

Do Pre-trained Models Benefit Equally in Continual Learning?

Kuan-Ying Lee

Yuanyi Zhong

Yu-Xiong Wang

University of Illinois at Urbana-Champaign

{kylee5, yuanyiz2, yxw}@illinois.edu

Abstract

Existing work on continual learning (CL) is primarily devoted to developing algorithms for models trained from scratch. Despite their encouraging performance on contrived benchmarks, these algorithms show dramatic performance drop in real-world scenarios. Therefore, this paper advocates the systematic introduction of pre-training to CL, which is a general recipe for transferring knowledge to downstream tasks but is substantially missing in the CL community. Our investigation reveals the multifaceted complexity of exploiting pre-trained models for CL, along three different axes: pre-trained models, CL algorithms, and CL scenarios. Perhaps most intriguingly, improvements in CL algorithms from pre-training are very inconsistent – an underperforming algorithm could become competitive and even state of the art, when all algorithms start from a pre-trained model. This indicates that the current paradigm, where all CL methods are compared in from-scratch training, is not well reflective of the true CL objective and desired progress. In addition, we make several other important observations, including that 1) CL algorithms that exert less regularization benefit more from a pre-trained model; and 2) a stronger pre-trained model such as CLIP does not guarantee a better improvement. Based on these findings, we introduce a simple yet effective baseline that employs minimum regularization and leverages the more beneficial pre-trained model, coupled with a two-stage training pipeline. We recommend including this strong baseline in the future development of CL algorithms, due to its demonstrated state-of-the-art performance. Our code is available at <https://github.com/eric11220/pretrained-models-in-CL>.

1. Introduction

Continual learning (CL) has gained increasing research momentum recently, due to the ever-changing nature of real-world data [1, 2, 12, 14, 21, 23, 32, 34, 41, 44]. Despite their encouraging performance, many notable CL al-

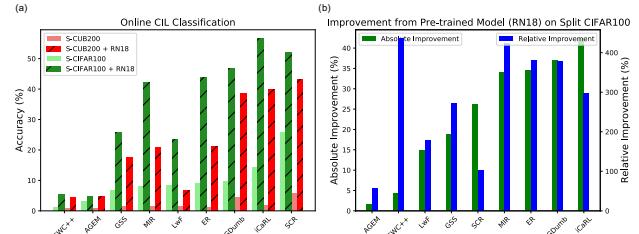


Figure 1. (a) CL algorithms trained from scratch fail on **Split CUB200**, a more complex dataset than **Split CIFAR100**, which necessitates the use of pre-trained models (denoted as ‘+ RN18’) that dramatically increase the accuracy of a wide spectrum of algorithms. (b) Different CL algorithms receive vastly different benefits from pre-trained models, and the superiority between algorithms changes. These findings suggest that it is critical for the community to develop CL algorithms with a pre-trained model and understand their behaviors. [Best viewed in color.]

gorithms were developed to work with a model trained from scratch. As one of the key objectives, this paper advocates the *systematic introduction of pre-training to CL*. This is rooted in the following two observed fundamental limitations of building CL algorithm on top of a from-scratch trained model, which fails to reflect the true progress in the CL research for real-world scenarios as shown in Fig. 1.

First, training from scratch does not reflect the actual performance, because if one were to apply a CL algorithm to real-world scenarios, it would be counter-intuitive not to build upon off-the-shelf pre-trained models given the large performance gap (Fig. 1). One might argue that applying all algorithms to a from-scratch trained model simplifies comparison between different algorithms. However, intriguingly, our study shows that *an underperforming algorithm could become competitive and even achieve state-of-the-art performance, when all algorithms start from a pre-trained model*. In particular, iCaRL [32], which shows mediocre performance in online class incremental learning (CIL) when trained from scratch, is comparable to or even outperforms SCR [26], when both are initialized from a ResNet18¹ pre-trained on ImageNet (accuracy in-

¹We refer to ResNet as RN throughout the paper.

crease from 14.26% to 56.64% for iCaRL vs. increase from 25.80% to 51.93% for SCR on Split CIFAR100 in Fig. 1 and Table 2). This potentially indicates that the efforts funneled into the development of CL algorithms could be in a less effective direction and are not well reflective of the actual progress in CL. Therefore, we should develop any future CL algorithms in the context of how we are going to use them in practice – starting from a pre-trained model.

Second, for many more realistic datasets with diverse visual concepts, data scarcity makes it impossible to train a CL learner from scratch [9] (as also shown in the results on Split CUB200 in Fig. 1. We believe that this is partially the reason why CL classification literature still heavily evaluates on contrived benchmarks such as Split MNIST and Split CIFAR [14, 23], as opposed to much more complex datasets typically used in offline learning.

Through our investigation, this paper reveals the *multi-faceted complexity* of exploiting pre-trained models for CL. As summarized in Table 1, we conduct the investigation along *three* different axes: different pre-trained models, different CL algorithms, and different CL scenarios. In particular, we analyze models pre-trained in either supervised or self-supervised fashion and from three distinct sources of supervision – curated labeled images, non-curated image-text pairs, and unlabeled images. These models cover supervised RN18/50 [19] trained on ImageNet classification [15], CLIP RN50 [31], and self-supervised RN50 trained with SimCLR [10], SwAV [6], or Barlow Twins [43].

We make several important observations. **1)** Benefits of a pre-trained model on different CL algorithms vary widely, as represented by the aforementioned comparison between iCaRL and SCR. **2)** As shown in Fig. 1, algorithms applying less regularization to the gradient (i.e., replay-based methods like ER [33]) seem to benefit the most from pre-trained models. **3)** Intriguingly, despite its impressive zero-shot capability, CLIP RN50 mostly *underperforms* ImageNet RN50. **4)** Self-supervised fine-tuning helps alleviate catastrophic forgetting. For example, fine-tuning SimCLR RN50 on the downstream dataset in a self-supervised fashion with the SimCLR loss demonstrates a huge reduction in forgetting, compared with supervised models (17.99% forgetting of SimCLR RN50 vs. 91.12% forgetting of supervised RN50). **5)** Iterating over data of a given task for multiple epochs as in class incremental learning (CIL) *does not necessarily* improve the performance over online CIL.

Based on these observations, we further propose a strong baseline by applying ER, which exerts minimum regularization (the second observation), on an ImageNet pre-trained model (the third observation). Coupled with a two-stage training pipeline [18] (Sec. 3.3), we show that such a simple baseline achieves state-of-the-art performance. We recommend including this strong baseline in the future development of CL algorithms.

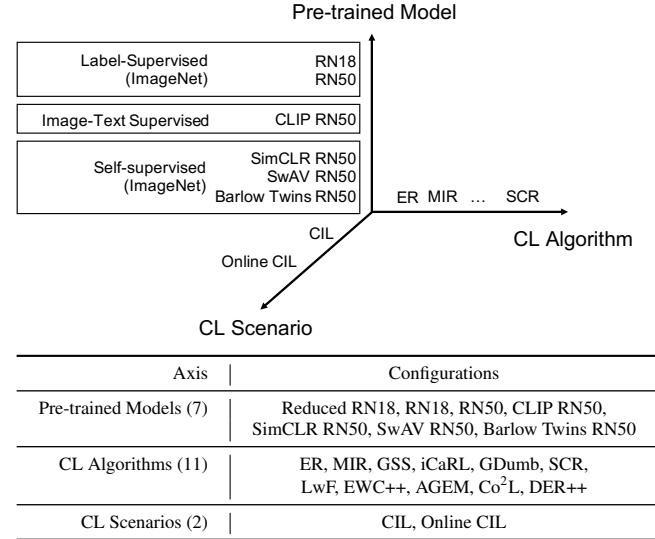


Table 1. We conduct the analyses of pre-trained models in CL by dissecting the space into three axes: 1) different pre-trained models, 2) different CL algorithms, and 3) different CL scenarios.

Our contributions are summarized as follows. 1) We show the necessity of pre-trained models on more complex CL datasets and the dramatic difference in their benefits on different CL algorithms, which may overturn the comparison results between algorithms. Therefore, we suggest the community consider pre-trained models when developing and evaluating new CL algorithms. 2) We show that replay-based CL algorithms seem to benefit more from a pre-trained model, compared with regularization-based counterparts. 3) We propose a simple yet strong baseline based on ER and ImageNet RN50, which achieves state-of-the-art performance for CL with pre-training.

2. Related Work

Continual Learning Scenarios. A large portion of CL literature focuses on incremental learning, which can be further divided into three different scenarios – task, domain, and class incremental learning [37]. Amongst, the most challenging scenario is class incremental learning (CIL), where the model has to predict all previously seen classes with a single head in the absence of task information. Most recent work [17, 20, 25] has investigated this setting.

However, being able to iterate over the entire data of a specific task for multiple epochs is not realistic [9, 30]. To this end, an online version of CIL is proposed [9, 14], where the model trains in an online fashion and thereby can only have access to each example once. In this work, we also mainly investigate pre-trained models in online CIL but also report results in CIL for several representative algorithms.

Continual Learning Methods. According to [13], continual learning approaches can be divided into three classes:

regularization, parameter isolation, and replay methods. Regularization methods [1, 21, 44] prevent the learned parameters from deviating too much to prevent forgetting. Parameter isolation methods counter forgetting completely by dedicating a non-overlapping set of parameters to each task [34, 41]. Replay methods either store previous instances [14, 17, 23, 32] or generate pseudo-instances [12, 35, 36] on the fly for replay to alleviate forgetting.

While the aforementioned approaches all show promising results in different CL scenarios, we specifically explore regularization and memory replay-based methods, given their popularity in recent literature. And we study the behaviors of pre-trained models on these methods.

Continual Learning with Pre-trained Models. While most of the CL work investigates training the learner from scratch [4, 14, 21, 23, 25, 32, 44], there is also some work that initializes the learner from a pre-trained model [3, 9, 11, 20, 28, 29]. They harness pre-trained models for reasons such as coping with data scarcity of the downstream task [9] and simulating prior knowledge of the continual learner [20]. However, they **do not** 1) systematically show the substantial benefits of pre-trained models over from-scratch trained models, 2) investigate different types of pre-trained models or fine-tuning strategies, or 3) investigate pre-trained models on different CL scenarios (incremental and online learning). Note that we claim no contribution to be the first to apply pre-trained models on CL, but rather study the aforementioned aspects *comprehensively*.

3. Methodology

We mainly focus on online class incremental learning (CIL), which is formally defined in Sec. 3.1. Next, we discuss various pre-trained models and how we leverage them (Sec. 3.2). In Sec. 3.3, we introduce the two-stage training pipeline that combines online training and offline training.

3.1. Problem Formulation

The most widely adopted continual learning scenarios are 1) task incremental learning, 2) domain incremental learning, and 3) class incremental learning (CIL). Amongst, CIL is the most challenging and draws the most attention, for its closer resemblance to real-world scenarios, where the model is required to make predictions on all classes seen so far with no task identifiers given (we refer interested readers to [37] for more details). In this paper, we focus on a more difficult scenario – *online CIL*, where the model can only have access to the data once unless with a replay buffer. In other words, the model **can not** iterate over the data of the current task for multiple epochs, which is common in CIL. The experiments of the remaining paper are based on online CIL, unless noted otherwise (we also evaluate CIL).

Formally, we define the problem as follows. C_{total} classes are split into N tasks, and each task t contains C_t

non-overlapping classes (e.g., CIFAR100 is divided into 20 tasks, with each task containing 5 unique classes). The model is presented with N tasks sequentially. Each task is a data stream $\{S_t | 0 < t \leq N\}$ that presents a mini-batch of samples to the model at a time. Each sample belongs to one of the C_t unique classes, which are non-overlapping to other tasks. We explore pre-trained models with $N = 20$ tasks and present the results in Sec. 5.2.

3.2. Fine-Tuning Strategy

When using a pre-trained model, we initialize the model with the pre-trained weights. Then, we fine-tune the model in either 1) supervised, or 2) self-supervised manner.

For the self-supervised fine-tuning, we experiment with the SimCLR pre-trained RN50 and fine-tune it with the SimCLR loss. Specifically, we leverage a replay buffer to store images and labels, which are then used to train the classifier on top of the fine-tuned feature representation at the end of each CL task. Note that we train the SimCLR feature with both images sampled from the memory and images from the data stream.

3.3. Two-Stage Pipeline

We combine the two-stage training pipeline proposed in [18] with pre-trained models to build a strong baseline for continual learning with pre-training.

The two-stage pipeline divides the learning process into two phases – a *streaming* phase where the learner (sampler) is exposed to the data stream, and an *offline* phase where the learner learns from the memory. Similar to another widely used method GDumb [30], the two-stage pipeline trains offline on samples in the memory after all data have been streamed. Specifically, we iterate over samples in the memory for 30 epochs after the streaming phase. However, GDumb performs no learning in the streaming phase and discards most of the data without learning from them, making it sub-optimal. By contrast, the two-stage pipeline improves over GDumb, through training the model on the data while storing them at the same time.

We found that this simple two-stage pipeline is particularly effective when coupled with pre-trained models. With the two-stage training, ER [33] could outperform the best-performing models, when all leverage a pre-trained model.

4. Experimental Setting

4.1. Datasets

We experiment on five datasets. Classes in each dataset are equally and randomly split into 20 tasks with no overlaps. The orderings are random, but kept the same across different experiments.

Split CIFAR100. We follow the suggested train and test split, where each category has 500 training and 100 test

data [22].

Split CUB200. CUB200 [40] contains 5594 training (~ 30 per category) and 6194 test images for 200 different bird species.

Split Mini-ImageNet. Mini-ImageNet [38] contains 100 subcategories from the original 1000 categories in ImageNet [15]. The 100 categories are divided into 64 training, 20 validation, and 16 test categories. We follow [26] and re-purpose all 100 categories for training and set aside 100 images per category for testing. In total, it contains 50k and 10k images for training and testing, respectively. Note the data overlap of Split Mini-ImageNet with ImageNet which is used to pre-train ResNets – this experiment is *purposefully* designed to evaluate the scenario where the dataset is extremely similar to the data used for pre-training.

Split FGVC-Aircraft. FGVC-Aircraft [27] contains 100 aircraft model variants, each with 100 images. We combine the suggested train and validation split for training, and the test split for testing. In total, there are 6667 training images (~ 66 images per category) and 3333 test images.

Split QuickDraw. QuickDraw [16] contains 50M drawings across 345 categories. We randomly select 100 categories, from which we randomly sample 100 images per category (50 for training and 50 for testing) to create a dataset of 10000 images in total.

4.2. CL Algorithms

We explore 11 different CL algorithms, including both regularization-based methods (in *Italics*) that require no memory and replay-based methods.

ER. Experience Replay [33] stores samples from the data stream with reservoir sampling and retrieves stored samples randomly.

MIR. Maximally Interfered Retrieval [2] performs a pseudo-update on the incoming data and retrieves samples stored in the memory whose loss increases most.

GSS. Gradient-based Sample Selection [4] attempts to diversify the gradients of the samples in the memory.

LwF. Learning without Forgetting [24] applies a knowledge distillation loss from the previous model when learning the new task.

iCaRL. Incremental Classifier and Representation Learning [32] stores a small set of images for different classes and learns features continually with distillation and classification losses.

EWC++. Elastic Weight Consolidation++ [8] is an online version of EWC [21] that constrains updates to important parameters for the previous tasks.

GDumb. Termed Greedy Sampler and Dumb Learner [30] randomly samples data to store during the streaming phase and only performs training when queried.

AGEM. Averaged Gradient Episodic Memory [9] constrains gradients with samples in the memory.

SCR. Supervised Contrastive Replay [26] applies supervised contrastive loss in addition to the cross-entropy loss.

DER++. Dark Experience Replay++ [5] matches the logits of the network throughout the past.

Co²L. Contrastive Continual Learning [7] learns and maintains features in a self-supervised fashion. Re-purposing Co²L for *online* CIL, we train the feature encoder on each task for one epoch. At the end of the streaming phase, we train the classifier using the stored buffer data for one epoch.

4.3. Pre-trained Models

Reduced RN18 (R-RN18). ResNet18 whose number of channels is reduced [25] compared with a standard one, which is used in the experiment of training from scratch.

ImageNet Pre-trained RN18, RN50. ResNets pre-trained on ImageNet [15].

CLIP Pre-trained RN50. ResNet50 pre-trained on the WebImageText dataset based on Contrastive Language–Image Pre-training (CLIP) [31].

SimCLR RN50. ResNet50 pre-trained on ImageNet with the SimCLR loss that brings closer features of different augmentations from the same image, while pushing apart those from different images [10].

SwAV RN50. ResNet50 pre-trained on ImageNet with the SwAV mechanism that predicts the cluster assignment of a view from the representation of another one [6].

Barlow Twins RN50. ResNet50 pre-trained on ImageNet with the Barlow Twins loss that encourages the correlation of two views from the same image to be one, while discouraging that of views from different images to be zero [43].

4.4. Implementation Details

Image Pre-processing and Architectures. We use the pre-trained ResNets provided in PyTorch and the CLIP Github repository². For SimCLR, SwAV, and Barlow Twins, we utilize pre-trained weights provided in Pytorch Lightning Bolts³. For different CL algorithms evaluated, we adopt the publicly available code from [26]⁴. The memory capacity for all replay-based methods is set to 1000.

When training from scratch, we do not perform any resizing and pre-processing for most of the datasets; for Split CUB200, we resize the short side of the images to 224, with a random crop for training and a center crop for testing. When initializing the continual learner from a pre-trained model, we resize the short side of the image to 224, followed by a random crop and a center crop for training and testing, respectively. Note that except random cropping, we do not apply other data augmentation. We normalize the images by the statistics computed during pre-training

²<https://github.com/openai/CLIP>

³<https://pytorch-lightning-bolts.readthedocs.io/en/stable>

⁴<https://github.com/RaptorMai/online-continual-learning>

Split CIFAR100

Model	ER [33]	MIR [2]	GSS [4]	<i>LwF</i> [24]	iCaRL [32]	<i>EWC++</i> [8]	GDumb [30]	AGEM [9]	SCR [26]
R-RN18	9.07 \pm 1.31	8.03 \pm 0.78	6.86 \pm 0.60	8.44 \pm 0.82	14.26 \pm 0.79	1.00 \pm 0.00	9.80 \pm 0.46	3.00 \pm 0.47	25.80\pm0.99
RN18	43.69 \pm 1.67	42.02 \pm 1.53	25.59 \pm 0.45	23.40 \pm 0.12	56.64\pm0.23	5.36 \pm 0.26	46.76 \pm 0.73	4.72 \pm 0.21	51.93 \pm 0.06
Δ	+34.62	+33.99	+18.73	+14.96	+42.38	+4.36	+36.96	+1.72	+26.13

Split CUB200

Model	ER	MIR	GSS	<i>LwF</i>	iCaRL	<i>EWC++</i>	GDumb	AGEM	SCR
R-RN18	1.24 \pm 0.11	1.44 \pm 0.08	1.46 \pm 0.22	1.47 \pm 0.11	1.82 \pm 0.24	0.80 \pm 0.20	4.49 \pm 0.56	0.67 \pm 0.12	5.64\pm0.75
RN18	21.05 \pm 1.07	20.95 \pm 0.66	17.65 \pm 0.45	6.79 \pm 0.36	39.95 \pm 1.43	4.47 \pm 0.10	38.63 \pm 0.44	4.59 \pm 0.30	43.03\pm1.80
Δ	+19.81	+19.51	+16.19	+5.32	+38.13	+3.67	+34.14	+3.92	+37.39

Split Mini-ImageNet

Model	ER	MIR	GSS	<i>LwF</i>	iCaRL	<i>EWC++</i>	GDumb	AGEM	SCR
R-RN18	8.56 \pm 0.24	8.00 \pm 0.82	6.74 \pm 0.15	7.58 \pm 0.65	11.61 \pm 0.78	1.00 \pm 0.00	7.01 \pm 0.40	3.04 \pm 0.21	33.87\pm1.84
RN18	56.91 \pm 0.54	54.96 \pm 0.46	25.74 \pm 4.53	20.41 \pm 0.99	72.40\pm0.52	4.79 \pm 0.14	40.00 \pm 0.37	5.23 \pm 0.41	67.94 \pm 0.11
Δ	+48.35	+46.96	+19.00	+12.83	+60.79	+3.79	+29.99	+2.19	+34.07

Table 2. **Accuracy in online CIL.** Different CL algorithms benefit from a pre-trained model very differently, and the comparison results between algorithms change when they are initialized from a pre-trained model. For instance, iCaRL outperforms SCR, the best-performing model when trained from scratch, on Split CIFAR100 (56.64 vs. 51.93) and Split Mini-ImageNet (72.40 vs. 67.94). This indicates that training from scratch does not serve as a fairground for comparison between different algorithms, in addition to its poor applicability to complex datasets. R-RN18 and RN18 stand for Reduced ResNet18 trained from scratch and ImageNet pre-trained ResNet18, respectively.

Model	ER [33]	MIR [2]	GSS [4]	<i>LwF</i> [24]	iCaRL [32]	<i>EWC++</i> [8]	GDumb [30]	AGEM [9]	SCR [26]	DER++ [5]	Co ² L [7]
R-RN18	9.07 \pm 1.31	8.03 \pm 0.78	6.86 \pm 0.60	8.44 \pm 0.82	14.26 \pm 0.79	1.00 \pm 0.00	9.80 \pm 0.46	3.00 \pm 0.47	25.80\pm0.99	15.72 \pm 1.33	2.31 \pm 0.64
RN18	43.69 \pm 1.67	42.02 \pm 1.53	25.59 \pm 0.45	23.40 \pm 0.12	56.64\pm0.23	5.36 \pm 0.26	46.76 \pm 0.73	4.72 \pm 0.21	51.93 \pm 0.06	44.42 \pm 1.29	5.68 \pm 3.19
RN50	50.88 \pm 0.84	50.20 \pm 2.80	31.53 \pm 3.37	26.68 \pm 0.97	59.20\pm0.33	3.47 \pm 1.42	57.37 \pm 0.21	4.49 \pm 0.27	56.22 \pm 0.42	49.37 \pm 1.36	8.57 \pm 0.57
CLIP	52.31 \pm 2.66	55.38\pm0.83	25.60 \pm 4.50	37.21 \pm 2.14	26.05\pm12.33	—*	55.10 \pm 0.22	17.22 \pm 2.52	30.93\pm5.44	53.01\pm0.18	1.12 \pm 0.16
SimCLR RN50	37.04 \pm 0.48	40.01 \pm 1.86	16.32 \pm 1.52	3.40\pm0.17	33.76 \pm 0.84	6.39 \pm 0.82	24.63 \pm 0.84	3.87 \pm 0.32	52.60\pm0.22	15.63 \pm 0.96	1.44 \pm 0.45
SwAV RN50	38.32 \pm 0.11	40.97 \pm 0.36	15.00 \pm 0.30	3.32\pm0.45	24.29 \pm 1.32	3.58 \pm 3.00	20.95 \pm 1.33	3.86 \pm 0.29	50.59\pm0.09	20.10 \pm 0.88	1.18 \pm 0.26
B.T. RN50	26.15 \pm 0.62	18.18 \pm 1.60	8.38 \pm 0.23	3.70\pm0.16	40.77 \pm 0.92	6.65 \pm 1.06	31.56 \pm 2.01	3.95 \pm 0.31	48.35\pm0.73	5.26 \pm 0.17	1.10 \pm 0.10

*EWC++ fails to train with losses of *nan*.

Table 3. **Accuracy of different pre-trained models** when fine-tuned in a supervised manner on Split CIFAR100 in online CIL. In most cases, RN pre-trained on ImageNet (RN50 vs. CLIP RN50) in a supervised fashion (RN50 vs. SimCLR, SwAV, and Barlow Twins RN50) brings the largest accuracy increase. Red numbers mark pre-trained accuracy that is within/below one std. of the from-scratch counterpart, which indicates potential negative impacts. Bold numbers indicate the best accuracy amongst all methods with a specific model (e.g., 25.80 of SCR is the best within R-RN18). R-RN18, RN18, RN50, and CLIP stand for Reduced ResNet18 trained from scratch, ImageNet pre-trained ResNet18 and ResNet50, and CLIP pre-trained ResNet50, respectively. B.T. stands for Barlow Twins. [Best viewed in color.]

on either ImageNet (RN18, RN50, SimCLR RN50, SwAV RN50, and Barlow Twins RN50) or CLIP data.

When training from scratch, we follow [26] and use a reduced ResNet18 (R-RN18) with fewer channels and a smaller first convolution kernel.

Classifier Initialization. For all models except CLIP RN50, we randomly initialize the classifier and set its output dimension to the number of classes in the target dataset. For CLIP RN50, we utilize its text encoder and follow the zero-shot classification guideline provided in the CLIP Github repository to generate the classifier. Specifically, we use the text encoder to encode text in the format of ‘A photo of XX,’ with XX being one of the categories. The weights of

the classifier are then initialized from these N text features, with N being the number of classes [31].

Hyper-parameter Selection. For the two-stage training, we do not perform an extensive hyper-parameter search on CL scenarios, but instead directly use hyper-parameters that work reasonably in offline training. Specifically, for supervised fine-tuning, we set learning rate to 0.1, 0.01, and 1e-6 for from-scratch trained Reduced ResNet18, pre-trained ResNets (RN18, RN50, SimCLR RN50, SwAV RN50, and Barlow Twins RN50), and CLIP RN50, respectively. For self-supervised fine-tuning with the SimCLR pre-trained RN50, we set the learning rate to 0.375. For CIL, we train the model for 5 epochs on each task.

For all methods, the batch size is set as 10. However, when fine-tuning SimCLR in a self-supervised fashion, we use a batch size of 64 (due to GPU memory constraints). All models except SimCLR are optimized via stochastic gradient descent (SGD) with no momentum and weight decay. For SimCLR, we follow [10] and apply layer-wise adaptive rate scaling (LARS) [42] on top of the plain SGD. A momentum of 0.9 and a weight decay of 1e-6 are applied.

For all results, we take the mean and standard deviation across three runs with different random seeds.

5. Experimental Result

We perform investigation on pre-trained models along three axes: 1) different CL algorithms, 2) different pre-trained models, and 3) different CL scenarios.

In Sec. 5.1, along the *CL algorithm* axis, we show varied impact of pre-trained models on different algorithms. In Sec. 5.2, along the *pre-trained model* axis, we examine difference between pre-trained models. In Sec. 5.3, along the *CL scenario* axis, we explore different CL scenarios. Finally, in Sec. 5.4, we integrate several observations and demonstrate the proposed strong baseline.

5.1. Comparison of Different CL Algorithms

Pre-trained models are essential for more complex CL datasets. From Table 2, one can observe that while many of the CL algorithms work fine on Split CIFAR100, they witness a severe performance drop when applied to Split CUB200, a more complex dataset. This clearly illustrates the huge benefit and the necessity of pre-trained models, if one were to deal with real-world CL problems that are likely even more difficult than Split CUB200. In fact, we would almost always consider initializing the learner from a pre-trained model in practice, reasonably assuming that the learner has some prior knowledge [20] to achieve better performance.

Varied benefits of pre-trained models on distinct CL algorithms. From Fig. 1, Table 2, and Table 3, clearly, the benefits of pre-trained models are not homogeneous across different CL algorithms. For example, while ER enjoys a 34.62% accuracy gain, LwF only has a 14.96% improvement, despite the similar accuracy when applied on a from-scratch trained model. Further if we sort the performance gains brought by ImageNet RN18 on Split CIFAR100 to different algorithms, one could observe that methods enjoying the most gains are ones that do not apply too much regularization (e.g., to gradients) during training. One such comparison is between ER and EWC++: ER exerts no additional loss other than the cross-entropy loss, while EWC++ applies a regularization loss based on the importance of network parameters, which is approximated by the diagonal of the Fisher information matrix.

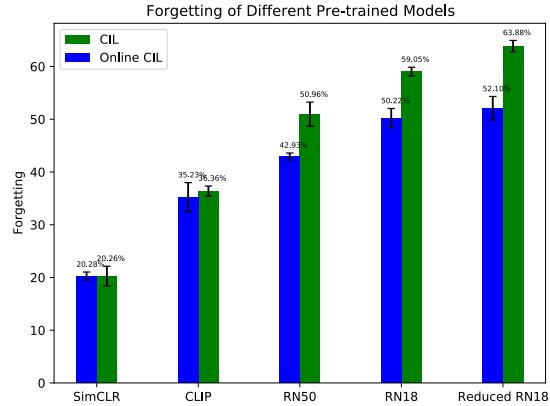


Figure 2. **Forgetting of different models.** Fine-tuning a pre-trained model in a self-supervised manner (with the SimCLR loss) largely decreases forgetting, compared with supervised fine-tuning. [Best viewed in color.]

Superiority inconsistency between from-scratch training and pre-training. The unequal gains also affect comparisons between different algorithms. For instance in Table 2, interestingly, while iCaRL only achieves 14.26% in accuracy, the comparison result with SCR (25.80%) is overturned when both are applied on ImageNet RN18 (56.64% of iCaRL vs. 51.93% of SCR).

Remark. Given the observations that pre-trained models should be applied in a more complex dataset and that the superiority of different algorithms does not hold when applied to a pre-trained model, future development and evaluation of CL algorithms should take pre-trained models into consideration.

5.2. Comparison of Different Pre-trained Models

Comparison on Forgetting. In Table 4, one could observe that CLIP RN50 enjoys consistently lower forgetting, compared with training from scratch (R-RN18). Yet, the other pre-trained models (also fine-tuned in a supervised manner) do not have the same advantage, regardless of being pre-trained in a supervised or self-supervised manner.

Table 5 and Fig. 2 show that self-supervised fine-tuning (SimCLR) has less forgetting than supervised fine-tuning. We identify two differences between self-supervised and supervised fine-tuning – 1) decoupled training of the feature and classifier, and 2) self-supervision. Our in-depth experiment in Sec. C of the Supplementary Material shows that the decoupled training likely helps mitigate forgetting.

CLIP vs. ImageNet Pre-trained RN50. From Table 6, surprisingly, despite its impressive zero-shot capability, CLIP RN50 is still worse than ImageNet RN50 in most cases. Based on Fig. 3, we conjecture that CLIP RN50 is better in few-shot regimes, while ImageNet RN50 prevails with plenty of data, which coincides with the CL scenarios considered in this work.

Model	ER [33]	MIR [2]	GSS [4]	LwF [24]	iCaRL [32]	EWC++ [8]	GDumb [30]	AGEM [9]	SCR [26]
R-RN18	52.10 \pm 2.22	50.51 \pm 2.34	50.69 \pm 0.23	8.83 \pm 1.31	12.27 \pm 0.34	2.19\pm0.28	13.59 \pm 0.14	59.23\pm1.47	21.00 \pm 0.79
RN18	50.22 \pm 1.82	51.52 \pm 1.77	66.85 \pm 0.98	-2.14 \pm 2.32	15.77 \pm 0.40	89.37 \pm 0.58	19.56 \pm 0.78	90.13 \pm 0.50	22.84 \pm 0.85
RN50	42.93 \pm 0.67	40.92 \pm 2.44	58.88 \pm 3.40	-3.42\pm1.18	14.65 \pm 0.65	87.36 \pm 2.76	17.02 \pm 0.48	90.26 \pm 0.40	27.52 \pm 0.20
CLIP	35.51\pm2.90	30.22\pm1.98	10.53\pm1.95	-1.81 \pm 0.93	9.73\pm0.17	—*	5.84\pm0.12	73.82 \pm 2.53	15.79 \pm 1.48
SimCLR RN50	46.88 \pm 0.90	40.12 \pm 1.80	59.12 \pm 2.47	-0.01 \pm 0.04	14.88 \pm 1.11	64.16 \pm 0.23	22.30 \pm 1.33	75.99 \pm 1.09	11.14\pm0.41
SwAV RN50	45.85 \pm 0.15	39.26 \pm 0.39	60.58 \pm 1.66	-0.22 \pm 0.13	13.39 \pm 0.72	48.83 \pm 14.62	24.55 \pm 1.26	74.66 \pm 1.59	12.20 \pm 0.48
B.T. RN50	60.57 \pm 0.55	56.80 \pm 2.05	63.71 \pm 1.38	0.25 \pm 0.24	16.15 \pm 0.42	74.19 \pm 0.40	18.13 \pm 2.20	78.31 \pm 1.40	14.50 \pm 0.63

*EWC++ fails to train with losses of *nan*.

Table 4. **Forgetting of different pre-trained models** when fine-tuned in a supervised manner on Split CIFAR100 in online CIL. **Bold** numbers indicate the least forgetting amongst all the models (column-wise comparison). CLIP shows consistently less forgetting compared with other ImageNet pre-trained ResNets. R-RN18, RN18, RN50, and CLIP stand for Reduced ResNet18 trained from scratch, ImageNet pre-trained ResNet18 and ResNet50, and CLIP pre-trained ResNet50, respectively. B.T. stands for Barlow Twins.

(a) Experience Replay (ER)					(b) Learning without Forgetting (LwF)					
Fine-Tuning	From-scratch		Supervised		Self-supervised	Fine-Tuning	From-scratch		Supervised	
	R-RN18	RN18	RN50	CLIP			R-RN18	RN18		
CIL	8.17 \pm 1.06 (63.88 \pm 1.07)	36.21 \pm 1.17 (59.05 \pm 0.82)	44.18 \pm 2.55 (50.96 \pm 2.29)	55.44 \pm 1.34 (36.36 \pm 0.98)	34.72 \pm 4.04 (20.26 \pm 1.87)	CIL	13.05 \pm 0.65 (8.33 \pm 4.35)	19.18 \pm 0.86 (-4.40 \pm 1.44)	17.82 \pm 1.83 (-3.50 \pm 2.10)	35.52 \pm 1.90 (-3.97 \pm 1.25)
Online CIL	9.07 \pm 1.31 (52.10 \pm 2.22)	43.69 \pm 1.67 (50.22 \pm 1.82)	50.88 \pm 0.84 (42.93 \pm 0.67)	52.79 \pm 1.91 (35.51 \pm 2.90)	33.39 \pm 0.42 (20.28 \pm 0.75)	Online CIL	8.44 \pm 0.82 (8.83 \pm 1.31)	23.40 \pm 0.12 (-2.14 \pm 2.32)	25.98 \pm 1.34 (-3.42 \pm 1.18)	37.73 \pm 1.19 (-1.81 \pm 0.93)

Table 5. **Accuracy (Forgetting) of different models** on Split CIFAR100. (a) *Self-supervised fine-tuning* (SimCLR) demonstrates a lower forgetting compared with supervised fine-tuning (20.26 vs. 50.96 of RN50 in CIL). (a)(b) CLIP, pre-trained with image-text pairs, shows less forgetting compared with ResNets pre-trained with curated ImageNet labels. Numbers outside/inside parentheses are accuracy/forgetting, respectively. R-RN18 and RN18 stand for Reduced ResNet18 and ImageNet pre-trained ResNet18, respectively.

S-CIFAR100	S-CUB200	S-QuickDraw	S-Aircraft	
RN50	57.68 \pm 0.23	48.05\pm0.87	55.19\pm0.41	34.23\pm0.98
CLIP	59.22\pm0.33	44.98 \pm 0.31	34.88 \pm 0.49	19.17 \pm 0.14

Table 6. **ImageNet RN50 vs. CLIP RN50** in online CIL. Despite the impressive zero-shot capability, CLIP underperforms ImageNet pre-trained RN50 in three out of the four datasets, suggesting that the better pre-trained model to consider is probably still ImageNet RN. S- stands for Split. RN50 and CLIP are ImageNet ResNet50 and CLIP ResNet50, respectively.

5.3. Comparison of Different CL Scenarios

From Table 5, one can observe that the performance of CIL is lower than online CIL. While seemingly counter-intuitive, this is likely because, in CIL, the model overfits too much to the current task, which worsens the forgetting. The same behavior is also observed in [39].

5.4. A Strong Baseline for CL with Pre-training

We have observed that 1) algorithms that apply less regularization tend to benefit more from pre-trained models, and 2) ResNets trained on ImageNet provide a better improvement on CL in general. Based on these insights, we propose a strong baseline for CL with pre-training. We combine the simplest ER that exerts no regularization during training, ImageNet RN50, and the two-stage training pipeline dis-

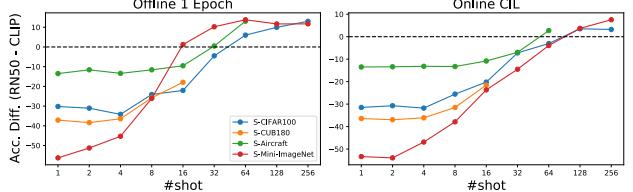


Figure 3. **Comparison between CLIP R50 and ImageNet RN50.** CLIP is better in the few-shot regime, while ImageNet RN50 prevails with plenty of data, as is the case for the CL scenarios considered in the paper. This explains the better performance when a CL learner is initialized from an ImageNet RN50 (cf. Table 6). S- stands for Split. [Best viewed in color.]

cussed in Sec. 3.3.

In Fig. 4 and Table 7, we compare this baseline with best-performing methods, iCaRL and SCR. We observe that with the additional offline training, the accuracy of ER increases from 50.83% to 65.35%. This demonstrates that a simple baseline ER, coupled with a pre-trained model, could turn a relatively naïve baseline into a strong baseline that even achieves the state-of-the-art result. Results of other methods with the two-stage pipeline are provided in the Supplementary Material, showing the *consistent* effectiveness in most cases. We hence recommend including this simple yet strong baseline for CL algorithm evaluation and comparison.

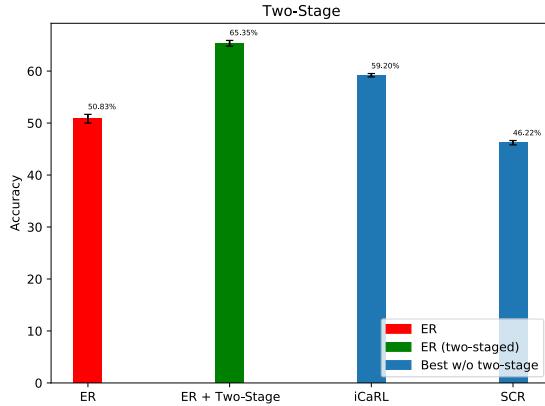


Figure 4. A simple second-staged offline training on memory data coupled with an ImageNet pre-trained ResNet50 turns a simple baseline into state of the art, suggesting the effectiveness of the proposed baseline. Note that SCR and iCaRL are the two best-performing methods when applied on the ImageNet pre-trained ResNet50. [Best viewed in color.]

Also, we note that the observation we made here is different from the one in [30]. In [30], for the from-scratch trained model, the authors found that simply fine-tuning in the offline phase without learning online achieves the best performance. However, we found that when initialized from a pre-trained model, learning in both the online and offline phases achieves the best performance. The discrepancy indicates potentially different training behaviors for CL between training from scratch and from a pre-trained model.

6. Limitation and Discussion

Although we have shown that pre-trained models bring a considerable amount of benefits to CL, they could limit the architecture choices. Also, it is still unclear how one should choose the optimal pre-trained model and fine-tuning strategy given a CL task. While we have shown that the optimality could depend on the downstream dataset and that ImageNet pre-trained ResNets are generally better, more research is required to have a better conclusion on this.

While the two-stage training pipeline seems promising, there are potential limitations which we have not addressed. First of all, in many real-world scenarios, there will be no ending to the data stream and the learner has to endlessly learn new concepts. This again reflects why we consider *online* CIL and not CIL, as in most real-world scenarios, it is not applicable to iterate the training data multiple times. Even if in a few scenarios where we do have the leisure to further train the model, there might be too many classes to fit into the memory. In such scenarios, it is questionable whether performing the additional fine-tuning on the limited classes in the memory is beneficial.

		ER	ER+2-stage	iCaRL	SCR
S-CIFAR100	RN18	43.69±1.67	58.59±0.31	56.64±0.23	51.93±0.06
	RN50	50.83±0.84	65.35±0.55	59.20±0.33	46.22±0.42
	CLIP	52.31±2.66	59.42±0.80	26.05±12.33	30.93±5.44
S-CUB200	RN18	21.05±1.07	52.32±0.98	39.95±1.43	43.03±1.80
	RN50	48.05±0.87	59.88±0.30	43.69±0.79	50.77±0.80
	CLIP	42.08±0.36	43.91±0.88	3.55±0.86	4.60±0.00

Table 7. **Two-stage training pipeline.** Building on top of a pre-trained model and the two-stage training pipeline, a simple ER baseline becomes a state-of-the-art approach. This further suggests that CL algorithms should be developed and evaluated with pre-trained models taken into consideration. S- stands for Split.

7. Conclusion

While most of existing continual learning algorithms have been developed in the context of from-scratch training, we show that this is not the most effective way to develop and evaluate continual learning algorithms. As supportive observations, our extensive empirical study reveals that 1) the pre-trained models are dramatically beneficial, making them of great necessity for real-world scenarios; and 2) the benefits of a pre-trained model on different CL algorithms are vastly different. The best algorithm when trained from scratch does not necessarily perform the best when coupled with a pre-trained model. These findings indicate that the current methodology of developing CL algorithms from scratch could be potentially less effective.

We also notice different behaviors between different pre-trained models. Somewhat surprisingly, despite being a powerful pre-trained model for its zero-shot capability, CLIP RN50 seems to underperform the supervised ImageNet RN50 in most CL cases. Further investigation into explaining different behaviors and developing model selection strategies is an interesting direction. Another observation is that fine-tuning a SimCLR pre-trained model in a self-supervised manner has significantly lower forgetting. The observation could be potentially leveraged to develop a hybrid method that incorporates both the self-supervised loss and the cross-entropy loss. By doing so, one could potentially enjoy the best of both worlds, reducing forgetting of supervised fine-tuning and increasing the accuracy of self-supervised fine-tuning.

Finally, based on the observations that algorithms exerting less regularization during training benefit more from a pre-trained model, and that ImageNet RN50 provides more benefit than the other pre-trained models, we propose a simple yet effective baseline that achieves state-of-the-art performance on multiple datasets.

Acknowledgement. This work was supported in part by NSF Grant 2106825, NIFA award 2020-67021-32799, the Jump ARCHES endowment through the Health Care Engineering Systems Center, the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign through the NCSA Fellows program, and the IBM-Illinois Discovery Accelerator Institute.

References

[1] R. Aljundi, F. Babiloni, M. Elhoseiny, M. Rohrbach, and T. Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *ECCV*, 2018. 1, 3

[2] R. Aljundi, E. Belilovsky, T. Tuytelaars, L. Charlin, M. Caccia, M. Lin, and L. Page-Caccia. Online continual learning with maximally interfered retrieval. In *NeurIPS*, 2019. 1, 4, 5, 7

[3] R. Aljundi, P. Chakravarty, and T. Tuytelaars. Expert gate: Lifelong learning with a network of experts. In *CVPR*, 2017. 3

[4] R. Aljundi, M. Lin, B. Goujaud, and Y. Bengio. Gradient based sample selection for online continual learning. In *NeurIPS*, 2019. 3, 4, 5, 7

[5] P. Buzzega, M. Boschini, A. Porrello, D. Abati, and S. Calderara. Dark experience for general continual learning: a strong, simple baseline. In *NeurIPS*, 2020. 4, 5

[6] M. Caron, I. Misra, J. Mairal, P. Goyal, P. Bojanowski, and A. Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *NeurIPS*, 2020. 2, 4

[7] H. Cha, J. Lee, and J. Shin. Co2L: Contrastive continual learning. In *ICCV*, 2021. 4, 5

[8] A. Chaudhry, P. K. Dokania, T. Ajanthan, and P. HS Torr. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *ECCV*, 2018. 4, 5, 7

[9] A. Chaudhry, M. Ranzato, M. Rohrbach, and M. Elhoseiny. Efficient lifelong learning with a-gem. In *ICLR*, 2019. 2, 3, 4, 5, 7

[10] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 2, 4, 6

[11] A. Chrysakis and Marie-Francine Moens. Online continual learning from imbalanced data. In *ICML*, 2020. 3

[12] Y. Cong, M. Zhao, J. Li, S. Wang, and L. Carin. GAN memory with no forgetting. In *NeurIPS*, 2020. 1, 3

[13] M. De Lange, R. Aljundi, M. Masana, S. Parisot, X. Jia, A. Leonardis, G. Slabaugh, and T. Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *TPAMI*, 44(7):3366–3385, 2021. 2

[14] M. De Lange and T. Tuytelaars. Continual prototype evolution: Learning online from non-stationary data streams. In *ICCV*, 2021. 1, 2, 3

[15] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 2, 4

[16] D. Ha and D. Eck. A neural representation of sketch drawings. In *ICLR*, 2018. 4

[17] T. L. Hayes, K. Kafle, R. Shrestha, M. Acharya, and C. Kanan. Remind your neural network to prevent catastrophic forgetting. In *ECCV*, 2020. 2, 3

[18] J. He, R. Mao, Z. Shao, and F. Zhu. Incremental learning in online scenario. In *CVPR*, 2020. 2, 3

[19] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 2

[20] G. Hocquet, O. Bichler, and D. Querlioz. Ova-inn: Continual learning with invertible neural networks. In *IJCNN*, 2020. 2, 3, 6

[21] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska, D. Hassabis, C. Clopath, D. Kumaran, and R. Hadsell. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 2017. 1, 3, 4

[22] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. 2009. 4

[23] S. Lee, J. Ha, D. Zhang, and G. Kim. A neural dirichlet process mixture model for task-free continual learning. In *ICLR*, 2020. 1, 2, 3

[24] Z. Li and D. Hoiem. Learning without forgetting. *TPAMI*, 40(12):2935–2947, 2017. 4, 5, 7

[25] D. Lopez-Paz and M. Ranzato. Gradient episodic memory for continual learning. In *NeurIPS*, 2017. 2, 3, 4

[26] Z. Mai, R. Li, H. Kim, and S. Sanner. Supervised contrastive replay: Revisiting the nearest class mean classifier in online class-incremental continual learning. In *CVPR*, 2021. 1, 4, 5, 7

[27] S. Maji, E. Rahtu, J. Kannala, M. Blaschko, and A. Vedaldi. Fine-grained visual classification of aircraft. Technical report, 2013. 4

[28] A. Mallya, D. Davis, and S. Lazebnik. Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In *ECCV*, 2018. 3

[29] D. Maltoni and V. Lomonaco. Continuous learning in single-incremental-task scenarios. *Neural Networks*, 116:56–73, 2019. 3

[30] A. Prabhu, P. HS Torr, and Puneet K. D. GDumb: A simple approach that questions our progress in continual learning. In *ECCV*, 2020. 2, 3, 4, 5, 7, 8

[31] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 2, 4, 5

[32] S.-A. Rebuffi, A. Kolesnikov, G. Sperl, and C. H. Lampert. iCaRL: Incremental classifier and representation learning. In *CVPR*, 2017. 1, 3, 4, 5, 7

[33] D. Rolnick, A. Ahuja, J. Schwarz, T. Lillicrap, and G. Wayne. Experience replay for continual learning. In *NeurIPS*, 2019. 2, 3, 4, 5, 7

[34] A. A. Rusu, N. C Rabinowitz, G. Desjardins, H. Soyer, J. Kirkpatrick, K. Kavukcuoglu, R. Pascanu, and R. Hadsell. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016. 1, 3

[35] H. Shin, J. K. Lee, J. Kim, and J. Kim. Continual learning with deep generative replay. In *NeurIPS*, 2017. 3

[36] G. M. van de Ven, H. T. Siegelmann, and A. S. Tolias. Brain-inspired replay for continual learning with artificial neural networks. *Nature communications*, 11(1):1–14, 2020. 3

[37] G. M. van de Ven and A. S. Tolias. Three scenarios for continual learning. *arXiv preprint arXiv:1904.07734*, 2019. 2, 3

[38] O. Vinyals, C. Blundell, T. Lillicrap, K. Kavukcuoglu, and D. Wierstra. Matching networks for one shot learning. In *NeurIPS*, 2016. 4

[39] M. Wallingford, A. Kusupati, K. Alizadeh-Vahid, A. Walsman, A. Kembhavi, and A. Farhadi. Fluid: A unified evaluation framework for flexible sequential data. 7

[40] P. Welinder, S. Branson, T. Mita, C. Wah, F. Schroff, S. Belongie, and P. Perona. Caltech-ucsd birds 200. 2010. 4

[41] J. Xu and Z. Zhu. Reinforced continual learning. In *NeurIPS*, 2018. 1, 3

[42] Y. Yang, G. Igor, and G. Boris. Large batch training of convolutional networks. *arXiv preprint arXiv:1708.03888*, 2017. 6

[43] J. Zbontar, L. Jing, I. Misra, Y. LeCun, and S. Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *ICML*, 2021. 2, 4

[44] F. Zenke, B. Poole, and S. Ganguli. Continual learning through synaptic intelligence. In *ICML*, 2017. 1, 3