

A Unified Continual Learning Framework with General Parameter-Efficient Tuning

Qiankun Gao¹ Chen Zhao^{†2} Yifan Sun³ Teng Xi³ Gang Zhang³ Bernard Ghanem² Jian Zhang^{†1}

¹Peking University Shenzhen Graduate School ²King Abdullah University of Science and Technology (KAUST) ³Baidu Inc.

gqk@stu.pku.edu.cn chen.zhao@kaust.edu.sa zhangjian.sz@pku.edu.cn

Abstract

The “pre-training \rightarrow downstream adaptation” presents both new opportunities and challenges for Continual Learning (CL). Although the recent state-of-the-art in CL is achieved through Parameter-Efficient-Tuning (PET) adaptation paradigm, only prompt has been explored, limiting its application to Transformers only. In this paper, we position prompting as one instantiation of PET, and propose a unified CL framework with general PET, dubbed as Learning-Accumulation-Ensemble (LAE). PET, e.g., using Adapter, LoRA, or Prefix, can adapt a pre-trained model to downstream tasks with fewer parameters and resources. Given a PET method, our LAE framework incorporates it for CL with three novel designs. 1) *Learning*: the pre-trained model adapts to the new task by tuning an online PET module, along with our adaptation speed calibration to align different PET modules, 2) *Accumulation*: the task-specific knowledge learned by the online PET module is accumulated into an offline PET module through momentum update, 3) *Ensemble*: During inference, we respectively construct two experts with online/offline PET modules (which are favored by the novel/historical tasks) for prediction ensemble. We show that LAE is compatible with a battery of PET methods and gains strong CL capability. For example, LAE with Adaptor PET surpasses the prior state-of-the-art by 1.3% and 3.6% in last-incremental accuracy on CIFAR100 and ImageNet-R datasets, respectively. Code is available at <https://github.com/gqk/LAE>.

1. Introduction

Continual Learning (CL) of new knowledge is an essential ability for AI models in the constantly changing world. However, neural networks often suffer from catastrophic forgetting [9, 39], in which previously learned knowledge is

[†] Corresponding authors.

This work was supported by Shenzhen General Research Project under Grant JCYJ20220531093215035, KAUST Office of Sponsored Research (OSR) under Award No. OSR-CRG2021-4648, and the SDAIA-KAUST Center of Excellence in Data Science and Artificial Intelligence.




Figure 1: **The pipeline of our LAE framework vs. prompt-pool approaches.** **Above:** Prompt-pool approaches, which query prompts from a pool of learnable prompts. **Below:** Proposed LAE, where an *online Parameter-Efficient Tuning (PET) module* attached to the pre-trained model to adapt a new task quickly, and an *offline PET module* accumulates the learned knowledge slowly. During inference, we use the ensemble of the predictions of the online and offline PET modules as the final prediction.

forgotten when the model incorporates novel information. Although many works have been devoted to reducing forgetting, such as dynamic networks [41, 50, 26, 18], regularization [28, 20, 52, 1], and memory replay [38, 15, 6, 30, 2], their performance still falls short of practical requirements.

Recently, pre-training and downstream adaptation techniques have opened up new opportunities and challenges for CL. Basically, these techniques [4, 36, 12, 11, 48] pre-train a deep model on large-scale data and then adapt the pre-trained model to novel tasks. We observe that downstream adaptation and CL are important for each other. On the one hand, in realistic AI systems, pre-trained models sometimes needs to be adapted to multiple downstream tasks sequentially, yielding the need of CL. On the other hand, recent efforts [47, 46, 45] show that the “pre-training \rightarrow downstream adaptation” techniques can boost CL performance.

Specifically, L2P [47], DualPrompt [46], and ESN [45] all use a popular adaptation technique named Parameter-Efficient-Tuning (PET). Generally, PET adapts pre-trained models to downstream tasks with much fewer learnable parameters, as well as fewer resources. Though these approaches have advanced the state-of-the-art in CL, they still have some limitations. 1) They are all constrained to a spe-

cific PET method, *i.e.*, prompt tuning, limiting their flexibility, considering that prompt can only cooperate with transformers and does not accommodate other network architectures. 2) Most of them rely on selecting task-specific parameters (the prompt tokens, in particular) for each individual task. The selection tends to be noisy with increasing task numbers and the task-specific prompts appear homogeneous, according to our investigation in the supplementary.

To circumvent these issues, this paper proposes Learning-Accumulation-Ensemble (LAE), a unified CL framework resort to the general Parameter-Efficient Tuning (PET). LAE is not restricted to Prompt, but can also utilize various other PET modules as shown in Fig. 2 (b). Given a PET method, our LAE directly reshapes it for CL with three steps, *i.e.*, learning, accumulation, and ensemble.

- **1) Learning with calibrated speed.** The pre-trained model adapts to the new task by tuning an online PET module. To accommodate various PET methods, a key challenge is that different PET modules have different *adaptation speeds* (for novel tasks), as well as different *forgetting speeds* (for historical tasks). In response, we design an adaptation calibration strategy, based on the gradient analysis for different PET modules. We empirically show that this calibration strategy aligns different PET against each other and is critical for LAE to be a unified framework.

- **2) Accumulation of multi-task knowledge.** After adapting the pre-trained model to a new task, the parameters in the online PET module are prone to the current novel task and may not fit historical tasks. Instead of memorizing multiple sets of PET modules and selecting some subsets (as in L2P and DualPrompt) for individual tasks, LAE accumulates all the knowledge of already-seen tasks into a single offline PET module through momentum update. This simple accumulation avoids noisy selection and is competent for *alleviating catastrophic forgetting, especially when the amount of learned tasks is large* (Figs. 3 and 4 in Sec. 5).

- **3) Ensemble of two expert models.** The online and offline PET modules respectively contain more novel and historical knowledge, therefore, two expert models constructed with them are correspondingly better at handling newer and older tasks. Instead of inference only using the online or offline expert model, we integrate the outputs of two expert models by an energy indicator (detailed in Sec. 4.2) to obtain the prediction for an inference sample from any learned task. This expert ensemble strategy helps our framework to achieve a more robust performance compared to inference using one of the expert models alone.

The contributions of this paper are summarized as follows:

- We thoroughly investigate the novel Continual Learning paradigm that constantly adapts a pre-trained model to novel tasks using general Parameter-Efficient Tuning (PET) methods, and propose a unified Learning-Accumulation-Ensemble (LAE) framework.

- Our LAE framework reshapes a given PET method into a competitive Memory-Free Continual Learning approach with three novel designs: Learning with calibrated speed, Accumulation of multi-task knowledge, and Ensemble of two expert models constructed with online and offline PET modules.
- We conduct extensive experiments on CIFAR100 and ImageNet-R benchmarks, on all of which, our LAE consistently achieves superior incremental performance than previous state-of-the-art approaches.

2. Related Works

Parameter-Efficient Tuning. As an efficient alternative to full fine-tuning, Adapter-Tuning [16] was first proposed to transfer large pre-trained Language models to downstream tasks. Inspired by textual prompting, Prompt-Tuning [25] and Prefix-Tuning [27] insert learnable tokens to adapt to the new task. More advanced methods [51, 7, 17, 33] achieved comparable or superior performance to full fine-tuning, and keep the same inference cost by merging the additional learnable parameters to the original pre-trained model. Following the step of successful Vision Transformers [5, 31], VPT [19] and AdapterFormer [3] have been proposed to solve visual transfer learning problems. Prompt-Tuning and Prefix-Tuning depend on the transformer architecture because they modify input or hidden tokens. Adapter and its variants are network architecture generalizable since they are new modules that can be implemented in forms compatible with pre-trained models. All types of Parameter-Efficient Tuning modules can be integrated into our LAE framework as long as they are suitable for the pre-trained model, but we focus on the representative Adapter [16], LoRA [17], and Prefix [27] in this paper.

Continual Learning. The central problem of Continual Learning (CL) is fighting catastrophic forgetting [9]. Memory-based approaches [38, 15, 6, 30, 2] save a subset of learned samples into a memory buffer and replay them when learning a new task. Memory-Free approaches do not rely on old samples that may raise privacy concerns, they dynamically expand the network or isolate parameters for different tasks [41, 50, 26, 18], regularize the network parameters that are important to learned tasks [28, 20, 52, 1], and replay generative or synthetic data [42, 49, 43, 8]. Conventional CL approaches learn tasks from scratch using a randomly initialized model, while pre-trained models have received little attention from CL researchers until recently. Two pioneering works [47, 46] introduce Prompt-Tuning to CL and achieve much higher incremental performance than previous approaches, demonstrating the advantage of using pre-trained models in CL. Side-Tuning [53] adopts a technique similar to Adapter-Tuning but requires the task identity of the inference sample. In this paper, we propose a

unified framework for Memory-Free CL that can incorporate various types of PET modules. Particularly, we focus on practical Class-Incremental Learning with the potential to extend our LAE to other CL scenarios in future work.

3. Preliminaries

3.1. Continual Learning Formulation

We focus on Continual Learning with incremental classes, *i.e.*, Class-Incremental Learning (CIL), where a model sequentially learns tasks $\mathcal{T} := \{\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_n\}$, the i^{th} task \mathcal{T}_i has $|\mathcal{T}_i|$ categories, the train set of \mathcal{T}_i is denoted as \mathcal{D}_i , and the categories are non-overlapping between tasks. The model $f(\cdot; \theta, \phi)$ predicts the category label $y \in \mathcal{Y}$ for a given sample $\mathbf{x} \in \mathcal{X}$ of the learned tasks, where \mathcal{Y} is all seen categories, θ and ϕ are parameters of feature extractor and classification head, respectively. In this paper, the feature extractor is a pre-trained model parameterized by θ_{pre} attached with the Parameter-Efficient Tuning module parameterized by θ_{pet} , and $\phi = \text{concatenate}(\phi_{old}, \phi_{new})$, where ϕ_{old} and ϕ_{new} are classifiers of all learned tasks $\mathcal{T}_{1:i}$ and the current learning task \mathcal{T}_i . Since θ_{pre} and ϕ_{old} are kept fixed during learning a new task, we may omit them for concise in the rest of this paper.

3.2. Parameter-Efficient Tuning Revisit

Parameter-Efficient Tuning (PET) keeps the pre-trained model frozen and tunes a small number of additional learnable parameters, called PET module in the paper. Below we revisit several representative PET modules, in which g is the module that PET attached to, \mathbf{e} and \mathbf{h} are input and output of the original g and \mathbf{h}' is output of g attached with PET.

Adapter [16] is a small module that can be inserted to any layer (*i.e.*, g) of the pre-trained model. As shown in Fig. 2 (b), the adapter is generally a residual block composed of a down-projection with parameters \mathbf{W}_{down} , a nonlinear activation function $\sigma(\cdot)$, and an up-projection with parameters \mathbf{W}_{up} . The two projections can be convolution [37] for CNN or linear [16] layers for Transformer architectures, respectively. We formulate the adapter as follows:

$$\mathbf{h}' = \mathbf{h} + \sigma(\mathbf{h} * \mathbf{W}_{down}) * \mathbf{W}_{up}, \quad (1)$$

where the $*$ is matrix multiplication or convolution operation, σ is the activation function. Alternatively, the adapter can also be parallel with g like a residual branch [53, 10]:

$$\mathbf{h}' = \mathbf{h} + \sigma(\mathbf{e} * \mathbf{W}_{down}) * \mathbf{W}_{up}. \quad (2)$$

LoRA [17] assumes the change of parameters is in a low-rank space when tuning the pre-trained model on a downstream task. For a linear layer with weight $\mathbf{W} \in \mathbb{R}^{d \times d'}$, the weight updates $\Delta \mathbf{W}$ can be decomposed into the multiplication of two small matrices:

$$\Delta \mathbf{W} = \mathbf{W}_{down} \mathbf{W}_{up}, \quad (3)$$

where $\mathbf{W}_{down} \in \mathbb{R}^{d \times r}$ and $\mathbf{W}_{up} \in \mathbb{R}^{r \times d'}$. For the convolution layer, the updates can be reshaped into the kernel shape. Finally, LoRA modifies the forward pass of the adapted layer into the following form:

$$\mathbf{h}' = \mathbf{h} + \mathbf{e} * (\mathbf{W}_{down} \mathbf{W}_{up}), \quad (4)$$

where $*$ is matrix multiplication or convolution operation, the bias and reshape operation are omitted for conciseness. Since LoRA adapts the weight of g , the weight updates can be merged into g to reduce the inference latency.

Prefix [27] and **Prompt** [25] are learnable tokens prepended to the input of a transformer block or keys and values of the attention module. Given two sets of prefix tokens $\mathbf{P}_k, \mathbf{P}_v \in \mathbb{R}^{l \times d}$ the attention module is modified as:

$$\mathbf{h}' = \text{Attn}(\mathbf{x} \mathbf{W}_q, [\mathbf{P}_k, \mathbf{e} \mathbf{W}_k], [\mathbf{P}_v, \mathbf{e} \mathbf{W}_v]), \quad (5)$$

where $[\cdot, \cdot]$ is concatenate, and Attn is defined as:

$$\text{Attn}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) := \text{softmax}\left(\frac{\mathbf{Q} \mathbf{K}^T}{\sqrt{d}}\right) \mathbf{V},$$

and the multi-head mechanism is omitted for conciseness.

In this paper, we follow [10] to add a learnable scale parameter s to the parallel Adapter (Eq. 2) and LoRA (Eq. 4) respectively to obtain:

$$\mathbf{h}' = \mathbf{h} + s \cdot \sigma(\mathbf{e} * \mathbf{W}_{down}) * \mathbf{W}_{up}, \quad (6)$$

$$\mathbf{h}' = \mathbf{h} + s \cdot \mathbf{e} * (\mathbf{W}_{down} \mathbf{W}_{up}). \quad (7)$$

The Eqs. (6) and (7) are general forms of Eqs. (2) and (4), and they degenerate to Eq. (2) and (4) when s is the constant 1. We use these PETs in Eqs. (6) and (7) rather than Eqs. (1) and (4) in our experiments (Sec. 5).

In addition to the three PET modules described above, there are many others, such as AdaptBias [7], Compacter [33], and AdapterFormer [3], and there will be new and superior PET methods in future as well. All of them can be applied to our CL framework (see Sec. 4.2), as long as they are compatible with the pre-trained model.

4. Methodology

4.1. Naive Baseline

We construct a baseline by leveraging pre-trained models and PET techniques, following the naive sequential fine-tuning (Seq-FT) that was usually considered as the lower bound of CIL. Intuitively, our baseline creates a PET module and attaches it to the pre-trained model, then sequentially learns tasks in the same way as Seq-FT but keeping the pre-trained model frozen. We chose the local cross-entropy loss (CE) rather than the global CE loss as the learning objective for Seq-FT and our baseline because local CE

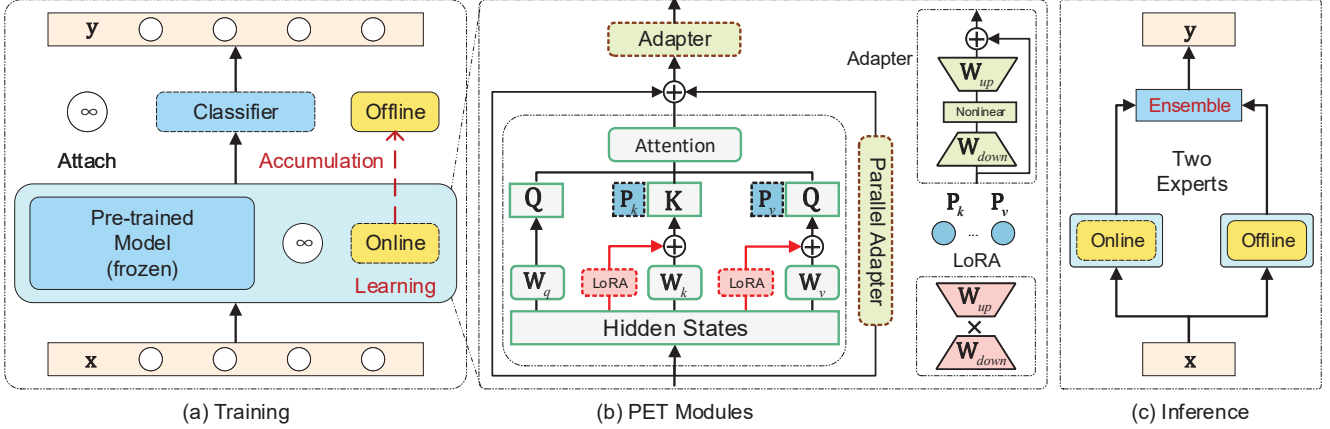


Figure 2: **Illustration of our LAE framework.** The left (a) is the training process, the right (c) is the inference flow, and the middle (b) lists some representative Parameter-Efficient Tuning (PET) modules attached to a transformer attention block, modules connected with dashed lines are optional and we use one of the PET modules in the experiments. There is no residual connection in Parallel Adapter. “Online” is a PET module to learn knowledge from the new task and “Offline” is a PET module to accumulate knowledge. The pre-trained model is omitted in the inference flow for concise.

empirically performs better than global CE when using a large pre-trained model [47, 46, 45]. The local CE is the standard CE computed on categories of the current task:

$$\mathcal{L} = \frac{1}{|\mathcal{D}_i|} \sum_{(\mathbf{x}, y) \in \mathcal{D}_i} \mathcal{L}_{ce}(\text{mask}(f(\mathbf{x}; \theta, \phi)), y), \quad (8)$$

where y is the ground truth label of the input \mathbf{x} in the current training set \mathcal{D}_i , $\text{mask}(\cdot)$ is a function that filters out the logits of old categories. The Eq. (8) falls back to global CE when $\text{mask}(\cdot)$ is removed. Although our baseline is very naive, the performance is comparable to the state-of-the-art DualPrompt when using the same Prefix module.

4.2. Proposed Framework

LAE builds upon our network architecture generalizable baseline and additionally introduces three novel designs, yielding a robust framework that can readily reshape any PET methods into a competitive Continual Learning approach. In the following, we will delve into these three key aspects of LAE: learning, accumulation, and ensemble.

Learning with calibrated speed. We observed that *PET modules vary in their speed for acquiring new knowledge, leading to disparities in performance.* In theory, adapting the PET module to a new task too quickly can lead to overfitting and result in worse catastrophic forgetting, whereas slower adaptation can maintain the model’s stability but limit its plasticity. We argue that aligning the adaptation speeds of different PET modules is crucial for transforming them into an efficient and robust CL approach in a unified way. To address this, we propose calibrating PET modules to align their adaptation speeds.

Moreover, as the same PET module θ_{pet} is shared by new and old tasks, the changes made to θ_{pet} for the new

task can cause forgetting for the old tasks. Therefore, slowing down the change in θ_{pet} , e.g., by reducing the learning rate (see supplementary), can help mitigate catastrophic forgetting. Kumar et al. [22] showcase how the linear probing followed by fine-tuning strategy effectively balances performance across out-of-distribution and in-distribution tasks. We employ a similar technique to calibrate the θ_{pet} ’s adaptation speed relative to ϕ_{new} . Specifically, at the beginning of the training, we only learn ϕ_{new} with θ_{pet} frozen; then after ϕ_{new} has sufficiently learned and the loss has significantly decreased, we jointly learn both ϕ_{new} and θ_{pet} .

According to the study by He et al. [10], the Prefix can be equivalently transformed into a similar form to the Adapter:

$$\mathbf{h}' \leftarrow (1 - \lambda(\mathbf{e}))\mathbf{h} + \lambda(\mathbf{e})\sigma(\mathbf{e}\mathbf{W}_1) \mathbf{W}_2, \quad (9)$$

in which $\mathbf{W}_1 = \mathbf{W}_q \mathbf{P}_k^\top$, $\mathbf{W}_2 = \mathbf{P}_v$, $\sigma = \text{softmax}$ and

$$\lambda(\mathbf{e}) = \frac{\sum_i \exp(\mathbf{e}\mathbf{W}_q \mathbf{P}_k^\top)_i}{\sum_i \exp(\mathbf{e}\mathbf{W}_q \mathbf{P}_k^\top)_i + \sum_j \exp(\mathbf{e}\mathbf{W}_q \mathbf{W}_k^\top \mathbf{C}^\top)_j}. \quad (10)$$

As \mathbf{P}_k usually contains much fewer tokens than input \mathbf{C} , $\lambda(\mathbf{e})$ is often a small positive number close to 0, which impacts the gradient of Prefix tokens \mathbf{P}_v :

$$\frac{\partial \mathcal{L}}{\partial \mathbf{P}_v} = \left(\frac{\partial \mathbf{h}'}{\partial \mathbf{P}_v} \right)^\top \frac{\partial \mathcal{L}}{\partial \mathbf{h}'} = \lambda(\mathbf{e}) (\sigma(\mathbf{e}\mathbf{W}_1))^\top \frac{\partial \mathcal{L}}{\partial \mathbf{h}'}. \quad (11)$$

Therefore, the gradient of \mathbf{P}_v is significantly smaller than \mathbf{W}_{up} of the corresponding Adapter parameterized by $\mathbf{W}_{down} = \mathbf{W}_1$ and $\mathbf{W}_{up} = \mathbf{W}_2$, and we can arrive at a similar conclusion regarding \mathbf{P}_k and \mathbf{W}_{down} . Then, we can easily observe that Prefix adapts to the new task much slower than Adapter. This is partly why the prompts for different

tasks in prior approaches [47, 46] are prone to be homogeneous (see supplementary). Here we align Prefix with Adapter by compensating its gradient by $\frac{1}{\lambda(\mathbf{e})}$ and adding two learnable scaling parameters s^k and s^v , calibrating the Prefix described by Eq. (9) into the following form:

$$\mathbf{h}' \leftarrow (1 - \lambda(\mathbf{e}))\mathbf{h} + \sigma(s^k \cdot \mathbf{e}\mathbf{W}_1)(s^v \cdot \mathbf{W}_2). \quad (12)$$

The adaptation speed of the calibrated Prefix is nearly equivalent to the Adapter depicted in Eq. (6) and elevates its performance to be on par with the Adapter. For other PET modules, we can also analyze them specifically and then calibrate their adaptation speeds to align with Adapter.

By aligning the adaptation speed of PET modules and calibrating their adaptation speed relative to the classifiers, our framework achieves a better and more consistent stability-plasticity balance with various PET modules.

Accumulation of multi-task knowledge. The PET module θ_{pet} is designed to continuously adapt to new tasks, making the model more proficient in dealing with novel tasks. However, this adaptation process can result in the model gradually forgetting how to handle older tasks. To address this issue, we propose to create an additional expert for older tasks to complement the expert for newer tasks, drawing inspiration from the Complementary Learning System [34, 23] of the human brain, which involves the hippocampus rapidly learning new knowledge and the neocortex integrating learned knowledge in an offline manner over time. We achieve this by duplicating the online PET module θ_{pet}^{on} (i.e., the θ_{pet} in the baseline) attached to the model as the offline PET module θ_{pet}^{off} after the model has learned the first task. The θ_{pet}^{off} slowly accumulates the learned knowledge when the model learns a new task by an accumulation function, and we empirically find the simple Exponential Moving Average (EMA) algorithm works well for our LAE:

$$\theta_{pet}^{off} \leftarrow \alpha \cdot \theta_{pet}^{off} + (1 - \alpha) \cdot \theta_{pet}^{on}, \quad (13)$$

where $\alpha \in (0, 1)$ is a large (i.e., close to 1) weight decay.

This way, the neocortex-like offline PET module gradually integrates the learned knowledge in a slow offline manner, while the hippocampus-like online PET module continues to rapidly learn new knowledge. Then, we can obtain two experts for newer tasks and older tasks with θ_{pet}^{on} and θ_{pet}^{off} , respectively. However, the task a sample belongs to is unknown during inference, we need to devise a method to effectively utilize both experts for inference.

Ensemble of two expert models. Two expert models constructed with θ_{pet}^{on} and θ_{pet}^{off} are respectively proficient at handling newer and older tasks. Instead of inference only using the online or offline expert model, we integrate their outputs to obtain the prediction for an inference sample.

A classifier can be viewed as an energy model when we define the unnormalized negative log probability as the energy function [24]. The optimization goal of the energy model is to minimize the energy of the model on the data distribution of its learning task. Previous research [29] has shown that the energy of an energy model trained on one data domain is generally very high on other data domains. Therefore, the Eq. (8) actually continuously minimizes the energy of the ϕ_{new} on the new task. Even if old data is not used during training, the energy of the old data on ϕ_{new} will be very high, as demonstrated by the recent work ESN [45].

Likewise, as the θ_{pet}^{on} and θ_{pet}^{off} respectively contain relatively more novel and historical knowledge, theoretically, the energy produced by θ_{pet}^{on} for the sample of the newer tasks should be smaller than that produced by θ_{pet}^{off} , and the vice versa for the sample of older tasks. Therefore, choosing the prediction result with the lowest energy as the final prediction of an inference sample seems like a simple but effective solution. However, in practice, we find that normalizing the energy produced by θ_{pet}^{on} and θ_{pet}^{off} before ensemble yields more robust results. Therefore, we adopt the following ensemble algorithm instead:

$$f_{ens}(\mathbf{o}^{on}, \mathbf{o}^{off}) := \max(\sigma(\mathbf{o}^{on}), \sigma(\mathbf{o}^{off})), \quad (14)$$

where σ is the softmax function, \mathbf{o}^{on} and \mathbf{o}^{off} are outputs of the online and offline expert models (i.e., $f(\cdot; \theta_{pet}^{on}, \phi)$ and $f(\cdot; \theta_{pet}^{off}, \phi)$) for an inference sample, respectively.

As illustrated in Fig. 2, in our LAE framework, the model learns a new task with θ_{pet}^{on} and accumulates the learned knowledge to θ_{pet}^{off} , the two experts favored by newer and older tasks are ensembled to get the final prediction for an inference sample. Our LAE can be applied to the pre-trained model in any network architecture as long as the PET modules are compatible with the model.

5. Experiment

5.1. Datasets and Evaluation Protocols

Our experiments use models pre-trained on the ImageNet21k [40] dataset without specified, and we follow prior works to train and evaluate the model on CIFAR100 [21] and ImageNet-R [14] benchmarks.

CIFAR100 is an extensively used dataset in prior continual learning (CL) works, containing 100 classes, each class with 500 training and 100 test images of size $32 \times 32 \times 3$.

ImageNet-R is first introduced to CL by Wang et al. [46], including 200 subcategories of ImageNet [40], but its samples are in different styles, such as cartoon, graffiti, and origami. There are also some hard examples from ImageNet that standard models, e.g., ResNet [13], fail to classify. The original dataset is split into the train set with 24000 samples and the test set with 6000 samples, and the number of training and testing samples varies between classes.

Table 1: Benchmark Results on CIFAR100. The PET modules are inserted into the first 5 transformer blocks of the standard ViT-B/16 pre-trained on the ImageNet21k dataset. The “5, 10, 20” indicate the size of PET modules.

Approach	PET Module	A_{10} (\uparrow)	\bar{A}_{10} (\uparrow)
Joint-FT	-	92.00 \pm 0.18	-
Seq-FT	-	77.61 \pm 0.37	85.82 \pm 0.86
Baseline	Adapter5	82.30 \pm 1.20	88.18 \pm 0.31
	Adapter10	81.76 \pm 1.21	87.84 \pm 0.49
	LoRA5	83.24 \pm 1.54	88.48 \pm 0.31
	LoRA10	82.55 \pm 1.61	88.35 \pm 0.48
	Prefix10	84.49 \pm 0.30	89.34 \pm 0.59
	Prefix20	84.44 \pm 0.75	89.46 \pm 0.42
L2P [47]	Prompt	82.57 \pm 0.42	86.95 \pm 0.68
DualPrompt [46]	Prefix20	84.27 \pm 0.41	88.92 \pm 0.78
ESN [45]	Prompt	84.18 \pm 0.08	88.49 \pm 0.64
LAE (Ours)	Adapter5	85.59 \pm 0.46	89.96 \pm 0.44
	Adapter10	85.33 \pm 0.20	89.77 \pm 0.50
	LoRA5	85.56 \pm 0.16	89.63 \pm 0.41
	LoRA10	85.37 \pm 0.39	89.87 \pm 0.50
	Prefix10	85.17 \pm 0.14	89.73 \pm 0.43
	Prefix20	85.25 \pm 0.66	89.71 \pm 0.42

We follow prior works to split the dataset into 10 tasks, and all tasks have the same number of classes, *i.e.*, 10 for CIFAR100 and 20 for ImageNet-R. We evaluate the model by the widely used incremental metrics: last incremental accuracy A_N and average incremental accuracy $\bar{A}_N = \frac{1}{N} \sum_{i=1}^N A_i$, where N is the total number of tasks (*i.e.*, 10), and A_i is formally defined as:

$$A_i = \frac{1}{|\mathcal{D}_{1:i}^{test}|} \sum_{(\mathbf{x}, y) \in \mathcal{D}_{1:i}^{test}} \mathbb{1}(\hat{y} = y), \quad (15)$$

where $\mathbb{1}(\cdot)$ is the indicator function that maps the boolean value to $\{0, 1\}$, $\mathcal{D}_{1:i}^{test}$ is the test set of all seen tasks so far, \hat{y} and y are predicted and ground truth labels of input \mathbf{x} . *We ran all experiments 3 times with different class orders and report the mean and standard deviation of these 3 runs.*

5.2. Implementation and Training Details

To make fair comparisons, we consider state-of-the-art approaches [47, 46] based on pre-trained models like our LAE and using the PyTorch code released by Jaeho Lee to conduct experiments. The joint fine-tuning (Joint-FT) and the naive sequential fine-tuning (Seq-FT) usually recognized as the upper and lower bounds of CIL are implemented in our codebase, referring to the code of Jaeho Lee. We also compare with recent work ESN [45], using its official PyTorch code. We chose three types of representative PET modules and two sizes per type for our baseline and LAE framework, where *the size denotes the down-projection dimension of the Adapter, the rank of LoRA, or the length of the Prefix* described in Sec. 3. We assume

<https://github.com/JH-LEE-KR>

Table 2: Benchmark Results on ImageNet-R. The PET modules are inserted into the first 5 transformer blocks of the standard ViT-B/16 pre-trained on the ImageNet21k dataset. The “5, 10, 20” indicate the size of PET modules.

Approach	PET Module	A_{10} (\uparrow)	\bar{A}_{10} (\uparrow)
Joint-FT	-	79.69 \pm 0.16	-
Seq-FT	-	40.42 \pm 3.28	59.39 \pm 2.36
Baseline	Adapter5	61.63 \pm 2.51	71.58 \pm 2.33
	Adapter10	57.08 \pm 3.67	68.58 \pm 2.89
	LoRA5	60.79 \pm 2.63	70.50 \pm 2.23
	LoRA10	57.62 \pm 3.74	68.21 \pm 2.85
	Prefix10	68.94 \pm 1.25	75.31 \pm 1.51
	Prefix20	68.99 \pm 0.98	75.38 \pm 1.43
L2P [47]	Prompt	63.91 \pm 1.60	69.27 \pm 2.25
DualPrompt [46]	Prefix20	68.99 \pm 0.08	74.21 \pm 1.13
ESN [45]	Prompt	62.61 \pm 0.96	68.58 \pm 1.64
LAE (Ours)	Adapter5	72.66 \pm 0.63	78.91 \pm 0.89
	Adapter10	72.45 \pm 0.81	79.07 \pm 0.88
	LoRA5	72.00 \pm 0.75	78.33 \pm 0.96
	LoRA10	71.83 \pm 0.57	78.24 \pm 0.90
	Prefix10	71.85 \pm 0.66	77.44 \pm 1.12
	Prefix20	72.05 \pm 0.66	77.55 \pm 1.00

that only a single PET module is attached to the pre-trained model in the previous discussion for convenience, in practice, multiple PET modules are inserted into the Attention blocks of Transformers or the convolution blocks of ConvNets in the shallow layers, following DualPrompt [46].

The training strategy of our baseline and LAE framework is the same as DualPrompt, *i.e.*, training the model with Adam optimizer for 5 and 50 epochs, and constant learning rate 0.03 and 0.005 based on batch size 256, for CIFAR100 and ImageNet-R, respectively. The EMA algorithm’s weight decay α defined in Eq. (13) is empirically set to 0.9999 in all experiments. The freezing epochs of PET modules are set to 3 and 30 for CIFAR100 and ImageNet-R, respectively. The data augmentation is consistent with that used in model pre-training. We train Joint-FT and Seq-FT with the recommended fine-tuning strategy of ViT [5], but the number of training epochs is the same as ours. More details can be found in the supplementary materials.

5.3. Benchmark Results

CIFAR100 benchmark results are present in Tab. 1. All approaches use the same ViT-B/16 [5] model pre-trained on the ImageNet21k [40] dataset. The numerical suffix of the PET module denotes its size (*i.e.*, down-projection dimension or length). L2P and DualPrompt are state-of-the-art approaches that adopt a pool to store Prompt or Prefix. However, the accuracy of their prompt selection gradually declines with the increase in the number of learning tasks and the prompts for different tasks appear homogeneous (see supplementary). Therefore, our baseline is very naive but achieves comparable performance to L2P and DualPrompt, and our LAE framework with all 6 PET modules consistently surpasses DualPrompt and ESN by about 1.5% in last

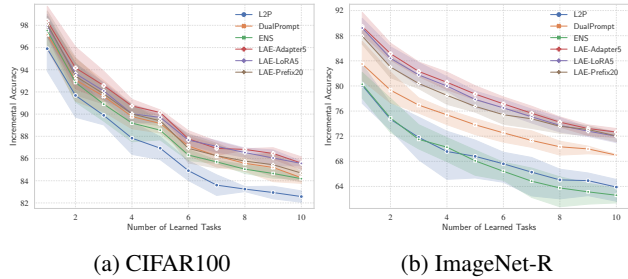


Figure 3: Task-by-Task Incremental Accuracy on two benchmarks. The lines show the task-by-task evaluation results of L2P [47], DualPrompt [46], ENS [45], and our LAE framework with different PET modules.

incremental accuracy A_{10} . Although PET modules have different performances in the baseline, they achieve better and same level performance in our LAE, mainly due to the calibration of adaptation speed. In particular, DualPrompt has 3-10x more learnable parameters than our LAE.

ImageNet-R benchmark is more difficult than CIFAR100, but it can better demonstrate the advantages of our LAE framework. From the results shown in Tab. 2, our baseline can only achieve comparable performance to DualPrompt when using Prefix. This is because the Adapter and LoRA adapt to a new task faster than Prefix, which is enlarged in the ImageNet-R dataset but successfully addressed by the adaption speed calibration of our LAE framework. Thus, we can see that our LAE achieves more than 3.5% performance improvement over DualPrompt in terms of the last incremental accuracy A_{10} , which is also scaled up compared to the easier CIFAR100 dataset. We can also observe that the size of the PET modules has little impact on performance in our LAE framework, and our LAE is more robust to the class order, while our baseline has a relatively large variance between different class orders.

The real-world CL is an endless procedure, and the performance of each learning phase is equally important to the AI system. So, we also plot the task-by-task incremental accuracy in Fig. 3a and 3b. We can observe that our LAE with all three types of PET modules performs better than L2P and DualPrompt at almost all learning phases. *Our LAE outperforms others by a wider margin in the 20-task experiments presented in the supplementary material, highlighting its ability to handle long-term CL scenarios.* In the supplementary, we compare our LAE with the contemporary CODA-Prompt [44] approach on ImageNet-R and DomainNet [35] datasets. The parameters and computation comparison can also be found in the supplementary.

5.4. Ablation Study

Our LAE consists of three main novel designs, *i.e.*, learning, accumulation, and ensemble, so we ablate on them and report the results in Tab. 3. The first and last rows are our baseline and LAE framework, respectively. The per-

Table 3: Ablation study on three key designs of our LAE framework: Learning (Lea.), Accumulation (Acc.), and Ensemble (Ens.). The experiments are conducted on the ImageNet-R dataset with Adapter10.

Lea.	Ens.	Acc.	A_{10} (\uparrow)	\bar{A}_{10} (\uparrow)
\times	\times	\times	57.08 ± 3.67	68.58 ± 2.89
\times	\checkmark	\times	63.36 ± 1.46	73.78 ± 1.54
\times	\checkmark	\checkmark	65.22 ± 2.56	75.71 ± 1.18
\checkmark	\times	\times	70.30 ± 1.48	77.25 ± 1.39
\checkmark	\checkmark	\times	71.37 ± 1.24	78.27 ± 1.30
\checkmark	\times	\checkmark	72.80 ± 0.81	78.63 ± 0.95
\checkmark	\checkmark	\checkmark	72.45 ± 0.81	79.07 ± 0.88




Figure 4: Ablation on inference with Online Adapter, Offline Adapter, and our Experts Ensemble. The three strategies correspond to rows 4, 6, and 7 in Table 3.

formance drops the most when removing our learning with calibrated speed, demonstrating that it contributes the most to our LAE. Accumulation and Ensemble are also important to our LAE, without them the last incremental accuracy decreases by 2.15%. The sixth row indicates our LAE inference with the Offline PET module only, whose A_{10} is even better than the inference by our expert ensemble. As illustrated in Fig. 4, inference by expert ensemble performs better than inference with the Online or Offline PET module alone in the earlier learning phases. However, as the number of learned tasks increases, the advantage of expert ensemble over the Offline PET module gradually decreases, partly due to the performance of the old tasks dominating the overall performance. Nonetheless, inference with expert ensemble yields more robust performance in most cases.

We also conducted ablation experiments on the calibration made to Prefix. We can see from Tab. 4 that both gradient compensation and learnable scaling parameters individually lead to significant improvements in performance. Moreover, when used together, the performance gain is approximately equal to the sum of the gains achieved by using each of them separately, indicating that their contributions to the performance are independent of each other.

5.5. Attach Position of PET Modules

Our LAE inserts PET modules directly into the first 5 Transformer blocks, following the DualPrompt. As shown

Table 4: Ablation study on the calibration of Prefix: gradient compensation and learnable scaling parameters. The experiments are conducted on ImageNet-R with Prefix20.

Compensation	Scale	A_{10} (\uparrow)	\bar{A}_{10} (\uparrow)
\times	\times	69.23 \pm 0.47	75.11 \pm 1.30
\checkmark	\times	70.53 \pm 0.55	76.16 \pm 1.16
\times	\checkmark	70.38 \pm 0.52	76.46 \pm 1.25
\checkmark	\checkmark	72.05 \pm 0.66	77.55 \pm 1.00




Figure 5: Ablation on the attached position of PET module (left) and the number of inserted blocks (right). There are 12 transformer blocks in the pre-trained model, “1-5” indicates attaching PET modules to the first 5 transformer blocks. The PET modules are inserted into the transformer starting from the first block in the right figure.

in Fig. 5 (left), inserting the PET modules in the shallowest position produces better results than inserting them in deeper positions, which is consistent with the observation in DualPrompt. Additionally, Fig. 5 (right) shows that inserting PET modules in the first 6 Transformer blocks achieves the best performance while inserting them in the first 5 Transformer blocks (*i.e.*, the default setting of our LAE) also results in nearly the same performance.

5.6. Results on Transformer variant and ConvNet

Prefix and Prompt are not flexible enough to be applied to ConvNets and Transformer variants, while our LAE is model architecture generalizable due to the ability to leverage various PET modules. We choose the Swin Transformer [31] and ConvNeXt [32] to validate our LAE.

Swin Transformer is a representative window-based Vision Transformer, but L2P and DualPrompt cannot be directly applied to it because the inserted tokens may disrupt the proper division of the windows. Therefore, we only compare our LAE with our baseline when using Swin Transformer, and report the results in Tab. 5. We can see that our LAE achieves better performance with Swin-B than with ViT-B/16, which is mainly due to the superior performance of Swin-B. Moreover, our LAE significantly improves the performance of our baseline on all two datasets.

ConvNeXt is a modern ConvNet for the 2020s that outperforms Swin Transformer by incorporating several novel designs into the standard ResNet [13]. Similarly, we compare

Table 5: The comparison between our LAE framework with our baseline on two datasets using the Swin-B model pre-trained on the ImageNet22k dataset. Adapters are inserted in the shallower 10 of 24 transformer blocks.

Approach	Dataset	A_{10} (\uparrow)	\bar{A}_{10} (\uparrow)
Baseline	CIFAR100	85.85 \pm 0.69	90.36 \pm 0.67
LAE (Ours)		86.52 \pm 0.38	90.58 \pm 0.61
Baseline	ImageNetR	71.81 \pm 1.09	78.91 \pm 1.41
LAE (Ours)		73.38 \pm 0.70	80.01 \pm 1.51

Table 6: The comparison between our LAE framework with our baseline on two datasets using the ConvNeXt-B model pre-trained on the ImageNet22k dataset. Adapters are inserted in the shallower 15 of 36 convolution blocks.

Approach	Dataset	A_{10} (\uparrow)	\bar{A}_{10} (\uparrow)
Baseline	CIFAR100	86.40 \pm 0.07	91.00 \pm 0.31
LAE (Ours)		87.01 \pm 0.28	91.18 \pm 0.34
Baseline	ImageNetR	76.35 \pm 1.36	82.60 \pm 1.79
LAE (Ours)		78.38 \pm 0.80	83.95 \pm 1.16

our LAE with the baseline in Tab. 6. The Adapter’s down and up projections are implemented using 1×1 convolution layers. Both our baseline and LAE achieve significantly better performance using ConvNeXt-B compared to using ViT-B/16 and Swin-B, highlighting LAE’s strengths beyond being limited to Transformers. Our LAE consistently improves the performance of our baseline on both datasets.

6. Conclusion

This paper thoroughly studied the novel Continual Learning (CL) paradigm that starts with a pre-trained model and continuously adapts the model to arriving tasks utilizing general Parameter-Efficient Tuning (PET) methods. We constructed a naive baseline that achieved performance comparable to the prior state-of-the-art approaches. We proposed the Learning-Accumulation-Ensemble (LAE) framework by introducing three novel designs to the baseline. Our LAE can convert any PET method into an efficient CL approach without accessing any old data. We conducted extensive experiments to validate the effectiveness of our LAE, and the results demonstrated that our LAE significantly outperforms the previous state-of-the-art approaches.

Limitations. There are still some limitations that need to be improved in the future, such as how to accumulate knowledge more efficiently and better ensemble expert models. Moreover, due to the lack of large-scale datasets that do not overlap with the pre-training dataset, our LAE has not been verified in CL scenarios with a larger number of tasks.

Overall, this paper provides a new solution for Memory-Free CL and offers some theoretical and experimental references for future research on this novel CL paradigm.

References

- [1] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [2] Lama Alssum, Juan Leon Alcazar, Meray Ramazanova, Chen Zhao, and Bernard Ghanem. Just a glimpse: Rethinking temporal information for video continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshop*, 2023.
- [3] Shoufa Chen, Chongjian Ge, Zhan Tong, Jiangliu Wang, Yibing Song, Jue Wang, and Ping Luo. Adaptformer: Adapting vision transformers for scalable visual recognition. *CoRR*, 2022.
- [4] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2019.
- [5] 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 *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021.
- [6] Arthur Douillard, Matthieu Cord, Charles Ollion, Thomas Robert, and Eduardo Valle. Podnet: Pooled outputs distillation for small-tasks incremental learning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.
- [7] Chin-Lun Fu, Zih-Ching Chen, Yun-Ru Lee, and Hung-yi Lee. Adapterbias: Parameter-efficient token-dependent representation shift for adapters in NLP tasks. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2022.
- [8] Qiankun Gao, Chen Zhao, Bernard Ghanem, and Jian Zhang. R-DFCIL: relation-guided representation learning for data-free class incremental learning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022.
- [9] Ian J Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*, 2013.
- [10] Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. Towards a unified view of parameter-efficient transfer learning. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2022.
- [11] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross B. Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [12] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [14] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, Dawn Song, Jacob Steinhardt, and Justin Gilmer. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [15] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [16] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for NLP. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2019.
- [17] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2022.
- [18] Steven C. Y. Hung, Cheng-Hao Tu, Cheng-En Wu, Chien-Hung Chen, Yi-Ming Chan, and Chu-Song Chen. Compacting, picking and growing for unforgetting continual learning. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [19] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge J. Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022.
- [20] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences (PNAS)*, 2017.
- [21] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *Technical Report*, 2009.
- [22] Ananya Kumar, Aditi Raghunathan, Robbie Matthew Jones, Tengyu Ma, and Percy Liang. Fine-tuning can distort pre-trained features and underperform out-of-distribution. In *Proceedings of the International Conference on Learning Representations (ICLR)*. OpenReview.net, 2022.
- [23] Dharshan Kumaran, Demis Hassabis, and James L McClelland. What learning systems do intelligent agents need? complementary learning systems theory updated. *Trends in cognitive sciences*, 2016.
- [24] Yann LeCun, Sumit Chopra, Raia Hadsell, M Ranzato, and Fugie Huang. A tutorial on energy-based learning. *Predicting structured data*, 2006.

- [25] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2021.
- [26] Xilai Li, Yingbo Zhou, Tianfu Wu, Richard Socher, and Caiming Xiong. Learn to grow: A continual structure learning framework for overcoming catastrophic forgetting. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2019.
- [27] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. In *Proceedings of the Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL-IJCNLP 2021)*, 2021.
- [28] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2017.
- [29] Weitang Liu, Xiaoyun Wang, John D. Owens, and Yixuan Li. Energy-based out-of-distribution detection. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [30] Yaoyao Liu, Bernt Schiele, and Qianru Sun. Adaptive aggregation networks for class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [31] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [32] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [33] Rabeeh Karimi Mahabadi, James Henderson, and Sebastian Ruder. Compacter: Efficient low-rank hypercomplex adapter layers. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [34] James L McClelland, Bruce L McNaughton, and Randall C O'Reilly. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological review*, 1995.
- [35] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1406–1415, 2019.
- [36] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018.
- [37] Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. Learning multiple visual domains with residual adapters. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [38] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [39] Anthony V. Robins. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connect. Sci.*, 1995.
- [40] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)*, 2015.
- [41] Andrei A. Rusu, Neil C. Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *CoRR*, 2016.
- [42] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [43] James Smith, Yen-Chang Hsu, Jonathan Balloch, Yilin Shen, Hongxia Jin, and Zsolt Kira. Always be dreaming: A new approach for data-free class-incremental learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [44] James Seale Smith, Leonid Karlinsky, Vyshnavi Gutta, Paola Cascante-Bonilla, Donghyun Kim, Assaf Arbelle, Rameswar Panda, Rogerio Feris, and Zsolt Kira. Coda-prompt: Continual decomposed attention-based prompting for rehearsal-free continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [45] Yabin Wang, Zhiheng Ma, Zhiwu Huang, Yaowei Wang, Zhou Su, and Xiaopeng Hong. Isolation and impartial aggregation: A paradigm of incremental learning without interference. 2023.
- [46] Zifeng Wang, Zizhao Zhang, Sayna Ebrahimi, Ruoxi Sun, Han Zhang, Chen-Yu Lee, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer G. Dy, and Tomas Pfister. Dualprompt: Complementary prompting for rehearsal-free continual learning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022.
- [47] Zifeng Wang, Zizhao Zhang, Chen-Yu Lee, Han Zhang, Ruoxi Sun, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer G. Dy, and Tomas Pfister. Learning to prompt for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [48] Teng Xi, Yifan Sun, Deli Yu, Bi Li, Nan Peng, Gang Zhang, Xinyu Zhang, Zhigang Wang, Jinwen Chen, Jian Wang, Lufei Liu, Haocheng Feng, Junyu Han, Jingtuo Liu, Errui Ding, and Jingdong Wang. UFO: unified feature optimization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022.
- [49] Hongxu Yin, Pavlo Molchanov, Jose M. Alvarez, Zhizhong Li, Arun Mallya, Derek Hoiem, Niraj K. Jha, and Jan Kautz. Dreaming to distill: Data-free knowledge transfer via deep-inversion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [50] Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong learning with dynamically expandable net-

works. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.

- [51] Elad Ben Zaken, Yoav Goldberg, and Shauli Ravfogel. Bit-fit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2022.
- [52] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2017.
- [53] Jeffrey O. Zhang, Alexander Sax, Amir Zamir, Leonidas J. Guibas, and Jitendra Malik. Side-tuning: A baseline for network adaptation via additive side networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.