fini, Moin Nabi, Elisa Ricci, and Karteek Alahari. A soft nearest-neighbor framework for continual semi-supervised learning. *arXiv preprint arXiv:2212.05102*, 2022. 6, 7, 8
- [28] Ronald Kemker and Christopher Kanan. Fearnert: Brain-inspired model for incremental learning. *ICLR2018*, 2018. 3
- [29] 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. *National Academy of Sciences*, 2017. 2, 6, 7, 8
- [30] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 6
- [31] Dharshan Kumaran, Demis Hassabis, and James L McClelland. What learning systems do intelligent agents need? complementary learning systems theory updated. *Trends in cognitive sciences*, 20(7):512–534, 2016. 1
- [32] Shuang Li, Yilun Du, Guido van de Ven, and Igor Mordatch. Energy-based models for continual learning. In *Conference on Lifelong Learning Agents*, pages 1–22. PMLR, 2022. 7, 8
- [33] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017. 2, 6, 7, 8
- [34] Xialei Liu, Marc Masana, Luis Herranz, Joost Van de Weijer, Antonio M Lopez, and Andrew D Bagdanov. Rotate your networks: Better weight consolidation and less catastrophic forgetting. In *International Conference on Pattern Recognition*, 2018. 2
- [35] Divyam Madaan, Jaehong Yoon, Yuanchun Li, Yunxin Liu, and Sung Ju Hwang. Representational continuity for unsupervised continual learning. In *International Conference on Learning Representations*, 2022. 3
- [36] Arun Mallya, Dillon Davis, and Svetlana Lazebnik. Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In *European Conference on Computer Vision*, 2018. 2
- [37] Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Conference on Computer Vision and Pattern Recognition*, 2018. 2
- [38] Marc Masana, Xialei Liu, Bartłomiej Twardowski, Mikel Menta, Andrew D Bagdanov, and Joost Van De Weijer. Class-incremental learning: survey and performance evaluation on image classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(5):5513–5533, 2022. 2, 6
- [39] Marc Masana, Tinne Tuytelaars, and Joost van de Weijer. Ternary feature masks: continual learning without any forgetting. *arXiv*, 2020. 2
- [40] James L McClelland, Bruce L McNaughton, and Randall C O’Reilly. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological review*, 102(3):419, 1995. 1
- [41] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier, 1989. 1
- [42] K L Navaneet, Soroush Abbasi Koohpayegani, Ajinkya Tejankar, and Hamed Pirsiavash. Simreg: Regression as a simple yet effective tool for self-supervised knowledge distillation. In *British Machine Vision Conference (BMVC)*, 2021. 6
- [43] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European conference on computer vision*, pages 69–84. Springer, 2016. 2
- [44] Randall C O’Reilly and Kenneth A Norman. Hippocampal and neocortical contributions to memory: Advances in the complementary learning systems framework. *Trends in cognitive sciences*, 6(12):505–510, 2002. 2
- [45] German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural networks*, 113:54–71, 2019. 2
- [46] German I Parisi, Jun Tani, Cornelius Weber, and Stefan Wermter. Lifelong learning of spatiotemporal representations with dual-memory recurrent self-organization. *Frontiers in neurorobotics*, page 78, 2018. 3
- [47] Quang Pham, Chenghao Liu, and Steven Hoi. Dualnet: Continual learning, fast and slow. *Advances in Neural Information Processing Systems*, 34:16131–16144, 2021. 2, 3
- [48] Ameya Prabhu, Philip HS Torr, and Puneet K Dokania. Gdumb: A simple approach that questions our progress in continual learning. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 524–540. Springer, 2020. 7
- [49] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Conference on Computer Vision and Pattern Recognition*, 2017. 2, 6, 7
- [50] David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience replay for continual learning. *Advances in Neural Information Processing Systems*, 32, 2019. 6, 7

- [51] Mert Bulent Sariyildiz, Yannis Kalantidis, Karteek Alahari, and Diane Larlus. No reason for no supervision: Improved generalization in supervised models. In *International Conference on Learning Representations*, 2023. 6
- [52] Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. In *International Conference on Machine Learning*, 2018. 2
- [53] Yujun Shi, Li Yuan, Yunpeng Chen, and Jiashi Feng. Continual learning via bit-level information preserving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16674–16683, 2021. 2
- [54] Stanford. Tiny imagenet challenge, cs231n course., CS231N. 6
- [55] Shixiang Tang, Dapeng Chen, Jinguo Zhu, Shijie Yu, and Wanli Ouyang. Layerwise optimization by gradient decomposition for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9634–9643, 2021. 2
- [56] Liyuan Wang, Kuo Yang, Chongxuan Li, Lanqing Hong, Zhenguo Li, and Jun Zhu. Ordisco: Effective and efficient usage of incremental unlabeled data for semi-supervised continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5383–5392, 2021. 2
- [57] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *Conference on Computer Vision and Pattern Recognition*, 2019. 2
- [58] Tete Xiao, Xiaolong Wang, Alexei A Efros, and Trevor Darrell. What should not be contrastive in contrastive learning. *ICLR*, 2021. 2
- [59] Shipeng Yan, Jiangwei Xie, and Xuming He. Der: Dynamically expandable representation for class incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3014–3023, 2021. 2
- [60] Lu Yu, Twardowski, Bartłomiej, Liu, Xialei, Luis Herranz, Kai Wang, Shangling Jui, and Joost van de Weijer. Semantic drift compensation for class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6982–6991, 2020. 8
- [61] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. *arXiv preprint arXiv:2103.03230*, 2021. 1, 2, 3
- [62] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, 2017. 2
- [63] Chen Zeno, Itay Golan, Elad Hoffer, and Daniel Soudry. Task-Agnostic Continual Learning Using Online Variational Bayes With Fixed-Point Updates. *Neural Computation*, 33(11):3139–3177, 10 2021. 7, 8
- [64] Mengyao Zhai, Lei Chen, and Greg Mori. Hyper-lifelonggan: Scalable lifelong learning for image conditioned generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2246–2255, 2021. 2
- [65] Junting Zhang, Jie Zhang, Shalini Ghosh, Dawei Li, Serafettin Tasci, Larry Heck, Heming Zhang, and C-C Jay Kuo. Class-incremental learning via deep model consolidation. In *Winter Conf. on App. of Computer Vision*, 2020. 2, 3
- [66] Fei Zhu, Xu-Yao Zhang, Chuang Wang, Fei Yin, and Cheng-Lin Liu. Prototype augmentation and self-supervision for incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5871–5880, 2021. 2
- [67] Simone Zini, Alex Gomez-Villa, Marco Buzzelli, Bartłomiej Twardowski, Andrew D Bagdanov, and Joost van de Weijer. Planckian jitter: countering the color-crippling effects of color jitter on self-supervised training. *ICLR*, 2023. 2