



# DLCFT: Deep Linear Continual Fine-Tuning for General Incremental Learning

Hyounguk Shon<sup>1\*</sup>, Janghyeon Lee<sup>2</sup>, Seung Hwan Kim<sup>2</sup>, and Junmo Kim<sup>1</sup>

<sup>1</sup> Korea Advanced Institute of Science and Technology

<sup>2</sup> LG AI Research

{hyounguk.shon, junmo.kim}@kaist.ac.kr

{janghyeon.lee, sh.kim}@lgresearch.ai

**Abstract.** Pre-trained representation is one of the key elements in the success of modern deep learning. However, existing works on continual learning methods have mostly focused on learning models incrementally from scratch. In this paper, we explore an alternative framework to incremental learning where we *continually fine-tune* the model from a pre-trained representation. Our method takes advantage of linearization technique of a pre-trained neural network for simple and effective continual learning. We show that this allows us to design a linear model where quadratic parameter regularization method is placed as the optimal continual learning policy, and at the same time enjoying the high performance of neural networks. We also show that the proposed algorithm enables parameter regularization methods to be applied to class-incremental problems. Additionally, we provide a theoretical reason why the existing parameter-space regularization algorithms such as EWC underperform on neural networks trained with cross-entropy loss. We show that the proposed method can prevent forgetting while achieving high continual fine-tuning performance on image classification tasks. To show that our method can be applied to general continual learning settings, we evaluate our method in data-incremental, task-incremental, and class-incremental learning problems.

**Keywords:** Continual Learning, Incremental Learning

## 1 Introduction

The ability to incrementally accumulate knowledge from a sequence of datasets is a crucial functionality that modern AI systems require. It is well known that deep neural networks suffer from significant performance degradation when the learning is done sequentially. Such phenomena is referred as catastrophic forgetting (CF), which continual learning aims to address.

Transfer learning is one of the key contributing elements in the recent success of deep learning across various applications from visual to linguistic tasks. When dealing with visual signals, neural nets are often pre-trained on large

---

\* Work done during an internship at LG AI Research.




Fig. 1: **Proposed continual fine-tuning framework.** The model is first trained on a large dataset (e.g., ImageNet) to obtain general representation. At incremental learning phase, the model is continually fine-tuned and accumulates knowledge by learning from the sequentially arriving batches of dataset.

datasets such as ImageNet [39] before training on the target task, which often brings significant performance boost. Although fine-tuning from a pre-trained representation is a standard practice adopted in a lot of modern deep learning applications, many of the existing approaches to continual learning assume a scenario where one needs to begin training from a randomly initialized model. In this work, we explore a practical alternative continual learning framework based on pre-training of neural network, coined *continual fine-tuning*.

Existing works on continual learning can be largely categorized into three groups: model regularization, data rehearsal, and parameter isolation [7]. Regularization-based methods aim to penalize the update either in function space [8, 24] or in parameter space [2, 9, 18, 20, 22, 23, 32, 38, 45]. Parameter isolation methods [29, 44] update the architecture in order to isolate the knowledge learned from each task in order to prevent forgetting. Rehearsal-based approaches [5, 26, 34, 36] replay examples from previous tasks either by storing samples to an external memory buffer or learning a generative model. The rehearsal method has shown to be effective at regularization with a small extra cost.

Parameter regularization methods based on Fisher information matrix aim to represent the source task objective by second-order Taylor approximation, which is typically too expensive due to the quadratic memory cost of the Hessian matrix. Naturally, existing works have focused on efficient representation of the matrix via diagonal approximation [18, 25] or Kronecker factorization [22, 38]. Our work is motivated by a question that has been relatively overlooked in prior works: “Does better Hessian approximation improve continual learning?”

Although parameter regularization approaches are founded on a principled framework that computes the importance of the weights for each task, they have shown relatively underwhelming performance compared to rehearsal-based approaches. One of the main roadblocks of the regularization methods is to address the problems coming from their high parameter dimension and non-linearity, which has left continual learning a particularly challenging problem to tackle. In this work, we show that it is possible to bring a significant boost to regularization methods by a simple modification on the loss function and reparametrization of the model. The remaining portion of the paper is organized as follows: In Section 2, we provide background on continual learning methods,

second-order derivative of neural networks, and linearization of neural networks. In Section 3.1, we provide reasoning on why existing parameter regularization methods underperform in mitigating catastrophic forgetting. In Section 3.2, we describe the proposed method, Deep Linear Continual Fine-tuning (DLCFT). In Section 3.3, we elaborate on how the proposed regularization method can be applied to the class-incremental problem where existing parameter regularization methods have been unable to be applied. In Section 4, we show the evaluation and analysis of the proposed method on data-/task-/class-incremental learning problems. Finally, in Section 5, we conclude the paper.

The main contributions of this work are summarized as follows:

- Instead of incrementally training a neural network from scratch, we propose an alternative approach to continual learning that leverages pre-trained representations, coined *continual fine-tuning*. We show that our approach introduces a novel method for simple and practical continual learning in deep learning.
- To continually adapt to a sequence of downstream tasks, we utilize pre-trained neural network through decomposing the model into nonlinear and linear components by linearization. We propose a learning algorithm that combines linearization and mean squared error loss that significantly boosts the effectiveness of quadratic weight regularization methods. Further, we provide justification on why linearization is the key component for a principled approach to optimal continual learning.
- Our method can be universally applied to various continual learning scenarios where new batches of data, task, or class are observed sequentially. Although data-incremental learning is an important open challenge to practical deep learning, relatively little attention has been given from the community. Notably, we show that our method can effectively learn in data-incremental scenario where batches of new data samples are observed sequentially. Additionally, we demonstrate our method in task-incremental and class-incremental learning scenarios with a small memory buffer.

## 2 Preliminaries

### 2.1 Continual learning

Existing approaches to continual learning can be largely grouped into three families: regularization-based, rehearsal-based, and parameter isolation methods. Regularization method incorporates additional training objectives to prevent model from changing too much during training. One can regularize the outputs of the model to mitigate forgetting. Learning without forgetting (LwF) [24] is one of the early works that utilizes target task data as a surrogate for the source task samples. This method shows reasonably good performance when the source task domains and target task domains are similar, however is less effective when the task domains are dissimilar.

Parameter regularization methods aim to estimate the importance of parameters and use that as the prior for the parameters during training of the subsequent tasks. This was first adopted by elastic weight consolidation (EWC) [18] where the authors proposed to use the diagonal entries of Fisher information matrix. Online-structured Laplace approximation (OSLA) [38] used Kronecker-factored approximate curvature (K-FAC) [30] to incorporate off-diagonals of the Fisher information matrix. With neural networks having intra-batch dependency due to batch normalization, Extended K-FAC (XK-FAC) [22] generalized the method to take intra-batch dependency into account.

Rehearsal-based methods maintain a buffer that stores a small number of samples, and replay the examples in the training of the subsequent tasks. One of the pioneering methods is iCaRL [36] which keeps track of per-class samples for class-incremental learning problems. Some other line of works such as GEM [26] and A-GEM [5] performs constrained optimization by projecting the gradients so that they do not interfere with the previous tasks. DER [4] proposed to distill from the logits that are sampled and stored throughout the training. Parameter isolation methods [3, 29, 40, 42] are based on dynamically allocating a set of parameters for each task so that the training does not interfere with each other.

## 2.2 Second-order derivatives of neural network

Computing and storing the Hessian matrix of neural networks is a difficult challenge due to their high parameter dimension. Hessian matrix of a probabilistic model can be approximated using Fisher information matrix (FIM) [33]. FIM can be interpreted as the second derivative of the KL-divergence between the model and target distribution, and can be efficiently estimated through Monte Carlo method.

Because the memory cost of the full Hessian matrix is quadratic to the size of the parameters, approximation or factorization technique is necessary to handle the matrix. One popular method is diagonal approximation of the Hessian, which neglects the influences of the off-diagonal components. Another approach that can consider off-diagonal influences is block-diagonal approximation, which considers the correlations among intra-layer parameters. K-FAC [12, 30] proposed to approximately factorize the block diagonals of a neural network’s Fisher information matrix into a Kronecker product of covariance matrices:

$$\begin{aligned}
 F^l &= \mathbb{E}_{y \sim p(y|x)} \left[ \nabla_{\theta} \log p(y|x) \nabla_{\theta} \log p(y|x)^{\top} \right] \\
 &= \mathbb{E}[(g \otimes a)(g \otimes a)^{\top}] = \mathbb{E}[(gg^{\top}) \otimes (aa^{\top})] \\
 &\approx \mathbb{E}[gg^{\top}] \otimes \mathbb{E}[aa^{\top}],
 \end{aligned} \tag{1}$$

where  $F^l$  is the FIM diagonal block of the  $l$ -th layer,  $g$  is the pre-activation gradient,  $a$  is the input activation, and  $\otimes$  indicates Kronecker product. K-FAC represents the curvature matrix in the form of a product between two factors, thereby reducing the memory cost from  $\mathcal{O}(M^2 \times N^2)$  to  $\mathcal{O}(M^2 + N^2)$ , where  $M$  and  $N$  are the sizes of the input and output units of the layer, respectively.

Eigenvalue-corrected K-FAC [11] proposed to re-compute the diagonals in the eigenbasis obtained from K-FAC, thereby correcting the eigenvalues of the factorized FIM. Trace-restricted K-FAC (TK-FAC) [10] is a trace-exact variant which corrects the norm of the factorized FIM by further tracking the trace of FIM.

### 2.3 Linearization of deep neural networks

Linear approximation in deep learning is a widely used concept as a means of estimating the proximal behavior of a neural network. Popular explanation methods such as Guided Backprop [43] and Grad-CAM [41] use the gradient with respect to the activations in order to estimate the network’s the sensitivity against the input. Such techniques are based on the first-order derivatives of neural networks with respect to the input.

The first-order derivatives with respect to the parameters, on the other hand, have lead to interesting insights about the training of neural networks. Recent works regarding the training dynamics of deep neural networks [16,21] have found that randomly initialized neural networks behave linearly throughout gradient descent training in the infinite-width regime. In [16], it is shown that such model can be described by a specific kernel function, coined neural tangent kernel, defined by the first-order derivatives of the neural network. Moreover, recent works [1,31] have empirically shown that the observation also apply to finite-sized neural networks when they are pre-trained, and showed the linearized network can be fine-tuned to achieve comparable performance to the nonlinear network.

## 3 Continual Fine-tuning

The goal of continual learning is to achieve the highest performance jointly on all tasks when the tasks arrive sequentially. Due to the sequential nature, the model is only allowed to observe a batch of data  $\mathcal{D}_t$  at each task  $t$ . The performance of continual learning can be seen as upper bounded by multi-task learning, where the model trains jointly on all tasks,

$$\min_{\theta} \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{(x,y) \in \mathcal{D}_t} [\lambda_t \mathcal{L}(f(x; \theta, t), y)], \quad (2)$$

where  $\theta$  is the vectorized parameter of the linearized model, and  $\lambda_t$  controls the stability–plasticity between the tasks. To learn all tasks from sequentially arriving batches, parameter regularization methods aim to capture the source task objective as a quadratic function of parameters, *i.e.*,

$$\begin{aligned} \frac{1}{t} \sum_{i=1}^t \mathbb{E}_{(x,y) \in \mathcal{D}_i} [\lambda_i \mathcal{L}(f(x; \theta, i), y)] &= C + \frac{1}{2}(\theta - \theta_t)^\top A(\theta - \theta_t) + \mathcal{O}(\theta^3) \\ &\simeq C + \frac{1}{2}(\theta - \theta_t)^\top A(\theta - \theta_t), \end{aligned} \quad (3)$$

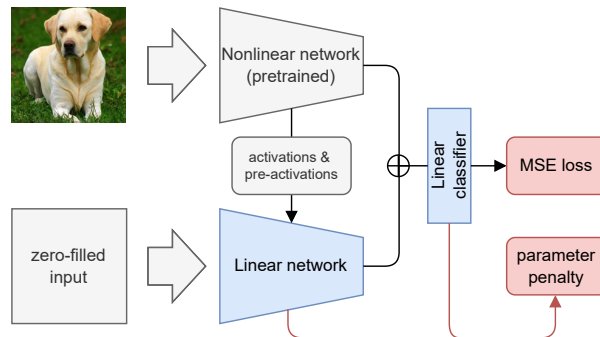


Fig. 2: **Architecture of our deep linear continual fine-tuning model.** Learnable modules are colored in blue. On forward pass, the image and a zero-filled tensor are passed into the nonlinear network (fixed) and linear network (learnable), respectively. The nonlinear feature and the linear residual feature are computed concurrently, which are then summed and passed to the linear classification layer. Notice that the output logits are strictly linear with respect to all learnable parameters. At training, MSE loss is used instead of cross-entropy. During continual learning, model is regularized by quadratic parameter penalty loss to prevent forgetting.

where  $A$  is the matrix whose each entry represents the importance of the corresponding parameter pair,  $C$  is a constant, and  $\theta_t$  is the model parameters after learning the  $t$ -th task.

Parameter-based continual learning methods aim to accurately estimate the importance matrix  $A$  in order to closely approximate the multi-task learning objective. One popular approach is to use second-order Taylor approximation, where  $A$  becomes the Hessian matrix of the source task objective function.

### 3.1 Understanding the problems of quadratic parameter regularization

In this section, we provide why existing parameter regularizations show underwhelming performance in continual training of neural networks. We particularly look into two sources of problems which we call *vanishing curvature* and *higher-order error*. First, we show that softmax cross-entropy loss causes curvature to vanish to zero, thereby losing ability to represent parameter importance. Then, we show why this problem is amplified by nonlinearity of neural network.

First, we investigate the behavior of FIM of a model that predicts categorical distribution and trained using cross-entropy loss. The second order derivatives of a probabilistic model is often estimated through FIM, which is the covariance matrix of the gradient with respect to the log-likelihood.

$$F_{\theta} = \mathbb{E}_{y \sim p_{\theta}(y|x)} \left[ \nabla_{\theta} \log p_{\theta}(y|x) \nabla_{\theta} \log p_{\theta}(y|x)^{\top} \right] \quad (4)$$

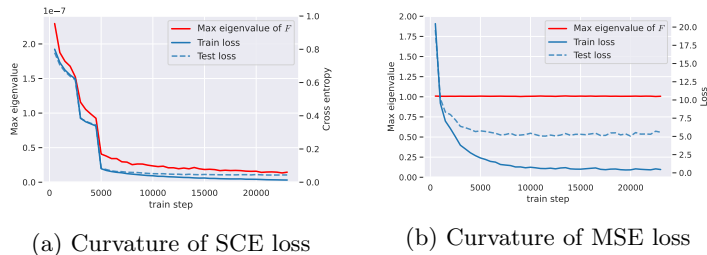


Fig. 3: An MLP is trained on MNIST using softmax cross-entropy (SCE) and mean squared error (MSE) losses, and the maximum eigenvalue of the Fisher information matrix of each loss function is traced and plotted in red. The eigenvalues of SCE loss diminish to zero as the model fits the data, whereas that of MSE loss remains constant.

It can be seen from Eq. (4) that as the model  $p_{\theta}(y|x)$  perfectly fits the target label, the sampled  $y$  becomes the target label with probability  $\rightarrow 1$  and  $\nabla_{\theta} \log p_{\theta}(y|x) \rightarrow 0$ , therefore  $F_{\theta} \rightarrow 0$ . This indicates that FIM loses its ability to represent parameter importance at near zero-loss optima. See Sec. A in the supplementary for detailed proof. We name this *vanishing curvature problem*. Additionally, notice that this behavior is caused by choosing to fit a categorical distribution, e.g., softmax layer. Fig. 3a shows that the maximum eigenvalue of FIM continues to vanish to zero as the negative log-likelihood approaches to zero loss, even after the test loss has converged. As a result, the parameters are under-regularized.

Secondly, quadratic approximation assumes that the objective function is a quadratic function with respect to the parameters. However, the true loss function is non-quadratic due to the cross-entropy loss and the non-linearity of neural networks. This introduces higher-order error terms to dominate in the second-order approximation. Combined with the vanishing curvature behavior, this amplifies the error of the loss approximation as the parameters are under-damped and the model can drift off the trust region of the local loss approximation which causes catastrophic forgetting.

### 3.2 Deep Linear Continual Fine-tuning (DLCFT)

To this end, we propose an alternative approach based on continual fine-tuning framework. Our change to the model is two-fold; To tackle the vanishing curvature problem, we replace the cross-entropy loss with MSE loss function. To resolve the non-convexity problem, we choose to *approximate the model* such that it has simple linear structure. The combination of changes to the model allows parameter-based regularization to be the optimal continual learning policy.

Firstly, to work around the vanishing curvature problem, we replace softmax cross-entropy loss for MSE loss which is a non-saturating, quadratic loss function.

$$\ell(f(x; \theta, t), y) = \frac{1}{2} \|\alpha \phi(y) - f(x; \theta, t)\|^2 \quad (5)$$

Here,  $\phi(\cdot)$  indicates the one-hot representation, and  $\alpha$  is a positive scaling constant which is fixed throughout all tasks. We followed [1, 14] and set  $\alpha = 15$ .

Secondly, to tackle the higher-order error problem, we linearize the neural network. We apply first-order Taylor approximation of a pre-trained neural network [1, 31] which decomposes the feature extractor into a frozen non-linear network and a trainable linear network. Instead of fine-tuning the full non-linear neural network, we train the linearized neural network, *i.e.*,

$$g_{lin}(x; \psi) = g(x; \psi_0) + D_\psi g(x; \psi_0) \cdot \psi, \quad (6)$$

where  $D_\psi g(x; \psi_0)$  is the Jacobian of the network evaluated at the pre-trained point  $\psi_0$ . Then we learn a linear classification layer using the linearized feature,

$$f(x; \theta, t) = w_t \cdot g_{lin}(x; \psi) + b_t, \quad (7)$$

where  $g(x; \psi_0)$  is the pretrained nonlinear feature extraction network and  $\theta = \{\psi, w_t, b_t\}$  corresponds to the parameters of the linearized model. For data-incremental learning setup, we train the linearized feature extraction network and a single linear classifier. We regularize both the feature extractor and the linear classifier  $\{\psi, w_t, b_t\}$ . For task-incremental learning setup, we use a shared linearized feature extraction network  $g_{lin}(x; \psi)$  along with a linear classifier assigned to each task. We regularize only the feature extractor  $\psi$  as each task-specific classifier does not interfere with each other. For class-incremental learning, we train the feature extractor and append output units to the classifier at the beginning of each task.

Finally, notice that when combined with linear model, this change makes the objective function fully quadratic with respect to the parameters. This enables us to accurately represent the objective function while allowing the model to be highly accurate and expressive. The difference that these changes bring is that because the model is linear with respect to its parameters and the loss function is mean-squared error, the objective becomes quadratic with respect to the parameters. Most notably, it follows that the quadratic parameter penalty is the optimal strategy to represent the source task objective for continual learning, *i.e.*,

$$\begin{aligned} \ell(f(x; \theta, t), y) &= \frac{1}{2} \|\alpha \phi(y) - f(x; \theta, t)\|^2 \\ &= \frac{1}{2} (\theta - \theta_t)^\top H_\theta (\theta - \theta_t) + C, \end{aligned} \quad (8)$$

where  $\theta_t$  is the trained parameters after the  $t$ -th task. We can apply any curvature approximation algorithm to efficiently store  $H_\theta$  in memory. *e.g.*, TK-FAC [10].



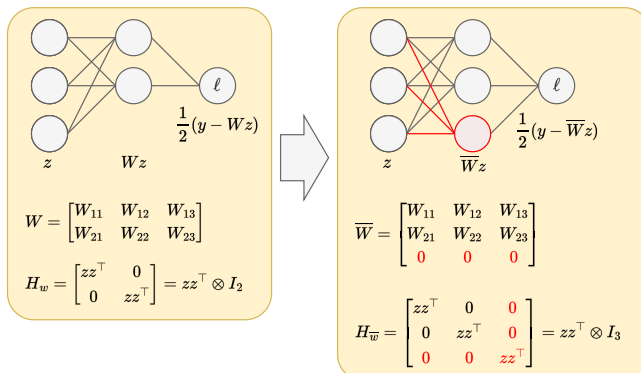


Fig. 4: **Curvature for classifier in class-incremental learning.** An example of a linear classification layer with two original output units and an appended unit. The appended output unit and its zero-initialized weights are outlined in red. Bias is omitted for simplicity. The corresponding updates to the weight matrix and the block diagonal of the Hessian matrix are marked in red.

The final objective function for continual learning is,

$$\mathcal{L}_{data} = \mathbb{E}_{(x,y) \sim \mathcal{D}_t} [\ell(f(x; \theta, t), y)] \quad (9)$$

$$\mathcal{L}_{reg} = \frac{1}{2}(\theta - \theta_{t-1})^\top H(\theta - \theta_{t-1}) \quad (10)$$

$$\mathcal{L} = \frac{1}{t}\mathcal{L}_{data} + \frac{t-1}{t}\mathcal{L}_{reg}, \quad (11)$$

where the target task objective and source task objective are weighted by 1 :  $(t - 1)$  for balancing, and  $H$  corresponds to the Hessian of the source task objectives.

### 3.3 Classifier regularization for class-incremental problem

In this section, we describe how the proposed method is extended to class-incremental learning problem. To the best of our knowledge, this is the first work that shows how parameter regularization can be reasonably applied to class-incremental learning problem.

In class-incremental learning, each task requires the model to learn a set of novel classes while the evaluation is done jointly over the current and previous tasks. Unlike task-incremental setup, task oracle is not provided at test time. At the beginning of each task, a set of output units that corresponds to the new classes is added to the classification layer. A key challenge to the problem is to apply correct regularization to the classification layer to prevent predictions from being biased towards more recent tasks. Memory-based method achieves this by replaying samples from a buffer. On the other hand, parameter regularization

methods require the curvature matrix of the source task loss to be defined for the parameters of the new output units. Here, we show that the curvature matrix is obtained without looking at the previous task data, but only from the existing curvature matrix. The key idea is that this is equivalent to the case where all the weights and biases of the unseen unit is set to zero. Note that this is only possible when MSE loss is used, whereas the new weights and biases diverge to infinity when SCE loss is used. Moreover, the curvature vanishes to zero when the weights and biases diverge.

Let us consider a linear classification layer with increasing output units for class-incremental setup. Adding a set of classes amounts to adding a set of corresponding output units to the weight matrix  $W$ . Let  $\overline{W}$  be the augmented weight matrix with the added output units, and  $y$  be the target. Then, for the loss function

$$\mathcal{L}(\overline{W}) = \mathbb{E} \left[ \frac{1}{2} \|y - \overline{W}z\|^2 \right], \quad (12)$$

the second derivative with respect to  $\overline{w} := \text{Vec}(\overline{W})$  is

$$\frac{\partial^2 \mathcal{L}}{\partial \overline{w} \partial \overline{w}^\top} = \frac{\partial}{\partial \overline{w}} \mathbb{E}[(y - \overline{W}z)z^\top] = \mathbb{E}[zz^\top \otimes I] = \mathbb{E}[zz^\top] \otimes I. \quad (13)$$

Because  $\mathbb{E}[zz^\top]$  has been already obtained through K-FAC regularization, we do not need additional computation or data to compute the second derivative of the appended weight.

We additionally employ a small buffer memory  $\mathcal{M}$  to replay samples of previous tasks. In previous works, combining replay with parameter regularization has not been a common practice due to underwhelming performance of curvature-based regularization. However in the proposed method, the only source of error is the approximation of the curvature matrix. Whereas in replay methods, the source of error is the subsampling of source task dataset. Therefore, we can combine the proposed parameter regularization with an additional replay loss to complement for the approximation error. The final regularization objective is,

$$\mathcal{L}_{reg} = \frac{\lambda}{2} \cdot (\theta - \theta_{t-1})^\top H(\theta - \theta_{t-1}) + (1 - \lambda) \cdot \mathbb{E}_{\mathcal{M}}[\ell(f(x; \theta, t), y)]. \quad (14)$$

## 4 Experiments

### 4.1 Evaluation methods and implementation details

**Evaluation settings.** We evaluate our method on three types of incremental learning (IL) problems: data-IL, task-IL, and class-IL.

**Models.** We use ResNet-18 [13] architecture for all benchmarks. For the linearized ResNet-18, we followed [1] and replace all ReLU nonlinearities with LeakyReLU [28]. We also followed [31] and folded the batch norm parameters into the convolution layers.

**Pre-training.** For the experiments using CIFAR-100 dataset [19], we use ImageNet32 [6] which consists of  $32 \times 32$  downsampled images of ImageNet-1k dataset [6, 39]. At pre-training phase, we train the model for 100 epochs using SGD optimizer with learning rate  $\eta = 10^{-1}$ , batch size = 256, weight decay =  $10^{-5}$ . We use the cosine annealing [27] learning rate schedule. For the experiments using MIT-67 dataset [35], we use the pretrained ResNet-18 model downloaded using TorchVision<sup>3</sup>, which trained on ImageNet-1k. To obtain the model with LeakyReLU nonlinearities, we replicate the scheme from [1] and fine-tune the downloaded model on ImageNet-1k for an additional epoch using SGD with learning rate  $\eta = 10^{-4}$ .

**Datasets.** For data-IL setting, we used Seq-CIFAR-100 dataset [19, 36] split into 10 and 100 tasks, each task having 5000 and 500 samples, respectively. Additionally, we used Seq-MIT-67 for large-resolution dataset, which is MIT-67 dataset [35] split into 4 tasks. For task-IL and class-IL settings, we used Seq-CIFAR-100 with 10 tasks each containing a disjoint set of classes.

**Curvature approximation.** For approximation of the Hessian matrix, we use K-FAC [12, 30] and TK-FAC [10].

**Data augmentation.** For Seq-CIFAR-100, we apply random crop with 4 pixels of zero padding, followed by random horizontal flip. For Seq-MIT-67, we first apply resizing to  $256 \times 256$  then apply random crop to  $224 \times 224$ . At test time, we resize to  $256 \times 256$  and apply center crop to  $224 \times 224$ .

**Training scheme and hyperparameters.** For training nonlinear models, we used softmax cross-entropy loss and SGD optimizer with initial learning rate  $10^{-3}$  and momentum 0.9. For training linearized models, we used MSE loss and Adam optimizer [17] with initial learning rate  $10^{-4}$  and  $(\beta_1, \beta_2) = (0.9, 0.999)$ . In data-IL and task-IL experiments, we enable batch normalization [15] at the first task only. For the loss used in class-IL, we set  $\lambda = 1/2$ .

**Other implementation details.** We used Nvidia RTX 3090 GPUs and PyTorch to conduct experiments. To add a pair of Kronecker-factored curvature matrices that each correspond to the source task and the target task, we take the weighted sum the factorized matrices by  $(t - 1) : t$  for each factors. For all experiments, we used weight decay rate of  $10^{-5}$ . For methods that uses buffer memory, we set the size of the buffer to 500 samples. We used reservoir sampling strategy to update the buffer.

## 4.2 Incremental learning benchmarks

**Data-incremental learning** We benchmark our method on data-IL setup in three different settings. Firstly, we tested on ten splits of Seq-CIFAR-100 training set to simulate a short sequence of data streams. Secondly, we tested on a hundred splits of Seq-CIFAR-100 training set to evaluate how the proposed method scales to a very long sequence. Finally, we tested on four splits of Seq-MIT-67 training to evaluate the methods on a high-resolution images.

<sup>3</sup> <https://github.com/pytorch/vision>

Table 1: **Experiment results for data-incremental learning.** ‘None’ indicates no regularization applied during incremental learning. The experiments are averaged over three runs. † denotes that K-FAC [12,30] is used for curvature approximation and ‡ denotes TK-FAC [10] is used.

Dataset	Seq-CIFAR-100		Seq-MIT-67
	Sequence length 10	100	4
None	78.74	71.83	63.48
LwF [24]	80.25	70.95	67.21
EWC [18]	78.88	72.61	63.68
MAS [2]	75.24	52.50	62.49
OSLA [38]	79.23	73.10	64.08
DLCFT <sup>†</sup> (Ours)	<b>81.95</b>	<b>75.92</b>	<b>70.55</b>
DLCFT <sup>‡</sup> (Ours)	<b>82.70</b>	<b>80.07</b>	<b>70.52</b>
Joint	83.57		74.40




Fig. 5: **Data-IL evaluated after each task.** Data-incremental learning using Seq-CIFAR-100. Here, we show the progress of the performance increase by evaluating after each task. Accuracy is evaluated on CIFAR-100 test set.

Table 1 shows performance in data-IL measured by final accuracy. We observed that Memory Aware Synapses (MAS) [2] fails to learn incrementally as it does not estimate accurate weight importance. On the other hand, our method performs much better than the baselines and achieves comparable performance to the joint training as it captures the curvature of the loss of the previous tasks accurately. Additionally, in Fig. 5, we show the performance trends plotted against increasing data-IL tasks. Notably, we observed that the proposed method significantly outperforms the baselines on the long sequence length setup.


**Task-/class-incremental learning** Table 2 shows performance comparison on task-incremental setup. The benchmark consists of 10 tasks with disjoint class categories obtained from CIFAR-100, each consisting of 10 classes. We observe that our method performs better than the baselines.

Table 2: **Experiment results for task-/class-incremental learning.** Experiments are conducted using Seq-CIFAR-100 with 10 tasks. TK-FAC [10] is used for our method.

Buffer size	Method	Task-IL	Class-IL
0	LwF [24]	92.16	-
	EWC [18]	77.44	-
	OSLA [38]	81.03	-
	DLCFT (Ours)	<b>95.79</b>	-
500	ER [37]	79.14	43.52
	DER [4]	91.47	58.07
	DER++ [4]	91.56	53.29
	DLCFT (Ours)	-	<b>59.98</b>



(a) Seq-CIFAR-100 (10 tasks)



(b) Seq-CIFAR-100 (100 tasks)

Fig. 6: **K-NN probing of the incrementally fine-tuned models.** The quality of the representation is evaluated using K-NN classification accuracy on test set. Evaluation is performed after training each task in data-IL. At each evaluation, all training samples were used to perform inference.

### 4.3 Evaluation of incrementally learned representations

Because we use pre-trained networks that has already learned transferable representation from a large labeled dataset, it is possible that the performance of the method is more attributed to the linear classification layer, rather than adjusting the features through the sequential tasks. To verify that our continual fine-tuning method does learn better representations through incremental learning, we use K-NN classifier to evaluate the quality of the representations learned through the continual fine-tuning process. The results are show in Fig. 6. The plot shows the discriminative performance of the fine-tuned feature by decoupling the linear classifier from the evaluation. We observed that the K-NN accuracy consistently and monotonically increases as the model observes more tasks.

### 4.4 Ablations

Table 3 shows the ablation study of the components of the proposed method performed on the data-incremental and class-incremental setting. The result

Table 3: **Ablations.** Experiments are conducted on data- and class-IL benchmarks using Seq-CIFAR-100 dataset. ‘Linear’ indicates linearized model, and ‘SCE’ indicates softmax cross-entropy loss.

Curvature	Linear	Loss	Data-IL		Class-IL
			10 tasks	100 tasks	10 tasks
K-FAC [30]	✗	SCE	79.23	73.51	44.93
	✓	MSE	<b>81.95</b>	<b>75.92</b>	<b>59.86</b>
TK-FAC [10]	✗	SCE	79.58	73.18	44.66
	✓	MSE	<b>82.70</b>	<b>80.07</b>	<b>59.98</b>

shows that jointly applying the linearization and MSE loss can significantly increase performance and mitigate forgetting. Note that between nonlinear networks trained with softmax cross-entropy, employing more accurate curvature approximation brings marginal performance difference. However, between linearized networks trained with MSE loss, adopting better curvature approximation brings significant performance gain. This is because in the former case, the effectiveness of better curvature is minimal due to the vanishing curvature and higher-order error problems.

## 5 Conclusion

In this paper, we have explored *continual fine-tuning*, which is a practical framework for incremental learning of deep neural networks. For this, we propose Deep Linear Continual Fine-tuning, which is a simple and effective continual learning algorithm using a pre-trained neural network. We provided theoretical reasons on why existing Hessian-based parameter regularization performs poorly with neural networks trained using softmax cross-entropy loss. We showed that a combination of model linearization technique and mean-squared error loss function allows the parameter regularization methods to closely match the optimal continual learning policy. We provided a principled approach to applying parameter regularization in class-incremental learning scenario, and showed that our method outperforms other baselines on data-/task-/class-incremental settings. Moreover, we show that our method can effectively accumulate knowledge over very long data-incremental tasks sequences.

**Acknowledgements** This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No. 2022-0-00951, Development of Uncertainty-Aware Agents Learning by Asking Questions).

## References

1. Achille, A., Golatkar, A., Ravichandran, A., Polito, M., Soatto, S.: Lqf: Linear quadratic fine-tuning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 15729–15739 (June 2021)
2. Aljundi, R., Babiloni, F., Elhoseiny, M., Rohrbach, M., Tuytelaars, T.: Memory aware synapses: Learning what (not) to forget. In: The European Conference on Computer Vision (ECCV) (September 2018)
3. Aljundi, R., Chakravarty, P., Tuytelaars, T.: Expert gate: Lifelong learning with a network of experts. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (July 2017)
4. Buzzega, P., Boschini, M., Porrello, A., Abati, D., Calderara, S.: Dark experience for general continual learning: a strong, simple baseline. *Advances in neural information processing systems* **33**, 15920–15930 (2020)
5. Chaudhry, A., Ranzato, M., Rohrbach, M., Elhoseiny, M.: Efficient lifelong learning with a-gem. In: ICLR (2019)
6. Chrabaszcz, P., Loshchilov, I., Hutter, F.: A downsampled variant of imagenet as an alternative to the cifar datasets. arXiv preprint arXiv:1707.08819 (2017)
7. Delange, M., Aljundi, R., Masana, M., Parisot, S., Jia, X., Leonardis, A., Slabaugh, G., Tuytelaars, T.: A continual learning survey: Defying forgetting in classification tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2021)
8. Dhar, P., Singh, R.V., Peng, K.C., Wu, Z., Chellappa, R.: Learning without memorizing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5138–5146 (2019)
9. Ebrahimi, S., Elhoseiny, M., Darrell, T., Rohrbach, M.: Uncertainty-guided continual learning with bayesian neural networks. In: International Conference on Learning Representations (2020), <https://openreview.net/forum?id=HklUCCVKDB>
10. Gao, K., Liu, X., Huang, Z., Wang, M., Wang, Z., Xu, D., Yu, F.: A trace-restricted kronecker-factored approximation to natural gradient. Proceedings of the AAAI Conference on Artificial Intelligence **35**(9), 7519–7527 (May 2021), <https://ojs.aaai.org/index.php/AAAI/article/view/16921>
11. George, T., Laurent, C., Bouthillier, X., Ballas, N., Vincent, P.: Fast approximate natural gradient descent in a kronecker factored eigenbasis. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 31. Curran Associates, Inc. (2018), <https://proceedings.neurips.cc/paper/2018/file/48000647b315f6f00f913caa757a70b3-Paper.pdf>
12. Grosse, R., Martens, J.: A Kronecker-factored approximate Fisher matrix for convolution layers. In: International Conference on Machine Learning. pp. 573–582 (2016)
13. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
14. Hui, L., Belkin, M.: Evaluation of neural architectures trained with square loss vs cross-entropy in classification tasks. arXiv preprint arXiv:2006.07322 (2020)
15. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: International conference on machine learning. pp. 448–456. PMLR (2015)
16. Jacot, A., Gabriel, F., Hongler, C.: Neural tangent kernel: Convergence and generalization in neural networks. In: Bengio, S., Wallach, H., Larochelle, H., Grauman,

- K., Cesa-Bianchi, N., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 31. Curran Associates, Inc. (2018), <https://proceedings.neurips.cc/paper/2018/file/5a4be1fa34e62bb8a6ec6b91d2462f5a-Paper.pdf>
17. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
  18. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A.A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., et al.: Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences* **114**(13), 3521–3526 (2017)
  19. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009)
  20. Kurle, R., Cseke, B., Klushyn, A., van der Smagt, P., Günnemann, S.: Continual learning with bayesian neural networks for non-stationary data. In: *International Conference on Learning Representations* (2019)
  21. Lee, J., Xiao, L., Schoenholz, S., Bahri, Y., Novak, R., Sohl-Dickstein, J., Pennington, J.: Wide neural networks of any depth evolve as linear models under gradient descent. *Advances in neural information processing systems* **32**, 8572–8583 (2019)
  22. Lee, J., Hong, H.G., Joo, D., Kim, J.: Continual learning with extended kronecker-factored approximate curvature. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9001–9010 (2020)
  23. Lee, S.W.L., Kim, J.H., Jun, J., Ha, J.W., Zhang, B.T.: Overcoming Catastrophic Forgetting by Incremental Moment Matching (IMM). In: *Advances In Neural Information Processing Systems 30* (2017)
  24. Li, Z., Hoiem, D.: Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence* **40**(12), 2935–2947 (2018)
  25. Liu, X., Masana, M., Herranz, L., Van de Weijer, J., Lopez, A.M., Bagdanov, A.D.: Rotate your networks: Better weight consolidation and less catastrophic forgetting. In: *2018 24th International Conference on Pattern Recognition (ICPR)*. pp. 2262–2268. IEEE (2018)
  26. Lopez-Paz, D., Ranzato, M.: Gradient episodic memory for continual learning. In: *NIPS* (2017)
  27. Loshchilov, I., Hutter, F.: Sgdr: Stochastic gradient descent with warm restarts. arXiv preprint arXiv:1608.03983 (2016)
  28. Maas, A., Hannun, A., Ng, A.: Rectifier nonlinearities improve neural network acoustic models. In: *Proceedings of the International Conference on Machine Learning*. vol. 30 (2013)
  29. Mallya, A., Lazebnik, S.: Packnet: Adding multiple tasks to a single network by iterative pruning. In: *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*. pp. 7765–7773 (2018)
  30. Martens, J., Grosse, R.: Optimizing neural networks with Kronecker-factored approximate curvature. In: *International conference on machine learning*. pp. 2408–2417 (2015)
  31. Mu, F., Liang, Y., Li, Y.: Gradients as features for deep representation learning. In: *International Conference on Learning Representations* (2020), <https://openreview.net/forum?id=BkeoaeHKDS>
  32. Nguyen, C.V., Li, Y., Bui, T.D., Turner, R.E.: Variational continual learning. In: *International Conference on Learning Representations* (2018)
  33. Pascanu, R., Bengio, Y.: Revisiting natural gradient for deep networks. arXiv preprint arXiv:1301.3584 (2013)



34. Prabhu, A., Torr, P., Dokania, P.: Gdumb: A simple approach that questions our progress in continual learning. In: The European Conference on Computer Vision (ECCV) (August 2020)
35. Quattoni, A., Torralba, A.: Recognizing indoor scenes. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition. pp. 413–420 (2009). <https://doi.org/10.1109/CVPR.2009.5206537>
36. Rebuffi, S.A., Kolesnikov, A., Sperl, G., Lampert, C.H.: icarl: Incremental classifier and representation learning. 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) pp. 5533–5542 (2017)
37. Riemer, M., Cases, I., Ajemian, R., Liu, M., Rish, I., Tu, Y., , Tesauro, G.: Learning to learn without forgetting by maximizing transfer and minimizing interference. In: International Conference on Learning Representations (2019), <https://openreview.net/forum?id=BigTShAct7>
38. Ritter, H., Botev, A., Barber, D.: Online structured laplace approximations for overcoming catastrophic forgetting. In: Advances in Neural Information Processing Systems. pp. 3738–3748 (2018)
39. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: ImageNet Large Scale Visual Recognition Challenge. International Journal of Computer Vision (IJCV) **115**(3), 211–252 (2015). <https://doi.org/10.1007/s11263-015-0816-y>
40. Rusu, A.A., Rabinowitz, N.C., Desjardins, G., Soyer, H., Kirkpatrick, J., Kavukcuoglu, K., Pascanu, R., Hadsell, R.: Progressive neural networks. arXiv preprint arXiv:1606.04671 (2016)
41. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D., et al.: Grad-cam: Visual explanations from deep networks via gradient-based localization. iccv (2016)
42. Serra, J., Suris, D., Miron, M., Karatzoglou, A.: Overcoming catastrophic forgetting with hard attention to the task. In: International Conference on Machine Learning. pp. 4548–4557. PMLR (2018)
43. Springenberg, J.T., Dosovitskiy, A., Brox, T., Riedmiller, M.: Striving for simplicity: The all convolutional net. arXiv preprint arXiv:1412.6806 (2014)
44. Yoon, J., Kim, S., Yang, E., Hwang, S.J.: Scalable and order-robust continual learning with additive parameter decomposition. In: International Conference on Learning Representations (2020), <https://openreview.net/forum?id=r1gdj2EKPb>
45. Zenke, F., Poole, B., Ganguli, S.: Continual learning through synaptic intelligence. In: International Conference on Machine Learning. pp. 3987–3995. PMLR (2017)