

SLCA: Slow Learner with Classifier Alignment for Continual Learning on a Pre-trained Model

Gengwei Zhang^{1*}, Liyuan Wang^{2*}, Guoliang Kang^{3,4}, Ling Chen¹, Yunchao Wei^{5,6}

¹ AAIL, University of Technology Sydney ² Tsinghua University

³ Beihang University ⁴ Zhongguancun Laboratory

⁵ Institute of Information Science, Beijing Jiaotong University

⁶ Beijing Key Laboratory of Advanced Information Science and Network

{zgw david, kgl.prml, wychao1987}@gmail.com; wly19@mail.tsinghua.org.cn; ling.chen@uts.edu.au

Abstract

The goal of continual learning is to improve the performance of recognition models in learning sequentially arrived data. Although most existing works are established on the premise of learning from scratch, growing efforts have been devoted to incorporating the benefits of pre-training. However, how to adaptively exploit the pre-trained knowledge for each incremental task while maintaining its generalizability remains an open question. In this work, we present an extensive analysis for continual learning on a pre-trained model (CLPM), and attribute the key challenge to a progressive overfitting problem. Observing that selectively reducing the learning rate can almost resolve this issue in the representation layer, we propose a simple but extremely effective approach named Slow Learner with Classifier Alignment (SLCA), which further improves the classification layer by modeling the class-wise distributions and aligning the classification layers in a post-hoc fashion. Across a variety of scenarios, our proposal provides substantial improvements for CLPM (e.g., up to 49.76%, 50.05%, 44.69% and 40.16% on Split CIFAR-100, Split ImageNet-R, Split CUB-200 and Split Cars-196, respectively), and thus outperforms state-of-the-art approaches by a large margin. Based on such a strong baseline, critical factors and promising directions are analyzed in-depth to facilitate subsequent research. Code has been made available at: <https://github.com/GengDavid/SLCA>.

1. Introduction

Continual learning aims to learn effectively from sequentially arrived data, behaving as if they were observed

simultaneously. Current efforts are mainly based on the premise of learning from scratch, attempting to mitigate catastrophic forgetting [27] of previously-learned knowledge when adapting to each incremental task. However, the success of large-scale pre-training has revolutionized the training paradigm of deep neural networks. The pre-training stage brings both strong knowledge transfer and robustness to catastrophic forgetting for downstream continual learning [43], which tends to be more significant as the scale of pre-training increases [31, 28]. Therefore, continual learning on a pre-trained model (CLPM) turns out to be an emerging direction and receives growing attention.

For CLPM, the pre-trained knowledge is usually expressed by the representation layer, adapted to a sequence of incremental tasks. There are two major strategies to leverage the pre-trained knowledge [43]: (1) fine-tune the representations, or (2) keep the representations fixed while learning a few additional parameters (e.g., adaptor [16], prompt [24], instruction [8], etc.). Although (2) is becoming dominant in natural language processing (NLP) [18], the choice of (1) and (2) remains an open question for continual learning in computer vision (CV). The recently proposed prompt-based approaches, such as L2P [46] and DualPrompt [45], followed the second strategy and reported to be *far superior* to the traditional continual learning baselines of fine-tuning the representation layer. On the other hand, since the large amount of pre-training data are typically unlabeled and may also arrive incrementally, it seems more reasonable to use self-supervised pre-training than supervised pre-training [5, 43], also regarding that (upstream) continual learning in a self-supervised manner is generally more robust to catastrophic forgetting [17, 26, 9].

In this work, however, we present an extensive analysis to *reconsider the current progress and technical route* for CLPM in CV. Specifically, CLPM poses a critical challenge for continual learning that the pre-trained knowledge

*Equal contribution.

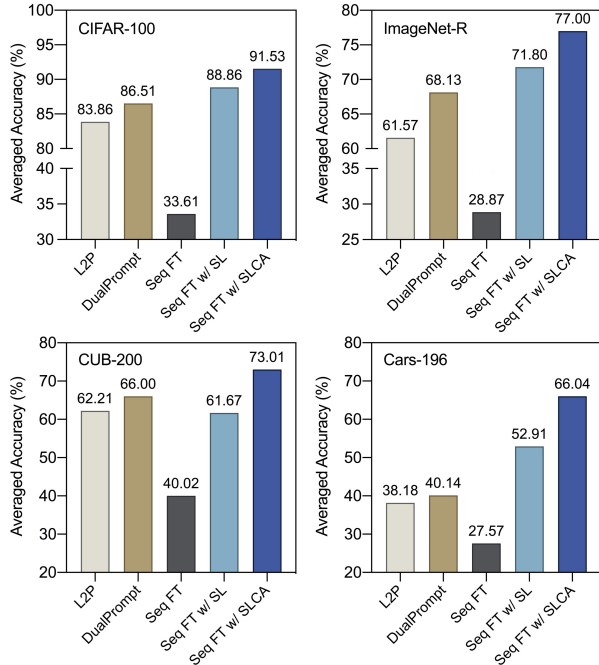


Figure 1. Performance of continual learning with supervised pre-training of ImageNet-21K. The proposed Slow Learner (SL) with Classifier Alignment (SLCA) enables sequential fine-tuning (Seq FT) to outperform prompt-based approaches such as L2P [46] and DualPrompt [45] by a large margin.

should be adaptively exploited for the current task while maintaining generalizability for future tasks. For traditional continual learning baselines that sequentially fine-tune (Seq FT) the entire model, a uniform learning rate is typically too large for the representation layer, leading to a *progressive overfitting* problem where the pre-trained representations overfit to each incremental task and gradually lose generalizability. This is the major cause of their inferiority to the recent prompt-based approaches. Such a phenomenon is further exacerbated by using the more realistic self-supervised pre-training, which generally requires larger updates of the pre-trained representations to adequately accommodate each incremental task.

To address the above challenges, we propose a simple but extremely effective approach named Slow Learner with Classifier Alignment (SLCA). First, we observe that the progressive overfitting problem of the representation layer can be almost avoided by selectively reducing its learning rate (*i.e.*, using a slow learner), which is sufficient to achieve a good trade-off between task specificity and generalizability in continually updating the pre-trained representations. Based on the proposed slow learner, we further improve the classification layer by modeling the class-wise distributions and aligning the classification predictions, so as to properly balance the relationships between different tasks.

Across a variety of continual learning benchmarks un-

der supervised or self-supervised pre-training, our proposal provides substantial improvements for CLPM in CV and dramatically fills the gap of current progress from the upper bound performance. For example, the naive Seq FT is improved by **49.76%**, **50.05%**, **44.69%** and **40.16%** on Split CIFAR-100, Split ImageNet-R, Split CUB-200 and Split Cars-196, respectively, thus outperforming the SOTA approaches by a large margin (summarized in Fig. 1). On Split CIFAR-100 and Split ImageNet-R, the performance gap is only less than **2%** for supervised pre-training and less than **4%** for self-supervised pre-training.

Our contributions include three aspects: (1) We present an extensive analysis of continual learning on a pre-trained model (CLPM), and demonstrate that the progressive overfitting problem is the key challenge for traditional continual learning baselines. (2) We propose a simple but effective approach to address this problem, which provides substantial improvements for CLPM and clearly outperforms the state-of-the-art approaches, serving as a strong baseline to re-evaluate the current progress and technical route. (3) Our results further identify critical factors and promising directions for CLPM, such as pre-training paradigm and downstream granularity, so as to facilitate subsequent research.

2. Related Work

Existing work on continual learning mainly focuses on sequential training of deep neural network(s) from scratch, seeking to effectively learn each new task without severely forgetting the old tasks. Representative strategies include regularization-based approaches [19, 1, 51, 25, 41, 7], which preserve the old model and selectively stabilize changes of parameters or predictions; replay-based approaches [49, 30, 3, 44, 40, 39], which approximate and recover the previously-learned data distributions; architecture-based approaches [34, 33, 42, 50], which allocate dedicated parameter sub-spaces for each incremental task; etc.

Recent works have increasingly explored the benefits of pre-training for continual learning. For example, the representations obtained from supervised pre-training have been shown to facilitate not only knowledge transfer but also robustness to catastrophic forgetting for downstream continual learning [31, 28, 57]. Also, learning a large number of base classes in the initial training phase allows class-incremental learning with small adaptations [48]. Inspired by the techniques of using pre-trained knowledge in NLP, L2P [46] employed an additional set of learnable parameters called “prompts” that dynamically instruct a pre-trained representation layer to learn incremental tasks. DualPrompt [45] extended this idea by attaching complementary prompts to the pre-trained representation layer for learning task-invariant and task-specific instructions. Such prompt-based approaches were reported to be far superior

to traditional continual learning baselines, which potentially challenges the current paradigm of using pre-trained knowledge in CV.

Since the large amount of training samples required to construct a pre-trained model are typically unlabeled and may also arrive incrementally, self-supervised pre-training emerges as a more preferable choice than supervised pre-training. Several recent works discovered that continual learning in a self-supervised manner suffers from less catastrophic forgetting [17, 26, 9]. Indeed, self-supervised paradigms have been shown to be better adapted to upstream continual learning [5]. However, the effectiveness of self-supervised pre-training for downstream continual learning remains to be investigated.

3. Continual Learning on a Pre-trained Model

In this section, we introduce the problem formulation of continual learning on a pre-trained model (CLPM), perform an extensive analysis, and then present our approach.

3.1. Problem Formulation

Let’s consider a neural network $M_\theta(\cdot) = h_{\theta_{cls}}(f_{\theta_{rps}}(\cdot))$ with parameters $\theta = \{\theta_{rps}, \theta_{cls}\}$ for classification tasks, consisting of a representation layer $f_{\theta_{rps}}(\cdot)$ that projects input images to feature representations, and a classification layer $h_{\theta_{cls}}(\cdot)$ that projects feature representations to output predictions. θ_{rps} is initialized on a pre-training dataset D_{pt} in a supervised or self-supervised manner (labels are not necessary for the latter). Then, M_θ needs to learn a sequence of incremental tasks from their training sets $D_t, t = 1, \dots, T$ and tries to perform well on their test sets. Following previous works of CLPM in CV [46, 45], we mainly focus on the class-incremental scenario of continual learning [37]. In details, $D_t = \bigcup_{c \in C_t} \{(x_{c,n}, y_{c,n})\}_{n=1}^{N_c}$ introduces a set of new classes C_t , where N_c denotes the number of training samples $(x_{c,n}, y_{c,n})$ for class c , and all the classes ever seen are evaluated without task labels.

To achieve this aim, the network needs to (1) effectively transfer the pre-trained knowledge to each incremental task while maintaining its generalizability for future tasks, and (2) properly balance learning plasticity of new tasks with memory stability of old tasks. A naive baseline is to sequentially fine-tune the entire model M_θ on each D_t , where $f_{\theta_{rps}}$ and $h_{\theta_{cls}}$ are updated in a similar speed (*i.e.*, using the same learning rate). However, due to the lack of D_{pt} and $D_{1:t-1} = \bigcup_{i=1}^{t-1} D_i$, the performance of M_θ is severely limited by a *progressive overfitting* problem in both aspects. Specifically, (1) the knowledge of D_{pt} is interfered by D_t , as θ_{rps} is continually updated to accommodate incremental tasks while the generalizability obtained from the pre-training stage is progressively lost; and (2) the knowledge of $D_{1:t-1}$ is interfered by D_t , as θ_{cls} (and θ_{rps}) catastrophically forgets the old tasks when learning each new task.

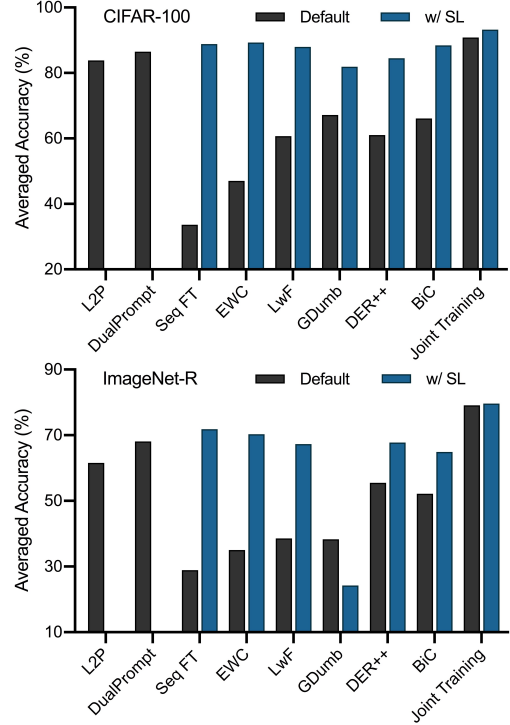


Figure 2. Slow Learner (SL) can greatly enhance continual learning performance on a pre-trained model. Here we adopt ImageNet-21K supervised pre-training for all baselines with default performance referenced from [46, 45], including prompt-based approaches (L2P [46] and DualPrompt [45]), regularization-based approaches (EWC [19] and LwF [25]), and replay-based approaches (GDumb [30], DER++ [3] and BiC [49]).

Most continual learning approaches [19, 25, 30, 49, 3] are established on the premise of learning from scratch and focus on improving the second aspect. We call them “traditional baselines”. To address the first aspect, the newly proposed prompt-based approaches [46, 45] fixed the representation layer and employed an additional set of learnable parameters to instruct the pre-trained model, which reported significantly better performance for CLPM than other traditional baselines that update the representation layer.

3.2. Slow Learner is (Almost) All You Need?

Implementation of Slow Learner: However, we argue that the performance of traditional baselines reported in [46, 45] is severely limited by using a uniform learning rate (*e.g.*, 0.005) for both θ_{rps} and θ_{cls} . Through extensive experiments, we observe that using a much smaller learning rate (0.0001) for θ_{rps} and a slightly larger learning rate (0.01) for θ_{cls} can greatly enhance the performance of traditional baselines (Fig. 2).¹ The reported performance of sequential fine-tuning (Seq FT) is improved by more than

¹A comprehensive analysis of the effect of learning rate can be found in the appendix.




Figure 3. Comparison of pre-training paradigms on ImageNet-1K. DeiT [36] is a strong supervised method for (pre-)training vision transformer, while MoCo v3 [4], MAE [13] and BEiT [2] are representative self-supervised methods. The pre-trained checkpoints are obtained from their official release.

40% for challenging continual learning benchmarks such as Split CIFAR-100 and Split ImageNet-R, respectively, thus clearly outperforming the prompt-based approaches such as L2P [46] and DualPrompt [45].² We call this simple but surprisingly effective strategy “Slow Learner (SL)”, corresponding to slowing down the updating speed of the representation layer. In contrast to previous works of using different learning rates for transfer learning [10, 53, 14], the benefit of our proposal is *specific to continual learning*, as the upper bound performance (*i.e.*, joint training) is only marginally improved by the SL while the performance gap between continual learning and joint training is greatly filled (*e.g.*, only 4.36% on Split CIFAR-100 and 7.80% on Split ImageNet-R for Seq FT w/ SL).

Effect of Pre-training Paradigm: We then evaluate the effect of pre-training paradigms (*i.e.*, supervised or self-supervised) on downstream continual learning. Here we focus on ImageNet-1K as the pre-training dataset, since most self-supervised methods only release checkpoints on it. Considering architectural consistency with previous works of CLPM [46, 45], we select representative self-supervised methods (*i.e.*, MoCo v3 [4], MAE [13] and BEiT [2]) that release checkpoints on ViT-B/16 in our comparisons. We further compare DeiT [36], a strong supervised method for

²Please note that the prompt-based approaches [46, 45] use a fixed representation layer, so our proposal is not applicable for them.




Figure 4. Similarity of the pre-trained representations (1) before and after joint training (left Y-axis, yellow dot), and (2) after joint training and after continual learning (right Y-axis, column). We adopt Centered Kernel Alignment (CKA) [20] as the similarity metric. *Best viewed in color.*

(pre-)training vision transformer. As shown in Fig. 3, self-supervised pre-training, while more realistic regarding labeling requirements and upstream continual learning, typically results in a larger performance gap between Seq FT and joint training than supervised pre-training. Similarly, the use of SL can effectively reduce this gap. Interestingly, the performance of Seq FT w/ SL for MoCo v3 [4] far exceeds that of the more recent MAE [13], although their joint training performance is comparable. This is possibly because the pre-trained representations of MoCo v3 [4] require remarkably smaller updates to learn all tasks well (Fig. 4, left Y-axis), thus alleviating the progressive overfitting problem. Meanwhile, the use of SL allows MoCo v3 [4] to learn representations much closer to that of the joint training (Fig. 4, right Y-axis). The above results suggest a new direction for the design of self-supervised paradigms, *i.e.*, how to effectively perform downstream continual learning and combine the advantages of SL.

Evaluation of Representation: So, why is the Slow Learner such effective, and what accounts for the remaining performance gap? We perform a linear probing experiment [13] to evaluate the performance of the representation layer. Specifically, after learning each incremental task (*e.g.*, 10 classes per task for 10 tasks in Split CIFAR-100) via Seq FT w/ SL, we fix the representation layer and employ an extra classification layer, called a linear probe, to learn all classes of the corresponding benchmark dataset (*e.g.*, a total of 100 classes in CIFAR-100 dataset). The performance of these linear probes is presented in Fig. 5, which tends to grow with learning more tasks, indicating that the representation layer is accumulating knowledge for better adaptation. After learning all incremental tasks, it can be clearly seen that using the continually-learned representation layer to jointly train an extra classifier for all classes can almost reach the




Figure 5. Linear probing results of Slow Learner. All experiments are based on ImageNet-21K supervised pre-training. We report the averaged accuracy of all classes in the corresponding benchmark dataset (e.g., a total of 100 classes in CIFAR-100 dataset). The dark red arrow represents the performance gap caused by a sub-optimal classification layer.

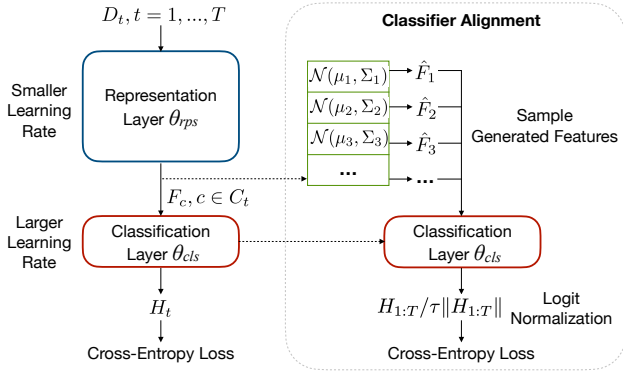


Figure 6. Slow Learner with Classifier Alignment (SLCA). H_t is the logit of predicting the current training data in D_t .

joint training performance of the entire model, and far outperform its counterpart with a continually-learned classifier (i.e., Seq FT w/ SL in Fig. 5). Therefore, the proposed SL can almost address the problem of the representation layer, yet the classification layer remains sub-optimal. In particular, the problem of classification layer becomes more severe for fine-grained continual learning benchmarks such as Split CUB-200 and Split Cars-196.

3.3. Slow Learner with Classifier Alignment

To further improve the classification layer, we propose to align the statistics of previously-learned classes in a *post-hoc* fashion (see Fig. 6 and Algorithm 1). Specifically, when learning each incremental task, only the parameters corresponding to the classes of the current task are trained together with the representation layer. After learning each task, we collect feature representations $F_c = [r_{c,1}, \dots, r_{c,N_c}]$ for each class $c \in C_t$ of the current task, where $r_{c,n} = f_{\theta_{rps}}(x_{c,n})$ and N_c denotes its amount. Instead of saving the extracted features F_c of training samples, CA preserves their mean $\mu_c \in \mathbb{R}^d$ and covariance $\Sigma_c \in \mathbb{R}^{d \times d}$ for each class c (d denotes the feature dimension). Since the use of pre-training provides well-distributed representations, each class tends to be single-peaked and can

be naturally modelled as a Gaussian $\mathcal{N}(\mu_c, \Sigma_c)$.

Whenever the model needs to be evaluated, the classification layers are further aligned as follows. We first sample generated features $\hat{F}_c = [\hat{r}_{c,1}, \dots, \hat{r}_{c,S_c}]$ from the distribution $\mathcal{N}(\mu_c, \Sigma_c)$ of each class $c \in C_{1:T}$ ever seen in $C_{1:T} = \bigcup_{i=1}^T C_i$, where S_c is the amount of generated features for each class ($S_c = 256$ in our experiments), and the task amount T can be any positive integer without being known in advance. Next, we apply a widely-used cross-entropy loss to adjust the classification layer $h_{\theta_{cls}}$ by feeding $\hat{F}_{1:T} = [\hat{F}_1, \dots, \hat{F}_{|C_{1:T}|}]$ as the input of the classification layer, where $|C_{1:T}|$ denotes the number of classes in $C_{1:T}$.

Algorithm 1 Slow Learner with Classifier Alignment (SLCA)

Inputs: Pre-training dataset D_{pt} ; training dataset D_t for task $t = 1, \dots, T$; network $M_\theta(\cdot) = h_{\theta_{cls}}(f_{\theta_{rps}}(\cdot))$ with parameters $\theta = \{\theta_{rps}, \theta_{cls}\}$; learning rates α for θ_{rps} and β for θ_{cls} ($\alpha < \beta$); temperature hyperparameter τ .

Initialization: Initialize θ_{rps} by pre-training on D_{pt} ; initialize θ_{cls} randomly.

- 1: # sequential tasks.
 - 2: **for** task $t = 1, \dots, T$ **do**
 - 3: # different learning rates for θ_{rps} and θ_{cls} .
 - 4: **while** not converged **do**
 - 5: Train M_θ with cross-entropy loss on D_t .
 - 6: **end while**
 - 7: Collect $F_c = [r_{c,1}, \dots, r_{c,N_c}]$ for $c \in C_t$.
 - 8: Save mean μ_c and covariance Σ_c of F_c for $c \in C_t$.
 - 9: **end for**
 - 10: # classifier alignment.
 - 11: Sample \hat{F}_c from $\mathcal{N}(\mu_c, \Sigma_c)$ for $c \in C_{1:T}$.
 - 12: **while** not converged **do**
 - 13: Compute logit $H_{1:T}$ and its magnitude $\|H_{1:T}\|$.
 - 14: Train $h_{\theta_{cls}}$ with normalized logit in Eqn. 1.
 - 15: **end while**
-

However, a prolonged training of the classification layer can lead to an overconfidence issue, which potentially impairs generalizability to the test set(s). To overcome this, we

Method	Memory-Free	Pre-trained	Split CIFAR-100		Split ImageNet-R	
			Last-Acc (%)	Inc-Acc (%)	Last-Acc (%)	Inc-Acc (%)
Joint-Training	-	IN21K-Sup	93.22±0.16	-	79.60±0.87	-
GDumb [30]		IN21K-Sup	81.92±0.15	89.46±0.94	24.23±0.35	43.48±0.49
DER++ [3]		IN21K-Sup	84.50±1.67	91.49±0.61	67.75±0.93	78.13±1.14
BiC [49]		IN21K-Sup	88.45±0.57	93.37±0.32	64.89±0.80	73.66±1.61
L2P [46]	✓	IN21K-Sup	82.76±1.17	88.48±0.83	66.49±0.40	72.83±0.56
DualPrompt [45]	✓	IN21K-Sup	85.56±0.33	90.33±0.33	68.50±0.52	72.59±0.24
EWC [19]	✓	IN21K-Sup	89.30±0.23	92.31±1.66	70.27±1.99	76.27±2.13
LwF [25]	✓	IN21K-Sup	87.99±0.05	92.13±1.16	67.29±1.67	74.47±1.48
Seq FT	✓	IN21K-Sup	88.86±0.83	92.01±1.71	71.80±1.45	76.84±1.26
SLCA (Ours)	✓	IN21K-Sup	91.53±0.28	94.09±0.87	77.00±0.33	81.17±0.64
Joint-Training	-	IN1K-Self	89.11±0.06	-	72.80±0.23	-
GDumb [30]		IN1K-Self	69.72±0.20	80.95±1.19	28.24±0.58	43.64±1.05
DER++ [3]		IN1K-Self	63.64±1.30	79.55±0.87	53.11±0.44	65.10±0.91
BiC [49]		IN1K-Self	80.57±0.86	89.39±0.33	57.36±2.68	68.07±0.22
EWC [19]	✓	IN1K-Self	81.62±0.34	87.56±0.97	64.50±0.36	70.37±0.41
LwF [25]	✓	IN1K-Self	77.94±1.00	86.90±0.90	60.74±0.30	68.55±0.65
Seq FT	✓	IN1K-Self	81.47±0.55	87.55±0.95	64.43±0.44	70.48±0.54
SLCA (Ours)	✓	IN1K-Self	85.27±0.08	89.51±1.04	68.07±0.21	73.04±0.56

Table 1. Experimental results for continual learning on Split CIFAR-100 and Split ImageNet-R. IN21K-Sup: supervised pre-training on ImageNet-21K. IN1K-Self: self-supervised pre-training on ImageNet-1K with MoCo v3 [4]. All other fine-tuning based methods are reproduced according to their officially-released codes with the proposed **Slow Learner** implemented.

draw inspirations from out-of-distribution (OOD) detection [47] and normalize the magnitude of network outputs when computing the cross-entropy. Let $H_{1:T} = h_{\theta_{cls}}(\hat{F}_{1:T}) = [h_{\theta_{cls}}(\hat{F}_1), \dots, h_{\theta_{cls}}(\hat{F}_{[C_{1:T}]})] := [l_1, \dots, l_{[C_{1:T}]}]$ denote the logit (*i.e.*, pre-softmax output) of $\hat{F}_{1:T}$, which can be re-written as the product of two components: $H_{1:T} = \|H_{1:T}\| \cdot \vec{H}_{1:T}$, where $\|\cdot\|$ denotes L_2 -norm. $\|H_{1:T}\| = \sqrt{\sum_{c \in C_{1:T}} \|l_c\|^2}$ represents the magnitude of $H_{1:T}$, and $\vec{H}_{1:T}$ represents its direction. Then we adopt a modified cross-entropy loss with logit normalization to perform classifier alignment:

$$\mathcal{L}(\theta_{cls}; \hat{F}_{1:T}) = -\log \frac{e^{l_y / (\tau \|H_{1:T}\|)}}{\sum_{c \in C_{1:T}} e^{l_c / (\tau \|H_{1:T}\|)}}, \quad (1)$$

where l_y denotes the y -th element of $H_{1:T}$ corresponding to the ground-truth label y . τ is a temperature hyperparameter. The intuition is that normalizing $H_{1:T}$ with an input-dependent constant $\tau \|H_{1:T}\|$ will not change the result of softmax prediction $\arg \max_{c \in C_{1:T}} (l_c)$, while forcing the magnitude $\|H_{1:T}\|$ before softmax becomes $\frac{1}{\tau}$ can make the criterion only adjust the direction $\vec{H}_{1:T}$ [47]. Therefore, the normalization in Eqn. 1 can alleviate the overconfidence issue in classifier alignment. In practice, we observe that the temperature hyperparameter is not sensitive and empirically find $\tau = 0.1$ to be a reasonable choice.

4. Experiments

In this section, we first briefly describe the experimental setups, and then present the experimental results.

4.1. Experimental Setups

Benchmark: Following L2P [46] and DualPrompt [45], we adopt pre-training from ImageNet-21K dataset [32], also known as the full ImageNet [6] consisting of 14,197,122 images with 21,841 classes. We also consider pre-training from ImageNet-1K dataset [23], a subset of ImageNet-21K introduced for the ILSVRC2012 visual recognition challenge, consisting of 1000-class images.

To evaluate the performance of downstream continual learning, we consider four representative benchmark datasets and randomly split each of them into 10 disjoint tasks: The first two follow previous works [46, 45] and are relatively coarse-grained in terms of classification, while the last two are relatively fine-grained. Specifically, CIFAR-100 dataset [22] consists of 100-class natural images with 500 training samples per class. ImageNet-R dataset [15] contains 200-class images, splitting into 24,000 and 6,000 images for training and testing (similar ratio for each class), respectively. Note that although the image categories of ImageNet-R are overlapped with ImageNet-21K, all images are out-of-distribution samples for the pre-train dataset, *i.e.*, hard examples from ImageNet or newly collected data of different styles. It requires considerable adaptations of the pre-trained model, therefore serving as a challenging benchmark for continual learning. CUB-200 dataset [38] includes 200-class bird images with around 60 images per class, 30 of which are used for training and the rest for testing. Cars-196 dataset [21] includes 196 types of car images, split into 8,144 and 8,040 images for training and testing (similar ratio for each class), respectively. We present the average accuracy of all classes after learning the last task, denoted as Last-Acc (equivalent to ‘‘Avg. Acc’’ in [46, 45]). We also

Method	Memory-Free	Pre-trained	Split CUB-200		Split Cars-196	
			Last-Acc (%)	Inc-Acc (%)	Last-Acc (%)	Inc-Acc (%)
Joint-Training	-	IN21K-Sup	88.00±0.34	-	80.31±0.13	-
GDumb [30]		IN21K-Sup	61.80±0.77	79.76±0.18	25.20±0.84	49.48±0.74
DER++ [3]		IN21K-Sup	77.42±0.71	87.61±0.09	60.41±1.76	75.04±0.57
BiC [49]		IN21K-Sup	81.91±2.59	89.29±1.57	63.10±5.71	73.75±2.37
L2P [46]	✓	IN21K-Sup	62.21±1.92	73.83±1.67	38.18±2.33	51.79±4.19
DualPrompt [45]	✓	IN21K-Sup	66.00±0.57	77.92±0.50	40.14±2.36	56.74±1.78
EWC [19]	✓	IN21K-Sup	68.32±2.64	79.95±2.28	52.50±3.18	64.01±3.25
LwF [25]	✓	IN21K-Sup	69.75±1.37	80.45±2.08	49.94±3.24	63.28±1.11
Seq FT	✓	IN21K-Sup	68.07±1.09	79.04±1.69	49.74±1.25	62.83±2.16
SLCA (Ours)	✓	IN21K-Sup	84.71±0.40	90.94±0.68	67.73±0.85	76.93±1.21
Joint-Training	-	IN1K-Self	79.55±0.04	-	74.52±0.09	-
GDumb [30]		IN1K-Self	45.29±0.97	66.86±0.63	20.95±0.42	45.40±0.66
DER++ [3]		IN1K-Self	61.47±0.32	77.15±0.61	50.64±0.70	67.64±0.45
BiC [49]		IN1K-Self	74.39±1.12	82.13±0.33	65.57±0.93	73.95±0.29
EWC [19]	✓	IN1K-Self	61.36±1.43	72.84±2.18	53.16±1.45	63.61±1.06
LwF [25]	✓	IN1K-Self	61.66±1.95	73.90±1.91	52.45±0.48	63.87±0.31
Seq FT	✓	IN1K-Self	61.67±1.37	73.25±1.83	52.91±1.61	63.32±1.31
SLCA (Ours)	✓	IN1K-Self	73.01±0.16	82.13±0.34	66.04±0.08	72.59±0.04

Table 2. Experimental results for continual learning on Split CUB-200 and Split Cars-196. IN21K-Sup: supervised pre-training on ImageNet-21K. IN1K-Self: self-supervised pre-training on ImageNet-1K with MoCo v3 [4]. The results of all fine-tuning based baselines are reproduced according to their officially-released codes with the proposed **Slow Learner** implemented.

compute the average accuracy of the classes ever seen after learning each incremental task and then present their average, denoted as Inc-Acc.

Implementation: Following previous works [46, 45], we adopt a pre-trained ViT-B/16 backbone for all baselines. In addition to supervised pre-training, we consider representative self-supervised paradigms that provide pre-trained checkpoints on ViT-B/16, *i.e.*, MoCo v3 [4], BEiT [2] and MAE [13]. For continual learning of downstream tasks, we follow the previous implementation that employs an Adam optimizer for L2P [46] and DualPrompt [45] while a SGD optimizer for other baselines, with the same batch size of 128. Our Slow Learner adopts a learning rate of 0.0001 for the representation layer and 0.01 for the classification layer, different from [46, 45] using 0.005 for the entire model.

Baseline: We adopt joint training as the upper bound performance and consider continual learning baselines with or without replaying old training samples. As for the former, a memory buffer of 1000 images is maintained, and we evaluate three representative replay-based approaches such as BiC [49], GDumb [30] and DER++ [3]. As for the latter, we evaluate representative regularization-based approaches such as EWC [19] and LwF [25], and prompt-based approaches such as L2P [46] and DualPrompt [45]. Note that sequential fine-tuning usually serves as the lower bound performance of continual learning, but we observe that simply adjusting the learning rate (*i.e.*, using the Slow Learner) makes it a surprisingly strong baseline.

4.2. Experimental Results

Overall Performance: All baseline approaches in Table 1, 2 are trained with our **Slow Learner** for a fair

comparison, except the prompt-based approaches that keep the representation layer fixed. For continual learning of relatively coarse-grained classification tasks, such as Split CIFAR-100 and Split ImageNet-R in Table 1 (also shown in Fig. 2), the SL can substantially enhance the performance of continual learning. With the help of Classifier Alignment (CA) and its Logit Normalization (LN), our approach clearly outperforms L2P [46] and DualPrompt [45], and almost reach the joint training upper bound (the performance gap is less than **2%** for supervised pre-training and **4%** for self-supervised pre-training). As for fine-grained continual learning benchmarks in Table 2, SLCA achieves the strongest performance for supervised pre-training, while for self-supervised pre-training the performance is comparable or slightly inferior to the SL version of BiC [49], which additionally employs some old training samples. Please note that the presented implementation of SLCA can be seen as adding CA+LN together with SL to the simplest sequential fine-tuning (Seq FT) baseline. In fact, the proposed SLCA can be naturally plug-and-play with other continual learning approaches. We leave it as a further work.

It is worth noting that different replay-based approaches (w/ SL) behave differently in CLPM. In general, BiC [49], which updated the entire model with the old training samples in a similar way to learning the new ones and added a bias correction layer to mitigate imbalance between old and new classes, receives the most substantial improvements. While GDumb [30], which simply used the old training samples to train a new model from scratch at test time, has difficulty to adapt the pre-trained model with limited training samples and thus performs the worst.

Ablation Study: We present an extensive ablation study

	Method	Pre-trained	Split CIFAR-100	Split ImageNet-R	Split CUB-200	Split Cars-196
Baseline	Seq FT [†]	IN21K-Sup	41.77 \pm 13.8	26.95 \pm 11.8	40.02 \pm 1.08	27.57 \pm 1.79
	w/ Fixed θ_{rps}	IN21K-Sup	63.75 \pm 0.67	34.64 \pm 14.3	60.44 \pm 1.80	24.51 \pm 6.90
Ours	w/ SL	IN21K-Sup	88.86 \pm 0.83	71.80 \pm 1.45	68.07 \pm 1.09	49.74 \pm 1.25
	w/ Fixed θ_{rps} +CA	IN21K-Sup	75.64 \pm 0.26	50.73 \pm 0.21	82.71 \pm 0.14	54.45 \pm 0.16
	w/ Fixed θ_{rps} +CA+LN	IN21K-Sup	75.62 \pm 0.21	51.83 \pm 0.34	83.65 \pm 0.18	53.43 \pm 0.09
	w/ SL+CA	IN21K-Sup	90.70 \pm 0.52	74.41 \pm 0.51	83.20 \pm 0.19	67.90 \pm 0.53
	w/ SL+CA+LN	IN21K-Sup	91.53 \pm 0.28	77.00 \pm 0.33	84.71 \pm 0.40	67.73 \pm 0.85
Baseline	Seq FT	IN1K-Self	27.99 \pm 5.16	45.84 \pm 4.19	45.35 \pm 1.38	35.96 \pm 2.04
	w/ Fixed θ_{rps}	IN1K-Self	77.30 \pm 0.56	51.97 \pm 0.17	55.54 \pm 1.55	43.16 \pm 0.12
Ours	w/ SL	IN1K-Self	81.47 \pm 0.55	64.43 \pm 0.44	61.67 \pm 1.37	52.91 \pm 1.61
	w/ Fixed θ_{rps} +CA	IN1K-Self	81.83 \pm 0.12	55.59 \pm 0.21	70.67 \pm 0.02	57.01 \pm 0.07
	w/ Fixed θ_{rps} +CA+LN	IN1K-Self	81.95 \pm 0.17	56.47 \pm 0.23	72.97 \pm 0.17	63.00 \pm 0.21
	w/ SL+CA	IN1K-Self	84.64 \pm 0.21	67.54 \pm 0.29	72.52 \pm 0.06	64.80 \pm 0.20
	w/ SL+CA+LN	IN1K-Self	85.27 \pm 0.08	68.07 \pm 0.21	73.01 \pm 0.16	66.04 \pm 0.08

Table 3. Ablation study. Here we present the Last-Acc (%) after continual learning of all classes. [†]The reproduced performance of Seq FT is slightly different from the reported one in [46, 45] due to the use of different random seeds. SL: Slow Learner; LN: Logit Normalization; CA: a naive implementation of Classifier Alignment without LN.

of our approach in Table 4.2. To demonstrate the necessity of the proposed Slow Learner (SL), we consider two baselines: (1) sequential fine-tuning (Seq FT), which adopts a uniform learning rate of 0.005; and (2) Seq FT with fixed θ_{rps} , which keeps the representation layer fixed and continually adjusts the classification layer. Seq FT with fixed θ_{rps} is generally superior to Seq FT while significantly inferior to Seq FT w/ SL, indicating the *necessity* of updating the representation layer (but with a properly reduced learning rate to mitigate the progressive overfitting problem).

We further validate the effectiveness of the proposed Classifier Alignment (CA) and its Logit Normalization (LN). In particular, as the fine-grained continual learning benchmarks severely exacerbate the problem of classification layer (Fig. 5), the benefits of SL are still significant but slightly reduced (*e.g.*, for ImageNet-21K supervised pre-training, the improvements of SL are **47.09%**, **44.85%**, **28.05%** and **22.17%** on Split CIFAR-100, Split ImageNet-R, Split CUB-200 and Split Cars-196, respectively), while the benefits of CA+LN are greatly enhanced (*e.g.*, the improvements of CA+LN are **2.67%**, **5.20%**, **19.64%** and **17.99%** on Split CIFAR-100, Split ImageNet-R, Split CUB-200 and Split Cars-196, respectively for IN21K-Sup pre-train). We also evaluate CA with fixed representations θ_{rps} , which gain consistent improvements as well. Please note that our CA or CA+LN is operated in a *post-hoc* fashion rather than aligning the classifier during representation learning [56, 11, 12], as the latter would even *worsen* the performance (*e.g.*, by **27.89%** on Split CIFAR-100).

Pre-training Paradigm and Downstream Granularity: Here we analyze two critical factors for CLPM identified in our results. Compared to supervised pre-training, self-supervised pre-training usually results in larger performance gaps between continual learning baselines and joint training (Fig. 3, Table 1, 2). Although our proposal can greatly improve the performance of continual learning, the

effectiveness varies with the choice of self-supervised methods (Fig. 3). This is because they differ in the magnitude of updates required to learn all tasks well (Fig. 4), resulting in different degrees of difficulty in refining the representation layer. Given that the large amount of data required for pre-training is typically unlabeled and incrementally collected, we suggest subsequent works to develop self-supervised pre-training paradigms that are more suitable for downstream continual learning, and potentially use this as a criterion to evaluate the progress of self-supervised learning. On the other hand, the performance gap is increased as downstream continual learning becomes more fine-grained (Table 1, 2, 4.2). This is primarily due to the sub-optimal classification layer (Fig. 5), which can be greatly improved by our classifier alignment strategy (Table 4.2).

Scalability: We further discuss the scalability of our method. First, the generated features are only used to align the output layer at test time rather than training the entire backbone, thus the *computation* is efficient and not accumulated in continual learning (*e.g.*, ranging from only 0.67% to 5% of the total running time for various benchmarks). Second, the covariance Σ_c can be further simplified as the variance $\sigma_c^2 \in \mathbb{R}^d$ with tolerable performance degradation (*e.g.*, up to only 0.61% on Split CIFAR-100 and 0.83% on Split ImageNet-R). In this way, the *storage* of μ_c and σ_c^2 for 100 classes corresponds to only 0.18% parameters of the ViT-B/16 backbone, which is clearly lightweight.

5. Conclusion

Continual learning on a pre-trained model (CLPM) requires effective transfer of pre-trained knowledge to each incremental task while maintaining its generalizability for future tasks. However, using a uniform learning rate to update the entire model makes traditional continual learning baselines fail to accommodate both objectives. Observing that using a slow learner can almost address this challeng-

ing issue in the representation layer, we further improve the classification layer through a classifier alignment strategy. Across a variety of upstream and downstream scenarios, our proposal enables the simplest sequential fine-tuning baseline to almost reach the joint training upper bound, far beyond the current state-of-the-art performance. Such a simple but extremely effective approach provides a strong criterion to *re-evaluate the current progress and technical route* for CLPM in CV. In addition, our empirical study demonstrates two critical factors for CLPM, such as pre-training paradigm and downstream granularity. Subsequent works could further explore these directions and develop more powerful strategies based on our proposal, so as to better leverage pre-trained knowledge for continual learning.

Discussion of Limitations. Although we argue that self-supervised pre-training enjoys an advantage of less forgetting in upstream continual learning and propose an effective approach to overcome its shortcomings in downstream continual learning, how to perform upstream and downstream continual learning together remains to be explored. Besides, our results are based on a pre-trained ViT backbone for a fair comparison with previous works [46, 45], without considering other representative architectures such as ResNet and other CL applications on downstream computer vision tasks such as object detection [35], image segmentation [54, 55, 52]. We leave the exploration of our work in the mentioned aspects to future works.

References

- [1] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, et al. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European Conference on Computer Vision*, pages 139–154, 2018.
- [2] Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*.
- [3] Pietro Buzzega, Matteo Boschini, Angelo Porrello, et al. Dark experience for general continual learning: a strong, simple baseline. In *Advances in Neural Information Processing Systems*, volume 33, pages 15920–15930, 2020.
- [4] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9640–9649, 2021.
- [5] Andrea Cossu, Tinne Tuytelaars, Antonio Carta, Lucia Pasaro, Vincenzo Lomonaco, and Davide Bacciu. Continual pre-training mitigates forgetting in language and vision. *arXiv preprint arXiv:2205.09357*, 2022.
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009.
- [7] Prithviraj Dhar, Rajat Vikram Singh, Kuan-Chuan Peng, et al. Learning without memorizing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5138–5146, 2019.
- [8] Avia Efrat and Omer Levy. The turking test: Can language models understand instructions? *arXiv preprint arXiv:2010.11982*, 2020.
- [9] Enrico Fini, Victor G Turrissi da Costa, Xavier Alameda-Pineda, et al. Self-supervised models are continual learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9621–9630, 2022.
- [10] Yunhui Guo, Honghui Shi, Abhishek Kumar, et al. Spottune: transfer learning through adaptive fine-tuning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4805–4814, 2019.
- [11] Tyler L Hayes, Kushal Kafle, Robik Shrestha, et al. Remind your neural network to prevent catastrophic forgetting. In *Proceedings of the European Conference on Computer Vision*, pages 466–483, 2020.
- [12] Tyler L Hayes and Christopher Kanan. Lifelong machine learning with deep streaming linear discriminant analysis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 220–221, 2020.
- [13] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, et al. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022.
- [14] Kaiming He, Ross Girshick, and Piotr Dollár. Rethinking imagenet pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4918–4927, 2019.
- [15] Dan Hendrycks, Steven Basart, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8340–8349, 2021.
- [16] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*, pages 2790–2799. PMLR, 2019.
- [17] Dapeng Hu, Qizhengqiu Lu, Lanqing Hong, Hailin Hu, Yifan Zhang, Zhenguo Li, Alfred Shen, and Jiashi Feng. How well self-supervised pre-training performs with streaming data? *arXiv preprint arXiv:2104.12081*, 2021.
- [18] Zixuan Ke and Bing Liu. Continual learning of natural language processing tasks: A survey. *arXiv preprint arXiv:2211.12701*, 2022.
- [19] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13):3521–3526, 2017.
- [20] Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. In *Proceedings of International Conference on Machine Learning*, pages 3519–3529. PMLR, 2019.
- [21] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 554–561, 2013.
- [22] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, 2009.

- [23] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, volume 25, pages 1097–1105, 2012.
- [24] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, 2021.
- [25] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12):2935–2947, 2017.
- [26] 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.
- [27] James L McClelland, Bruce L McNaughton, and Randall C O’Reilly. Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102(3):419, 1995.
- [28] Sanket Vaibhav Mehta, Darshan Patil, Sarath Chandar, and Emma Strubell. An empirical investigation of the role of pre-training in lifelong learning. *arXiv preprint arXiv:2112.09153*, 2021.
- [29] Xingchao Peng et al. Moment matching for multi-source domain adaptation. In *ICCV*, 2019.
- [30] Ameya Prabhu, Philip HS Torr, and Puneet K Dokania. Gdumb: A simple approach that questions our progress in continual learning. In *Proceedings of European Conference on Computer Vision*, pages 524–540, 2020.
- [31] Vinay Venkatesh Ramasesh, Aitor Lewkowycz, and Ethan Dyer. Effect of scale on catastrophic forgetting in neural networks. In *Proceedings of the International Conference on Learning Representations*, 2021.
- [32] Tal Ridnik, Emanuel Ben-Baruch, Asaf Noy, and Lihi Zelnik-Manor. Imagenet-21k pretraining for the masses. *arXiv preprint arXiv:2104.10972*, 2021.
- [33] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016.
- [34] Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. In *Proceedings of International Conference on Machine Learning*, pages 4548–4557, 2018.
- [35] Konstantin Shmelkov, Cordelia Schmid, and Karteek Alahari. Incremental learning of object detectors without catastrophic forgetting. In *Proceedings of the IEEE international conference on computer vision*, pages 3400–3409, 2017.
- [36] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *Proceedings of International Conference on Machine Learning*, pages 10347–10357. PMLR, 2021.
- [37] Gido M van de Ven and Andreas S Tolias. Three scenarios for continual learning. *arXiv preprint arXiv:1904.07734*.
- [38] Catherine Wah, Steve Branson, Peter Welinder, et al. The caltech-ucsd birds-200-2011 dataset. 2011.
- [39] Liyuan Wang, Bo Lei, Qian Li, Hang Su, Jun Zhu, and Yi Zhong. Triple-memory networks: A brain-inspired method for continual learning. *IEEE Transactions on Neural Networks and Learning Systems*, 2021.
- [40] Liyuan Wang, Kuo Yang, Chongxuan Li, Lanqing Hong, Zhenguo Li, and Jun Zhu. Ordisco: Effective and efficient usage of incremental unlabeled data for semi-supervised continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5383–5392, 2021.
- [41] Liyuan Wang, Mingtian Zhang, Zhongfan Jia, Qian Li, Chenglong Bao, Kaisheng Ma, Jun Zhu, and Yi Zhong. Afec: Active forgetting of negative transfer in continual learning. In *Advances in Neural Information Processing Systems*, volume 34, 2021.
- [42] Liyuan Wang, Xingxing Zhang, Qian Li, Jun Zhu, and Yi Zhong. Coscl: Cooperation of small continual learners is stronger than a big one. In *European Conference on Computer Vision*, pages 254–271. Springer, 2022.
- [43] Liyuan Wang, Xingxing Zhang, Hang Su, and Jun Zhu. A comprehensive survey of continual learning: Theory, method and application. *arXiv preprint arXiv:2302.00487*, 2023.
- [44] Liyuan Wang, Xingxing Zhang, Kuo Yang, Longhui Yu, Chongxuan Li, Lanqing Hong, Shifeng Zhang, Zhenguo Li, Yi Zhong, and Jun Zhu. Memory replay with data compression for continual learning. In *Proceedings of the International Conference on Learning Representations*, 2021.
- [45] Zifeng Wang, Zizhao Zhang, Sayna Ebrahimi, Ruoxi Sun, Han Zhang, Chen-Yu Lee, Xiaoqi Ren, et al. Dual-prompt: Complementary prompting for rehearsal-free continual learning. *arXiv preprint arXiv:2204.04799*, 2022.
- [46] Zifeng Wang, Zizhao Zhang, Chen-Yu Lee, et al. Learning to prompt for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 139–149, 2022.
- [47] Hongxin Wei, Renchunzi Xie, Hao Cheng, Lei Feng, Bo An, and Yixuan Li. Mitigating neural network overconfidence with logit normalization. *arXiv preprint arXiv:2205.09310*, 2022.
- [48] Tz-Ying Wu, Gurumurthy Swaminathan, Zhizhong Li, et al. Class-incremental learning with strong pre-trained models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9601–9610, 2022.
- [49] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, et al. Large scale incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 374–382, 2019.
- [50] Binbin Yang, Xincheng Deng, Han Shi, Changlin Li, Gengwei Zhang, Hang Xu, Shen Zhao, Liang Lin, and Xiaodan Liang. Continual object detection via prototypical task correlation guided gating mechanism. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9255–9264, 2022.
- [51] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *Proceedings of the International Conference on Machine Learning*, pages 3987–3995, 2017.

- [52] Gengwei Zhang, Guoliang Kang, Yi Yang, and Yunchao Wei. Few-shot segmentation via cycle-consistent transformer. *Advances in Neural Information Processing Systems*, 34:21984–21996, 2021.
- [53] Tianyi Zhang, Felix Wu, Arzoo Katiyar, Kilian Q Weinberger, and Yoav Artzi. Revisiting few-sample bert fine-tuning. *arXiv preprint arXiv:2006.05987*, 2020.
- [54] Zekang Zhang, Guangyu Gao, Zhiyuan Fang, Jianbo Jiao, and Yunchao Wei. Mining unseen classes via regional objectness: A simple baseline for incremental segmentation. *NeurIPS*, 35, 2022.
- [55] Zekang Zhang, Guangyu Gao, Jianbo Jiao, Chi Harold Liu, and Yunchao Wei. Coinseg: Contrast inter- and intra- class representations for incremental segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023.
- [56] Fei Zhu, Xu-Yao Zhang, Chuang Wang, Fei Yin, and Cheng-Lin Liu. Prototype augmentation and self-supervision for incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5871–5880, 2021.
- [57] Hongguang Zhu, Yunchao Wei, Xiaodan Liang, Chunjie Zhang, and Yao Zhao. Ctp: Towards vision-language continual pretraining via compatible momentum contrast and topology preservation, 2023.

A. More Details and Results.

Implementation Details. All baselines follow an implementation similar to the one described in [46, 45]. Specifically, we adopt a pre-trained ViT-B/16 backbone. We use an Adam optimizer for prompting-based approaches that keep the representation layer fixed, while a SGD optimizer for other baselines that update the entire model, with the same batch size of 128. The original implementation of [46, 45] adopts a constant learning rate of 0.005 for all baselines, while our slow learner using 0.0001 for the representation layer and 0.01 for the classification layer. In practice, we observe that supervised pre-training usually converges faster than self-supervised pre-training in downstream continual learning. Therefore, for supervised pre-training, we train all baselines for 20 epochs on Split CIFAR-100 and 50 epochs on other benchmarks. For self-supervised pre-training, we train all baselines for 90 epochs on all benchmarks.

Extended Analysis. In this section, we provide extended results to support the main claims in our paper. First, we present the CKA similarity of pre-trained representation (1) before and after learning downstream tasks in Fig. 7, and (2) after joint training and after continual learning in Fig. 8.

Results on Additional Dataset. Except CIFAR-100, CUB-200-2011, ImageNet-R and Cars-196, we further consider a subset of DomainNet with 345-class sketch images (for short, Sketch-345). Our SLCA delivers consistently strong performance as shown in Table 4.

Combine with other methods. In the main text, the efficacy of SL has been widely validated by combining it with all baseline methods. We have further validated the efficacy of CA, presenting representative non-replay and replay methods on IN21K-Sup as shown in Table 5.

Method	Sketch-345, IN21K-Sup		Sketch-345, IN1K-Self	
	Last-Acc (%)	Inc-Acc (%)	Last-Acc (%)	Inc-Acc (%)
Joint-Training	72.18 \pm 0.03	-	66.04 \pm 0.07	-
Seq FT	40.40 \pm 14.87	46.91 \pm 24.25	12.98 \pm 4.09	38.80 \pm 5.49
w/ SL	63.41 \pm 0.53	71.24 \pm 0.67	56.94 \pm 0.05	66.07 \pm 0.38
w/ SL+CA	64.92 \pm 0.81	72.69 \pm 0.57	59.88 \pm 0.06	67.99 \pm 0.54

Table 4. Results on Sketch-345, a subset of DomainNet dataset [29].

Method	CIFAR-100	ImageNet-R	CUB-200	Cars-196
EWC	47.01 \pm 0.29	35.00 \pm 0.43	51.28 \pm 2.37	47.02 \pm 3.90
EWC w/ SL	89.30 \pm 0.23	70.27 \pm 1.99	81.62 \pm 0.34	64.50 \pm 0.36
EWC w/ SL+CA	90.61 \pm 0.17	71.48 \pm 0.31	84.29 \pm 0.37	69.61 \pm 0.29
BiC	66.11 \pm 1.76	52.14 \pm 1.08	78.69 \pm 1.97	55.03 \pm 3.27
BiC w/ SL	88.45 \pm 0.57	64.89 \pm 0.80	81.91 \pm 2.59	63.10 \pm 5.71
BiC w/ SL+CA	91.57 \pm 0.13	74.49 \pm 0.08	86.82 \pm 0.69	73.90 \pm 0.38

Table 5. Ablations for CA combining with EWC and BiC.




Figure 7. CKA similarity of pre-trained representations before and after learning downstream tasks.




Figure 8. CKA similarity of pre-trained representations after joint training and after continual learning.

Benchmark	Pre-trained	0.005 [†]	0.001	0.0001	0.00001	0.000001	Fixed θ_{rps}
Split CIFAR-100	IN21K-Sup	44.77 ±13.8	83.04±1.46	88.86±0.83	88.81±0.46	85.11±0.42	63.75±0.67
Split ImageNet-R	IN21K-Sup	26.95 ±11.8	70.38±0.80	71.80±1.45	62.64±2.35	53.57±4.33	34.64±14.3
Split CUB-200	IN21K-Sup	40.02 ±1.08	60.02±1.24	68.07±1.09	66.58±3.93	64.38±3.36	60.44±1.80
Split Cars-196	IN21K-Sup	27.57 ±1.79	15.74±26.3	49.74±1.25	30.66±9.01	24.85±7.90	24.51±6.90
Split CIFAR-100	IN1K-Self	27.99 ±5.16	81.49±0.75	81.47±0.55	81.57±0.14	78.61±0.29	77.30±0.56
Split ImageNet-R	IN1K-Self	45.84 ±4.19	68.72±0.48	64.43±0.44	59.19±0.33	54.54±0.32	51.97±0.17
Split CUB-200	IN1K-Self	45.35 ±1.38	68.58±1.16	61.67±1.37	56.46±1.86	55.10±2.13	55.54±1.55
Split Cars-196	IN1K-Self	35.96 ±2.04	58.39±2.31	52.91±1.61	43.64±0.73	41.74±0.23	43.16±0.12

Table 6. Continual learning performance with different learning rates of the representation layer. Here we present the Last-Acc (%) after continual learning of all classes. IN21K-Sup: supervised pre-training on ImageNet-21K. IN1K-Self: self-supervised pre-training on ImageNet-1K with MoCo v3 [4]. The column labeled by [†] uses the same learning rate of 0.005 for the entire model, while the others use a learning rate of 0.01 for the classification layer.