

This CVPR workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Alleviating Representational Shift for Continual Fine-tuning

Shibo Jie Zhi-Hong Deng^{*} Ziheng Li School of Artificial Intelligence, Peking University

{parsley, zhdeng, liziheng}@pku.edu.cn

Abstract

We study a practical setting of continual learning: finetuning on a pre-trained model continually. Previous work has found that, when training on new tasks, the features (penultimate layer representations) of previous data will change, called representational shift. Besides the shift of features, we reveal that the intermediate layers' representational shift (IRS) also matters since it disrupts batch normalization, which is another crucial cause of catastrophic forgetting. Motivated by this, we propose ConFiT, a fine-tuning method incorporating two components, crossconvolution batch normalization (Xconv BN) and hierarchical fine-tuning. Xconv BN maintains pre-convolution running means instead of post-convolution, and recovers post-convolution ones before testing, which corrects the inaccurate estimates of means under IRS. Hierarchical finetuning leverages a multi-stage strategy to fine-tune the pretrained network, preventing massive changes in Conv layers and thus alleviating IRS. Experimental results on four datasets show that our method remarkably outperforms several state-of-the-art methods with lower storage overhead. Code: http://github.com/JieShibo/ConFiT.

1. Introduction

An *artificial neural network* (ANN) well-trained on a task suffers from *catastrophic forgetting* when learning a new task: its performance on already learned tasks drops dramatically after learning a new one [14, 28]. However, in a real-world scenario where re-training on past data may be expensive or infeasible, we expect ANNs to be able to learn continually, acquiring new knowledge without interfering with previously learned skills. This learning paradigm is referred to as *continual learning*.

Meanwhile, *pre-trained models* have already shown overwhelming results in many fields of deep learning, such as *computer vision* and *neural language processing* (NLP) [10,32,45]. By pre-training an ANN on a large-scale dataset

Pre-train | Fine-tune



Figure 1. Continual Fine-tuning on a pre-trained model in multihead (task-incremental) scenario.

and then fine-tuning all or part of the network parameters on the downstream dataset, the general knowledge obtained from large-scale datasets can be transferred to downstream tasks. Due to the effectiveness of pre-training & fine-tuning paradigm, using a pre-trained model as initialization has already been adopted as a generic approach in continual learning of NLP [4, 6, 15, 40].

Whereas in vision, most previous works of continual learning prefer to initialize the network randomly. Other works either directly leverage the pre-trained model as a frozen feature extractor, upon which shallower classifiers or task-specific masks are continually learned [19, 27, 38]; or investigate the performance of their methods when fine-tuning continually on hard fine-grained datasets [1, 7, 41, 43, 44] and/or tough scenarios [26, 43, 44]. Since our ultimate aim is to deploy a high-performance continual learning system in the real world, it is sensible to leverage the pre-trained models and their general knowledge. Therefore, how to fine-tune a network without catastrophic forgetting given its pre-trained parameters is a key issue for designing effective methods in continual learning.

So, what is the crucial cause of catastrophic forgetting when fine-tuning continually? We here follow the multihead setting adopted by some previous works [7, 20, 36], under which each task monopolizes a specific output fullyconnected (FC) layer and the other layers are shared among

^{*}Corresponding author



Figure 2. Fine-tuning (left) and ConFiT (right). After fine-tuning on new tasks, the true moments of previous tasks' intermediate representations are shifted, and the running moment estimates of BN are biased towards new tasks. ConFiT alleviates IRS via hierarchical fine-tuning and corrects running means via Xconv BN.

tasks, as shown in Fig. 1. Ideally, the penultimate layer representations of previous tasks should be kept unchanged when learning new tasks. Otherwise, the distorted representations may not correspond with the output layer, because the output layers of previous tasks are not updated. We refer to this phenomenon as *representational shift*. The shift of the penultimate layer representations has already been discussed by [44], but we will step further here: the representational shift of other intermediate layers in the network is also noteworthy, since it significantly disrupts *Batch Normalization* (BN) [18].

BN is a widely used technique in deep backbone networks such as ResNet [16]. While testing, BN normalizes the intermediate representations using running estimates (means and variances) of the global moments calculated during training. In continual learning, two problems arise when we directly employ BN in networks. First, the running moments will be inevitably biased towards new tasks. Second, even if we keep the running moments unchanged, the representational shift of intermediate layers (i.e., intermediate representational shift, IRS) still leads to changes in the distribution of intermediate representations. Due to these two problems, the BN becomes uncontrollable in continual learning, and thus we cannot normalize intermediate representations accurately when testing on previous tasks. As a result, the model's performance on previous tasks is degraded, which causes catastrophic forgetting.

In this paper, we propose *ConFiT*, a method to address these issues on both sides via two novel techniques, *crossconvolution batch normalization (Xconv BN)* and *hierarchical fine-tuning*, as illustrated in Fig. 2. Xconv BN is designed on the basis of BN, but can provide more accurate running estimates under IRS. For BNs placed after convolutional (Conv) layers, we calculate the pre-convolution running means before the Conv layers during training, and use them to recover the post-convolution running means. Hierarchical fine-tuning prevents massive changes in Conv layers by fine-tuning the network with a multi-stage strategy and thus alleviates IRS.

In summary, our contributions are as follows:

· We reveal that intermediate representational shift

(IRS) deserves to be noticed since it disrupts BN and consequently leads to forgetting. As far as we know, no previous work has presented this issue.

- We propose ConFiT, a novel fine-tuning method that both alleviates IRS and corrects BN estimates.
- Experimental results show that the proposed method remarkably promotes the performance of fine-tuning and outperforms several state-of-the-art baselines with lower additional storage overhead.

2. Related Work

2.1. Continual Learning

Continual learning methods focus on alleviating catastrophic forgetting, which can be categorized into three groups [23]. *Regularization-based* methods measure the importance of parameters in previous tasks, and then impose a regularization term on the loss to prevent important parameters from significant deviation [1, 7, 8, 20, 41, 46]. *Replay-based* methods maintain a limited memory space to store samples from previous tasks, and jointly train them with new data to overcome forgetting [2, 3, 5, 24, 31, 34, 35, 42]. *Parameter-isolation* methods offer distinct parameters for each task, so as to isolate interference among tasks [13, 27, 33, 36].

Pre-trained models such as BERT [10] and GPT2 [32] have already been used in NLP tasks of continual learning [6, 15, 40]. However, in vision tasks, most recent works prefer to start from random initialization, or use a frozen pre-trained model as a feature extractor, which simplifies the problem by only learning the shallow classifier continually. There are only a few works that have made limited attempts to fine-tune pre-trained vision models. For instance, [1, 7, 41] investigate the fine-tuning performance of their methods on hard datasets (e.g., CUB200 [39] and CORe50 [25]), and [26,43,44] use pre-trained initialization under a more difficult class-incremental setting where new classes and patterns emerge over time. Whereas in this paper, we directly study the forgetting on pre-trained model and suggest strategies for fine-tuning continually.

2.2. Batch Normalization

Batch normalization (BN) is widely used in deep neural networks to accelerate training and reduce the sensitivity to initialization [18]. To be specific, in a convolutional neural network, the activation a has a shape of (B, C, H, W), standing for batch size, channel size, height and width, respectively. During training, BN calculates moments as:

$$\boldsymbol{\mu}_{BN} = \frac{1}{BHW} \sum_{b=1}^{B} \sum_{h=1}^{H} \sum_{w=1}^{W} \boldsymbol{a}_{b,:,h,w}$$
(1)

$$\sigma_{BN}^2 = \frac{1}{BHW} \sum_{b=1}^{B} \sum_{h=1}^{H} \sum_{w=1}^{W} \left(a_{b,:,h,w} - \mu_{BN} \right)^2 \quad (2)$$

where $\mu_{BN}, \sigma_{BN}^2 \in \mathbb{R}^C$. Meanwhile, BN maintains momentum-updated running moments to approximate global moments:

$$\boldsymbol{\mu}_{r} \leftarrow \boldsymbol{\mu}_{r} + \eta \left(\boldsymbol{\mu}_{BN} - \boldsymbol{\mu}_{r} \right), \quad \boldsymbol{\sigma}_{r}^{2} \leftarrow \boldsymbol{\sigma}_{r}^{2} + \eta \left(\boldsymbol{\sigma}_{BN}^{2} - \boldsymbol{\sigma}_{r}^{2} \right)$$
(3)

The normalized activation is:

$$a' = \gamma \left(\frac{a-\mu}{\sqrt{\sigma^2 + \epsilon}}\right) + \beta$$
 (4)

where $\boldsymbol{\mu} = \boldsymbol{\mu}_{BN}, \boldsymbol{\sigma} = \boldsymbol{\sigma}_{BN}$ in training, and $\boldsymbol{\mu} = \boldsymbol{\mu}_r, \boldsymbol{\sigma} = \boldsymbol{\sigma}_r$ in testing. $\boldsymbol{\gamma}, \boldsymbol{\beta} \in \mathbb{R}^C$ are trainable affine parameters. All operations are broadcast alone the axes of B, H and W.

An obvious weakness of BN is that the training may fail due to unstable μ_{BN} and σ_{BN}^2 when data batches are small and/or non-i.i.d., especially when the model is confronted with a continual unstable data stream. [26] suggests replacing BN with Batch Renormalization (BRN) [17] so as to bypass this drawback. But this weakness will not be exposed in task-incremental learning where the data batches are i.i.d. within each single task.

Another flaw of BN is that the moments used in testing are estimated with training data, which may lead to trouble when the estimates do not match the true distribution of testing data. This inconsistency becomes more severe in continual learning, which we will discuss comprehensive in Sec. 3.1. For this reason, a novel normalization method has been proposed by a recent study [30] to replace BN in online continual learning. In this paper, we will improve BN in another way to address this issue, which only differs from BN in testing without modifying its normalization behavior in training.

3. Methodology

3.1. Preliminaries & Motivation

In this paper, we consider a fundamental setting where task boundaries are available. The model is supposed to learn a sequence of tasks: $\mathcal{T} = \{\mathcal{T}_1, \mathcal{T}_2, ..., \mathcal{T}_T\}$. At each step t, the model can only have access to training set of $\mathcal{T}_t = \{(\mathbf{x}_i^{tr}, y_i^{tr}, t)\}$. After training, the model will be tested on test set of all learned tasks $\{(\mathbf{x}_i^{te}, y_i^{te}, j) | j = 1, ..., T\}$. Following multi-head setting in previous works, we denote the feature extractor of the model as a mapping f, and the classifier head of task \mathcal{T}_t as another mapping h_t . Then at step t, without any continual learning approaches, we train the network by minimizing the empirical risk $\mathcal{L}_t = \frac{1}{|\mathcal{T}_t|} \sum_{i=1}^{|\mathcal{T}_t|} \mathcal{L}(h_t \circ f(\mathbf{x}_i^{tr}), y_i^{tr})$. As the training goes on, f is evolving all the time, whereas h_t is updated only at step t. This is to say that at t' > t, the representations outputted by f have changed and may no longer correspond with the classifier h_t . We refer to this phenomenon as representational shift of the penultimate layer.

To generalize this phenomenon to other layers, we now consider a more complicated case: a network with BN. We decompose the feature extractor f into $f_{BN} \circ f_1$. Since f_{BN} maintains running moments for testing, at t' > t, the running moments will be flooded by moments of the newest task, and thus the intermediate $f_1(x)$ for $x \in T_t$ will not be normalized correctly during testing.

A trivial solution to this issue is to store specific moments for each task respectively. However, as discussed above, even if the running moments of f_{BN} are taskspecific, the f_1 is still evolving all the time, which also results in IRS — the distribution of intermediate representation $f_1(x)$ for $x \in T_t$ has changed, so that the running moments of previous task T_t have been no longer representative of this distribution. This IRS intensifies the inconsistency between training and testing of BN, and thus leads to catastrophic forgetting.

We now introduce the two components of ConFiT.

3.2. Cross-convolution Batch Normalization

Since BN usually follows a Conv layer, we now consider an intermediate fragment of network composed of a Conv layer and a BN layer: $f_{BN} \circ f_{Conv}$. Firstly, for μ_r , σ_r , β , and γ of f_{BN} , we store specific ones for each task to prevent them from leaning to new tasks. For a previous task, if input activation a of this fragment is fixed, then the IRS disrupting f_{BN} is caused by the change of f_{Conv} . We start our analysis with an interesting relation between the input's and output's mean of a Conv layer. For convenience, we first declare two following operators:

Definition 1 (Average Pooling) Suppose activation a has a shape of (B, C, H, W), then AvgPool(a) has a shape of (C), in which

$$AvgPool(a)_{c} = \frac{1}{BHW} \sum_{b'=1}^{B} \sum_{h'=1}^{H} \sum_{w'=1}^{W} a_{b',c,h',w'}$$
(5)

Definition 2 (Dimension-Preserving Average Pooling)

Suppose activation **a** has a shape of (B, C, H, W), then $AvgPool_{DP}(\mathbf{a})$ also has a shape of (B, C, H, W), in which

$$AvgPool_{DP}(\boldsymbol{a})_{b,c,h,w} = AvgPool(\boldsymbol{a})_c \tag{6}$$

 $AvgPool_{DP}$ is dimension-preserving because it does not change the shape of input as AvgPool does. Note that if we broadcast AvgPool(a) along the axes of B, H and W, we



Figure 3. BN (left) and Xconv BN (right). Moments in blue and red circles are used for training and testing, respectively. μ_r of BN is inaccurate since Conv layer has changed, so Xconv BN calculates the post-convolution μ_r with stable pre-convolution μ'_r before testing.

will get $AvgPool_{DP}(a)$. With these operators, we show an invariance of the Conv layer's output:

Proposition 1 Suppose $f_{Conv}(\cdot)$ denotes a 2D-Conv layer with stride = 1 and padding = kernel_size - 1. Then,

$$AvgPool(f_{Conv}(\boldsymbol{a})) = AvgPool(f_{Conv}(AvgPool_{DP}(\boldsymbol{a})))$$
(7)

Proof and other variants when $stride \neq 1$ or $padding \neq kernel_size - 1$ are in Appendix.

The f_{BN} following f_{Conv} calculates mean as:

$$\boldsymbol{\mu}_{BN} = AvgPool(f_{Conv}(\boldsymbol{a})) \tag{8}$$

From Prop. 1 we know that this mean can also be calculated as:

$$\boldsymbol{\mu}_{BN} = AvgPool(f_{Conv}(AvgPool_{DP}(\boldsymbol{a}))$$
(9)

This is to say that, if we assume the distribution of input *a* is certain, we can always get the accurate post-convolution μ_{BN} in testing by only storing its preconvolution mean AvgPool(a) during training, no matter how the weight of f_{Conv} changes later. However, if we directly store μ_{BN} instead (or use moving average μ_r to approximate global μ_{BN} when batch learning) as BN does, the changed Conv layer will make the stored training μ_{BN} (or μ_r) unequal to true testing μ_{BN} .

On the other hand, if f_{BN} normalizes the mean accurately, its outputs have a certain mean vector — equal to its β parameter. So the pre-convolution mean of the next Conv layer is also certain — exactly this β (if there are no other layers between them). Then as our analysis above, the pre-convolution mean of the next Conv layer will also be certain. Inductively, all BN layers will normalize the means accurately when testing.

But in practice, there are also nonlinear activation functions between a Conv layer and the BN layer ahead of it, so we cannot directly use β of the previous BN as the preconvolution mean. Instead, the proposed Xconv BN maintains a momentum-updated pre-convolution running mean μ'_r instead of f_{BN} 's μ_r . μ'_r is updated using batch preconvolution means $\mu_{Pre} = AvgPool(a)$:

$$\boldsymbol{\mu'}_r \leftarrow \boldsymbol{\mu'}_r + \eta \left(\boldsymbol{\mu}_{Pre} - \boldsymbol{\mu'}_r \right) \tag{10}$$

Before testing, we uses it to calculate the approximate postconvolution global mean of this Conv layer:

$$\boldsymbol{\mu}_{r} = AvgPool(f_{Conv}(Broadcast(\boldsymbol{\mu'}_{r})))$$
(11)

As for σ_r , we will empirically prove that it is not necessary to correct it in Sec. 4.2.3, so we here retain the postconvolution σ_r of BN in Xconv BN. The post-convolution σ_r and recovered post-convolution μ_r are used for normalization in testing, as demonstrated in Fig. 3. Besides, we also store task-specific μ'_r like f_{BN} 's μ_r . Note that Xconv BN just modifies the running moments of BN, so Xconv BN still uses post-convolution μ_{BN} and σ_{BN} in training just as BN does.

3.3. Hierarchical Fine-tuning

Xconv BN adjusts the running means to match the shifted representations, i.e., adapts f_{BN} to changed f_{Conv} , rather than correcting the intermediate representational shift itself, which remains an issue.

Firstly, we give an intuition to prevent IRS during finetuning. In a common fine-tuning manner, the classification head is randomly initialized and jointly trained with the pretrained feature extractor. However, since the random head may be far from its optimal value, the loss will be large at the beginning, which also misleads the feature extractor and forces it to change a lot. If we just fine-tune on a single task, it will not cause a problem. But if we fine-tune continually, the massively changed feature extractor may lead to IRS on previous tasks.

To formulate this intuition, we now analyze a simple overparametrized model with linear feature extractor following previous work [22, 37]. Consider a regression task $y = v^{\top} B x$, in which $v \in \mathbb{R}^k$ is classifier and $B \in \mathbb{R}^{k \times d}$ is feature extractor, and the number of inputs is n which satisfies $1 \le k < n < d$. We use v_t^* and B_t^* to denote a zero-loss minimum point of task \mathcal{T}_t . The MSE loss of task \mathcal{T}_t is denoted as $\mathcal{L}_t(B, v)$. For a previous task $\mathcal{T}_{t'}$, since its data could be arbitrary, we consider the worst case in which $\mathcal{L}_{t'}(B, v) = \max_{\|x\| \le 1} \left(v^{\top} B x - v_{t'}^{*\top} B_{t'}^* x\right)^2$. For simplicity, we first analyze the loss of previous tasks in singlehead setting, i.e., the model is initialized with (B_{t-1}^*, v_{t-1}^*) and finally updated to (B_t^*, v_t^*) . We here give a proposition which is a deduction of the theorem in [22].



Figure 4. Hierarchical fine-tuning. Colored layers are trainable and dashed layers are task-specific.

Proposition 2 Let $X = \{x | x \in \mathcal{T}_t\}$ be the training data in \mathcal{T}_t , S and \mathcal{R} be the orthogonal basis of SpanSpace(X)and $RowSpace(\mathbf{B}_{t-1}^*)$ respectively, and $(\mathbf{B}_*, \mathbf{v}_*)$ be the optimal parameters on both current task \mathcal{T}_t and previous task $\mathcal{T}_{t'}$. If $\sigma_k (\mathcal{R}^\top S^\perp) > 0$, after fine-tuning on \mathcal{T}_t , the loss on $\mathcal{T}_{t'}$ is lower bounded as

$$\sqrt{\mathcal{L}_{t'}\left(\boldsymbol{B}_{t}^{*},\boldsymbol{v}_{t}^{*}\right)} \geq \frac{\sigma_{k}\left(\mathcal{R}^{\top}\mathcal{S}^{\perp}\right)}{\sqrt{k}} \frac{\min(\phi,\phi^{2}/\|\boldsymbol{B}_{*}\boldsymbol{v}_{*}\|_{2})}{\left(1+\|\boldsymbol{B}_{*}\boldsymbol{v}_{*}\|_{2}\right)^{2}} - \epsilon$$
(12)

where σ_k denotes the k-th largest singular value, $\phi^2 = \left| \left(\boldsymbol{v}_{t-1}^* \boldsymbol{v}_* \right)^2 - \left(\boldsymbol{v}_*^\top \boldsymbol{v}_* \right)^2 \right|$ is the alignment between \boldsymbol{v}_{t-1}^* and \boldsymbol{v}_* , and $\epsilon = \min_{\boldsymbol{U}} \left\| \boldsymbol{B}_{t-1}^* - \boldsymbol{U} \boldsymbol{B}_* \right\|_2^2$ is the distance between \boldsymbol{B}_{t-1}^* and \boldsymbol{B}_* under a rotation.

This proposition can be easily generalized to multi-head setting, which will be introduced in Appendix. This proposition indicates that the loss of previous tasks is lowerbounded by the alignment between the initial and optimal classification head. Moreover, the bound is tighter when the feature extractor is better, such as pre-trained as we use. It indicates that it is important to use a "good" initialization for the classification head when using a pre-trained feature extractor.

Along this line of thinking, we propose to fine-tune the classification head first (a.k.a. linear probing) before fine-tuning the entire network, which provides a well-initialized head. To be more specific, when task \mathcal{T}_t arrives, we freeze the feature extractor B_{t-1}^* and only fine-tune the head to v_t^{lp} . Then the model is initialized with (B_{t-1}^*, v_t^{lp}) and fine-tune to the optimal point (B_t^*, v_t^*) on \mathcal{T}_t . We show that if pre-trained B_0^* is "good enough", fine-tuning on new task with this proposed strategy will not degrade the performance on previous tasks in multi-head setting.

Proposition 3 If every task $\mathcal{T}_{t'}$ with $t' \leq t$ satisfies: (i) while training on this task, v is initialized with $v_{t'}^{lp}$, and (ii)

there exists v_0 such that $\mathcal{L}_{t'}(B_0^*, v_0) = \mathcal{L}_{t'}(B_{t'}^*, v_{t'}^*)$ (i.e., B_0^* is good enough), then for all task $\mathcal{T}_{t'}$ with $t' \leq t$,

$$\mathcal{L}_{t'}(\boldsymbol{B}_{t}^{*}, \boldsymbol{v}_{t'}^{*}) = \mathcal{L}_{t'}(\boldsymbol{B}_{t'}^{*}, \boldsymbol{v}_{t'}^{*})$$
(13)

Generalized from the overparametrized model, we can fine-tune ordinary deep networks with a similar strategy. Since Xconv BN stores task-specific moment estimates μ , σ and affine parameters γ , β for each task, we can also regard a Xconv BN layer as a "head" of a network fragment $f_{BN} \circ f_{Conv}$, and hope that the Xconv BN layers are well-initialized as well, so as to reduce changes of f_{Conv} . Consequently, we also fine-tune the Xconv BN layers before fine-tuning the whole network, which finally raises a three-step fine-tuning strategy illustrated in Fig. 4: when a new task arrives, we first fine-tune the classification head, and then the head as well as Xconv BNs, and finally all parameters.

4. Experiments

4.1. Setup

4.1.1 Datasets

We evaluate our method on four datasets: CIFAR100 [21], CUB200 [39], Caltech101 [12], and Flowers102 [29]. CI-FAR100 is the most widely used dataset in continual learning, which we also adopt in our experiments. However, CIFAR100 is somewhat easy for a pre-trained model, so we also use three harder datasets: CUB200, which is finegrained; Caltech101, which is unbalanced; and Flowers102, both fine-grained and unbalanced.

We randomly divide CIFAR100 into 20 tasks where each has 5 classes, and divide CUB200 into 10 tasks where each has 20 classes. For Caltech101 and Flowers102, 100 classes of each dataset are randomly selected and evenly split into 10 tasks. We do not involve Imagenet [9] as a benchmark dataset since it has already been leveraged for pre-training in our experiments. For a fair and clear comparison, the task partitioning of each dataset, as well as its train-test split, is fixed in all runs.

4.1.2 Baselines

Since our method does not need episodic memory to store samples and thus is orthogonal to replay-based methods, we compare our method with baselines that also do without episodic memory, including regularization-based methods: EWC [20], SI [46], RWalk [8], MAS [1], CPR [7], and AFEC [41]; and a parameter-isolation method: CCLL [36]. We also report two reference results: single-task learning (STL), where each task is fine-tuned on an individual pretrained network, and linear probing (LP), where the pretrained feature extractor is frozen. Note that STL is not a continual learning method.

Methods	CIFAR100 (T = 20)		CUB200 (T = 10)		Caltech101 (T = 10)		Flowers102 $(T = 10)$	
	ACC ↑	$\mathrm{FGT}\downarrow$	ACC \uparrow	$FGT\downarrow$	ACC \uparrow	$\mathrm{FGT}\downarrow$	ACC \uparrow	$\mathrm{FGT}\downarrow$
Fine-tuning	$87.43_{\pm 0.78}$	$7.11_{\pm 0.92}$	$71.76_{\pm 0.58}$	$14.76_{\pm 0.66}$	$71.89_{\pm 0.86}$	$17.24_{\pm 1.11}$	$78.06_{\pm 1.29}$	$10.93_{\pm 1.16}$
EWC [20]	$88.04_{\pm 0.59}$	$6.40_{\pm0.58}$	$79.30_{\pm0.43}$	$7.20_{\pm0.52}$	$76.10_{\pm 1.12}$	$12.35_{\pm 1.46}$	$79.94_{\pm0.68}$	$8.56_{\pm0.56}$
SI [46]	$88.20_{\pm 0.44}$	$6.42_{\pm0.46}$	$77.94_{\pm 0.53}$	$8.42_{\pm 0.51}$	$76.26_{\pm 1.32}$	$12.28_{\pm 1.10}$	$79.61_{\pm 0.62}$	$8.77_{\pm 0.36}$
RWalk [8]	$87.75_{\pm 0.74}$	$6.82_{\pm0.78}$	$77.27_{\pm 0.64}$	$9.24_{\pm 0.64}$	$76.14_{\pm 1.15}$	$12.51_{\pm 1.03}$	$79.48_{\pm0.61}$	$9.06_{\pm 0.68}$
MAS [1]	$90.59_{\pm 0.35}$	$3.61_{\pm0.28}$	$83.64_{\pm 0.39}$	$2.42_{\pm0.44}$	$81.76_{\pm 1.60}$	$6.18_{\pm0.96}$	$80.90_{\pm 0.82}$	$6.60_{\pm 0.76}$
CPR [7]	$91.17_{\pm 0.21}$	$2.89_{\pm0.16}$	$83.64_{\pm 0.36}$	$2.49_{\pm0.36}$	$81.85_{\pm 1.47}$	$5.68_{\pm0.86}$	$80.97_{\pm 0.90}$	$6.53_{\pm0.92}$
AFEC [41]	$90.68_{\pm 0.11}$	$3.55_{\pm0.04}$	$83.70_{\pm 0.31}$	$2.60_{\pm0.35}$	$82.55_{\pm 0.89}$	$5.34_{\pm0.72}$	$81.26_{\pm 0.64}$	$6.38_{\pm0.57}$
ConFiT (Ours)	$\textbf{92.02}_{\pm 0.21}$	$\textbf{2.72}_{\pm 0.29}$	$\textbf{87.43}_{\pm 0.16}$	$1.31_{\pm0.38}$	$\textbf{88.73}_{\pm 0.34}$	$0.67_{\pm 0.27}$	$\textbf{86.99}_{\pm 0.25}$	$\textbf{1.79}_{\pm 0.49}$
CCLL [36]	77.59	0.00	44.66	0.00	63.18	0.00	64.13	0.00
LP	$91.48_{\pm 0.12}$	0.00	$86.03_{\pm 0.08}$	0.00	$77.07_{\pm 0.40}$	0.00	$71.14_{\pm 0.72}$	0.00
STL	$93.50_{\pm0.11}$	0.00	$86.60_{\pm 0.15}$	0.00	$86.65_{\pm 0.69}$	0.00	$85.96_{\pm 0.50}$	0.00

Table 1. Experimental results on benchmark datasets. Methods in the upper part permit update of shared parameters among tasks. For all methods except CCLL, we report results averaged over 5 runs and their 95% confidence intervals. \uparrow and \downarrow stand for higher is better and lower is better, respectively.



Figure 5. Average Accuracy on all already learned tasks after learning each task. We report results averaged over 5 runs.

4.1.3 Implementation Details

In all experiments, including all baselines, LP, STL, and ConFiT, we use a ResNet18 pre-trained on ImageNet as initialization. For all regularization-based baselines and our ConFiT, we use a SGD optimizer with a learning rate of 0.01, and train the network with a mini-batch size of $\{128,$ 32, 32, 32} for {10, 50, 20, 10} epochs on each task of {CIFAR100, CUB200, Caltech101, Flowers102}, respectively. For CPR and AFEC which are plug-in approaches, we report the performance of MAS-CPR and MAS-AFEC, because MAS performs the best among {EWC, SI, RWalk, MAS}. As for our hierarchical fine-tuning, we allocate 20%, 30% and 50% of the total epochs to the three stages, respectively. For parameter-isolation CCLL, we follow the setting in its original paper that trains the network for 150 epochs using SGD with an initial learning rate of 0.01, and multiplies the learning rate by 0.1 at 50, 100 and 125 epochs. All experiments are in multi-head setting, i.e., taskincremental scenario.

4.1.4 Evaluation Metrics

We measure the performance with two metrics, as in previous works [7, 8]: Average Accuracy (ACC) and Average Forgetting (FGT). We denote the accuracy on task T_i after training the model on task T_j as a_{ij} , and task T_T is the final task. Then ACC and FGT can be calculated as follows:

$$ACC = \frac{1}{T} \sum_{i=1}^{T} a_{iT} \tag{14}$$

$$FGT = \frac{1}{T-1} \sum_{i=1}^{T-1} \max_{i \le j \le T} (a_{ij} - a_{iT})$$
(15)

4.2. Results & Discussion

4.2.1 Comparisons with Baselines

The main results are shown inTab. 1. MAS performs better than EWC, SI, and RWalk on all datasets, which is reasonable since MAS will not underestimate the importance of parameters for pre-trained knowledge, as [1] indicates. Plug-in methods CPR and AFEC do not provide much improvement on MAS in this setting. CCLL performs poorly since the added calibration modules significantly change the structure of the pre-trained network and thus disrupt the pretrained knowledge.

On all the four datasets, ConFiT has higher ACC than all continual learning baselines, and lower FGT than all



Figure 6. Additional parameters to be stored for each method. ConFiT uses orders of magnitude fewer additional parameters than other baselines.

Ablation (Cases	Flowers102	Caltech101		
Fine-tuning	BN	$78.06{\pm}1.29$	$71.89{\pm}0.86$		
Hierarchical FT	BN	$83.39{\pm}0.43$	$85.12{\pm}0.34$		
Fine-tuning	Xconv BN	$81.26{\pm}0.72$	$82.09{\pm}0.59$		
Hierarchical FT	Xconv BN	$86.99{\pm}0.25$	$\textbf{88.73}{\pm}\textbf{0.34}$		

Table 2. Ablation study for ConFiT. We report ACC on Flowers102 and Caltech101.

regularization-based methods. Furthermore, ConFiT outperforms STL on CUB200, Caltech101, and Flowers102. Since STL uses an individual network for each task, it can achieve zero forgetting, but will not benefit from task-level knowledge transfer. It indicates that ConFiT also has the ability to share knowledge among tasks. On CUB200 and CIFAR100, LP even outperforms many baselines and has good results close to STL. This reflects the great ability of the pre-trained model — it has accumulated almost enough knowledge during pre-training. Whereas ConFiT still performs better than LP, which indicates the necessity to obtain more task-specific knowledge via fine-tuning.

In Fig. 6 we illustrate that ConFiT uses orders of magnitude fewer additional parameters than other baselines. Methods like EWC need to store regularization weights for each parameter (11.2M for ResNet18). In addition to these, AFEC needs to store an extra network and its regularization weights. CCLL stores parameters of calibration modules (0.17M/task). ConFiT, on the other hand, only stores Xconv BN parameters for each task (0.02M/task). It is worth noting that 0.02M/task is an extremely tiny requirement for continual learning, since the episodic memory will take up about 0.15M/task even if only a single $224 \times 224 \times 3$ image/task is stored for replay-based methods.

4.2.2 Ablation Study

We provide an ablation study on Flowers102 and Caltech101. As shown in Tab. 2, both hierarchical fine-tuning and Xconv BN can boost the average accuracy significantly. ConFiT combines both of them and achieves the best performance.



Figure 7. Δ_1 and $\Delta_2 - \Delta_0$ of the 20 Conv layers in ResNet18 (three of which are in shortcuts) on Flowers102. Layer 0 is closest to input. Lower Δ_1 and $\Delta_2 - \Delta_0$ are better.

Moreover, to verify our motivation — ConFiT alleviates IRS and corrects the BN running estimates, we record these statistics for output of each Conv layer:

- μ_{te}^i : the true mean on test set of \mathcal{T}_1 after training on task \mathcal{T}_i ;
- μ_r^i : the post-convolution running mean of BN/XConv BN of \mathcal{T}_1 after training on \mathcal{T}_i .

Note that for BN the running mean is shared across tasks, whereas for Xconv BN it is specific for T_1 . All the means are calculated along the axes of *BHW* as BN does. Based on them, we calculate the following values on Flowers102:

- $\Delta_1 = \|\boldsymbol{\mu}_{te}^1 \boldsymbol{\mu}_{te}^{10}\|_2$: the shift of the mean of \mathcal{T}_1 's intermediate representations;
- $\Delta_2 = \|\boldsymbol{\mu}_r^{10} \boldsymbol{\mu}_{te}^{10}\|_2$: the gap between the running mean and true mean of \mathcal{T}_1 when the training is over.

Note that the running estimates are not absolutely precise even with Xconv BN, since Xconv BN inherits the intrinsic inconsistency between running and true moments of BN. For a clear comparison, we also calculate this intrinsic inconsistency as

• $\Delta_0 = \|\boldsymbol{\mu}_r^1 - \boldsymbol{\mu}_{te}^1\|_2$: the gap between the running mean and true mean of \mathcal{T}_1 when the network has only been trained on \mathcal{T}_1 , which is a lower bound of Δ_2 .

We conduct experiments on ablation cases where hierarchical fine-tuning and Xconv BN are not used, and report the Δ_1 and $\Delta_2 - \Delta_0$ of each layer.



Figure 8. Results when using transductive moment estimates. "R-" and "T-" stand for "running" and "transductive", respectively. The R-mean has been corrected by Xconv BN.

Fig. 7 shows that, for the majority of layers, IRS has been alleviated by hierarchical fine-tuning (lower Δ_1), and error of estimation has been reduced by Xconv BN (lower $\Delta_2 - \Delta_0$). Moreover, the IRS of the final Conv layer's output also decreases via hierarchical fine-tuning, which means the representational shift of features is also reduced. We also find that Xconv BN is more effective on the shallow layers (close to input), where the estimation error of Xconv BN is almost 0. This is because the estimation error may accumulate layer by layer, which makes the inputs' means of deep layers not be fixed, potentially violating the assumption in Sec. 3.2.

4.2.3 Is Running Variance Correction Necessary?

Xconv BN maintains pre-convolution running means instead of post-convolution, but it still retains the postconvolution running variances like BN. A straightforward reason for using post-convolution running variances is that the variance is a nonlinear secondary moment, so it is difficult to find a feasible mapping between pre- and postconvolution variances under a linear Conv layer.

To figure out whether running variance correction is necessary, we explore the best case of correction: using true moments to normalize. We refer to this ideal setting as "transductive" since the model has access to the whole test set to determine parameters. As shown in Fig. 8, if the mean and variance are both transductive, the normalization is accurate, so the model outperforms ConFiT as expected, whose mean and variance are both running. If only the mean or only the variance is transductive, the performance is also improved. However, the improvement brought about by transductive variance is somewhat limited (less than 1%) compared with the large improvement that ConFiT has already provided, not to mention that the transductive results are just upper bounds in the ideal cases. Therefore, we argue that running variance correction is insignificant for Xconv BN.

4.2.4 Performance on Deeper Networks

Since our experiments are conducted on ResNet18, we here deploy ConFiT on two other networks: ResNet34 and



Figure 9. ACC on ResNet34 and ResNet50.

ResNet50 [16], which are much deeper than ResNet18. These networks also use BN after Conv layers, so our Con-FiT can be directly applied. Since deeper networks converge more slowly, we train each network for 50 epochs on Caltech101 and Flowers102, more than that for ResNet18.

As shown in Fig. 9, when directly fine-tuning the network, both ResNet34 and ResNet50 perform poorly, and the deeper ResNet50 even suffers from more forgetting. ConFiT boosts the performance of both ResNet34 and ResNet50, in which the ACC of ResNet50 with ConFiT is higher than that of ResNet34. This illustrates that ConFiT has great scalability and can be applied to deeper networks according to actual demand.

5. Limitations

Firstly, the proposed hierarchical fine-tuning is based on a strong assumption that the feature extractor has already been good enough before continual fine-tuning. This assumption makes it hard to generalize this method to a randomly initialized network, or to cases when the tasks are too complicated for the pre-trained model to handle well.

Secondly, Xconv BN needs to store a small number of task-specific parameters, so the task identifiers are necessary in testing for selecting the parameters of Xconv BN. Consequently, Xconv BN cannot be used for class-incremental learning and task-agnostic scenarios directly. Xconv BN can also not be applied to models without BN, such as Vision Transformers [11].

6. Conclusion

In this paper, we focus on fine-tuning continually on pretrained models. We reveal a crucial cause of catastrophic forgetting — the IRS distorts BN. From this perspective, we propose ConFiT, which corrects running means via Xconv BN and alleviates IRS via hierarchical fine-tuning. These two components are proven to be effective both empirically and theoretically. Our ConFiT achieves superior performance, surpassing current SOTA regularization-based and parameter-isolation methods on several datasets.

References

- Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *ECCV*, 2018. 1, 2, 5, 6
- [2] Rahaf Aljundi, Eugene Belilovsky, Tinne Tuytelaars, Laurent Charlin, Massimo Caccia, Min Lin, and Lucas Page-Caccia. Online continual learning with maximal interfered retrieval. In *NeurIPS*, 2019. 2
- [3] Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient based sample selection for online continual learning. In *NeurIPS*, 2019. 2
- [4] Magdalena Biesialska, Katarzyna Biesialska, and Marta R. Costa-jussà. Continual lifelong learning in natural language processing: A survey. In Donia Scott, Núria Bel, and Chengqing Zong, editors, COLING, 2020. 1
- [5] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: a strong, simple baseline. In *NeurIPS*, 2020. 2
- [6] Pengfei Cao, Yubo Chen, Jun Zhao, and Taifeng Wang. Incremental event detection via knowledge consolidation networks. In *EMNLP*, 2020. 1, 2
- [7] Sungmin Cha, Hsiang Hsu, Taebaek Hwang, Flávio P. Calmon, and Taesup Moon. CPR: classifier-projection regularization for continual learning. In *ICLR*, 2021. 1, 2, 5, 6
- [8] Arslan Chaudhry, Puneet Kumar Dokania, Thalaiyasingam Ajanthan, and Philip H. S. Torr. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In ECCV, 2018. 2, 5, 6
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 5
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, 2019.
 1, 2
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 8
- [12] Li Fei-Fei, Robert Fergus, and Pietro Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. *Comput. Vis. Image Underst.*, 2007. 5
- [13] Chrisantha Fernando, Dylan Banarse, Charles Blundell, Yori Zwols, David Ha, Andrei A. Rusu, Alexander Pritzel, and Daan Wierstra. Pathnet: Evolution channels gradient descent in super neural networks. *CoRR*, abs/1701.08734, 2017. 2
- [14] Ian J Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. In *ICLR*, 2014. 1
- [15] Xu Han, Yi Dai, Tianyu Gao, Yankai Lin, Zhiyuan Liu, Peng Li, Maosong Sun, and Jie Zhou. Continual relation learning

via episodic memory activation and reconsolidation. In ACL, 2020. 1, 2

- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In CVPR, 2016. 2, 8
- [17] Sergey Ioffe. Batch renormalization: Towards reducing minibatch dependence in batch-normalized models. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *NIPS*, 2017. 3
- [18] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015. 2
- [19] Ronald Kemker and Christopher Kanan. Fearnet: Braininspired model for incremental learning. In *ICLR*, 2018. 1
- [20] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, et al. Overcoming catastrophic forgetting in neural networks. *PNAS*, 2017. 1, 2, 5, 6
- [21] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. *Computer Science Department, University of Toronto, Tech. Rep*, 1, 2009. 5
- [22] Ananya Kumar, Aditi Raghunathan, Robbie Matthew Jones, Tengyu Ma, and Percy Liang. Fine-tuning can distort pretrained features and underperform out-of-distribution. In *ICLR*, 2022. 4
- [23] Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Alevs. Leonardis, Gregory G. Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *IEEE transactions on pattern* analysis and machine intelligence, 2021. 2
- [24] Matthias De Lange and Tinne Tuytelaars. Continual prototype evolution: Learning online from non-stationary data streams. *CoRR*, abs/2009.00919, 2020. 2
- [25] Vincenzo Lomonaco and Davide Maltoni. Core50: a new dataset and benchmark for continuous object recognition. In *CoRL*, 2017. 2
- [26] Vincenzo Lomonaco, Davide Maltoni, and Lorenzo Pellegrini. Rehearsal-free continual learning over small non-i.i.d. batches. In *CVPR Workshops*, 2020. 1, 2, 3
- [27] Arun Mallya, Dillon Davis, and Svetlana Lazebnik. Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *ECCV*, 2018. 1, 2
- [28] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*. Elsevier, 1989. 1
- [29] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *ICVGIP*, 2008. 5
- [30] Quang Pham, Chenghao Liu, and Steven HOI. Continual normalization: Rethinking batch normalization for online continual learning. In *ICLR*, 2022. **3**
- [31] Ameya Prabhu, Philip H. S. Torr, and Puneet K. Dokania. Gdumb: A simple approach that questions our progress in continual learning. In *ECCV*, 2020. 2

- [32] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2018. 1, 2
- [33] Jathushan Rajasegaran, Munawar Hayat, Salman H. Khan, Fahad Shahbaz Khan, and Ling Shao. Random path selection for continual learning. In *NeurIPS*, 2019. 2
- [34] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In CVPR, 2017. 2
- [35] Anthony V. Robins. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connect. Sci.*, 7(2), 1995. 2
- [36] Pravendra Singh, Vinay Kumar Verma, Pratik Mazumder, Lawrence Carin, and Piyush Rai. Calibrating cnns for lifelong learning. In *NeurIPS*, 2020. 1, 2, 5, 6
- [37] Nilesh Tripuraneni, Michael I. Jordan, and Chi Jin. On the theory of transfer learning: The importance of task diversity. In *NeurIPS*, 2020. 4
- [38] Gido M. van de Ven, Zhe Li, and Andreas S. Tolias. Classincremental learning with generative classifiers. In CVPR Workshops, 2021. 1
- [39] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The caltech-ucsd birds-200-2011 dataset. Technical report, California Institute of Technology, 2011. 2, 5
- [40] Hong Wang, Wenhan Xiong, Mo Yu, Xiaoxiao Guo, Shiyu Chang, and William Yang Wang. Sentence embedding alignment for lifelong relation extraction. In *NAACL-HLT*, 2019.
 1, 2
- [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 *NeurIPS*, 2021. 1, 2, 5, 6
- [42] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *CVPR*, 2019. 2
- [43] Haiyan Yin, peng yang, and Ping Li. Mitigating forgetting in online continual learning with neuron calibration. In *NeurIPS*, 2021. 1, 2
- [44] Lu Yu, Bartlomiej Twardowski, Xialei Liu, Luis Herranz, Kai Wang, Yongmei Cheng, Shangling Jui, and Joost van de Weijer. Semantic drift compensation for class-incremental learning. In *CVPR*, 2020. 1, 2
- [45] Matthew D. Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *ECCV*, 2014. 1
- [46] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *ICML*, 2017. 2, 5, 6