

Self-Supervised Models are Continual Learners

Enrico Fini^{*1,2} Victor G. Turrissi da Costa^{*1} Xavier Alameda-Pineda²
Elisa Ricci^{1,3} Karteek Alahari² Julien Mairal²

¹ University of Trento ² Inria[†] ³ Fondazione Bruno Kessler

Abstract

Self-supervised models have been shown to produce comparable or better visual representations than their supervised counterparts when trained offline on unlabeled data at scale. However, their efficacy is catastrophically reduced in a Continual Learning (CL) scenario where data is presented to the model sequentially. In this paper, we show that self-supervised loss functions can be seamlessly converted into distillation mechanisms for CL by adding a predictor network that maps the current state of the representations to their past state. This enables us to devise a framework for Continual self-supervised visual representation Learning that (i) significantly improves the quality of the learned representations, (ii) is compatible with several state-of-the-art self-supervised objectives, and (iii) needs little to no hyperparameter tuning. We demonstrate the effectiveness of our approach empirically by training six popular self-supervised models in various CL settings. Code: github.com/DonkeyShot21/cassle.

1. Introduction

During the last few years, self-supervised learning (SSL) has become the most popular paradigm for unsupervised visual representation learning [3, 7, 8, 13, 14, 25, 27, 54]. Indeed, under certain assumptions (*e.g.*, offline training with large amounts of data and resources), SSL methods are able to extract representations that match the quality of representations obtained with supervised learning, without requiring annotations. However, these assumptions do not always hold in real-world scenarios, *e.g.*, when new unlabeled data are made available progressively over time. In fact, in order to integrate new knowledge into the model, training needs to be repeated on the whole dataset, which is impractical, expensive, and sometimes even impossible when old data is not available. This issue is exacerbated by the fact that SSL models are notoriously computationally expensive to train.

^{*}Enrico Fini and Victor G. Turrissi da Costa contributed equally.

[†]Univ. Grenoble Alpes, CNRS, Grenoble INP, LJK, 38000 Grenoble, France.




Figure 1. Linear evaluation accuracy of representations learned with different self-supervised methods on class-incremental CIFAR100 and ImageNet100. In blue the accuracy of SSL fine-tuning, in green the improvement brought by CaSSLe. The red dashed line is the accuracy attained by supervised fine-tuning.

Continual learning (CL) studies the ability of neural networks to learn tasks sequentially. Prior art in the field focuses on mitigating catastrophic forgetting [17, 22, 24, 37]. Common benchmarks in the CL literature evaluate the discriminative performance of classifiers learned *with supervision* from non-stationary distributions. In this paper, we tackle the same forgetting phenomenon in the context of SSL. Unsupervised representation learning is indeed appealing for sequential learning since it does not require human annotations, which are particularly hard to obtain when new data is generated on-the-fly. This setup, called Continual Self-Supervised Learning (CSSL), is surprisingly under-investigated in the literature.

In this work, we propose CaSSLe, a simple and effective framework for CSSL of visual representations based on the intuition that SSL models are intrinsically capable of learning continually, and that SSL losses can be seamlessly converted into distillation losses. Our key idea is to train the

current model to predict past representations with a prediction head, thus encouraging it to remember past knowledge. CaSSLe has several favourable features: (i) it is compatible with popular state-of-the-art SSL loss functions and architectures, (ii) it is simple to implement, and (iii) it does not require any additional hyperparameter tuning with respect to the original SSL method. Our experiments demonstrate that SSL methods trained continually with CaSSLe significantly outperform all the related methods (CSSL baselines and several methods adapted from supervised CL).

We also perform a comprehensive analysis of the behavior of six popular SSL methods in diverse CL settings (*i.e.*, class, data, and domain incremental). We provide empirical results on small (CIFAR100), medium (ImageNet100), and large (DomainNet) scale datasets. Our study sheds new light on interesting properties of SSL methods that emerge when learning continually. Among other findings, we discover that, in the class-incremental setting, SSL methods typically approach or outperform supervised learning (see Fig.1), while this is not generally true for other settings (data-incremental and domain-incremental) where supervised learning still shows a sizeable advantage.

2. Related Work

Self-Supervised Learning. Recent SSL approaches have shown performance comparable to their supervised learning equivalents [3, 7, 8, 13, 14, 25, 27, 54]. In a nutshell, most of these methods use image augmentation techniques to generate correlated views (positives) from a sample, and then learn a model that is invariant to these augmentations by enforcing the network to output similar representations for the positives. Initially, contrastive learning, based on instance discrimination [52] using noise-contrastive estimation [26, 38], was a popular strategy [13, 27]. However, this learning paradigm requires large batch sizes or memory banks. A few methods that use a negative-free cosine similarity loss [15, 25] have addressed such issues.

Concurrently, clustering-based methods (SwAV [7], DeepCluster v2 [6, 7] and DINO [8]) have also been proposed. They do not operate on the features directly, and instead compare positives through a cross-entropy loss using cluster prototypes as a proxy. Redundancy reduction-based methods have also been popular [3, 20, 54]. Among them, BarlowTwins [54] considers an objective function measuring the cross-correlation matrix between the features, and VicReg [3] uses a mix of variance, invariance and covariance regularizations. Methods such as [19] have explored the use of nearest-neighbour retrieval and divide and conquer [49]. However, none of these works studied the ability of SSL methods to learn continually and adaptively.

Continual Learning. A plethora of methods have been developed to counteract catastrophic forgetting [2, 4, 9–12, 18, 21, 30, 31, 33, 35, 39, 41, 43–47, 51, 55]. Following [17],

these works can be organized into three macro-categories: replay-based [4, 12, 35, 39, 43, 44], regularization-based [2, 9–11, 18, 21, 30, 31, 33, 47, 51, 55], and parameter isolation methods [45, 46]. All these works evaluate the effectiveness of CL methods using a linear classifier learned sequentially over time. However, this evaluation does not reflect an important aspect, *i.e.*, the internal dynamics of the hidden representations. Moreover, most CL methods tend to rely on supervision in order to mitigate catastrophic forgetting. A few of them can be adapted for the unsupervised setting, although their effectiveness is greatly reduced (see discussion in Sec. 5, Sec. 6 and the supplementary material).

Works such as [1, 42, 48] laid the foundations of unsupervised CL, but their studies are severely limited to digit-like datasets, *e.g.*, MNIST and Omniglot, and the proposed methods are unfit for large-scale scenarios. Recently, [5, 23] explored self-supervised pretraining for supervised continual learning with online and few-shot tasks, and [10] presented a supervised contrastive CL approach. Two concurrent works [34, 36] have also attempted to address CSSL recently. The former [34] extends [10] to the unsupervised setting, but is specifically designed for contrastive SSL, such as [13, 27], and lacks generalizability to other popular SSL paradigms. The latter [36] is also limited as it only shows small-scale experiments in the class-incremental setting and considers just two SSL methods. In contrast, we present a general framework for CSSL with superior performance, conduct large-scale experiments on three challenging settings, thereby presenting a deeper analysis of CSSL.

3. Preliminaries

Self-Supervised Learning. The training procedure of several state-of-the-art SSL methods [3, 7, 8, 13, 19, 25, 27, 54] can be summarized as follows. Given an image \mathbf{x} in a batch sampled from a distribution \mathcal{D} , two correlated views \mathbf{x}^A and \mathbf{x}^B are extracted by applying stochastic image augmentations, such as random cropping, color jittering and horizontal flipping. View \mathbf{x}^A is fed to an encoder $f_\theta = f_p \circ f_b$, which is parametrized by θ and has a backbone f_b and a projection head f_p , that extracts feature representations $\mathbf{z}^A = f_\theta(\mathbf{x}^A)$. Similarly, \mathbf{x}^B is forwarded into the same networks, or possibly copies thereof, updated with exponential moving average (EMA), to obtain the representation \mathbf{z}^B . A loss function \mathcal{L}_{SSL} is applied to these representations to learn the parameters θ as follows:

$$\operatorname{argmin}_{\theta} \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [\mathcal{L}_{SSL}(\mathbf{z}^A, \mathbf{z}^B)]. \quad (1)$$

More details on the implementation of \mathcal{L}_{SSL} are provided in Sec. 5.1 and Tab. 1. This procedure turns out to be extremely powerful at extracting visual representations from large unlabeled datasets. The intuition behind the success of these models is that they learn to be invariant to augmentations. Importantly, augmentations are hand-crafted in a way

that the two views x^A and x^B contain roughly the same semantics as x , but their overall appearance (geometry, colors, resolution, *etc.*) is different. This forces the model map images with the same semantics to similar regions of the feature space. Interestingly, these augmentations are much stronger, *i.e.*, they distort the image more, than augmentations commonly used to train supervised models.

Continual Learning. The CL problem focuses on training models such as deep neural networks from non-stationary data distributions. More formally, this involves a network $f'_{\theta'} = f'_c \circ f'_b$ with parameters θ' , backbone f'_b and a classifier f'_c , that learns from an ordered set of tasks $\{1, \dots, T\}$, each exhibiting a different data distribution \mathcal{D}_t . Usually, an image x sampled i.i.d. from \mathcal{D}_t is processed by f' that predicts a probability distribution \mathbf{p} over the set of classes \mathcal{Y}_t . The objective is to find parameters θ' such as:

$$\operatorname{argmin}_{\theta'} \sum_{t=1}^T \mathbb{E}_{(x,y) \sim \mathcal{D}_t} [\mathcal{L}_{CL}(\mathbf{p}, \mathbf{y})], \quad (2)$$

where, in most cases, \mathcal{L}_{CL} is the cross-entropy loss. However, during task t , the previous data distribution \mathcal{D}_{t-1} is not available and therefore Eq. (2) cannot be minimized directly. Current research focuses on approximating θ' using indirect approaches. Some of them [18, 33] are based on knowledge distillation [29], *i.e.*, transferring knowledge from one network to another by forcing them to produce the same outputs. We will discuss the applicability of distillation methods in CSSL in Sec. 5.

4. Continual Self-Supervised Learning

In this paper, we tackle the problem of Continual Self-Supervised Learning as an extension of both SSL and CL. In practice, a CSSL experiment starts with the first task, where the model is trained as per the specific self-supervised method that it implements, with no difference from offline training. Subsequent tasks are then presented to the model sequentially, and the data from the previous tasks are discarded. No labels are provided during this training phase. For the sake of simplicity and since we are exploring a new, challenging setting, we assume task boundaries to be provided to the model. More formally, the CSSL objective is to learn a strong feature extractor that is invariant to augmentations on all tasks. Following the notation introduced in Sec. 3, we define:

$$\operatorname{argmin}_{\theta} \sum_{t=1}^T \mathbb{E}_{x \sim \mathcal{D}_t} [\mathcal{L}_{SSL}(z^A, z^B)]. \quad (3)$$

Note the absence of labels \mathbf{y} when sampling from \mathcal{D}_t , the summation over the set of tasks inherited from Eq. (2) and the SSL loss function in Eq. (1). The expectation is approximated using stochastic gradient descent on minibatches.

Evaluation. After each task, it is possible (for evaluation purposes) to train a linear classifier on top of the obtained backbone f_b . With this linear classifier we report accuracy on the test set. This protocol is compatible with standard CL metrics, as shown in Sec. 6.1. We explore three CSSL settings in our work.

► **Class-incremental:** each task t is represented by a dataset $D_t \sim \mathcal{D}_t$ containing images that belong to a set of classes \mathcal{Y}_t such that $\mathcal{Y}_t \cap \mathcal{Y}_s = \emptyset$ for each other task $s \neq t$. Note that the class labels are only used for splitting the dataset and they are unknown to the model. In practice, the set of classes in the dataset are shuffled and then partitioned into T tasks. Each task contains the same number of classes.

► **Data-incremental:** each task t contains a set of images D_t such that $D_t \cap D_s = \emptyset$ for each other task $s \neq t$. No additional constraints are imposed on the classes. In practice, the whole dataset is shuffled and then partitioned into T tasks. Each task can potentially contain all the classes.

► **Domain-incremental:** each task t contains a set of images D_t drawn from a different domain. We assume that the set of classes \mathcal{Y}_t in each dataset remains the same for all tasks but the data distribution changes, as if the data were collected from different sources.

5. The CaSSLe Framework

We now introduce ‘‘CaSSLe’’, our framework for continual self-supervised learning of visual representations and detail its compatibility with several SSL methods.

Distillation in CSSL. From a supervised CL perspective, the concept of invariance is interesting. Here, we would like to learn representations of previously-learned semantic concepts that are invariant to the state of the model’s parameters. Indeed, this idea was investigated in prior works [18, 30] that leverage knowledge distillation for CL. However, such approaches are only mildly effective in a CSSL scenario, as we show in Sec. 6. We believe this is due to CSSL being fundamentally different from supervised CL. In CSSL, we aim to extract the best possible representations that can be subsequently reused in a variety of tasks, and maximize the linear separability of features at the end of the CL phase. Thus, the linear classifier does not benefit much from the stability of the representations. Also, forcing the representations not to change may prevent the model from learning new concepts. This is especially critical for SSL methods for two reasons: (i) the performance of the models improve substantially with longer training, implying that the representations continue to get refined, and (ii) they exhibit different losses and feature normalizations that might interfere with distillation and vice-versa (*e.g.*, BarlowTwins uses standardization while [18, 30] use l_2 -normalization). Nonetheless, the features still need to be informative of previous tasks to maximize the separability of the old distribution but the current state might be too different from the pre-

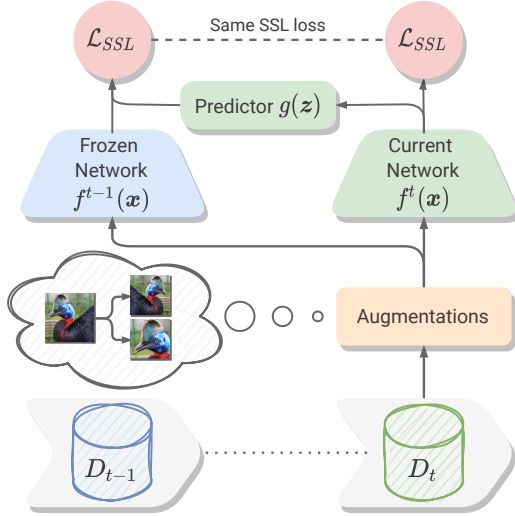


Figure 2. Overview of the CaSSLe framework.

vious one making comparing representations complicated.

Distillation through SSL losses. Our framework, shown in Fig. 2, is based on the following ideas: (i) a predictor network that maps the current state of the representations to their past state, by leveraging a distillation through time strategy that satisfies both stability and plasticity principles, and (ii) a family of adaptable distillation losses inherited from the SSL literature that solves the issue of having different objectives interfering with each other.

When a new task is received, we start by making a copy of the current model. This copy does not require gradient computation and will not be updated. We call this the *frozen encoder* f^{t-1} . As soon as an image $x \in D_t$ is available we apply our stochastic image augmentations and extract its features $z = f^t(x)$. In addition, we also use the frozen encoder to extract another feature vector $\bar{z} = f^{t-1}(x)$. Now, our goal is to ensure that z contains at least as much information as (and ideally more than) \bar{z} . Instead of enforcing the two feature vectors to be similar, and hence discouraging the new model from learning new concepts, we propose to use a predictor network g to project the representations from the new feature space to the old one. If the predictor is able to perfectly map from one space to the other, then it implies that z is at least as powerful as \bar{z} .

We are now ready to perform distillation, but which is the most appropriate distillation loss? Since we want the representations produced by g to be invariant to the state of the model, we propose to use the same SSL loss used to simulate invariance to augmentations. Empirically, we verify that this choice reduces interference and minimizes the need for hyperparameter tuning. We can hence write a generic distillation loss by reusing the definition of \mathcal{L}_{SSL} :

$$\mathcal{L}_D(z, \bar{z}) = \mathcal{L}_{SSL}(g(z), \bar{z}). \quad (4)$$

Note that \bar{z} is always detached from the computational graph, such that the frozen encoder does not receive any gradient, and the gradient only flows through the predictor g , as prescribed in [15]. On the one hand, if training converges and \mathcal{L}_D is minimized, the features predicted by g will likely be quasi-invariant to the state of the model, which satisfies the stability principle. On the other hand, the current encoder is less bound to its previous state, hence representations z can be more plastic. The loss can be extended to multiple views by applying it to both representations, *i.e.*, $\mathcal{L}_D(z^A, \bar{z}^A) + \mathcal{L}_D(z^B, \bar{z}^B)$, and also swapped distillation, *e.g.*, $\mathcal{L}_D(z^A, \bar{z}^B)$ and vice-versa (see ablation in Tab. 6).

The final loss of an SSL method trained continually with the CaSSLe framework is given by:

$$\begin{aligned} \mathcal{L} &= \mathcal{L}_{SSL}(z^A, z^B) + \mathcal{L}_D(z^A, \bar{z}^A) \\ &= \mathcal{L}_{SSL}(z^A, z^B) + \mathcal{L}_{SSL}(g(z^A), \bar{z}^A). \end{aligned} \quad (5)$$

This loss can be made symmetric by applying it to both the views (swapping A and B in Eq. (5)) and it can also be easily adapted for multi-crop [7]. Note that we do not use any hyperparameter to weight the importance of the distillation loss with respect to the SSL loss.

5.1. Compatibility of SSL methods with CaSSLe

The main difference among SSL methods is the loss function that they use. Following the notation defined in Sec. 3, and the loss functions in Tab. 1, we now detail if and how SSL losses can be used in our CaSSLe framework. Full derivation of distillation losses is deferred to the supplementary material.

InfoNCE-based methods [13, 27] perform instance discrimination, where positive samples help to build invariance to augmentations. The negatives prevent the model from falling into degenerate solutions. The InfoNCE (*a.k.a.* contrastive) loss can be written as in Eq. (6), where subscript i is the index of a generic sample in the batch, sim is the cosine similarity and $\eta(i)$ is the set of negatives for sample i in the current batch. Distilling knowledge with this loss is equivalent to performing instance discrimination of current task samples but in the feature space learnt in the past. Thus, the predictor g learns to project samples from the present to the past space to maximize the distance with the negative samples, and the similarity with itself in the past.

MSE-based approaches [15, 25] enforce consistency among positive samples and ignore the negatives. BYOL [25] uses a momentum encoder and SimSiam [15] performs a stop gradient operation to avoid degenerate solutions. Since the representations are l_2 -normalized, their loss (Eq. 7) can be rewritten as the negative cosine similarity: $-\text{sim}(q^A, z^B) = -\frac{q^A}{\|q^A\|_2} \cdot \frac{z^B}{\|z^B\|_2}$, where $q^A = h(z^A)$ and h is a prediction head. The gradient is backpropagated only

Table 1. Overview of state-of-the-art SSL methods and losses. In all tables, highlight colors are coded according to the type of loss.

Methods	Loss	Equation
SimCLR [13] MoCo [27] NNCLR [19]	InfoNCE	$-\log \frac{\exp(\text{sim}(z_i^A, z_i^B)/\tau)}{\sum_{z_j \in \eta(i)} \exp(\text{sim}(z_i^A, z_j)/\tau)}$ (6)
BYOL [25] SimSiam [15] VICReg [3]	MSE	$- q^A - z^B _2^2$ (7)
SwAV [7] DCV2 [7] DINO [8]	Cross-entropy	$-\sum_d \alpha_d^B \log \frac{\exp(\text{sim}(z^A, c_d)/\tau)}{\sum_k \exp(\text{sim}(z^A, c_k)/\tau)}$ (8)
Barlow Twins [54] VICReg [3]	Cross-correlation	$\sum_u (1 - C_{uv})^2 + \lambda \sum_u \sum_{v \neq u} C_{uv}^2$ (9)

through the representations of the first argumentation. A special case of this family of methods is VICReg [3], which uses a combination of multiple losses, where MSE acts as invariance term. Features are not l_2 -normalized in VICReg and its predictor is the identity function. In our framework, this loss encourages the model to predict the past state of representations without additional regularization.

Cross-entropy-based. Instead of simply enforcing invariance of the representations to augmentations, cluster prototypes $\mathbf{C} = \{c_1, \dots, c_K\}$ are used as a proxy in these approaches, so that the model learns to predict invariant cluster assignments. Slight variations of this idea result in different methods: SwAV [7], DeepClusterV2 [7] and DINO [8]. Once a probability distribution over the prototypes is predicted, the cross-entropy loss (Eq. 8) is used to compare the two views. Features and cluster prototypes c are l_2 -normalized. The assignments α^B can be calculated in several ways, e.g., k-means in DeepCluster, Sinkhorn-Knopp in SwAV and EMA in DINO. When employed as a distillation loss, cross-entropy encourages g to predict the assignments generated by the frozen encoder with a set of frozen prototypes: $\alpha^B = \frac{\exp(\text{sim}(z^B, c_d^{t-1})/\tau)}{\sum_k \exp(\text{sim}(z^B, c_k^{t-1})/\tau)}$, where $\mathbf{C}^{t-1} = \{c_1^{t-1}, \dots, c_K^{t-1}\}$.

Cross-correlation-based. These methods use a different approach based on decorrelating the components of the feature space, e.g., Barlow Twins [54], VICReg [3] and W-MSE [20]. For our analysis, we will mainly focus on Barlow Twins’ implementation of this objective. Extensions to VICReg are left for future work. The cross-correlation based objective function is shown in Eq. 9, where λ is an hyperparameter to control the importance of the first and the second terms of the loss, and $C_{uv} = \frac{\sum_i z_{i,u}^A z_{i,v}^B}{\sqrt{\sum_i (z_{i,u}^A)^2} \cdot \sqrt{\sum_i (z_{i,v}^B)^2}}$ is the value of position (u, v) of the cross-correlation matrix computed between the representations of the views along the batch dimension. Note that the representations here are mean centered along the batch dimension, such that each unit has mean output zero over the batch. Performing distillation with this loss has the addi-

tional effect of decorrelating the dimensions of the predicted features $g(z^A)$.

6. Experiments

6.1. Experimental Protocol

Evaluation Metrics. Following previous work [35], we propose the following metrics to evaluate the quality of the representations extracted by our CSSL model:

► **Linear Evaluation Accuracy:** accuracy of a classifier trained on top of the backbone f_b on all tasks (or a subset, e.g., 10% of the data) or a downstream task. For class-incremental and data-incremental, we use the task-agnostic setting, meaning that at evaluation time we do not assume to know the task ID. For the domain-incremental setting, we perform both task-aware and task-agnostic evaluations (the latter is discussed in the supplementary material). To calculate the average accuracy we compute $A = \frac{1}{T} \sum_{i=1}^T A_{T,i}$, where $A_{j,k}$ is the linear evaluation accuracy of the model on task k after observing the last sample from task j .

► **Forgetting:** a common metric in the CL literature, it quantifies how much information the model has forgotten about previous tasks. It is formally defined as: $F = \frac{1}{T-1} \sum_{i=1}^{T-1} \max_{t \in \{1, \dots, T\}} (A_{t,i} - A_{T,i})$.

► **Forward Transfer:** measures how much the representations that we learned so far are helpful in learning new tasks, namely: $FT = \frac{1}{T-1} \sum_{i=2}^T A_{i-1,i} - R_i$ where R_i is the linear evaluation accuracy of a random network on task i .

Datasets. We perform experiments on 3 datasets: CIFAR100 [32] (class-incremental), a 100-class dataset with 60k 32x32 colour images; ImageNet100 [50] (class- and data-incremental), 100-class subset of the ILSVRC2012 dataset with $\approx 130k$ images in high resolution (resized to 224x224); DomainNet [40] (domain-incremental), a 345-class dataset containing roughly 600k high-resolution images (resized to 224x224) divided into 6 domains. We experiment with 5 tasks for the class- and data-incremental settings and with 6 tasks (one for each domain in DomainNet) in the case of domain-incremental. The supplementary material presents additional results with different number of tasks. For the domain-incremental setting, we order the domains in decreasing number of images.

Implementation details. The SSL methods are adapted from solo-learn [16], an established SSL library, which is the main code base for all our experiments. The number of epochs per task is as follows: 500 for CIFAR100, 400 for ImageNet100, 200 for DomainNet. The backbone f_b is a ResNet18 [28], with batch size 256. We use LARS [53] for all our experiments. The offline version of each method, that serves as an upper bound, is trained for the same number of epochs as the continual counterpart for a fair comparison. All the results for offline upper bounds are ob-

Table 2. Comparison with state-of-the-art CL methods on CIFAR100 (5 tasks, class-incremental) using linear evaluation top-1 accuracy, forgetting and forward transfer.

Strategy	SimCLR			Barlow Twins			BYOL		
	A (↑)	F (↓)	T (↑)	A (↑)	F (↓)	T (↑)	A (↑)	F (↓)	T (↑)
Fine-tuning	48.9	1.0	33.5	54.3	0.4	39.2	52.7	0.1	35.9
EWC [31]	53.6	0.0	33.3	56.7	0.2	39.1	56.4	0.0	39.9
ER [44]	50.3	0.1	32.7	54.6	3.0	39.4	54.7	0.4	36.3
DER [4]	50.7	0.4	33.2	55.3	2.5	39.6	54.8	1.1	36.7
LUMP [36]	52.3	0.3	34.5	57.8	0.3	41.0	56.4	0.2	37.9
Less-Forget [30]	52.5	0.2	33.8	56.4	0.2	40.1	58.6	0.2	41.1
POD [18]	51.3	0.1	33.8	55.9	0.3	40.3	57.9	0.0	41.1
CaSSLe	58.3	0.2	36.4	60.4	0.4	42.2	62.2	0.0	43.6
Offline	65.8	-	-	70.9	-	-	70.5	-	-

Table 3. Comparison with Lin *et al.* [34] on CIFAR100 (2 and 5 tasks, class-incremental setting). MoCoV2+ is an updated version of MoCoV2 that uses a symmetric loss. The difference between the two is $\approx 1\%$ at convergence [15].

Strategy	Method	2 Tasks	5 Tasks
Lin <i>et al.</i> [34]	SimCLR	55.7	-
	MoCoV2	56.1	53.8
CaSSLe	SimCLR	61.8	58.3
	MoCoV2+	63.3	59.5

tained using the checkpoints provided in [16]. For some SSL methods, it was necessary to slightly increase the learning rate over the values provided by [16] in order for the methods to fully converge in the CSSL setting. Although tuning the hyperparameters might be beneficial in some settings, we do **not** perform any hyperparameter tuning for CaSSLe. We also neither change the parameters of the SSL methods, nor use a weight for the distillation loss (as per Eq. (5)).

Baselines. Most of the CL methods require labels which makes them unsuitable for CSSL. However, a few works can be adapted for our setting with minimal changes. We choose baselines from three categories [17]: prior-focused regularization (EWC [31]), data-focused regularization (POD [18], Less-Forget [30]), and rehearsal-based replay (ER [44], DER [4]) methods. We also compare with two concurrent works that propose approaches for CSSL (LUMP [36], Lin *et al.* [34]). Finally, we do not consider methods based on VAEs [1, 42], since they have been shown to yield poor performance on large scale. Details on how the baselines are selected, implemented and tuned for CSSL can be found in the supplementary material.

6.2. Results

Comparison with the state of the art. In Tab. 2 we report comparison with CL baselines and fine-tuning in composition with three SSL methods: SimCLR, Barlow Twins and BYOL. We select these three methods for the following reasons: (i) they feature different losses (InfoNCE, Cross-correlation and MSE), (ii) they exhibit different feature normalizations (l_2 , standardization and mean centering), and (iii) they use different techniques to avoid collapse (neg-




Figure 3. Evolution of top-1 linear evaluation accuracy over tasks on CIFAR100 (5 tasks, class-incremental).

atives, redundancy reduction, momentum encoder). The comparison is performed on class-incremental CIFAR100 with 5 tasks. Offline learning results are reported as upper bound.

First, we notice that CaSSLe produces better representations than all the other strategies, outperforming them by large margins with all SSL methods in terms of top-1 accuracy. Moreover, our framework also shows better forward transfer, meaning that its features are easier to generalize to other tasks (also evident in Tab. 8). CaSSLe appears to reduce catastrophic forgetting with respect to fine-tuning, and is comparable to other methods. In general, SSL methods already have low forgetting with respect to supervised learning on CIFAR100 (see Tab. 4) and therefore there is little margin for improvement. However, on higher resolution images (ImageNet100) CaSSLe actually achieves remarkable results in the mitigation of catastrophic forgetting.

Replay-based methods (ER, DER) clearly do not help against forgetting in CSSL. We found two reasons for this failure. First, in supervised CL, replay-based methods benefit from storing labels, which contain a lot of information about previous tasks and enable the retraining of the linear classifier on old classes. This is not the case in CSSL, where labels are unavailable. Second, SSL models need more training epochs to converge, which means that samples in the buffer are also replayed many more times. This causes severe overfitting on these exemplars, defeating the purpose of the replay buffer. LUMP mitigates this effect by augmenting the buffer using mixup but does not reach too far, surpassing other baselines only with Barlow Twins. EWC holds up surprisingly well, outperforming more recent methods, meaning that the importance of the weights can be calculated accurately with the self-supervised loss. Distillation methods (POD, Less-Forget) show good performance. However, they use l_2 -normalization in their loss, causing loss of information when coupled with Barlow Twins, which decreases accuracy.

Fig. 3 shows the evolution of top-1 linear evaluation accuracy over the whole training trajectory on class-incremental CIFAR100 with 5 tasks. CaSSLe outperforms the other methods, and keeps improving throughout the sequence. We found BYOL to be unstable when simply fine-tuning the model. CaSSLe, EWC and Less-Forget mitigate this instability completely. On the other hand, LUMP first

Table 4. Linear evaluation top-1 accuracy on class-incremental CIFAR100 and ImageNet100 with 5 tasks. CaSSLe is compared to fine-tuning, offline and supervised learning.

Method	Strategy	CIFAR100			ImageNet100		
		A (\uparrow)	F (\downarrow)	T (\uparrow)	A (\uparrow)	F (\downarrow)	T (\uparrow)
Barlow Twins	Fine-tuning	54.3	0.4	39.2	63.1	10.7	44.4
	CaSSLe	60.4	0.4	42.2	68.2	1.3	47.9
	Offline	70.9	-	-	80.4	-	-
SwAV	Fine-tuning	55.5	0.0	32.8	64.4	4.3	42.8
	CaSSLe	57.8	0.0	34.5	66.0	0.2	43.6
	Offline	64.9	-	-	74.3	-	-
BYOL	Fine-tuning	52.7	0.1	35.9	66.0	2.9	43.2
	CaSSLe	62.2	0.0	42.2	66.4	1.1	46.6
	Offline	70.5	-	-	80.3	-	-
VICReg	Fine-tuning	51.5	0.9	36.4	61.3	7.9	42.0
	CaSSLe	53.6	0.2	41.1	64.8	4.3	45.3
	Offline	68.5	-	-	79.4	-	-
MoCoV2+	Fine-tuning	47.3	0.2	33.4	62.0	8.4	41.6
	CaSSLe	59.5	0.0	39.6	68.8	1.5	46.8
	Offline	69.9	-	-	79.3	-	-
SimCLR	Fine-tuning	48.9	1.0	33.5	61.5	8.1	40.3
	CaSSLe	58.3	0.2	36.4	68.0	2.2	45.8
	Offline	65.8	-	-	77.5	-	-
Supervised	Fine-tuning	54.1	6.8	36.5	63.1	5.6	42.5
	Offline	75.6	-	-	81.9	-	-

drops slightly and then recovers. We believe this is due to some instability introduced by the mixup regularization, to which the model takes time to adapt.

In Tab. 3 we also compare with Lin *et al.* [34] on class-incremental CIFAR100. Although our method is not specifically designed for contrastive learning, it substantially outperforms Lin *et al.* with 2 and 5 tasks. It is worth noting that MoCoV2+ is slightly better than MoCoV2 ($\approx 1\%$ difference), whereas our gains are much larger ($\approx 7\%$).

Ablation study. We ablate the most critical design choices we adopt in CaSSLe: (i) distillation without swapped views, and (ii) the presence of a prediction head g . These results are reported in Tab. 6. Our full framework clearly outperforms its variants with swapped views and without predictor. This validates our hypothesis that a predictor to map new features to the old feature space is crucial. The result that swapping views does not help is likely due to the frozen encoder not being invariant to the current task.

Class-incremental. In Tab. 4 we report a study of CSSL with 6 SSL methods in composition with the CaSSLe framework on class-incremental CIFAR100 and ImageNet100. Fine-tuning and Offline SSL results are reported as lower and upper bounds. The accuracy of supervised learning is also reported. CaSSLe always improves with respect to fine-tuning. In particular, our framework produces higher forward transfer and lower forgetting, especially on ImageNet100, where methods tend to forget more. Notably, CaSSLe outperforms supervised fine-tuning, ex-

Table 5. Training 5 times longer on 1/5 of the data vs. training continually w/ and w/o CaSSLe on ImageNet100 (5 tasks, class- and data-incremental). **Bold** is best, underlined is second best.

Setting	Method	Fine-tune	Offline 1/5	CaSSLe
Class-inc.	SimCLR	61.5	<u>63.1</u>	68.0
	Barlow Twins	63.1	<u>63.5</u>	68.2
	BYOL	<u>66.0</u>	60.6	66.4
Data-inc.	SimCLR	68.9	67.2	72.1
	Barlow Twins	<u>71.3</u>	70.2	74.9
	BYOL	74.0	66.7	<u>73.3</u>

Table 6. Ablation study of design choices in CaSSLe.

Strategy	Method	Swap	No pred.	Ours
CaSSLe	SimCLR	49.3	52.6	58.3
	Barlow Twins	57.4	57.3	60.4
	BYOL	52.0	58.6	62.2

cept when coupled with VICReg on CIFAR100. On average, SSL methods trained continually with CaSSLe improve by 6.8% on CIFAR100 and 4% on ImageNet100.

Data-incremental. Tab. 7 presents results for linear evaluation top-1 accuracy on ImageNet100 with 5 tasks in a data-incremental scenario. While no SSL method is better than supervised fine-tuning, Barlow Twins coupled with CaSSLe is competitive. CaSSLe improves performance in all cases by 2% on average, except for BYOL. This is likely due to the fact that in the data-incremental scenario remembering past knowledge is less important than in other scenarios, and BYOL already has a momentum encoder that provides some information about the past. This hypothesis is validated by the fact that MoCoV2+ (that uses a momentum encoder) improves less than SimCLR when coupled with CaSSLe. We believe that, by tuning the EMA schedule, improvement could also be achieved for BYOL. In addition, BYOL already shows impressive performance with fine-tuning, outperforming all the other methods by more than 2%. Interestingly, SwAV comes closest to its offline upper bound, with only a 3% decrease in performance when coupled with CaSSLe.

Domain-incremental. We also examine the capability of CaSSLe to learn continually when the domain from which the data is drawn changes. Tab. 7 shows the average top-1 accuracy of a linear classifier trained on top of the frozen feature extractor on all domains separately (domain-aware). Domain-agnostic evaluation and results for each domain are presented in the supplementary material. Again, CaSSLe improves every method by 4.4% on average, showing that our distillation strategy is robust to domain shift, and although the data distribution is really different, information transfer is still performed. Interestingly, most of the methods, when trained with CaSSLe get very close to their offline accuracy.

Long training vs continual training. We also analyze the following question: is it worth training continually or is it better to train for longer on a small dataset? This

Table 7. Linear evaluation accuracy on ImageNet100 (5 tasks, data-incremental) and DomainNet (6 tasks, domain-incremental).

Method	Strategy	ImageNet100 (Data-inc.)	DomainNet (Domain-inc.)
Barlow Twins	Fine-tuning	71.3	50.3
	CaSSLe	74.9	55.5
	Offline	80.4	57.2
SwAV	Fine-tuning	70.8	49.6
	Knowledge	71.3	54.3
	Offline	74.3	54.6
BYOL	Fine-tuning	74.0	50.6
	CaSSLe	73.3	55.1
	Offline	80.3	56.6
VICReg	Fine-tuning	70.2	49.3
	CaSSLe	72.3	52.9
	Offline	79.4	56.7
MoCoV2+	Fine-tuning	69.5	43.2
	CaSSLe	71.9	46.7
	Offline	78.2	53.7
SimCLR	Fine-tuning	68.9	45.1
	CaSSLe	72.1	50.0
	Offline	77.5	52.6
Supervised	Fine-tuning	75.9	55.9
	Offline	81.9	66.4

depends on two factors: (i) the SSL method, and (ii) the CSSL setting. For SimCLR and Barlow Twins in the class-incremental setting it seems to be better to train offline on 1/5 of the classes instead of training continually with 5 tasks. In this setting, offline BYOL seems to suffer from instability, ending up lower than fine-tuning. On the other hand, on the data-incremental setting, fine-tuning outperforms longer training, especially for BYOL, which also outperforms CaSSLe (as explained previously). Apart from this exception, CaSSLe always produces better representations than other strategies, making it the go-to option.

Downstream and semi-supervised. In Tab. 8, we present the downstream performance of CaSSLe compared with fine-tuning when trained on ImageNet100 and evaluated on DomainNet (Real). Barlow Twins, SwAV and BYOL show higher performance than the supervised model, even when considering a fine-tuning strategy. This is probably due to the fact that SSL methods tend to learn more general features than their supervised counterparts. CaSSLe improves performance on all the SSL methods, making them surpass the supervised baseline. Lastly, when compared with fine-tuning, CaSSLe improves the performance of SSL methods by 3.4% on average. Tab. 9 contains the top-1 accuracy on ImageNet100 when training a linear classifier on a frozen backbone with limited amount of labels (10% and 1%). First, we can observe that no SSL method with fine-tune surpasses the performance of supervised learning. When using CaSSLe, MoCoV2+ outperforms supervised with 10% labels and, in general, Barlow Twins and MoCoV2+ work best in both semi-supervised settings. CaSSLe

Table 8. Downstream performance with different SSL methods trained on Imagenet-100 and evaluated on DomainNet (Real).

Strategy	Barlow Twins	SwAV	BYOL	VICReg	MoCoV2+	SimCLR	Supervised
Fine-tune	56.2	55.9	55.0	54.0	52.4	51.6	54.3
CaSSLe	60.3	56.9	56.9	56.3	58.7	56.5	

Table 9. Top-1 linear accuracy on Imagenet-100 with different SSL methods, semi-supervised setting with 10% and 1% of labels.

Percentage	Strategy	Barlow Twins	SwAV	BYOL	VICReg	MoCoV2+	SimCLR	Supervised
10%	Fine-tune	56.6	57.6	55.7	53.6	54.9	52.5	60.8
	CaSSLe	60.3	58.2	56.5	56.5	61.7	58.9	
1%	Fine-tune	42.6	42.5	42.3	40.4	40.9	39.7	48.1
	CaSSLe	47.0	43.1	43.4	43.2	47.8	46.8	

improves all SSL methods when compared with fine-tuning.

7. Conclusion

In this work, we study Continual Self-Supervised Learning (CSSL), the problem of learning a set of tasks without labels continually. We make two important contributions for the SSL and CL communities: (i) we present CaSSLe, a simple and effective framework for CSSL that shows how SSL methods and losses can be seamlessly reused to learn continually, and (ii) we perform a comprehensive analysis of CSSL, leading to the emergence of interesting properties of SSL methods.

Limitations. Although CaSSLe shows exciting performance, it has some limitations. First, it is applicable in settings where task boundaries are provided. Second, our framework increases the amount of computational resources needed for training by roughly 30%, both in terms of memory and time. Finally, CaSSLe does not perform clustering, meaning that it is unable to directly learn a mapping from data to latent classes, and thus needs either a linear classifier trained with supervision, or some clustering algorithm.

Broader impact. The capabilities of supervised CL agents are bounded by the need for human-produced annotations. CSSL models can potentially improve without the need for human supervision. This facilitates the creation of powerful AIs that may be used for malicious purposes such as discrimination and surveillance. Also, since in CSSL the data is supposed to come from a non-curved stream, the model may be affected by biases in the data. This is problematic because biases are then be transferred to downstream tasks.

Acknowledgements. This work was supported by the European Institute of Innovation & Technology (EIT) and the H2020 EU project SPRING, funded by the European Commission under GA 871245. It was carried out under the ‘‘Vision and Learning joint Laboratory’’ between FBK and UNITN. Kartek Alahari was funded by the ANR grant AVENUE (ANR-18-CE23-0011). Julien Mairal was funded by the ERC grant number 714381 (SOLARIS project) and by ANR 3IA MIAI@Grenoble Alpes (ANR-19-P3IA-0003). Xavier Alameda-Pineda was funded by the ARN grant ML3RI (ANR-19-CE33-0008-01). This project was granted access to the HPC resources of IDRIS under the allocation 2021-[AD011013084] made by GENCI.

References

- [1] Alessandro Achille, Tom Eccles, Loic Matthey, Christopher P Burgess, Nick Watters, Alexander Lerchner, and Irina Higgins. Life-long disentangled representation learning with cross-domain latent homologies. *NeurIPS*, 2018. 2, 6
- [2] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *ECCV*, 2018. 2
- [3] Adrien Bardes, Jean Ponce, and Yann LeCun. Vireg: Variance-invariance-covariance regularization for self-supervised learning. *arXiv preprint arXiv:2105.04906*, 2021. 1, 2, 5
- [4] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: a strong, simple baseline. In *NeurIPS*, 2020. 2, 6
- [5] Lucas Caccia and Joelle Pineau. Special: Self-supervised pretraining for continual learning. *arXiv preprint arXiv:2106.09065*, 2021. 2
- [6] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *ECCV*, pages 132–149, 2018. 2
- [7] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *NeurIPS*, 2020. 1, 2, 4, 5
- [8] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. *ICCV*, 2021. 1, 2, 5
- [9] Francisco M Castro, Manuel J Marin-Jimenez, Nicolas Guil, Cordelia Schmid, and Karteek Alahari. End-to-end incremental learning. In *ECCV*, 2018. 2
- [10] Hyuntak Cha, Jaeho Lee, and Jinwoo Shin. Co2l: Contrastive continual learning. In *CVPR*, pages 9516–9525, 2021. 2
- [11] Arslan Chaudhry, Puneet K Dokania, Thalaiyasingam Ajanthan, and Philip HS Torr. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *ECCV*, 2018. 2
- [12] Arslan Chaudhry, Marc’ Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a gem. In *ICLR*, 2019. 2
- [13] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 1, 2, 4, 5
- [14] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. 1, 2
- [15] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *ICCV*, 2021. 2, 4, 5, 6
- [16] Victor G. Turrissi da Costa, Enrico Fini, Moin Nabi, Nicu Sebe, and Elisa Ricci. Solo-learn: A library of self-supervised methods for visual representation learning. *arXiv 2108.01775*, 2021. 5, 6
- [17] Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Ales Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. Continual learning: A comparative study on how to defy forgetting in classification tasks. *IEEE TPAMI*, 2(6), 2019. 1, 2, 6
- [18] Arthur Douillard, Matthieu Cord, Charles Ollion, Thomas Robert, and Eduardo Valle. Podnet: Pooled outputs distillation for small-tasks incremental learning. In *ECCV*, 2020. 2, 3, 6
- [19] Debidatta Dwibedi, Yusuf Aytar, Jonathan Tompson, Pierre Sermanet, and Andrew Zisserman. With a little help from my friends: Nearest-neighbor contrastive learning of visual representations. *ICCV*, 2021. 2, 5
- [20] Aleksandr Ermolov, Aliaksandr Siarohin, Enver Sangineto, and Nicu Sebe. Whitening for self-supervised representation learning. In *ICML*, 2021. 2, 5
- [21] Enrico Fini, Stéphane Lathuilière, Enver Sangineto, Moin Nabi, and Elisa Ricci. Online continual learning under extreme memory constraints. In *ECCV*, 2020. 2
- [22] Robert M. French. Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, 3(4):128–135, 1999. 1
- [23] Jhair Gallardo, Tyler L Hayes, and Christopher Kanan. Self-supervised training enhances online continual learning. *arXiv preprint arXiv:2103.14010*, 2021. 2
- [24] I. J. Goodfellow, M. Mirza, D. Xiao, A. Courville, and Y. Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*, 2013. 1
- [25] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent: A new approach to self-supervised learning. *NeurIPS*, 2020. 1, 2, 4, 5
- [26] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *AISTATS*, 2010. 2
- [27] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 1, 2, 4, 5
- [28] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 5
- [29] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 3
- [30] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *CVPR*, pages 831–839, 2019. 2, 3, 6
- [31] 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. *Proc. of the national academy of sciences*, 2017. 2, 6
- [32] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, Univ. Toronto, 2009. 5
- [33] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE TPAMI*, 2017. 2, 3

- [34] Zhiwei Lin, Yongtao Wang, and Hongxiang Lin. Continual contrastive self-supervised learning for image classification. *arXiv preprint arXiv:2107.01776*, 2021. [2](#), [6](#), [7](#)
- [35] David Lopez-Paz and Marc-Aurelio Ranzato. Gradient episodic memory for continual learning. In *NeurIPS*, 2017. [2](#), [5](#)
- [36] Divyam Madaan, Jaehong Yoon, Yuanchun Li, Yunxin Liu, and Sung Ju Hwang. Rethinking the representational continuity: Towards unsupervised continual learning. *arXiv preprint arXiv:2110.06976*, 2021. [2](#), [6](#)
- [37] 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](#)
- [38] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. [2](#)
- [39] Oleksiy Ostapenko, Mihai Puscas, Tassilo Klein, Patrick Jah-nichen, and Moin Nabi. Learning to remember: A synaptic plasticity driven framework for continual learning. In *CVPR*, 2019. [2](#)
- [40] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *ICCV*, 2019. [5](#)
- [41] Ameya Prabhu, Philip HS Torr, and Puneet K Dokania. Gdumb: A simple approach that questions our progress in continual learning. In *ECCV*, pages 524–540. Springer, 2020. [2](#)
- [42] Dushyant Rao, Francesco Visin, Andrei A Rusu, Yee Whye Teh, Razvan Pascanu, and Raia Hadsell. Continual unsuper-vised representation learning. *NeurIPS*, 2019. [2](#), [6](#)
- [43] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *CVPR*, 2017. [2](#)
- [44] Anthony Robins. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connection Science*, 1995. [2](#), [6](#)
- [45] A. A. Rusu, N. C. Rabinowitz, G. Desjardins, H. Soyer, J. Kirkpatrick, K. Kavukcuoglu, R. Pascanu, and R. Hadsell. Progressive neural networks. *arXiv 1606.04671*, 2016. [2](#)
- [46] Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. *ICML*, 2018. [2](#)
- [47] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In *NeurIPS*, 2017. [2](#)
- [48] James Smith, Cameron Taylor, Seth Baer, and Constantine Dovrolis. Unsupervised progressive learning and the stam architecture. *arXiv preprint arXiv:1904.02021*, 2019. [2](#)
- [49] Yonglong Tian, Olivier J Henaff, and Aaron van den Oord. Divide and contrast: Self-supervised learning from uncu-rated data. *arXiv preprint arXiv:2105.08054*, 2021. [2](#)
- [50] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Con-
trastive multiview coding. In *ECCV*, pages 776–794. Springer, 2020. [5](#)
- [51] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale in-cremental learning. In *CVPR*, 2019. [2](#)
- [52] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, 2018. [2](#)
- [53] Yang You, Igor Gitman, and Boris Ginsburg. Large batch training of convolutional networks. *arXiv preprint arXiv:1708.03888*, 2017. [5](#)
- [54] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. *ICML*, 2021. [1](#), [2](#), [5](#)
- [55] Friedeman Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *ICML*, 2017. [2](#)