

Revisiting Supervision for Continual Representation Learning

Daniel Marczak^{*1,2}, Sebastian Cygert^{1,3},
Tomasz Trzcinski^{1,2,4}, and Bartłomiej Twardowski^{1,5,6}

¹ IDEAS NCBR

² Warsaw University of Technology

³ Gdańsk University of Technology

⁴ Tooploox

⁵ Autonomous University of Barcelona

⁶ Computer Vision Center

Abstract. In the field of continual learning, models are designed to learn tasks one after the other. While most research has centered on supervised continual learning, there is a growing interest in unsupervised continual learning, which makes use of the vast amounts of unlabeled data. Recent studies have highlighted the strengths of unsupervised methods, particularly self-supervised learning, in providing robust representations. The improved transferability of those representations built with self-supervised methods is often associated with the role played by the multi-layer perceptron projector. In this work, we depart from this observation and reexamine the role of supervision in continual representation learning. We reckon that additional information, such as human annotations, should not deteriorate the quality of representations. Our findings show that supervised models when enhanced with a multi-layer perceptron head, can outperform self-supervised models in continual representation learning. This highlights the importance of the multi-layer perceptron projector in shaping feature transferability across a sequence of tasks in continual learning. The code is available on github.

Keywords: Continual Learning · Representation Learning

1 Introduction

In continual learning (CL), the goal of the model is to learn new tasks sequentially. Most of the works focus on supervised continual learning (SCL) for image classification where the learner is provided with labeled training data and the metric of interest is accuracy on all the tasks seen so far. More recently, unsupervised continual learning (UCL) gained more attention [10, 11, 25]. UCL considers the problem of learning robust and general representations on a sequence of tasks, without accessing the data labels. Effective UCL methods would allow

* Corresponding author, email: daniel.marczak.dokt@pw.edu.pl




Figure 1: In a two-task continual learning scenario, supervised learning (SL) results in representations that perform well on the second task but poorly on the first task due to high forgetting. On the other hand, representations trained with self-supervised learning (SSL) have higher first-task performance but they underperform on the second task. We show that simple modifications to supervised learning (SL+MLP) yield representations that are superior on the first task and on par with SL on the second task. We report average task-aware k-NN accuracy on 6 different 2-task combinations of CIFAR100, CIFAR10 and SVHN datasets (3 runs for each scenario).

the utilization of vast amounts of unlabeled data emerging on a daily basis and continually improve existing models.

A number of recent works study continual learning from a representation learning perspective and show that unsupervised approaches build more robust representations when trained continually [7, 25]. More specifically, [25] shows that self-supervised learning (SSL) methods build representations that are more robust to forgetting than supervised learning (SL). [7] notice that training SimCLR [4] have advantageous properties for continual learning compared to supervised training. However, it is still counter-intuitive that access to more information (labels) results in worse representations in continual learning.

One of the potential reasons is the transferability gap between supervised and unsupervised learning. It was believed that the superior transferability of unsupervised learning can be attributed to a special design of contrastive loss [15, 46] or lack of annotations during training [8, 36]. However, recent works [35, 41] identify that a multi-layer perceptron (MLP) projector commonly used in SSL [4, 5, 12, 45] is a crucial component that improves transferability of SSL models. Following that finding [35, 41] use an MLP projector to improve transferability of supervised learning and achieve state-of-the-art transfer learning performance, surpassing unsupervised methods.

In this work, encouraged by these advancements in improving the transferability of supervised models, we revisit supervision for continual representation learning. We argue that additional information (human annotations) should not hurt the quality of representations in continual learning, as suggested by [25]. Motivated by the latest study on transferability of representations in self-supervised and supervised learning, we aim to improve transferability between tasks in continual learning. We are the first to show that supervised models can contin-

ually learn representations of higher quality than self-supervised models when trained with a simple MLP head (see Figure 1). We identify the crucial role of an MLP projector in representation learning through the perspective of feature transferability, forgetting, and retention for continually trained models.

The main contributions of this paper are as follows:

- We empirically show that SL equipped with a simple MLP projector can learn higher-quality representations than SSL methods in continual finetuning scenarios in both in-distribution and transfer learning scenarios
- We show that the use of the MLP projector can be coupled with several continual learning methods, further improving their performance.
- We shed light on the reasons behind the strong performance of supervised learning with MLP projector: better transferability, lower forgetting, and increasing diversity of representations.

2 Related Work

Self-supervised Learning (SSL). Learning effective visual representations without annotations is a long-standing problem that aims at leveraging large volumes of unlabeled data. Recent SSL methods show impressive performance, matching or even exceeding the performance of their supervised equivalents [3–5, 12, 45]. The majority of these techniques rely on image augmentation methods to produce multiple views for a given sample. They train a model to be insensitive to these augmentations by ensuring that the network generates similar representations for the views of the same image and different representations for views of other images. In this work, we use BarlowTwins [45] which considers an objective function measuring the cross-correlation matrix between the features and SimCLR [4] which uses contrastive learning based on noise-contrastive estimation. A number of studies [1, 5, 16, 45] show that an MLP projector between the encoder and the loss function is a crucial component to prevent the collapse of the representations and improve their transferability.

Transferable representations. [41] seeks to understand the transferability gap between unsupervised (SSL) and supervised pretraining. They found out that adding a projection network (which is commonly used in SSL) boosts the transferability of the supervised models’ features. This was further explored in [35] and it was shown that it is possible to build representations that are good for both the source and the downstream tasks. In this work, we revisit those findings in the context of models learned on a sequence of tasks. Contrary to the transfer learning literature [8, 35], which usually focuses on the downstream task performance, we evaluate the model on all tasks during the sequential training. This allows us to gain more insight into learned representations, i.e. representation forgetting.

Supervised Continual Learning (SCL). SCL aims to create systems that can acquire the ability to solve novel tasks using new annotated data while retaining the knowledge acquired from previously learned tasks [31]. A popular formulation of CL is class-incremental learning (CIL) [27, 39] where each task introduces unseen classes that will not occur in the following tasks. In an exemplar-free setting, the

model is not allowed to store any samples from previous tasks which might be important in situations where privacy concerns apply and such a setting remains a great challenge [37]. A popular strategy is *feature distillation* [23, 43] which minimizes representational changes in subsequent learning stages by enforcing consistent output between the current model and the one trained in the previous task. Despite the recent progress in SCL, [19] highlight the fact that state-of-the-art SCL methods focus on eliminating bias and forgetting on a level of last classification layer. As a result they fail to improve the feature extractor during the continual training which is a main goal of this paper.

Unsupervised Continual Learning (UCL). Despite the success of SSL methods, they are designed to learn from large static datasets. UCL methods aim to overcome this issue and allow the models to learn from an ever-changing stream of data without excessive memory requirements. Recent works [10, 11, 25] apply SSL in the UCL setting and claim their superior results for continual representation learning. Most successful methods apply feature distillation through learnable non-linear projector: CaSSLe [10] distills features outputted by the projector while PFR [11] distills the features outputted by the backbone. UCL models are evaluated by measuring their representation strength through linear probing or k-nearest neighbors (k-NN) and this paper follows this evaluation protocol.

3 Experimental Setup

Datasets. We utilize four different datasets: CIFAR10 [22] (C10), CIFAR100 [22] (C100), SVHN [28] and ImageNet100 [38] (IN100), 100-class subset of the ILSVRC2012 dataset [34] with $\approx 130k$ images in high resolution (resized to 224×224). We consider popular settings in continual learning: CIFAR10/5, CIFAR100/5, CIFAR100/20 and ImageNet100/5 sequences, where D/N denotes that dataset D is split into N tasks with an equal number of classes in each task without overlapping ones. To gain further insight, we construct multiple two-task settings where we investigate representation strength and stability. We denote task shift with " \rightarrow ", e.g. sequence $A \rightarrow B$ means that the model was trained on two tasks, the first one was dataset A and the second one was dataset B . We consider C10 \rightarrow C100 and C100 \rightarrow C10 scenarios as having low distribution shifts, while C10 \rightarrow SVHN and SVHN \rightarrow C10 scenarios involve higher distribution shifts. We also perform transfer learning experiments on a set of diverse array of datasets: Food101 (Food) [2], Oxford-IIIT Pets (Pets) [32], Oxford Flowers-102 (Flowers) [29], Caltech101 (Caltech) [9], Stanford Cars (Cars) [21], FGVC-Aircraft (Aircrafts) [26], Describable Textures (DTD) [6] and Caltech-Birds-200 (Birds) [40].

Methods. We use the following supervised methods: (1) SL - the standard approach of training a model with linear classification head with a cross-entropy loss function [27]. (2) SL+MLP - SL with MLP projector added between the backbone and a linear head that is discarded at test-time (see Figure 1), (3) t-ReX [35], and (4) SupCon [17]. Note that SL is the only method that does not utilize an additional MLP projector during training. For SSL approaches we

choose BarlowTwins [45] and SimCLR [4]. Results denoted as SSL were obtained using BarlowTwins. We use ResNet-18 [13] as a feature extractor network for all the experiments. For CL strategies we use LwF [23], CaSSLe [10] and PFR [11]. Note that we do not include most of the approaches designed for class-incremental learning because most of them fail to improve their feature extractor during continual training [19].

Training. We use the code repository from CaSSLe [10] and we follow their training procedure. We train SSL models for 500 epochs per task using SGD optimizer with momentum with batch size 256 and cosine learning rate schedule. We adapt the procedure to SL by reducing the number of epochs to 100 per task and the batch size to 64. We tune the learning rate on CIFAR100/5 for each method. We use augmentations from SimCLR [4] for SSL and SupCon and augmentations proposed in [41] for SL approaches (SL, SL+MLP and t-ReX). Note that, unless stated otherwise, we investigate continual finetuning scenario and we do not employ any methods for continual learning nor replay buffer.

Evaluation. We use k-NN classifier to evaluate the quality of representations following [10,25] and Nearest Mean Classifier (NMC) as in [33,44] to evaluate the stability of representations. We use CKA [20] to measure the similarity between representations of two models. Moreover, we use forgetting (F) and forward transfer (FT) commonly used in continual learning [24]. We also measure task exclusion difference EXC [14] to evaluate the level of retention of task-specific features. We use subscripts to indicate the evaluation dataset, e.g. Acc_{C10} means "accuracy on C10 dataset". We report means and standard deviations computed across 3 runs unless stated otherwise.

4 Main Results

This section presents the experimental results of continual representation learning. In Section 4.1 we present our main results showing that supervised models can outperform self-supervised models in continual representation learning. In Section 4.2 we perform continual transfer learning evaluation that further supports this claim. In Section 4.3 we combine different models with CL strategies to investigate the synergy between them. Then we follow with an extensive analysis that sheds light on the reasons for improved performance. Section 5.1 investigates the quality of representations, including forgetting, task exclusion comparison, similarity, and forward transfer. Section 5.2 presents a spectral analysis of representations. In Section ??, we evaluate the stability of the representations learned by different methods under various distribution shifts. Finally, in Section 5.3 we present an ablation study on impact of different components of the projector as well as data efficiency and robustness to label noise.

4.1 Continual representation learning

Figure 2 presents our main finding. Namely, we show that supervised models can build stronger representations than self-supervised models under continual finetuning, contrary to previous beliefs [25]. We identify that the key component to improving the performance of supervised models is an additional MLP projector used during training and discarded afterward - without it, SL significantly underperforms compared to SSL.

Figure 3 presents the performance of SL, SSL and SL+MLP after each task. We identify two factors contributing to superior results of SL+MLP. Firstly, we observe that the performance of supervised models after the initial task is largely improved by the addition of the MLP projector, resulting in accuracy close to SSL models. In order to achieve good task-agnostic accuracy on the whole dataset (seen and unseen classes), the model trained on a single task needs to perform well on unseen data. Therefore, we attribute the advantage of SL+MLP to the increased transferability of representations induced by MLP projector, which is in line with [35, 41].

Secondly, we notice that SL+MLP is the only method able to incrementally accumulate knowledge and consistently improve performance. This observation is in line with the increasing diversity of features presented in Section 5.2.

Table 1 presents extended results including multiple SL and SSL approaches in continual finetuning. We observe that all the supervised methods equipped with the projector significantly outperform simple SL. SL+MLP, t-ReX, and SupCon achieve much higher results than SL in all the finetuning experiments. What is worth noting is the fact that all these methods were trained with different




Figure 2: SL finetuning underperforms compared to SSL. However, when equipped with the MLP projector it consistently outperforms SSL. We report the difference in k-NN accuracy (%) between supervised approaches and SSL.




Figure 3: SL+MLP: (1) achieves strong performance after the initial task compared to SL which indicates that it produces representations that are transferable to the unseen tasks; (2) is the only method that is able to accumulate knowledge learned on a sequence of tasks. We report task-agnostic k-NN accuracy after each task on the whole dataset (notice that yet unseen tasks are also included in the evaluation).

Table 1: Supervised methods that utilize MLP projector largely outperform vanilla SL. They also outperform SSL on most of the datasets. We report k-NN accuracy of the learned representations. The best result in **bold** and the second best underlined.

Method	C10/5	C100/5	C100/20	IN100/5
SL	59.8±1.8	45.3±0.7	23.1±0.2	40.4±1.2
SL+MLP	65.9±0.7	61.9±0.5	<u>47.1±0.7</u>	62.4±0.4
t-ReX	69.3±1.1	<u>59.2±0.6</u>	50.8±0.1	59.2±0.6
SupCon	60.4±0.6	49.4±0.3	30.0±0.7	57.6±0.6
BarlowTwins	76.2±1.2	54.1±0.3	40.0±0.8	57.0±0.4
SimCLR	<u>72.4±1.3</u>	48.9±0.4	33.4±0.5	54.7±0.4

supervised losses: SL+MLP uses cross-entropy, t-ReX uses cosine softmax cross-entropy and SupCon uses supervised contrastive loss. However, they all utilize the MLP projector and all outperform vanilla SL and SSL on most of the datasets.

4.2 Continual transfer learning

In this Section we evaluate the continually trained feature extractors on a set of diverse downstream classification tasks described in Section 3: Food, Pets, Flowers, Caltech, Cars, Aircrafts, DTD and Birds. We evaluate the performance on downstream datasets after each task in the continual learning sequence and present the results in Figure 4. We observe that SL+MLP outperforms the competitors on all but one dataset. Interestingly, the pattern of results is similar to the *in-distribution* results presented in Figure 3: (1) SL+MLP is on-par with SSL after the first task and significantly outperforms SL, and (2) improves its performance when trained continually. These results highlight the superior transferability of the representations learned by the proposed SL+MLP approach.

4.3 Synergy with CL methods

In this Section, we investigate our findings combined with different CL approaches. More precisely, we combine both supervised and self-supervised methods with existing CL methods to verify the synergy between them.

Firstly, we evaluate the resulting combinations *in-distribution* as in Section 4.1 and present the results in Table 2. We observe that most of the combinations of methods and CL strategies are outperforming simple finetuning (the only exceptions are some combinations of LwF with supervised methods). Moreover, the conclusions from simple finetuning scenarios still hold – supervised methods equipped with MLP projectors outperform vanilla SL and SSL. However, with the help of CL methods, the performance gap is much smaller.

We also evaluate transfer learning performance as in Section 4.2 and present the results in Table 3. We observe that CL methods bring an improvement over finetuning in almost all experiments. Similarly to the in-distribution experiments, the best results are obtained by the supervised methods that utilize MLP projector combined with one of CL strategies.

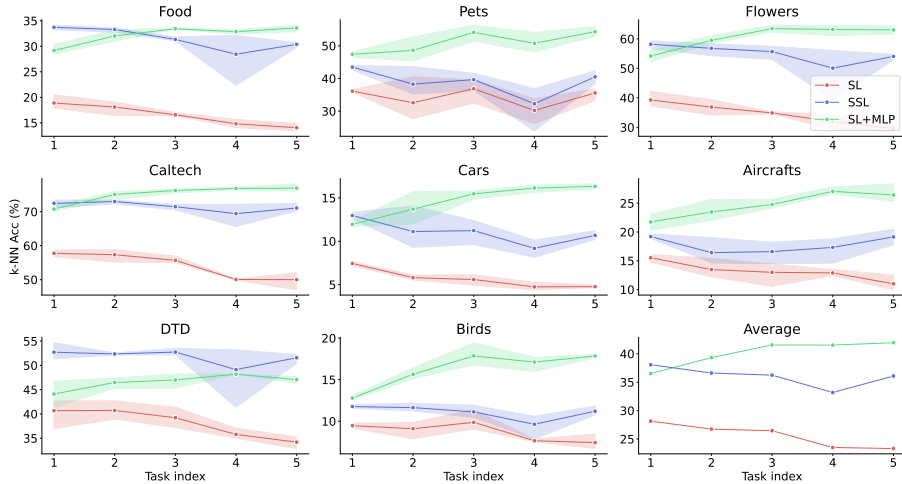


Figure 4: Representations learned by SL+MLP are more transferable than those learned by SL and SSL. They also improve when trained on new tasks. We present the results of the models trained continually on ImageNet/5 and evaluated after each task. We report k-NN accuracy (%) on a set of 8 diverse downstream classification tasks and an average performance.

To sum up, models that achieve the best transfer learning performance are those: (1) trained in a supervised way (2) with the use of the MLP projector and (3) coupled with CL strategy based on temporal learnable projection, namely CaSSLe or PFR. These conclusions are coherent with the previous Section where a similar set of models performed the best.

5 Analysis

5.1 Quality of representations

Forgetting. In Table 4 we observe high representation forgetting for SL, significantly lower for SSL, and the lowest for SL equipped with MLP projector. Low forgetting is an important factor contributing to the superior performance of SL+MLP.

Task exclusion difference. In the two-task sequence *EXC* answers the question: *what is the performance gap between the model trained on B and a model trained on a sequence $A \rightarrow B$ when evaluated on A?* Results from Table 4 show that SL achieves small positive *EXC* meaning that it forgets most features specific to the initial task (but not all of them). SL+MLP achieves the highest *EXC* which shows that it is able to successfully retain a large portion of task-specific features. Surprisingly, SSL exhibits negative *EXC*. It means that it is better to train SSL model from scratch on another task than to finetune the model pretrained on the task of interest.

Table 2: The effect of MLP projector compounds with the positive effect of CL strategies leading to combinations that outperform SSL on most datasets. We report k-NN accuracy of the learned representations. The best result in each group in **bold** and the second best underlined.

Method	CL strategy	C10/5		C100/5		C100/20		IN100/5	
		Acc	Δ	Acc	Δ	Acc	Δ	Acc	Δ
SUPERVISED CONTINUAL LEARNING									
SL	Finetune	59.8±1.8		45.3±0.7		23.1±0.2		40.4±1.2	
	LwF	69.6±1.1	+9.8	62.9±0.1	+17.6	51.5±0.2	+28.4	67.0±0.1	+26.6
	PFR	71.0±2.0	+11.2	<u>63.3±0.2</u>	+18.0	52.3±1.2	+29.2	65.0±0.3	+24.6
SL+MLP	Finetune	65.9±0.7		61.9±0.5		47.1±0.7		62.4±0.4	
	LwF	72.6±3.4	+6.7	58.7±0.2	-3.2	51.9±0.1	+4.8	60.4±0.2	-2.0
	PFR	<u>76.3±1.0</u>	+10.4	63.6±0.2	+1.7	54.5±0.2	+7.4	65.2±0.1	+2.8
t-ReX	Finetune	69.3±1.1		59.2±0.6		50.8±0.1		59.2±0.6	
	LwF	74.5±0.7	+5.2	58.3±0.4	-0.9	50.4±0.1	-0.4	58.6±1.0	-0.6
	PFR	75.9±1.2	+6.6	60.9±0.5	+1.7	<u>53.4±0.3</u>	+2.6	63.9±0.6	+4.7
SupCon	Finetune	60.4±0.6		49.4±0.3		30.0±0.7		57.6±0.6	
	CaSSLe	75.1±0.4	+14.7	61.1±0.2	+11.7	49.2±1.2	+19.2	70.4±0.6	+12.8
	PFR	78.1±1.0	+17.7	57.0±0.2	+7.6	51.2±0.8	+21.2	<u>68.0±0.7</u>	+10.4
UNSUPERVISED CONTINUAL LEARNING									
BarlowTwins	Finetune	76.2±1.2		54.1±0.3		40.0±0.8		57.0±0.4	
	CaSSLe	80.9±0.2	+4.7	58.6±0.6	+4.5	<u>49.3±0.1</u>	+9.3	64.9±0.1	+7.9
	PFR	78.8±0.6	+2.6	<u>57.2±0.2</u>	+3.1	46.0±0.7	+6.0	<u>61.1±0.2</u>	+4.1
SimCLR	Finetune	72.4±1.3		48.9±0.4		33.4±0.5		54.7±0.4	
	CaSSLe	<u>80.6±0.5</u>	+8.2	55.9±0.5	+7.0	48.2±0.4	+14.8	59.3±0.5	+4.6
	PFR	79.2±0.7	+6.8	53.8±0.3	+4.9	49.4±0.1	+16.0	57.7±0.2	+3.0

CKA similarity. In Table 4 we report CKA similarity between the models trained on C10 and the rest of the models. We observe that usage of MLP head in SL increases CKA between the C10 model and other models. Moreover, in the case of SL+MLP, the models pretrained on C10 and finetuned on another task have higher similarity to C10 models than the models trained on another dataset from scratch. This is not necessarily the case for SL models. SSL models have the highest CKA scores, however, they usually underperform compared to SL+MLP. This suggests that SSL produces similar features when trained on different datasets but their discriminative power for a classification task is worse than those learned with SL+MLP.

Positive and negative forward transfer. We present the results of the forward transfer evaluation in Table 5. All the methods benefit from pretraining on CIFAR100 which is semantically close to CIFAR10. However, pretraining on semantically distant SVHN hinders the performance of SL but it hardly influences the performance of SSL and SL+MLP.

5.2 Spectra of representations

To gain further insight into the properties of continually trained representations, we analyze the spectrum of their covariance matrix. We follow the procedure

Table 3: SL+MLP combined with PFR achieves the best average performance on transfer learning tasks. CL methods also improve the performance of SL and SSL but they achieve worse overall performance. Best results in each group (finetuning and continual learning methods) in **bold**.

Method	Food	Pets	Flowers	Caltech	Cars	Aircrafts	DTD	Birds	Avg
FINETUNING									
SL	14.0±0.6	35.6±1.6	29.6±1.0	50.0±2.1	4.8±0.2	11.0±1.1	34.2±1.0	7.4±0.8	23.3
SSL	30.4±0.4	40.5±1.6	54.0±0.8	71.1±1.0	10.7±0.4	19.1±1.1	51.6±0.8	11.2±0.5	36.1
SL+MLP	33.6±0.5	54.3±1.3	63.0±1.1	76.9±0.9	16.3±0.2	26.4±1.4	47.1±0.5	17.8±0.2	41.9
CONTINUAL LEARNING METHODS									
SL+LwF	30.2±0.4	57.8±0.8	57.9±0.7	72.7±1.2	13.6±0.3	22.2±0.3	46.6±0.5	17.1±0.4	39.8
SL+PFR	30.4±0.6	57.4±0.3	58.5±0.7	72.6±0.2	14.0±0.2	21.9±0.6	46.3±1.1	16.3±0.2	39.7
SSL+PFR	35.1±0.3	43.4±1.0	59.9±1.1	74.6±0.4	12.2±0.5	19.8±0.3	53.9±0.7	12.4±0.4	38.9
SSL+CaSSLe	39.1±0.2	48.5±0.7	65.1±0.8	75.5±0.4	14.2±0.6	21.1±0.9	56.4±0.5	13.5±0.2	41.7
SL+MLP+LwF	34.7±0.4	55.4±0.9	61.9±1.3	75.3±0.4	15.1±0.1	24.5±0.5	47.3±1.0	16.2±0.2	41.3
SL+MLP+PFR	38.0±0.2	58.8±0.7	67.8±0.1	79.3±0.2	17.5±0.3	26.2±0.3	49.2±0.4	18.1±0.2	44.4

Table 4: We observe high representation forgetting for SL, significantly lower for SSL, and the lowest for SL equipped with MLP projector. SL is able to preserve a small fraction of task-specific features while SL+MLP can retain much more, based on their *EXC* scores. Surprisingly, SSL achieves negative *EXC* meaning that pretraining on a given task hurts the performance on this task after the finetuning. We report evaluation on CIFAR10 dataset. The best value between SL, SSL, and SL-MLP in **bold**.

Training sequence	SL				SSL				SL+MLP			
	Acc ↑	F ↓	EXC ↑	CKA ↑	Acc ↑	F ↓	EXC ↑	CKA ↑	Acc ↑	F ↓	EXC ↑	CKA ↑
C10	93.7±0.1	-	-	-	88.8±0.1	-	-	-	93.3±0.1	-	-	-
C100	79.4±0.3	-	-	0.46	80.8±0.1	-	-	0.56	84.5±0.4	-	-	0.49
C10→C100	81.3±0.6	12.4±0.6	1.9±0.3	0.50	79.1±0.2	9.7±0.3	-1.8±0.2	0.52	88.8±0.2	4.5±0.3	4.3±0.6	0.57
SVHN	27.3±0.2	-	-	0.05	58.6±1.2	-	-	0.27	56.3±0.2	-	-	0.20
C10→SVHN	29.6±1.1	64.1±1.1	2.3±1.3	0.05	54.9±0.7	33.8±0.7	-3.7±1.9	0.25	62.7±0.8	30.6±0.8	6.4±1.0	0.25

from [16]. We gather the representations of the validation set and compute the covariance matrix of the representations, C . We perform singular value decomposition of the covariance matrix $C = USV^T$, where $S = \text{diag}(\sigma^k)$ and σ^k is k -th singular value of C . We call the representations *diverse* when a large number of principal directions (independent features) is needed to explain most of the variance. Figure 5 presents how singular value spectra change after each task for different training methods and different sequences of tasks.

Representation collapse. Figure 5 reveals that supervised learning exhibits signs of neural collapse [30] - a large fraction of variance is described by a few dimensions roughly equal to the number of classes in the training set. This is an undesirable property in continual representation learning as the representations should be more versatile and useful not only for current but also for past and future tasks. SSL, on the other hand, learns a more diverse set of features resulting

Table 5: All methods benefit from pretraining on C100 which is semantically close to C10. However, pretraining on semantically distant SVHN hinders the performance of SL. Best method in **bold**.

Method	C10	C100→C10		SVHN→C10	
	$Acc_{C10} \uparrow$	$Acc_{C10} \uparrow$	$FT_{C10} \uparrow$	$Acc_{C10} \uparrow$	$FT_{C10} \uparrow$
SL	93.7±0.1	94.3±0.1	0.6±0.0	93.3±0.1	-0.4±0.1
SSL	88.8±0.1	89.2±0.1	0.5±0.2	88.5±0.1	-0.3±0.2
SL+MLP	93.3±0.1	94.3±0.1	1.0±0.0	93.2±0.2	-0.1±0.1

Table 6: Impact of projector architecture on different methods. Projector architectures are arranged in order of increasing number of parameters. We report k-NN accuracy (Acc) and average improvement over the method without projector (Δ). Best results for each method in **bold**.

Method	Projector architecture	CIFAR10/5		CIFAR100/5	
		Acc	Δ	Acc	Δ
SL	None	59.8±1.8		45.3±0.7	
	MLPP [41]	65.9±0.7	+6.1	61.9±0.5	+16.6
	t-ReX [35]	67.3±0.1	+7.5	58.3±0.2	+13.0
t-ReX	None	60.0±1.2		36.7±1.1	
	MLPP [41]	66.7±1.4	+6.7	58.1±0.2	+21.4
	t-ReX [35]	69.3±1.1	+9.3	59.2±0.6	+22.5
SupCon	None	58.7±0.4		23.7±0.5	
	Linear	61.0±0.7	+2.3	37.7±0.6	+14.0
	SupCon [17]	60.4±0.6	+1.7	49.4±0.3	+25.7
	MLPP [41]	66.2±0.7	+7.5	54.3±0.5	+30.6
	t-ReX [35]	63.5±0.2	+4.8	58.1±0.1	+34.4

in a flatter singular values spectrum. In our experiments adding MLP to SL prevents neural collapse and results in features’ properties more similar to SSL. **Evolution of representations.** An important property of representations learned in continual learning is the change in their diversity: the diversity that increases after each task is desired. In Figure 5 we can observe that for SL, the diversity of the features usually decreases, except for C10→C100 where the increase is caused by a higher number of training classes [30]. For SSL, the diversity increases in the five-task scenario and remains close to constant for two-task settings. SL+MLP is able to improve the diversity of the representations consistently across all the presented scenarios suggesting its superiority in continual representation learning. It may be related to its ability to effectively accumulate knowledge when trained on a sequence of tasks, as presented in Figure 3.

5.3 Ablation study

Architecture of the projector. In Table 6 we examine the impact of different projector architectures on supervised learning methods. All the methods achieve the worst performance when not using a projector. Moreover, all the methods

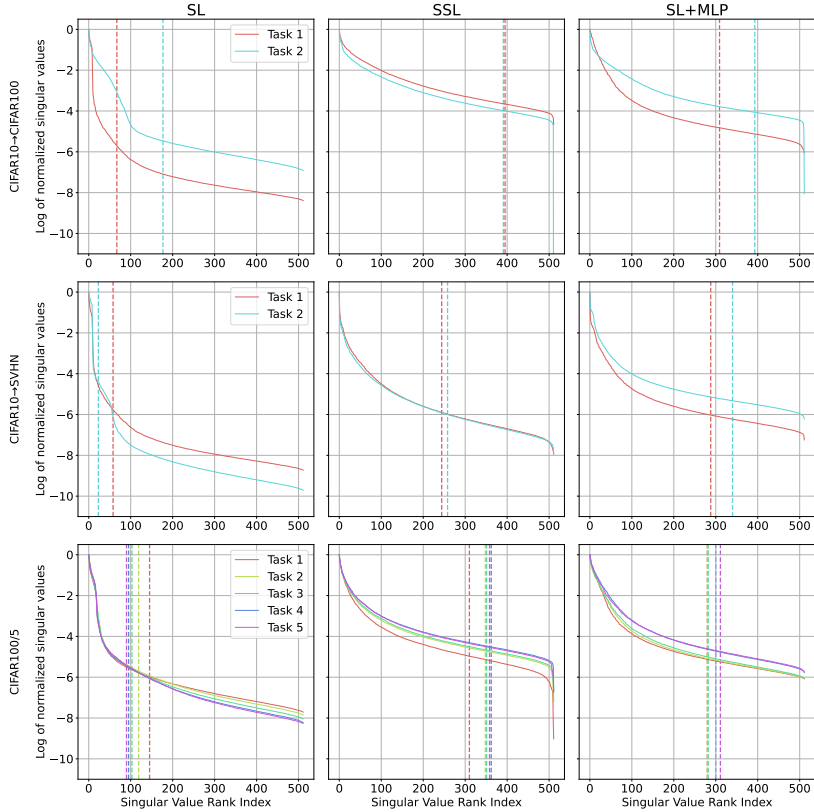


Figure 5: Singular value spectra of 512-dimensional representation space. Representations learned with SL+MLP (right) exhibit desirable properties from the continual learning point of view: (1) they consist of a more diverse set of features (contrary to SL, left); (2) they improve feature diversity when learning new tasks consistently across all the presented settings. Singular values are ordered descending, are normalized by σ^1 (the largest singular value) and the scale is logarithmic. Vertical dashed lines denote 95% of the variance explained. Intuitively, it indicates how many relevant independent features the representation contains.

achieve similarly high accuracy when with a suitable projector. These results highlight the importance of MLP projector for supervised continual representation learning and diminish the importance of other factors (*e.g.* loss function) on the final performance.

To gain a deeper insight into the impact of different components of the projector we systematically investigate the impact of the depth and width of the projector. We follow the MLPP [41] architecture for the projector, using a linear layer, batch normalization (BN), and ReLU activation, which we call a *block* with a hidden dimension d_h . This block is followed by an output linear

Figure 6: One MLP block brings significant improvements and the next blocks have minor effect. Impact of projector depth and width on CIFAR100/5. Note that d_h does not apply to linear projectors.

$n =$	0 (linear)		1		2		3	
$h_o =$	512	2048	512	2048	512	2048	512	2048
$d_h = 512$			60.06	60.88	60.61	60.45	60.59	60.93
$d_h = 1024$			60.61	60.50	61.04	60.61	60.84	61.25
$d_h = 2048$	46.74	47.45	60.54	60.88	61.42	61.05	60.97	61.00
$d_h = 4096$			61.46	61.10	60.95	61.46	61.09	60.91

Figure 7: BatchNorm seems to have the highest positive impact on the final performance. Impact of block components on CIFAR100/5.

Input Lin.	✗	✗	✓	✓
ReLU	✗	✓	✗	✓
BN ✗	46.74	49.06	47.59	53.87
✓	57.33	59.15	58.47	61.46




Figure 8: SL+MLP outperforms SSL and SL approaches: (left) trained on full dataset even with only 30% of data available and (right) with 40% label noise. We present the accuracy (%) after the final task on CIFAR100/5 for different data fractions (left) and varying label noise level (right).

layer with a hidden dimension d_o , and then a classification layer. The projector contains n blocks, where $n = 0$ means a linear projector. We present the results in Table 6. We observe that the width and depth of the projector have a very limited impact on the final performance as long as the projector has at least one block. Therefore, we further decompose the projector starting from an MLP with a single block ($d_h = 4096$, $d_o = 512$) and remove basic components. We present the results in Table 7. We observe that BN has the highest positive impact on performance. However, using BN without the linear layer and ReLU lags 4 p.p. behind the performance of the full block, highlighting the importance of all the components.

Data efficiency. SL with MLP projector outperforms SSL in continual representation learning when having access to the same amount of data. However, in real-world scenarios, SSL is able to utilize vast amounts of unlabeled data while SL needs costly data annotations. Therefore, we examine how SL behaves when trained on a fraction of data available for SSL in each task to simulate limited access to labeled data. We restrict the data and select a fraction of it (equal for each task). We present the results in Figure 8 (left). SL+MLP surpasses SimCLR with less than 20% of data and BarlowTwins with about 30% of data.

Robustness to mislabelling. We evaluate the robustness of our claims against noisy labels following previous study [18] and randomly change the labels of the fraction of samples in the train dataset. Figure 8 (right) presents the results

on CIFAR100/5. We observe that SL+MLP outperforms SSL even with 40% label noise. It validates the efficacy of SL+MLP in the presence of imperfect annotations.

6 Discussion and limitations

Although supervised learning with the MLP projection head seems to be more effective in continual representation learning, it comes at a price. SL requires mundane image labeling of the whole dataset which can be costly and impractical at scale. Self-supervised learning, on the other hand, is not dependent on image annotations and, therefore, can operate on a vast amount of unlabeled data.

However, SSL faces its own limitations. Firstly, most SSL approaches depend on strong image augmentations and learn representations that are invariant to them [3, 4, 45]. This can hinder the performance on the downstream tasks which require attention to the traits that it has been trained to be invariant to [42]. Moreover, SSL usually requires longer training which increases computational requirements in comparison to SL. It is also worth noting that both SL+MLP and SSL introduce additional costs to the model during the training, as both introduce MLP projector that requires more computational requirements. However, at test time every method operates at the same number of parameters, as we discard MLP projectors after training.

Furthermore, it’s worth re-emphasizing that this work focuses on continual representation learning. While we utilize data from previous tasks to construct k-NN and nearest mean classifiers for evaluating learned representations, our primary objective is not centered on the continual approach to the downstream task (classification). We are not delving into class-incremental learning, a prevalent continual learning setting. Nonetheless, our analysis of representation strength and stability can offer valuable insights into continual learning dynamics, potentially aiding in the creation of more effective algorithms for continual downstream task solutions.

7 Conclusions

In this work, we are first to show that supervised learning can significantly outperform self-supervised learning in continual representation learning. We achieve it by equipping SL with a simple MLP projector discarded after the training, following the common practice from SSL. We show that SL+MLP can be successfully coupled with several continual learning strategies, further improving the performance. Finally, we shed some light on the reasons for improved performance when using MLP with SL: better transferability, lower forgetting, and higher diversity of learned features.

Acknowledgments

Daniel Marczak is supported by National Centre of Science (NCN, Poland) Grant No. 2021/43/O/ST6/02482. This research was partially funded by National Science Centre, Poland, grant no 2020/39/B/ST6/01511, grant no 2022/45/B/ST6/02817, and grant no 2023/51/D/ST6/02846. Bartłomiej Twardowski acknowledges the grant RYC2021-032765-I. This paper has been supported by the Horizon Europe Programme (HORIZON-CL4-2022-HUMAN-02) under the project "ELIAS: European Lighthouse of AI for Sustainability", GA no. 101120237. We gratefully acknowledge Polish high-performance computing infrastructure PLGrid (HPC Center: ACK Cyfronet AGH) for providing computer facilities and support within computational grant no. PLG/2023/016393.

References

1. Bordes, F., Balestriero, R., Garrido, Q., Bardes, A., Vincent, P.: Guillotine regularization: Why removing layers is needed to improve generalization in self-supervised learning. *TMLR* (2023)
2. Bossard, L., Guillaumin, M., Van Gool, L.: Food-101 – mining discriminative components with random forests. In: *ECCV* (2014)
3. Caron, M., Touvron, H., Misra, I., Jegou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. *ICCV* (2021)
4. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.E.: A simple framework for contrastive learning of visual representations. In: *ICML* (2020)
5. Chen, X., He, K.: Exploring simple siamese representation learning. *CVPR* (2020)
6. Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., Vedaldi, A.: Describing textures in the wild. In: *CVPR* (2014)
7. Davari, M.R., Asadi, N., Mudur, S., Aljundi, R., Belilovsky, E.: Probing representation forgetting in supervised and unsupervised continual learning. In: *CVPR* (2022)
8. Ericsson, L., Gouk, H., Hospedales, T.M.: How well do self-supervised models transfer? *CVPR* (2020)
9. Fei-Fei, L., Fergus, R., Perona, P.: One-shot learning of object categories. *IEEE TPAMI* (2006)
10. Fini, E., da Costa, V.G.T., Alameda-Pineda, X., Ricci, E., Alahari, K., Mairal, J.: Self-supervised models are continual learners. In: *CVPR* (2022)
11. Gomez-Villa, A., Twardowski, B., Yu, L., Bagdanov, A.D., van de Weijer, J.: Continually learning self-supervised representations with projected functional regularization. *CVPR Workshops* (2021)
12. Grill, J.B., Strub, F., Althç'e, F., Tallec, C., Richemond, P.H., Buchatskaya, E., Doersch, C., Pires, B., Guo, Z., Azar, M.G., Piot, B., Kavukcuoglu, K., Munos, R., Valko, M.: Bootstrap your own latent: A new approach to self-supervised learning. *NeurIPS* (2020)
13. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *CVPR* (2016)
14. Hess, T., Verwimp, E., van de Ven, G.M., Tuytelaars, T.: Knowledge accumulation in continually learned representations and the issue of feature forgetting. *TMLR* (2024)

15. Islam, A., Chen, C.F., Panda, R., Karlinsky, L., Radke, R., Feris, R.: A broad study on the transferability of visual representations with contrastive learning. *ICCV* (2021)
16. Jing, L., Vincent, P., LeCun, Y., Tian, Y.: Understanding dimensional collapse in contrastive self-supervised learning. In: *ICLR* (2022)
17. Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., Krishnan, D.: Supervised contrastive learning. *NeurIPS* (2020)
18. Kim, C.D., Jeong, J., Moon, S., Kim, G.: Continual learning on noisy data streams via self-purified replay. In: *ICCV* (2021)
19. Kim, D., Han, B.: On the stability-plasticity dilemma of class-incremental learning. *CVPR* (2023)
20. Kornblith, S., Norouzi, M., Lee, H., Hinton, G.E.: Similarity of neural network representations revisited. *ICML* (2019)
21. Krause, J., Stark, M., Deng, J., Fei-Fei, L.: 3d object representations for fine-grained categorization. In: *4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13)* (2013)
22. Krizhevsky, A.: Learning multiple layers of features from tiny images. University of Toronto (2009)
23. Li, Z., Hoiem, D.: Learning without forgetting. *IEEE TPAMI* (2018)
24. Lopez-Paz, D., Ranzato, M.: Gradient episodic memory for continual learning. *NeurIPS* (2017)
25. Madaan, D., Yoon, J., Li, Y., Liu, Y., Hwang, S.J.: Representational continuity for unsupervised continual learning. In: *ICLR* (2022)
26. Maji, S., Rahtu, E., Kannala, J., Blaschko, M., Vedaldi, A.: Fine-grained visual classification of aircraft (2013)
27. Masana, M., Liu, X., Twardowski, B., Menta, M., Bagdanov, A.D., van de Weijer, J.: Class-incremental learning: Survey and performance evaluation on image classification. *IEEE TPAMI* (2023)
28. Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., Ng, A.Y.: Reading digits in natural images with unsupervised feature learning (2011)
29. Nilsback, M.E., Zisserman, A.: Automated flower classification over a large number of classes. In: *Indian Conference on Computer Vision, Graphics and Image Processing* (2008)
30. Pappayan, V., Han, X.Y., Donoho, D.L.: Prevalence of neural collapse during the terminal phase of deep learning training. *PNAS* (2020)
31. Parisi, G.I., Kemker, R., Part, J.L., Kanan, C., Wermter, S.: Continual lifelong learning with neural networks: A review. *Neural Networks* (2019)
32. Parkhi, O.M., Vedaldi, A., Zisserman, A., Jawahar, C.V.: Cats and dogs. In: *CVPR* (2012)
33. Rebuffi, S., Kolesnikov, A., Sperl, G., Lampert, C.H.: icarl: Incremental classifier and representation learning. In: *CVPR* (2017)
34. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: ImageNet Large Scale Visual Recognition Challenge. *IJCV* (2015)
35. Sariyildiz, M.B., Kalantidis, Y., Alahari, K., Larlus, D.: No reason for no supervision: Improved generalization in supervised models. In: *ICLR* (2023)
36. Sariyildiz, M.B., Kalantidis, Y., Larlus, D., Karteek, A.: Concept generalization in visual representation learning. *ICCV* (2021)
37. Smith, J.S., Tian, J., Halbe, S., Hsu, Y., Kira, Z.: A closer look at rehearsal-free continual learning. In: *CVPR Workshops* (2023)

38. Tian, Y., Krishnan, D., Isola, P.: Contrastive multiview coding. ECCV (2019)
39. van de Ven, G.M., Tuytelaars, T., Tolias, A.S.: Three types of incremental learning. Nature Machine Intelligence (2022)
40. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The caltech-ucsd birds-200-2011 dataset (2011)
41. Wang, Y., Tang, S., Zhu, F., Bai, L., Zhao, R., Qi, D., Ouyang, W.: Revisiting the transferability of supervised pretraining: an mlp perspective. CVPR (2021)
42. Xiao, T., Wang, X., Efros, A.A., Darrell, T.: What should not be contrastive in contrastive learning. In: ICLR (2021)
43. Yan, S., Xie, J., He, X.: DER: dynamically expandable representation for class incremental learning. In: CVPR (2021)
44. Yu, L., Twardowski, B., Liu, X., Herranz, L., Wang, K., Cheng, Y., Jui, S., van de Weijer, J.: Semantic drift compensation for class-incremental learning. In: CVPR (2020)
45. Zbontar, J., Jing, L., Misra, I., LeCun, Y., Deny, S.: Barlow Twins: Self-Supervised Learning via Redundancy Reduction. ICML (2021)
46. Zhao, N., Wu, Z., Lau, R.W.H., Lin, S.: What makes instance discrimination good for transfer learning? ICLR (2020)