

InfLoRA: Interference-Free Low-Rank Adaptation for Continual Learning

Yan-Shuo Liang and Wu-Jun Li*

National Key Laboratory for Novel Software Technology,
Department of Computer Science and Technology, Nanjing University, P. R. China

liangys@smail.nju.edu.cn, liwujun@nju.edu.cn

Abstract

Continual learning requires the model to learn multiple tasks sequentially. In continual learning, the model should possess the ability to maintain its performance on old tasks (stability) and the ability to adapt to new tasks continuously (plasticity). Recently, parameter-efficient fine-tuning (PEFT), which involves freezing a pre-trained model and injecting a small number of learnable parameters to adapt to downstream tasks, has gained increasing popularity in continual learning. Although existing continual learning methods based on PEFT have demonstrated superior performance compared to those not based on PEFT, most of them do not consider how to eliminate the interference of the new task on the old tasks, which inhibits the model from making a good trade-off between stability and plasticity. In this work, we propose a new PEFT method, called interference-free low-rank adaptation (InfLoRA), for continual learning. InfLoRA injects a small number of parameters to reparameterize the pre-trained weights and shows that fine-tuning these injected parameters is equivalent to fine-tuning the pre-trained weights within a subspace. Furthermore, InfLoRA designs this subspace to eliminate the interference of the new task on the old tasks, making a good trade-off between stability and plasticity. Experimental results show that InfLoRA outperforms existing state-of-the-art continual learning methods on multiple datasets. Code is available at <https://github.com/liangyanshuo/InfLoRA>.

1. Introduction

Continual learning requires the model to learn multiple tasks sequentially [33]. To achieve continual learning, the model should possess two essential abilities, including the ability to keep its performance on the old tasks (stability) and the ability to adapt to the new tasks continuously (plasticity) [33]. Furthermore, two different scenar-

ios are often considered in continual learning, including task-incremental scenario [32] and class-incremental scenario [41]. Task-incremental scenario allows the model to get task identities during inference. On the contrary, class-incremental scenario does not allow the model to get task identities during inference, making the model learn to distinguish all the classes across all the tasks.

Recently, parameter-efficient fine-tuning (PEFT) [15, 16, 18], which involves freezing a pre-trained model and injecting a small number of learnable parameters to adapt to downstream tasks, has gained increasing popularity in continual learning [12, 38, 44], especially in the class-incremental scenario. More specifically, existing continual learning methods based on PEFT [21, 43] inject the learnable parameters into a pre-trained model using some popular PEFT methods such as prompt-tuning [25] or low-rank adaptation (LoRA) [16]. Subsequently, these methods freeze the pre-trained weights and sequentially fine-tune the injected parameters on multiple tasks throughout the continual learning process.

Although continual learning methods based on PEFT have demonstrated superior performance compared to those not based on PEFT [44], most of them do not consider how to eliminate the interference of the new task on the old tasks, which inhibits the model from making a good trade-off between stability and plasticity. Specifically, when learning a new task, existing continual learning methods based on PEFT either reuse the previously learned parameters to adapt to the new task [12, 44] or randomly expand some parameters first and then adapt to the new task [38, 42, 43]. During this process, the interference of the new task on the old tasks exists due to the shared parameters between new and old tasks, which means fine-tuning a pre-trained model on a new task may interfere with the model’s performance on the old tasks. As a result, it is hard for the model to make a good trade-off between stability and plasticity.

In this work, we propose a new PEFT method, called interference-free low-rank adaptation (InfLoRA), for continual learning. The contributions of this work are listed as follows:

*Wu-Jun Li is the corresponding author.

- InfLoRA injects a small number of parameters to reparameterize the pre-trained weights and shows that fine-tuning these injected parameters is equivalent to fine-tuning the pre-trained weights within a subspace.
- InfLoRA designs this subspace to eliminate the interference of the new task on the old tasks, making a good trade-off between stability and plasticity.
- Experimental results show that InfLoRA outperforms existing state-of-the-art continual learning methods on multiple datasets.

2. Related Work and Preliminaries

2.1. Related Work

Parameter-Efficient Fine-Tuning Parameter-efficient fine-tuning (PEFT) methods freeze a pre-trained model and inject a small number of learnable parameters to adapt to downstream tasks. In this way, PEFT methods reduce the inefficiency of full fine-tuning methods which fine-tune all the parameters of a pre-trained model to learn downstream tasks. For example, Adapter [15] adds small modules in different layers of Transformers and only tunes these added modules to adapt to downstream tasks. Prompt-tuning [25] and Prefix-tuning [27] insert a set of learnable tokens into the input of the Transformer layers and only tune these tokens to adapt to downstream tasks. Low-rank adaptation (LoRA) [16] reparameterizes the pre-trained weights with low-rank branches and only tunes these branches to adapt to the downstream tasks. Although these methods tune much fewer learnable parameters than full fine-tuning, they always show comparable or even superior performance compared with full fine-tuning [11, 16, 31, 45]. Early PEFT methods focus on natural language processing (NLP). Recently, PEFT methods have also been proposed for computer vision (CV). For example, visual prompt tuning (VPT) [18] and AdapterFormer [6] apply prompt-tuning and Adapter techniques to CV tasks, respectively. Both of them exhibit comparable performance to full fine-tuning.

Continual Learning Early continual learning was usually considered in the context of learning from scratch. Three types of continual learning methods are proposed, including regularization-based methods [1, 20, 23, 46], memory-based methods [2, 3, 7, 28, 39], and expansion-based methods [17, 26, 35]. Regularization-based methods employ a penalty loss (regularization) to prevent important parameters of old tasks from changing too much. Memory-based methods maintain a memory buffer to store information about old tasks. Expansion-based methods dynamically expand the model’s architecture for each new task.

Recently, with the advancements of pre-trained models [9, 10, 13], using pre-trained models for continual learning has gained increasing popularity. Some continual learning methods fully fine-tune the pre-trained models [4, 49],

which has been shown to be inefficient. Other methods explore PEFT methods in continual learning. For instance, some existing continual learning methods [21, 38, 43, 44] introduce prompt-tuning in continual learning, achieving much higher performance than previous methods that learn from scratch, especially in the class-incremental scenario. The method in [12] introduces a framework in continual learning that can be combined with many existing PEFT methods, such as prompt-tuning, LoRA and Adapter. However, all these methods do not consider how to eliminate the interference of the new task on the old tasks, which inhibits the model from making a good trade-off between stability and plasticity.

2.2. Preliminaries

We first introduce low-rank adaptation (LoRA) [16], a popular PEFT method related to our method. Then, we give the problem definition for continual learning.

Low-Rank Adaptation LoRA [16] is one of the most popular PEFT methods. It assumes that the changes of parameters lie in a low-rank space when the model is fully fine-tuned on a downstream task. Specifically, for a linear layer with the input dimension d_I and the output dimension d_O , we represent its weight with $\mathbf{W}^{d_O \times d_I}$. Then, LoRA reparameterizes the pre-trained weight \mathbf{W} by expanding a branch with two matrices, $\mathbf{A} \in \mathbb{R}^{d_O \times r}$ and $\mathbf{B} \in \mathbb{R}^{r \times d_I}$. Typically, r is much smaller than the input dimension d_I and output dimension d_O , making \mathbf{A} a dimensionality increasing matrix and \mathbf{B} a dimensionality reduction matrix. Finally, LoRA modifies the forward propagation in this linear layer as $e = \mathbf{W}h + \mathbf{A}Bh$. Here, h and e denote the input and output of this layer, respectively. LoRA initializes \mathbf{A} as $\mathbf{0}$ and initializes \mathbf{B} using Gaussian distribution. During the learning of the downstream tasks, LoRA freezes the pre-trained weight \mathbf{W} and only fine-tunes the parameters \mathbf{A} and \mathbf{B} .

Problem Definition In continual learning, there is a sequence of tasks with different distributions. We define the task sequence as $\mathcal{D} = \{\mathcal{D}_1, \dots, \mathcal{D}_T\}$, where the t -th task $\mathcal{D}_t = \{(\mathbf{x}_{i,t}, y_{i,t})\}_{i=1}^{n_t}$. Here, $\mathbf{x}_{i,t}$ denotes an input sample and $y_{i,t}$ denotes its label. The objective of continual learning is to train a model sequentially on these tasks and ensure that the model performs well on all of them.

We follow existing continual learning methods [43, 44] based on PEFT and assume the model is a pre-trained Vision Transformer (ViT) [10]. Specifically, assume the model is $h_{\Phi}(f_{\Theta}(\cdot))$ where $h_{\Phi}(\cdot)$ is the classifier with parameters Φ and $f_{\Theta}(\cdot)$ is the pre-trained ViT backbone with pre-trained parameters Θ . Similar to existing work [43], our focus is primarily on the class-incremental scenario, where task identities are unknown during inference. Furthermore, we concentrate on the exemplar-free setting [43, 51], where no historical data can be fetched for rehearsal.




Figure 1. (a) The architecture of our InfLoRA in a certain linear layer of a Transformer. During the learning of the t -th task, the pre-trained weight and all the old branches are frozen, and only A_t is fine-tuned. (b) The pipeline of designing dimensionality reduction matrix B_t .

3. Methodology

Figure 1 (a) illustrates the architecture of our InfLoRA within a linear layer. Before learning the t -th new task, our InfLoRA expands a LoRA-like branch, which includes a dimensionality reduction matrix $B_t \in \mathbb{R}^{r \times d_i}$ and a dimensionality increasing matrix $A_t \in \mathbb{R}^{d_o \times r}$. Then, the forward propagation of this linear layer is modified as

$$e = Wh + \sum_{j=1}^t A_j B_j h = W_{t-1} h + A_t B_t h = W_t h. \quad (1)$$

Here, $W_t = W_{t-1} + A_t B_t = W + \sum_{i=1}^t A_i B_i$. Similar to LoRA, our InfLoRA also initializes dimensionality increasing matrix A_t as $\mathbf{0}$. However, different from LoRA, which employs Gaussian distribution to initialize the dimensionality reduction matrix B , our InfLoRA designs the dimensionality reduction matrix B_t before learning the t -th task. During the learning of the t -th task, InfLoRA fine-tunes A_t to learn the new task while keeping the pre-trained weight W , all the old branches and the matrix B_t frozen. After learning the t -th tasks, for any given test sample belonging to the learned tasks, the model uses W_t and (1) to infer its label. This design ensures that our method is compatible with the class-incremental scenario where task identities are unknown during inference.

In the following subsections, we first build the relationship between our InfLoRA and the method that fine-tunes the pre-trained weight. Specifically, we show that fine-tuning parameters A_t is equivalent to fine-tuning the pre-trained weights W within a subspace spanned by the rows of B_t . Note that B_t is designed before learning the t -th task, making this subspace pre-designed. Then, building upon this relationship, we introduce how our InfLoRA designs this subspace to eliminate the interference of the new

task on the old tasks and make a good trade-off between stability and plasticity.

3.1. Relationship between InfLoRA and Fine-Tuning the Pre-Trained Weight

When the t -th task arrives and our method has expanded a new branch, the forward propagation in this layer can be represented by (1). At this time, we can prove the following proposition:

Proposition 1. *When learning the t -th task with forward propagation represented by (1), fine-tuning A_t is equivalent to fine-tuning the pre-trained weight W within the subspace $\text{span}\{\mathbf{b}_1^t, \dots, \mathbf{b}_r^t\}$. Here, \mathbf{b}_i^t ($1 \leq i \leq r$) denotes the i -th row vector of B_t .*

Proof. When tuning the pre-trained weight W to learn the t -th task, we can compute the gradient of W based on the chain rule:

$$\frac{\partial \mathcal{L}}{\partial W} = \frac{\partial \mathcal{L}}{\partial e} \frac{\partial e}{\partial W} = \frac{\partial \mathcal{L}}{\partial e} h^T. \quad (2)$$

Here, \mathcal{L} denotes the loss function. At this time, the change of W can be denoted as $\Delta W = -\alpha \frac{\partial \mathcal{L}}{\partial W}$, where α is the learning rate. Then, we can compute the change of the composed matrix $W_t = W + \sum_{j=1}^t A_j B_j$:

$$\begin{aligned} \Delta_W W_t &= [W + \Delta W + \sum_{j=1}^t A_j B_j] - (W + \sum_{j=1}^t A_j B_j) \\ &= \Delta W = -\alpha \frac{\partial \mathcal{L}}{\partial W_t} = -\alpha \frac{\partial \mathcal{L}}{\partial e} h^T \end{aligned} \quad (3)$$

Here, we use $\Delta_W W_t$ to denote the change of the composed matrix W_t causing by the change of W .

Similarly, when tuning the expanded weight A_t , we can get the gradient of A_t based on the chain rule:

$$\frac{\partial \mathcal{L}}{\partial A_t} = \frac{\partial \mathcal{L}}{\partial e} \frac{\partial e}{\partial A_t} = \frac{\partial \mathcal{L}}{\partial e} h^T B_t^T. \quad (4)$$

At this time, the change of A_t can be denoted as $\Delta A_t = -\alpha \frac{\partial \mathcal{L}}{\partial A_t}$. Then, we can compute the change of the composed matrix $W_t = W_{t-1} + A_t B_t$:

$$\begin{aligned} \Delta_{A_t} W_t &= [W_{t-1} + (A_t + \Delta A_t) B_t] - (W_{t-1} + A_t B_t) \\ &= \Delta A_t B_t = -\alpha \frac{\partial \mathcal{L}}{\partial A_t} B_t = -\alpha \frac{\partial \mathcal{L}}{\partial e} h^T B_t^T B_t \\ &= \Delta_W W_t B_t^T B_t \end{aligned} \quad (5)$$

Here, we use $\Delta_{A_t} W_t$ to denote the change of the composed matrix W_t causing by the change of A_t . The fourth equation in (5) holds because of (4). The fifth equation in (5) holds because of (2). (5) shows that $\Delta_{A_t} W_t$ is equal to $\Delta_W W_t$ multiplying a projection matrix $B_t^T B_t$. Since $B_t^T B_t$ projects each row vector of $\Delta_W W_t$ into the subspace $\text{span}\{b_1^t, \dots, b_r^t\}$, Proposition 1 holds. \square

Proposition 1 has demonstrated that using our InfLoRA to train the model is equivalent to directly fine-tuning the pre-trained weight W within the subspace $\text{span}\{b_1^t, \dots, b_r^t\}$. Therefore, before learning the t -th task, we can design matrix B_t such that learning the t -th task in the subspace $\text{span}\{b_1^t, \dots, b_r^t\}$ will not interfere with the performance of the model on the old tasks.

3.2. Eliminating the Interference of the New Task on the Old Tasks

We first introduce the desired characteristics that InfLoRA aims to let the subspace $\text{span}\{b_1^t, \dots, b_r^t\}$ have. With these characteristics, InfLoRA can eliminate the interference of the new task on the old tasks and make a good trade-off between stability and plasticity. Then, we introduce how to design dimensionality reduction matrix B_t so that subspace $\text{span}\{b_1^t, \dots, b_r^t\}$ has these characteristics.

3.2.1 Desired Characteristics

First, InfLoRA aims to make the subspace $\text{span}\{b_1^t, \dots, b_r^t\}$ orthogonal to the gradients of all the old tasks. In this way, according to Proposition 1, the update of InfLoRA, which can be represented as $\Delta_{A_t} W_t$, will also be orthogonal to the gradient of the old tasks. Note that the idea of making the update for the new task orthogonal to the gradient of the old tasks to eliminate the interference of the new task on the old tasks has been proposed in many existing continual learning methods [30, 36]. However, all these existing methods are designed for continual learning from scratch, involving updating all parameters of the model, which is incompatible with the setting in PEFT. On the contrary, our

method is a PEFT method, which only tunes the parameters in A_t .

Besides eliminating the interference of new tasks on old tasks, our InfLoRA further makes the subspace $\text{span}\{b_1^t, \dots, b_r^t\}$ lie in a subspace that the gradient of the new task lies in to make a good trade-off between stability and plasticity. Specifically, existing work [19] has shown that during fine-tuning, the weight increments of pre-trained ViT exhibit redundancy in terms of weight rank. Therefore, the gradients of the new task lie in a low-dimensional subspace. Our method makes $\text{span}\{b_1^t, \dots, b_r^t\}$ not only orthogonal to the gradient of the old tasks but also lie in the subspace in which the gradients of the new task t lie. By doing so, our method makes the model’s focus on the new task when eliminating the interference of the new task on the old tasks, thereby making a good trade-off between stability and plasticity. Section 4.2 verifies the effectiveness of these two characteristics.

3.2.2 Designing Dimensionality Reduction Matrix

InfLoRA first approximates the gradient space of the new task and old tasks. Here, we use \mathcal{N}_t to represent the gradient space of the new task approximated by InfLoRA. Similarly, we use \mathcal{M}_t to represent the gradient space of previous $t - 1$ old tasks approximated by InfLoRA. We also use \mathcal{M}_t^\perp to denote the residual gradient space, which is orthogonal to the space \mathcal{M}_t . Then, in order to meet the characteristics described in Section 3.2.1, InfLoRA ensures that each row of B_t lies in $\mathcal{N}_t \cap \mathcal{M}_t^\perp$. In other words, InfLoRA makes $\text{span}\{b_1^t, \dots, b_r^t\} \subseteq \mathcal{N}_t \cap \mathcal{M}_t^\perp$.

Existing works [29, 36] have shown that the gradient update of the linear layer lies in the span of the inputs. Please refer to supplementary material for a detailed explanation of this proposition. Therefore, InfLoRA uses the input matrix of the new task t to approximate the gradient space of the new task. Specifically, InfLoRA computes the input matrix $H_t = [h_1^t, \dots, h_n^t]$, with each column of H_t representing an input vector of the t -th task. Then, InfLoRA considers \mathcal{N}_t as the subspace spanned by the columns of matrix H_t .

However, InfLoRA cannot use the input matrix of the old tasks to approximate the gradient space of the old tasks since the data from the old tasks is not available when the model learns the new tasks. Instead, existing methods such as gradient projection memory (GPM) [36] and dual gradient projection memory (DualGPM) [29] can learn a matrix to preserve information about the gradients of the old tasks. InfLoRA incorporates DualGPM to preserve gradient information. With the assistance of DualGPM, the model can learn either a matrix $M_t \in \mathbb{R}^{d_I \times k_t}$ or a matrix $M_t^\perp \in \mathbb{R}^{d_I \times (d_I - k_t)}$. Here, the columns of M_t contribute to the orthonormal bases of \mathcal{M}_t and the columns of M_t^\perp contribute to the orthonormal bases of \mathcal{M}_t^\perp . k_t denotes the

dimension of \mathcal{M}_t . For detailed information of how DualGPM maintains orthonormal bases \mathbf{M}_t or \mathbf{M}_t^\perp , please refer to supplementary material or the original paper [29].

After approximating the gradient space of the new task and old tasks, InfLoRA gets the component of \mathcal{N}_t which lies in \mathcal{M}_t^\perp . Specifically, when the model maintains \mathcal{M}_t , InfLoRA performs the operation

$$\hat{\mathbf{H}}_t = \mathbf{H}_t - \mathbf{M}_t \mathbf{M}_t^T \mathbf{H}_t. \quad (6)$$

Similarly, when the model maintains \mathcal{M}_t^\perp , InfLoRA performs the operation

$$\hat{\mathbf{H}}_t = \mathbf{M}_t^\perp (\mathbf{M}_t^\perp)^T \mathbf{H}_t. \quad (7)$$

Note that when $t = 1$, \mathcal{M}_t is a null space and $\hat{\mathbf{H}}_t = \mathbf{H}_t$. Obviously, each column of $\hat{\mathbf{H}}_t$ lies in $\mathcal{N}_t \cap \mathcal{M}_t^\perp$. However, since $(\hat{\mathbf{H}}_t)^T \in \mathbb{R}^{n \times d_I}$ and $\mathbf{B}_t \in \mathbb{R}^{r \times d_I}$ have different shapes, InfLoRA can not directly define \mathbf{B}_t as $(\hat{\mathbf{H}}_t)^T$. Note that $n \gg r$, InfLoRA uses the principal components of $(\hat{\mathbf{H}}_t)^T$ to set \mathbf{B}_t . Specifically, singular value decomposition (SVD) is performed on $(\hat{\mathbf{H}}_t)^T = \mathbf{V}_t \Sigma_t \mathbf{U}_t$. Then, InfLoRA designs \mathbf{B}_t by

$$\mathbf{B}_t = (\mathbf{U}_t)_r. \quad (8)$$

Here, $(\mathbf{U}_t)_r$ denotes the rows of \mathbf{U}_t corresponding to the top- r singular values. Figure 1 (b) illustrates the pipeline of designing matrix \mathbf{B}_t .

Note that DualGPM expands subspace \mathcal{M}_t and reduces subspace \mathcal{M}_t^\perp when the number of tasks increases. Since InfLoRA constrains the update of the model within the subspace $\mathcal{N}_t \cap \mathcal{M}_t^\perp \subseteq \mathcal{M}_t^\perp$, the space for learning the new task reduces when the number of tasks increases. However, by adjusting the approximation error of the gradient for the old tasks, DualGPM can expand \mathcal{M}_t slowly and reduce \mathcal{M}_t^\perp slowly. Therefore, the constraints imposed by InfLoRA do not excessively affect the model’s learning of new tasks. Please refer to supplementary material for a detailed explanation.

3.3. Whole Process of InfLoRA

Algorithm 1 outlines the whole process of InfLoRA in continual learning. When the t -th new task arrives, InfLoRA first designs \mathbf{B}_t through (8) and expands a new branch. Then, InfLoRA learns the t -th task by fine-tuning the newly expanded branch. Please note that, based on empirical findings from existing methods [12, 38], we employ the local cross-entropy (CE) loss as the learning objective, as it usually performs better than the global CE loss in continual learning methods based on PEFT. The local CE is the CE loss constrained to the classes of the current new task, which can be denoted as

$$\mathcal{L}(\mathcal{D}_t) = \frac{1}{|\mathcal{D}_t|} \sum_{(\mathbf{x}, y) \in \mathcal{D}_t} \mathcal{L}_{ce}(\text{mask}(h_\Phi(f_\Theta(\mathbf{x}))), y). \quad (9)$$

Algorithm 1 InfLoRA for Continual Learning

- 1: **Input:** The data of different tasks $\{\mathcal{D}_t\}_{t=1}^T$, a pre-trained ViT model $f_\Theta(\cdot)$.
 - 2: **Output:** Network $f_\Theta(\cdot)$ with learned parameters \mathbf{W}_t .
 - 3: **for** t in $1 : T$ **do**
 - 4: Design \mathbf{B}_t through (8);
 - 5: Expand a new branch for the t -th task;
 - 6: **for** \mathcal{B}_t sampled from \mathcal{D}_t **do**
 - 7: Compute the loss $\mathcal{L}(f_\Theta(\mathcal{B}_t))$ through (9) and update the parameters;
 - 8: **end for**
 - 9: Preserve the information about the gradient of the t -th task through DualGPM;
 - 10: **end for**
-

Here, $\text{mask}(\cdot)$ is a function that filters out the logits of the old classes and \mathcal{L}_{ce} denotes the standard CE loss. After learning the t -th new task, InfLoRA follows the DualGPM to preserve the information about the gradient of the t -th task.

Note that the branch corresponding to the t -th task will be frozen once the model has learned the t -th task. Since the expanded branches are linear transformations, we can integrate the old branches into the pre-trained weight to reduce the expanded parameters. Specifically, after learning the first task, InfLoRA integrates the first branch into the pre-trained weight and obtains the weight $\mathbf{W}_1 = \mathbf{W} + \mathbf{A}_1 \mathbf{B}_1$. Before learning the t -th new task ($t > 1$), InfLoRA maintains the weight \mathbf{W}_{t-1} . After learning the t -th task, InfLoRA integrates the t -th branch into \mathbf{W}_{t-1} and obtains $\mathbf{W}_t = \mathbf{W}_{t-1} + \mathbf{A}_t \mathbf{B}_t$. In this way, the parameters in \mathbf{A}_t and \mathbf{B}_t do not need to be maintained in the learning of subsequent tasks. Therefore, during the whole learning process, the number of parameters expanded by InfLoRA equals the number of parameters in a single branch. Since a single branch contains $(d_I + d_O)r$ parameters, the number of parameters expanded by InfLoRA is $(d_I + d_O)r$ all the time.

4. Experiments

4.1. Experimental Settings

Datasets and Evaluation Metric Similar to existing continual learning methods [12, 44] based on PEFT, we use ImageNet-R [14], CIFAR100 [24], and DomainNet [34] to train and evaluate the models. Imagenet-R is generated through artistic processing of 200 classes from ImageNet [8]. This dataset is introduced to continual learning by existing work [43] and has become a standard benchmark for continual learning methods based on PEFT. CIFAR100 is a dataset commonly used in existing continual learning works. DomainNet contains 345 classes and is introduced by some existing works [38, 42] for continual learning. Following existing continual learning work [38],

Table 1. Results (%) on ImageNet-R. Results are included for 5 tasks, 10 tasks, and 20 tasks. We report results averaged over 5 trials.

Tasks	5		10		20	
Method	ACC_5 (\uparrow)	\overline{ACC}_5 (\uparrow)	ACC_{10} (\uparrow)	\overline{ACC}_{10} (\uparrow)	ACC_{20} (\uparrow)	\overline{ACC}_{20} (\uparrow)
<i>joint</i>	81.14 ± 0.34	-	81.14 ± 0.34	-	81.14 ± 0.34	-
<i>sequential</i>	58.74 ± 1.28	72.91 ± 0.28	46.07 ± 1.15	62.91 ± 0.68	34.62 ± 0.85	51.15 ± 1.50
L2P [44]	64.13 ± 0.78	68.66 ± 0.41	62.54 ± 0.24	67.98 ± 0.27	57.92 ± 0.28	64.57 ± 0.29
DualPrompt [43]	67.88 ± 0.17	71.16 ± 0.31	65.41 ± 0.52	69.39 ± 0.43	61.00 ± 0.72	65.80 ± 0.67
CODA-P [38]	73.09 ± 0.21	76.91 ± 0.21	71.47 ± 0.35	75.82 ± 0.29	67.28 ± 0.30	72.34 ± 0.17
C-LoRA [37]	75.85 ± 0.31	78.85 ± 0.34	71.89 ± 0.45	75.33 ± 0.28	65.71 ± 0.60	70.63 ± 0.85
LAE [12]	73.84 ± 0.14	77.29 ± 0.45	71.70 ± 0.39	76.71 ± 0.10	66.98 ± 0.35	73.72 ± 0.05
InfLoRA-b5	75.28 ± 0.01	78.95 ± 0.08	74.13 ± 0.18	78.54 ± 0.14	68.41 ± 0.29	74.00 ± 0.50
InfLoRA	77.52 ± 0.37	82.01 ± 0.12	75.65 ± 0.14	80.82 ± 0.24	71.01 ± 0.45	77.28 ± 0.45

we split ImageNet-R into 5, 10, and 20 tasks, with each task containing 40, 20, and 10 classes. We split CIFAR100 into 10 tasks, and each task constrains 10 classes. We split DomainNet into 5 tasks, and each task contains 69 classes.

Following existing continual learning methods [12, 44], we evaluate the performance of the model through two popular metrics, including the final accuracy ACC_T and the averaged accuracy $\overline{ACC}_T = \frac{1}{T} \sum_{i=1}^T ACC_i$, where T denotes the total number of tasks and ACC_i is defined as

$$ACC_i = \frac{1}{i} \sum_{j=1}^i a_{i,j}. \quad (10)$$

Here, $a_{i,j}$ denotes the accuracy of the j -th task once the model has learned the i -th task.

Baselines We compare our InfLoRA with state-of-the-art continual learning methods based on PEFT, including learn to prompt (L2P) [44], DualPrompt [43], continual decomposed attention-based prompt (CODA-P) [38], learning accumulation ensemble (LAE) [12], continual low-rank adaptation (C-LoRA) [37]. For LAE, we implement it with LoRA [16]. Following existing works [12, 38], we also include two methods without continual learning, *joint* and *sequential*, in the comparison. Here, *joint* denotes the method that learns all the tasks jointly, while *sequential* denotes the method that learns all the tasks sequentially without any operation to overcome the forgetting of the model. The accuracy of *joint* can be treated as the accuracy upper-bound and the accuracy of *sequential* can be treated as the accuracy lower-bound.

Architecture and Training Details We follow existing works [12, 43] to perform experiments. Specifically, we use the ViT-B/16 backbone [10] supervised pre-trained on ImageNet 21K as the pre-trained model.

For all the methods, we follow existing works [12, 38, 44] and use the Adam [22] optimizer with running averages of gradient and its square ($\beta_1 = 0.9$, $\beta_2 = 0.999$). Each task is trained for 50 epochs on ImageNet-R, 20 epochs on CIFAR100 and 5 epochs on DomainNet. The batch size

is set to 128 for all the experiments. Since our InfLoRA shares a similar architecture to LoRA, we follow existing work [12] and insert the architecture of our InfLoRA in the key and value of the attention module. Furthermore, existing method DualPrompt [43] treats the inserted blocks as hyperparameters and searches for the best positions for their prompts. On the contrary, we insert the architecture of InfLoRA for all the Transformer blocks to avoid searching. We also implement a variant of our method, which inserts the bottom 5 Transformer blocks like existing methods DualPrompt and CODA-P. We call this variant InfLoRA-b5. As for the hyperparameter r , we determine its value through a grid search on a validation dataset.

4.2. Experimental Results

Accuracy Table 1 shows the results of different methods on ImageNet-R with a different number of tasks. Table 2 shows the results of different methods on CIFAR100 and DomainNet. We can find that our methods InfLoRA and InfLoRA-b5 outperform existing continual learning methods.

Figure 2 shows the variation of the accuracy of different continual learning methods on ImageNet-R and CIFAR100. We can find that our method outperforms existing methods not only at the end of the learning but also throughout the whole learning process. This indicates that our InfLoRA eliminates the interference of the new task on the old tasks and thus the accuracy of our method decreases slower compared to other methods.

Analysis of Expanded Parameters Figure 3 shows the number of expanded parameters and the accuracy of different methods on ImageNet-R and CIFAR100. For L2P, DualPrompt and CODA-P, their expanded parameters are included in the added prompts and corresponding key. For LAE, its expanded parameters are the inserted LoRA modules and an additional copy. For C-LoRA, its expanded parameters are inserted LoRA modules. For our method, the expanded parameters are B_t and A_t . The details of computing the number of expanded parameters for different meth-

Table 2. Results (%) on CIFAR100 and DomainNet. We report results over 5 trials.

Tasks	CIFAR100		DomainNet	
	ACC_{10} (\uparrow)	\overline{ACC}_{10} (\uparrow)	ACC_5 (\uparrow)	\overline{ACC}_5 (\uparrow)
<i>joint</i>	91.92 ± 0.05	-	77.72 ± 0.04	-
<i>sequential</i>	62.18 ± 3.59	80.42 ± 0.23	53.44 ± 1.21	69.09 ± 0.33
L2P [44]	82.48 ± 0.20	87.64 ± 0.25	70.16 ± 0.05	75.60 ± 0.03
DualPrompt [43]	84.42 ± 0.30	90.06 ± 0.07	72.14 ± 0.05	77.71 ± 0.06
CODA-P [38]	86.62 ± 0.11	91.08 ± 0.28	73.23 ± 0.13	78.72 ± 0.07
C-LoRA [37]	82.97 ± 0.47	88.81 ± 0.34	69.34 ± 0.13	75.25 ± 0.11
LAE [12]	84.15 ± 0.10	89.84 ± 0.03	66.85 ± 0.40	75.01 ± 0.17
InfLoRA-b5	87.06 ± 0.25	91.59 ± 0.13	73.26 ± 0.50	78.82 ± 0.34
InfLoRA	86.51 ± 0.73	91.70 ± 0.32	74.53 ± 0.23	79.57 ± 0.57

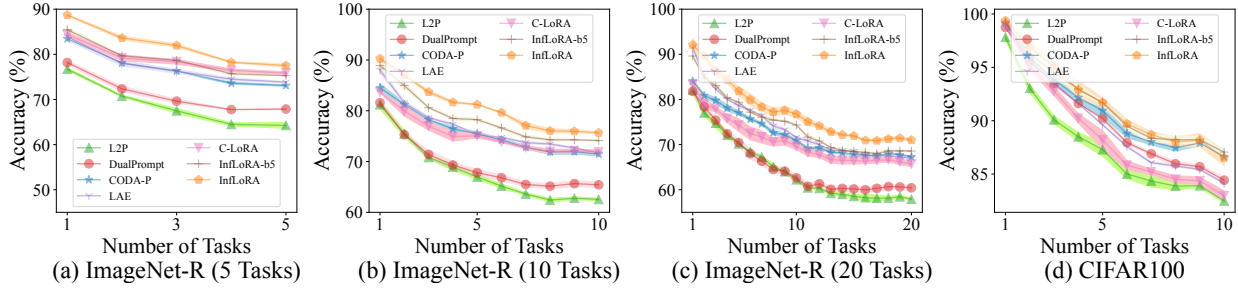


Figure 2. Variation of the performance of different methods during the learning of ImageNet-R and CIFAR100.

ods are given in supplementary material. We can find that CODA-P and C-LoRA expand much more parameters than other methods. Furthermore, our methods InfLoRA and InfLoRA-b5 expand comparable parameters to L2P, DualPrompt and LAE but perform better than these methods.

Ablation Study We perform experiment to verify the effectiveness of designing dimensionality reduction matrix B_t by (8). Specifically, we explore three different variants for designing B_t . The first variant designs B_t randomly using Gaussian distribution. We call this variant ‘Random $\rightarrow B_t$ ’. The second variant discards the operation in (6) or (7) and directly sets $\hat{H}_t = H_t$. Through this way, this variant ensures that each row of B_t lies in \mathcal{N}_t while ignoring \mathcal{M}_t^\perp . We call this variant ‘ $\mathcal{N}_t \rightarrow B_t$ ’. The third variant does not compute the input matrix but initializes H_t using a Gaussian distribution before applying the operation in (6) or (7). In this way, this variant ensures that each row of B_t lies in \mathcal{M}_t^\perp while ignoring \mathcal{N}_t . We call this variant ‘ $\mathcal{M}_t^\perp \rightarrow B_t$ ’. Since our method focuses both \mathcal{M}_t^\perp and \mathcal{N}_t , we use $\mathcal{N}_t \cap \mathcal{M}_t^\perp \rightarrow B_t$ to represent our method.

Table 3 shows the results of our method and its variants. We can find that all these variants fail to perform as well as our method. To further demonstrate the performance of different variants, we show the relative accuracy of different tasks after the model learns them all in Figure 4. Here, relative accuracy is the accuracy of different variants minus the accuracy of our InfLoRA. Note that the last task is the new task, and the other tasks are old tasks in Figure 4. As we can see, ‘Random $\rightarrow B_t$ ’ and ‘ $\mathcal{N}_t \rightarrow B_t$ ’ outper-

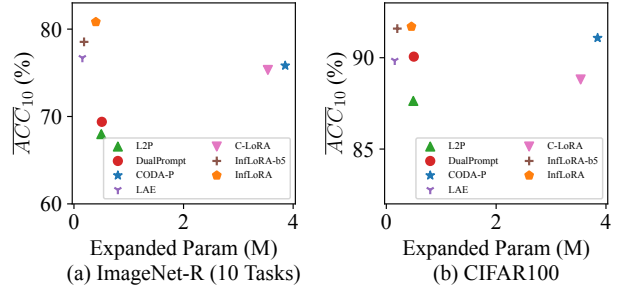


Figure 3. Variation of the performance of different methods during the learning of ImageNet-R and CIFAR100.

form ‘ $\mathcal{M}_t^\perp \rightarrow B_t$ ’ on the new task but shows much lower accuracy than ‘ $\mathcal{M}_t^\perp \rightarrow B_t$ ’ and our InfLoRA on the old tasks. This means these two variants fail to eliminate the interference of the new task on the old tasks, making the model suffer from low stability. On the contrary, ‘ $\mathcal{M}_t^\perp \rightarrow B_t$ ’ shows the lowest performance on the new task. This means ‘ $\mathcal{M}_t^\perp \rightarrow B_t$ ’ ignores the plasticity of the model. Our method outperforms all the variants on most of the tasks. This shows that our method can eliminate the interference of the new task on the old tasks and make a better trade-off between stability and plasticity than these variants.

Varying the Pre-Trained Model We also follow the existing method [40] and perform experiments using a ViT-B/16 pre-trained with two different self-supervised methods, including DINO [5] and iBOT [50]. All experimental settings, except for the choice of the pre-trained model, are kept consistent with the details outlined in Section 4.1.

Table 3. Results of different variants on ImageNet-R with a different number of tasks.

Tasks	5		10		20	
	$ACC_5 (\uparrow)$	$\overline{ACC}_5 (\uparrow)$	$ACC_{10} (\uparrow)$	$\overline{ACC}_{10} (\uparrow)$	$ACC_{20} (\uparrow)$	$\overline{ACC}_{20} (\uparrow)$
Random $\rightarrow \mathbf{B}_t$	72.49 ± 0.38	79.40 ± 0.29	67.38 ± 0.41	76.62 ± 0.06	56.17 ± 0.29	69.24 ± 0.35
$\mathcal{N}_t \rightarrow \mathbf{B}_t$	67.01 ± 0.11	76.09 ± 0.04	57.91 ± 0.30	70.23 ± 0.59	40.73 ± 0.29	59.68 ± 0.52
$\mathcal{M}_t^\perp \rightarrow \mathbf{B}_t$	75.94 ± 0.53	80.69 ± 0.27	74.61 ± 0.62	79.67 ± 0.27	68.79 ± 0.42	75.74 ± 0.26
$\mathcal{N}_t \cap \mathcal{M}_t^\perp \rightarrow \mathbf{B}_t$ (InfLoRA)	77.52 ± 0.37	82.01 ± 0.12	75.65 ± 0.14	80.82 ± 0.24	71.01 ± 0.45	77.28 ± 0.45




Figure 4. Relative accuracy of different tasks. Relative accuracy is the accuracy of different variants minus the accuracy of InfLoRA.

Table 4. Results (%) of different methods on ImageNet-R (10 tasks) using various self-supervised pre-trained models. Here, DINO-1k and iBOT-1k indicate that the ViT is pre-trained on ImageNet-1k using these respective methods.

	Method	$ACC_{10} (\uparrow)$	$\overline{ACC}_{10} (\uparrow)$
DINO-1k	L2P [44]	56.71 ± 0.12	63.59 ± 0.21
	DualPrompt [43]	60.23 ± 0.42	66.57 ± 0.25
	CODA-P [38]	64.02 ± 0.68	71.50 ± 0.42
	C-LoRA [37]	63.07 ± 0.36	68.09 ± 0.41
	LAE [12]	61.03 ± 0.27	69.89 ± 0.15
	InfLoRA-b5	66.16 ± 0.14	73.01 ± 0.17
	InfLoRA	68.31 ± 0.28	76.15 ± 0.05
iBOT-1k	L2P [44]	60.80 ± 0.35	66.58 ± 0.28
	DualPrompt [43]	63.78 ± 0.38	68.88 ± 0.16
	CODA-P [38]	68.02 ± 0.48	74.28 ± 0.47
	C-LoRA [37]	68.60 ± 0.07	73.47 ± 0.28
	LAE [12]	64.14 ± 0.29	72.59 ± 0.22
	InfLoRA-b5	69.72 ± 0.44	76.11 ± 0.13
	InfLoRA	71.84 ± 0.09	78.29 ± 0.09

Table 4 shows the results of different methods on ImageNet-R when using various pre-trained models. Comparing these results to those in Table 1, we can find that the performance of all methods utilizing self-supervised pre-trained models is lower than the performance of the corresponding methods using supervised pre-trained models. However, our methods still outperform all other methods.

Combining with Classifier Alignment Slow learner with classifier alignment (SLCA) [48] utilizes feature statistics to align classifiers, demonstrating superior performance compared to methods without aligned classifiers. Our InfLoRA can be combined with classifier alignment (CA) to get better performance. Specifically, after learning the t -th task

Table 5. Results (%) of different methods on ImageNet-R (10 tasks) and CIFAR100 using classifier alignment (CA) technique.

	Method	$ACC_{10} (\uparrow)$	$\overline{ACC}_{10} (\uparrow)$
CIFAR100	SLCA [48]	91.06 ± 0.24	93.65 ± 0.19
	InfLoRA+CA	91.59 ± 0.08	94.39 ± 0.05
ImageNet-R	SLCA [48]	77.34 ± 0.25	81.35 ± 0.16
	InfLoRA+CA	79.78 ± 0.25	83.38 ± 0.19

with parameters \mathbf{A}_t and \mathbf{B}_t and loss (9), we collect features $\mathbf{F}_t = \{\mathbf{r}_{i,t}\}_{i=1}^{n_t}$ of the t -th task. Here, $\mathbf{r}_{i,t} = f(\mathbf{x}_{i,t})$ denotes the features extracted by backbone $f_{\Theta}(\cdot)$. Then, mean and covariance of features for each class are computed and saved. After that, for each class c the model has seen during continual learning, S samples are sampled from Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}_c, \boldsymbol{\Sigma}_c)$. Here, $\boldsymbol{\mu}_c$ and covariance $\boldsymbol{\Sigma}_c$ denote the mean and covariance of the class c . Finally, we align the classifier using standard cross-entropy and these samples. The details of this experiment are given in supplementary material.

Table 5 shows that our method InfLoRA+CA outperforms SLCA. Note that SLCA tunes all the parameters of the model while our method InfLoRA only tunes the parameters in \mathbf{A}_t . Therefore, our InfLoRA+CA is much more efficient than SLCA.

5. Conclusion

In this work, we propose a new method, called interference-free low-rank adaptation (InfLoRA), for continual learning. InfLoRA injects a small number of parameters to reparameterize the pre-trained weights and shows that fine-tuning these injected parameters is equivalent to fine-tuning the pre-trained weights within a subspace. Furthermore, InfLoRA designs this subspace to eliminate the interference of the new task on the old tasks, making a good trade-off between stability and plasticity. Experimental results show that InfLoRA outperforms existing state-of-the-art continual learning methods on multiple datasets.

Acknowledgment

This work is supported by NSFC (No.62192783), National Key R&D Program of China (No.2020YFA0713901), and Fundamental Research Funds for the Central Universities (No.020214380108).

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*, pages 139–154, 2018. [2](#)
- [2] Rahaf Aljundi, Eugene Belilovsky, Tinne Tuytelaars, Laurent Charlin, Massimo Caccia, Min Lin, and Lucas Page-Caccia. Online continual learning with maximal interfered retrieval. In *Advances in Neural Information Processing Systems*, pages 11849–11860, 2019. [2](#)
- [3] Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient based sample selection for online continual learning. In *Advances in Neural Information Processing Systems*, pages 11816–11825, 2019. [2](#)
- [4] Matteo Boschini, Lorenzo Bonicelli, Angelo Porrello, Giovanni Bellitto, Matteo Pennisi, Simone Palazzo, Concetto Spampinato, and Simone Calderara. Transfer without forgetting. In *Proceedings of the European Conference on Computer Vision*, pages 692–709, 2022. [2](#)
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9650–9660, 2021. [7](#)
- [6] Shoufa Chen, Chongjian Ge, Zhan Tong, Jiangliu Wang, Yibing Song, Jue Wang, and Ping Luo. Adaptformer: Adapting vision transformers for scalable visual recognition. *Advances in Neural Information Processing Systems*, pages 16664–16678, 2022. [2](#)
- [7] Aristotelis Chrysakis and Marie-Francine Moens. Online continual learning from imbalanced data. In *Proceedings of the International Conference on Machine Learning*, pages 1952–1961, 2020. [2](#)
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. [5](#)
- [9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, pages 4171–4186, 2019. [2](#)
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. [2](#), [6](#)
- [11] 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 *Findings of the Association for Computational Linguistics*, pages 2608–2621, 2022. [2](#)
- [12] Qiankun Gao, Chen Zhao, Yifan Sun, Teng Xi, Gang Zhang, Bernard Ghanem, and Jian Zhang. A unified continual learning framework with general parameter-efficient tuning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11449–11459, 2023. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#), [3](#)
- [13] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. [2](#)
- [14] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8340–8349, 2021. [5](#)
- [15] 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*, pages 2790–2799, 2019. [1](#), [2](#)
- [16] Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. [1](#), [2](#), [6](#)
- [17] 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 *Advances in Neural Information Processing Systems*, pages 13647–13657, 2019. [2](#)
- [18] 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*, pages 709–727, 2022. [1](#), [2](#)
- [19] Shibo Jie and Zhi-Hong Deng. Fact: Factor-tuning for lightweight adaptation on vision transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1060–1068, 2023. [4](#)
- [20] Sangwon Jung, Hongjoon Ahn, Sungmin Cha, and Taesup Moon. Continual learning with node-importance based adaptive group sparse regularization. *Advances in Neural Information Processing Systems*, pages 3647–3658, 2020. [2](#)
- [21] Muhammad Gul Zain Ali Khan, Muhammad Ferjad Naeem, Luc Van Gool, Didier Stricker, Federico Tombari, and Muhammad Zeshan Afzal. Introducing language guidance in prompt-based continual learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11463–11473, 2023. [1](#), [2](#)
- [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [6](#)
- [23] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neu-

- ral networks. *Proceedings of the National Academy of Sciences*, pages 3521–3526, 2017. 2
- [24] A Krizhevsky. Learning multiple layers of features from tiny images. *Master’s thesis, University of Tront*, 2009. 5
- [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*, pages 3045–3059, 2021. 1, 2
- [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*, pages 3925–3934, 2019. 2
- [27] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 4582–4597, 2021. 2
- [28] Yan-Shuo Liang and Wu-Jun Li. Loss decoupling for task-agnostic continual learning. In *Advances in Neural Information Processing Systems*, 2023. 2
- [29] Yan-Shuo Liang and Wu-Jun Li. Adaptive plasticity improvement for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7816–7825, 2023. 4, 5, 2, 3
- [30] Sen Lin, Li Yang, Deliang Fan, and Junshan Zhang. Trgp: Trust region gradient projection for continual learning. In *International Conference on Learning Representations*, 2022. 4
- [31] Rabeeh Karimi Mahabadi, James Henderson, and Sebastian Ruder. Compacter: Efficient low-rank hypercomplex adapter layers. In *Advances in Neural Information Processing Systems*, pages 1022–1035, 2021. 2
- [32] Marc Masana, Joost Van de Weijer, Bartłomiej Twardowski, et al. On the importance of cross-task features for class-incremental learning. *arXiv preprint arXiv:2106.11930*, 2021. 1
- [33] German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, pages 54–71, 2019. 1
- [34] 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*, pages 1406–1415, 2019. 5
- [35] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016. 2
- [36] Gobinda Saha, Isha Garg, and Kaushik Roy. Gradient projection memory for continual learning. In *International Conference on Learning Representations*, 2021. 4
- [37] James Seale Smith, Yen-Chang Hsu, Lingyu Zhang, Ting Hua, Zsolt Kira, Yilin Shen, and Hongxia Jin. Continual diffusion: Continual customization of text-to-image diffusion with c-lora. *CoRR*, 2023. 6, 7, 8, 3, 5
- [38] 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*, pages 11909–11919, 2023. 1, 2, 5, 6, 7, 8, 3
- [39] Qing Sun, Fan Lyu, Fanhua Shang, Wei Feng, and Liang Wan. Exploring example influence in continual learning. *Advances in Neural Information Processing Systems*, pages 27075–27086, 2022. 2
- [40] Liyuan Wang, Jingyi Xie, Xingxing Zhang, Mingyi Huang, Hang Su, and Jun Zhu. Hierarchical decomposition of prompt-based continual learning: Rethinking obscured sub-optimality. *arXiv preprint arXiv:2310.07234*, 2023. 7, 3, 4, 5
- [41] Liyuan Wang, Xingxing Zhang, Hang Su, and Jun Zhu. A comprehensive survey of continual learning: Theory, method and application. *arXiv preprint arXiv:2302.00487*, 2023. 1
- [42] Yabin Wang, Zhiwu Huang, and Xiaopeng Hong. S-prompts learning with pre-trained transformers: An occam’s razor for domain incremental learning. *Advances in Neural Information Processing Systems*, pages 5682–5695, 2022. 1, 5
- [43] 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*, pages 631–648, 2022. 1, 2, 5, 6, 7, 8, 3
- [44] Zifeng Wang, Zizhao Zhang, Chen-Yu Lee, Han Zhang, Ruoxi Sun, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, and Tomas Pfister. Learning to prompt for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 139–149, 2022. 1, 2, 5, 6, 7, 8, 3
- [45] 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 (Short Papers)*, pages 1–9, 2022. 2
- [46] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *Proceedings of the International Conference on Machine Learning*, pages 3987–3995, 2017. 2
- [47] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, pages 107–115, 2021.
- [48] Gengwei Zhang, Liyuan Wang, Guoliang Kang, Ling Chen, and Yunchao Wei. SLCA: slow learner with classifier alignment for continual learning on a pre-trained model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19091–19101, 2023. 8
- [49] Zangwei Zheng, Mingyuan Ma, Kai Wang, Ziheng Qin, Xianguyu Yue, and Yang You. Preventing zero-shot transfer degradation in continual learning of vision-language models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19068–19079, 2023. 2

- [50] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. Image bert pre-training with online tokenizer. In *International Conference on Learning Representations*, 2022. 7
- [51] Fei Zhu, Xu-Yao Zhang, Chuang Wang, Fei Yin, and Cheng-Lin Liu. Prototype augmentation and self-supervision for incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5871–5880, 2021. 2

InfLoRA: Interference-Free Low-Rank Adaptation for Continual Learning

Supplementary Material

A. Details of GPM and DualGPM

GPM and DualGPM are established on the fact that the gradient updates lie in the span of input data points [47].

For a linear layer, we denote its forward propagation as

$$\mathbf{e} = \mathbf{W}\mathbf{h} + \mathbf{b}, \quad (11)$$

$\mathbf{W} \in \mathbb{R}^{d_I \times d_O}$, $\mathbf{h} \in \mathbb{R}^{d_I}$, and $\mathbf{e} \in \mathbb{R}^{d_O}$. d_I and d_O denote input and output dimension, respectively. We further denote the loss function as \mathcal{L} . Through the chain rule, we can get the gradient of \mathbf{W} :

$$\frac{\partial \mathcal{L}}{\partial \mathbf{W}} = \frac{\partial \mathcal{L}}{\partial \mathbf{e}} \frac{\partial \mathbf{e}}{\partial \mathbf{W}} = \frac{\partial \mathcal{L}}{\partial \mathbf{e}} \mathbf{h}^T = \begin{bmatrix} a_1 \mathbf{h}^T, \\ a_2 \mathbf{h}^T, \\ \dots, \\ a_{d_O} \mathbf{h}^T \end{bmatrix}. \quad (12)$$

$[a_1, a_2, \dots, a_{d_O}]^T$ denotes the vector $\frac{\partial \mathcal{L}}{\partial \mathbf{e}}$. Through (12), we can find that each column of $\frac{\partial \mathcal{L}}{\partial \mathbf{W}}$ can be represented as input \mathbf{h} multiplied by a real value a_k ($1 \leq k \leq d_O$). Therefore, in the linear layer, each column of the gradient $\frac{\partial \mathcal{L}}{\partial \mathbf{W}}$ lies in the span of input.

A.1. Gradient Projection Memory

GPM learns a subspace \mathcal{M}_t with orthogonal bases \mathbf{M}_t to approximate the gradient space of the old tasks. Here, the columns of \mathbf{M}_t contribute a set of orthogonal bases in \mathcal{M}_t . GPM expands the bases of \mathcal{M}_t to the bases of \mathcal{M}_{t+1} after learning the t -th new task. Specifically, GPM computes the inputs matrix \mathbf{H}_t such that each column of \mathbf{H}_t represents an input of this layer. Then, the part of \mathbf{H}_t that has already in \mathcal{M}_t is removed by

$$\hat{\mathbf{H}}_t = \mathbf{H}_t - \mathbf{M}_t(\mathbf{M}_t)^T \mathbf{H}_t = \mathbf{H}_t - \mathbf{H}_{t,proj}. \quad (13)$$

Please note that when $t = 1$, $\dim(\mathcal{M}_t) = 0$ and hence $\mathbf{H}_{t,proj}$ is a zero matrix. After that, singular value decomposition (SVD) is performed on $\hat{\mathbf{H}}_t = \hat{\mathbf{U}}\hat{\Sigma}\hat{\mathbf{V}}^T$. Then, u new orthogonal bases are chosen from the columns of $\hat{\mathbf{U}}$ for a minimum of u satisfying the following criteria for given threshold ϵ_{th} :

$$\|(\hat{\mathbf{H}}_t)_u\|_F^2 + \|\mathbf{H}_{t,proj}\|_F^2 \geq \epsilon_{th} \|\mathbf{H}_t\|_F^2. \quad (14)$$

Here, $(\hat{\mathbf{H}}_t)_u = [\mathbf{u}_1, \dots, \mathbf{u}_u]$ denotes the components of $\hat{\mathbf{H}}_t$ that correspond to top- u singular values. Then, subspace \mathcal{M}_{t+1} is obtained with the bases $\mathbf{M}_{t+1} = [\mathbf{M}_t, \mathbf{u}_1, \dots, \mathbf{u}_u]$.

A.2. Dual Gradient Projection Memory

Different from GPM that learns a subspace \mathcal{M}_t with orthogonal bases \mathbf{M}_t to approximate the gradient space of the old tasks, DualGPM either learns a subspace \mathcal{M}_t with orthogonal bases \mathbf{M}_t to approximate the gradient of the old tasks, or learns a subspace \mathcal{M}_t^\perp with orthogonal bases \mathbf{M}_t^\perp to approximate orthogonal complement of the gradient space of the old tasks.

DualGPM decides whether to keep \mathbf{M}_t or \mathbf{M}_t^\perp in memory according to $\dim(\mathcal{M}_t)$ and $\dim(\mathcal{M}_t^\perp)$. Specifically, during the learning of the first several tasks, $\dim(\mathcal{M}_t) \leq \dim(\mathcal{M}_t^\perp)$. At this time, DualGPM maintains \mathbf{M}_t , and expands \mathbf{M}_t to \mathbf{M}_{t+1} after each task. When $\dim(\mathcal{M}_t)$ increases and exceeds $\dim(\mathcal{M}_t^\perp)$, DualGPM obtains \mathbf{M}_t^\perp through some transformations on \mathbf{M}_t . After that, DualGPM only maintains \mathbf{M}_t^\perp in memory, and reduces \mathbf{M}_t^\perp to \mathbf{M}_{t+1}^\perp after each task. Through this way, the number of bases kept for each layer is $\min\{\dim(\mathcal{M}_t), \dim(\mathcal{M}_t^\perp)\}$.

There are three key problems in DualGPM: expanding the bases of \mathcal{M}_t , obtaining the bases of \mathcal{M}_t^\perp through the bases of \mathcal{M}_t , and reducing the bases of \mathcal{M}_t^\perp .

Expanding the Bases of \mathcal{M}_t The expansion of \mathcal{M}_t is the same as that in GPM.

Transforming \mathcal{M}_t to \mathcal{M}_t^\perp DualGPM transforms \mathcal{M}_t to \mathcal{M}_t^\perp by performing SVD to the matrix \mathbf{M}_t . Specifically, let $\mathbf{M}_t = \mathbf{U}\Sigma\mathbf{V}^T$, the column vectors of \mathbf{U} which correspond to the zero singular values form a set of orthogonal bases of \mathcal{M}_t^\perp . Please refer to the paper of DualGPM [29] for this explanation.

Reducing the Bases of \mathcal{M}_t^\perp DualGPM reduces space \mathcal{M}_t^\perp by removing the part of \mathcal{M}_t^\perp which contains the gradient of the t -th task. Specifically, DualGPM first computes the input matrix \mathbf{R}_t . Then, the part of \mathbf{R}_t which lies in \mathcal{M}_t^\perp can be computed through

$$\hat{\mathbf{R}}_t^\perp = \mathbf{M}_t^\perp (\mathbf{M}_t^\perp)^T \mathbf{R}_t = \mathbf{R}_{t,proj}^\perp. \quad (15)$$

After that, SVD is performed on $\hat{\mathbf{R}}_t^\perp = \hat{\mathbf{U}}^\perp \hat{\Sigma}^\perp (\hat{\mathbf{V}}^\perp)^T$. Then, k new orthogonal bases are chosen from the columns of $\hat{\mathbf{U}}^\perp$ for a maximum of k satisfying the following criteria for the given threshold ϵ_{th} (the same as ϵ_{th} in (14)):

$$\|(\hat{\mathbf{R}}_t^\perp)_k\|_F^2 \leq (1 - \epsilon_{th}) \|\mathbf{R}_t\|_F^2. \quad (16)$$

Let $\mathbf{Z} = (\hat{\mathbf{R}}_t^\perp)_k = [\mathbf{u}_1^\perp, \dots, \mathbf{u}_k^\perp]$, $\mathcal{Z} = \text{span}\{\mathbf{u}_1^\perp, \dots, \mathbf{u}_k^\perp\}$. Here, \mathcal{Z} is the subspace of \mathcal{M}_t^\perp that contains the gradient of




Figure 5. Change of the dimension of subspace \mathcal{M}_t^\perp throughout the whole learning process.

the t -th task. DualGPM removes \mathcal{Z} from \mathcal{M}_t^\perp to get \mathcal{M}_{t+1}^\perp . Specifically, let $\hat{M}_t^\perp = M_t^\perp - \mathbf{Z}(\mathbf{Z}^T M_t^\perp)$. DualGPM performs the second SVD on $\hat{M}_t^\perp = \tilde{U}^\perp \tilde{\Sigma}^\perp (\tilde{V}^\perp)^T$. The columns of \tilde{U}^\perp which correspond to the non-zero singular values form the bases M_{t+1}^\perp . Please refer to the paper of DualGPM [29] for this explanation.

A.3. Approximation Error in DualGPM

DualGPM either learns a subspace \mathcal{M}_t to approximate the gradient space of the old tasks or learns a subspace \mathcal{M}_t^\perp to represent the orthogonal complement of the gradient of the old tasks. From Seciton A.2, we can find that the approximation error is related to the hyperparameter ϵ_{th} in (14) and (16). Specifically, when the value of ϵ_{th} in (14) and (16) increases, the approximation error decreases. As a result, the dimension of subspace \mathcal{M}_t becomes larger, while the dimension of \mathcal{M}_t^\perp becomes smaller. Note that our InfLoRA constrains the update of the model to lie within the subspace $\mathcal{N}_t \cap \mathcal{M}_t^\perp \subseteq \mathcal{M}_t^\perp$. Therefore, we can adjust the value of ϵ_{th} to adjust the space for learning the new task. Here, for all the experiments, we set

$$\epsilon_{th} = \epsilon + \frac{(1 - \epsilon) * t}{T}, \quad (17)$$

where t denotes the task id and T denotes the total number of tasks. In other words, we gradually increase the value of ϵ_{th} as the number of tasks increases throughout the whole learning process. Table 6 shows the setting of ϵ in our InfLoRA.

Figure 5 illustrates the variation of the dimension of the subspace \mathcal{M}_t^\perp in different Transformer layers of ViT-B/16. We can find that the dimension of the subspace \mathcal{M}_t^\perp in different Transformer layers of ViT-B/16 is always much larger than zero, which means the space for learning the new task always exists throughout the whole learning process.

B. More Experimental Details

B.1. Training Details

For all the methods in all the experiments except for the comparison with SLCA, the batch size is set to 128 to fol-

low many existing continual learning methods based on PEFT [38, 40]. Hyperparameters for different methods are selected based on the experimental settings in existing works [12, 38, 44] or through hyperparameter search. For example, Adam is used as the optimizer with running averages of gradient and its square ($\beta_1 = 0.9, \beta_2 = 0.999$). The learning rate is searched among $[5e-4, 1e-3, 2e-3, 1e-2]$ for all the methods through the validation sets we split from the training sets. For the hyperparameter r in our InfLoRA, we search it among $[1, 5, 10, 20, 30]$ through the validation sets we split from the training sets. Table 6 shows the hyperparameters of different methods.

When compared with SLCA, our method is combined with classifier alignment (CA). At this time, we follow SLCA to train the expanded LoRA branches and classifiers using the SGD optimizer. Each task is trained for 50 epochs on ImageNet-R, 20 epochs on CIFAR100 and 5 epochs on DomainNet. The batch size is set to 128.

B.2. Expanded Parameters

For L2P [44], the expanded parameters consist of the inserted prompts and their corresponding keys. Let d denote the embedding dimension, e denote the prompt length, p denote the number of prompts, and l denote the number of layers in which prompts are inserted. To compute the total number of expanded parameters, the formula used is $dlp(e + 1)$.

For DualPrompt [43], the expanded parameters also consist of the inserted prompts and corresponding keys. However, DualPrompt contains expert prompts and shared prompts. Let d denote the embedding dimension, T denote the number of tasks, e_E denote the expert prompt length, e_S denote the shared prompt length, l_E denote the number of layers in which expert prompts are inserted and l_S denote the number of layers in which shared prompts are inserted. To compute the total number of expanded parameters, the formula used is $d[Tl_E(e_E + 1) + e_S l_S]$.

For CODA-Prompt [38], the expanded parameters consist of the inserted prompts, corresponding keys and attention parameters. Let d denote the embedding dimension, e denote the prompt length, p denote the number of prompts, and l denote the number of layers in which prompts are inserted. To compute the total number of expanded parameters, the formula used is $dlp(e + 2)$.

For LAE [12], we implement it with LoRA. Therefore, the expanded parameters in this method consist of the inserted LoRA modules and the corresponding ensemble modules. Let d denote the embedding dimension, r denote the rank, and l denote the number of layers in which LoRA modules are inserted. Since LAE inserts LoRA modules into key and value projection in multi-head attention, the number of expanded parameters is $8ldr$.

For C-LoRA [37], the expanded parameters in this

Table 6. List of hyper-parameters for different methods. The meaning of different hyperparameters is given in Section B.2. The hyperparameter ϵ in InfLoRA is explained in Section A.3

Methods	Hyper-Parameters
L2P	lr: 0.001 (ImageNet-R, DomainNet, CIFAR100) l: 1 (ImageNet-R, DomainNet, CIFAR100) p: 30 (ImageNet-R, DomainNet, CIFAR100) e: 20 (ImageNet-R, DomainNet, CIFAR100)
DualPrompt	lr: 0.001 (ImageNet-R, DomainNet, CIFAR100) l _E : 3 (ImageNet-R, DomainNet, CIFAR100) l _S : 2 (ImageNet-R, DomainNet, CIFAR100) e _E : 20 (ImageNet-R, DomainNet, CIFAR100) e _S : 6 (ImageNet-R, DomainNet, CIFAR100)
CODA-P	lr: 0.001 (ImageNet-R, DomainNet, CIFAR100) l: 5 (ImageNet-R, DomainNet, CIFAR100) p: 100 (ImageNet-R, DomainNet, CIFAR100) e: 8 (ImageNet-R, DomainNet, CIFAR100)
LAE	lr: 0.001 (ImageNet-R, DomainNet, CIFAR100) r: 5 (ImageNet-R, DomainNet, CIFAR100)
C-LoRA	lr: 0.001 (ImageNet-R, DomainNet, CIFAR100) r: 64 (ImageNet-R, DomainNet, CIFAR100) λ : 0.5 (ImageNet-R, DomainNet, CIFAR100)
InfLoRA-b5	lr: 0.001 (CIFAR100), 0.0005 (ImageNet-R, DomainNet) r: 10 (ImageNet-R, CIFAR100), 20 (DomainNet) ϵ : 0.99 (ImageNet-R), 0.95 (CIFAR100, DomainNet)
InfLoRA	lr: 0.0005 (ImageNet-R, DomainNet, CIFAR100) r: 10 (ImageNet-R, DomainNet, CIFAR100) ϵ : 0.98 (ImageNet-R), 0.95 (CIFAR100, DomainNet)

method consist of the inserted LoRA modules. Let d denote the embedding dimension, r denote the rank, and l denote the number of layers in which LoRA modules are inserted. Since C-LoRA inserts LoRA modules into query, key and value projection in multi-head attention, the number of expanded parameters is $6ldr$.

For our methods, since we integrate the branches of the old tasks when the model learns a new task, the number of expanded parameters equals the number of parameters in a single branch. Let d denote the embedding dimension, r denote the rank, and l denote the number of layers in which our InfLoRA modules are inserted. Since we also insert InfLoRA modules into key and value projection in multi-head attention, the number of expanded parameters is $4ldr$.

C. More Experimental Results

C.1. Compare with More Methods

We compare with SeqLoRA, which initials LoRA modules and finetunes these modules on multiple tasks sequentially without any operation to overcome forgetting. The results are given in Table 7, Table 8 and Table 9. We can find that

our method outperforms this method.

A recent continual learning PEFT method hierarchical decomposition prompt (HiDe-Prompt) [40] proposes to perform continual learning hierarchically. This method maintains a set of task-specific prompts for each task and contains two stages during training and inference. Specifically, given an input sample, Hide-Prompt infers the prompt index and then uses the corresponding prompt to infer its label. We also compare our method with this method, and the results are also given in Table 7, Table 8 and Table 9. We can find that our method outperforms this method. Furthermore, this method shows comparable performance to our method in terms of final accuracy ACC_T on ImageNet-R. However, there is a notable gap between this method and our method in terms of averaged accuracy \overline{ACC}_T . Note that averaged accuracy \overline{ACC}_T is more important than final accuracy ACC_T since \overline{ACC}_T represents the performance of the model over the whole learning process.

C.2. Hyperparameter Analysis

We perform the hyperparameter analysis for our method InfLoRA. There are two specific hyperparameters in our

Table 7. The comparison between our InfLoRA and more methods on ImageNet-R.

Tasks	5		10		20	
	ACC_5 (\uparrow)	\overline{ACC}_5 (\uparrow)	ACC_{10} (\uparrow)	\overline{ACC}_{10} (\uparrow)	ACC_{20} (\uparrow)	\overline{ACC}_{20} (\uparrow)
SeqLoRA	70.96 \pm 0.25	79.14 \pm 0.32	64.32 \pm 0.09	74.78 \pm 0.29	56.98 \pm 0.29	69.29 \pm 0.26
HiDe-Prompt [40]	76.82 \pm 0.91	77.19 \pm 0.34	75.06 \pm 0.12	76.60 \pm 0.01	66.88 \pm 1.29	76.71 \pm 0.23
InfLoRA	77.52 \pm 0.37	82.01 \pm 0.12	75.65 \pm 0.14	80.82 \pm 0.24	71.01 \pm 0.45	77.28 \pm 0.45

Table 8. The comparison between our InfLoRA and more methods on DomainNet.

	ACC_5 (\uparrow)	\overline{ACC}_5 (\uparrow)
SeqLoRA	71.69 \pm 0.13	78.68 \pm 0.12
HiDe-Prompt [40]	71.48 \pm 0.10	76.15 \pm 0.05
InfLoRA	74.53 \pm 0.23	79.57 \pm 0.57

Table 9. Results (%) of different methods on ImageNet-R (10 tasks) using various self-supervised pre-trained models. Here, DINO-1k and iBOT-1k indicate that the ViT is pre-trained on ImageNet-1k using these respective methods.

	Method	ACC_{10} (\uparrow)	\overline{ACC}_{10} (\uparrow)
DINO-1k	SeqLoRA	60.67 \pm 0.11	66.29 \pm 0.21
	HiDe-Prompt [40]	68.11 \pm 0.18	71.70 \pm 0.01
	InfLoRA	68.31 \pm 0.28	76.15 \pm 0.05
iBOT-1k	SeqLoRA	66.87 \pm 0.40	71.80 \pm 0.28
	HiDe-Prompt [40]	71.33 \pm 0.21	73.62 \pm 0.13
	InfLoRA	71.84 \pm 0.09	78.29 \pm 0.09




Figure 6. (a) Analysis of the hyperparameter r . (b) Analysis of the hyperparameter ϵ .

method InfLoRA. The first hyperparameter is r , which controls the expanded parameters in InfLoRA. The second hyperparameter is ϵ , which is not the specific hyperparameter of our InfLoRA but the hyperparameter introduced by DualGPM. This hyperparameter controls the component maintained in the matrix M_t .

Figure 6 shows the results of our method with different values of r or ϵ . We can find that the performance of InfLoRA increases first and then decreases with the increase of r and ϵ .

Table 10. Results of DomainNet for domain incremental setting.

Method	ACC_6 (\uparrow)	\overline{ACC}_6 (\uparrow)
L2P [44]	34.15 \pm 0.10	49.84 \pm 0.03
DualPrompt [43]	35.24 \pm 0.12	48.44 \pm 0.13
CODA-P [38]	56.89 \pm 0.04	57.56 \pm 0.03
C-LoRA [37]	44.96 \pm 0.01	52.95 \pm 0.08
InfLoRA	68.44 \pm 0.04	67.46 \pm 0.03




Figure 7. The time of inferring one task for different methods.

C.3. Domain Incremental Setting

InfLoRA can be extended to the domain incremental setting. Specifically, DomainNet contains six domains and InfLoRA can learn on these domains sequentially. Table 10 shows that InfLoRA outperforms other baselines.

C.4. Inference Efficiency

Existing methods often involve multiple forward propagations through the pre-trained backbone. Specifically, prompt-based continual learning methods, including L2P, DualPrompt, and CODA-P, require an extra forward propagation to generate instance-specific prompts. LAE requires an extra forward propagation for ensembling. In contrast, our InfLoRA only requires a single forward propagation through the pre-trained backbone. Figure 7 provides a comparison of the time consumed by different methods during inference. We can find that our method consistently outperforms existing methods in terms of time efficiency.