

Adapt Before Continual Learning

Aojun Lu¹, Tao Feng², Hangjie Yuan³, Chunhui Ding¹, Yanan Sun^{1*}

¹College of Computer Science, Sichuan University, Chengdu, China

²Department of Computer Science and Technology, Tsinghua University, Beijing, China

³College of Computer Science and Technology, Zhejiang University, Hangzhou, China

Abstract

Continual Learning (CL) seeks to enable neural networks to incrementally acquire new knowledge (plasticity) while retaining existing knowledge (stability). Although pre-trained models (PTMs) have provided a strong foundation for CL, existing approaches face a fundamental challenge in balancing these two competing objectives. Current methods typically address stability by freezing the PTM backbone, which severely limits the model’s plasticity, particularly when incoming data distribution diverges largely from the pre-training data. Alternatively, sequentially fine-tuning the entire PTM can adapt to new knowledge but often leads to catastrophic forgetting, highlighting the critical stability-plasticity trade-off in PTM-based CL. To address this limitation, we propose Adapting PTMs before the core CL process (ACL), a novel framework that introduces a plug-and-play adaptation phase prior to learning each new task. During this phase, ACL refines the PTM backbone by aligning embeddings with their original class prototypes while distancing them from irrelevant classes. This mechanism theoretically and empirically demonstrates desirable balance between stability and plasticity, significantly improving CL performance across benchmarks and integrated methods.

Code — https://github.com/byyx666/ACL_code

Introduction

In open-world scenarios, data typically arrives in a streaming fashion, necessitating a machine learning paradigm capable of incrementally acquiring new knowledge while retaining previous information, known as Continual Learning (CL) (Wang et al. 2023; Zhou et al. 2024b; Bian et al. 2024; Feng et al. 2025). Effective CL critically hinges on a neural network’s ability to balance *plasticity*, which enables the learning of new concepts, and *stability*, which ensures the retention of previously acquired knowledge. However, an overemphasis on stability can hinder the model’s adaptability to new information, whereas excessive plasticity may lead to *catastrophic forgetting* of prior knowledge (Goodfellow et al. 2013). This fundamental conflict is known as the *stability-plasticity dilemma* (Grossberg 2013; Lu et al. 2025), which remains a central challenge in CL research.

*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

The advent of powerful Pre-Trained Models (PTMs) has significantly reshaped the machine learning domain, spurring considerable interest in their application to CL (Zhou et al. 2024b). PTMs, trained on large-scale datasets, exhibit strong generalization capabilities, providing a strong foundation for downstream CL tasks. A dominant strategy in PTM-based CL, therefore, involves freezing the PTM backbone to preserve this foundational knowledge while training lightweight modules (*e.g.*, prompts or adapters) tailored to new tasks (Wang et al. 2022b,a; Tan et al. 2024; Sun et al. 2024). While such parameter-efficient approaches excel at preserving PTMs’ generalizable knowledge for stability, they may trade off plasticity, as the model cannot sufficiently adapt its learned representations to new tasks.

A critical challenge emerges from this trade-off: the pre-trained feature space, though broadly general, is not always optimal for the discriminative requirements of every downstream task (Zhou et al. 2022; Hendrycks et al. 2021b). When the incoming data distribution diverges significantly from the pre-training data, the highly stable but rigid feature representations can become a bottleneck, limiting the model’s ability to learn new concepts effectively (Zhou et al. 2024a). This limitation may result in diminished plasticity and suboptimal CL performance. On the other hand, fully fine-tuning the PTM for each task may enhance plasticity but risks degrading its generalizable knowledge, leading to catastrophic forgetting (Kumar et al. 2022). This underscores a more nuanced dilemma in PTM-based CL: *how to realign the PTM’s feature space for new tasks without destabilizing the foundational knowledge crucial for stability.*

To address this pivotal challenge, we introduce Adapting PTMs before the core CL process (ACL), a novel framework that performs a plug-and-play adaptation of the PTM’s feature space before learning each new task. By adapting the PTM’s feature space to align better with incremental data, ACL provides a stronger and more task-relevant foundation for existing PTM-based CL methods to build upon. Specifically, during the adaptation phase, the PTM backbone is fine-tuned by encouraging its output embeddings to move closer to their respective original class prototypes while simultaneously distancing them from other class prototypes. This straightforward yet effective adaptation mechanism ensures that the PTM’s feature space is realigned for new tasks without destabilizing the foundational knowledge critical for




Figure 1: Performance comparison on ImageNet-A-Inc20 between the frozen PTM and the PTM adapted using our ACL. Plasticity: the average of the optimal accuracy of each task during CL; Stability: the average forgetting across previous tasks after learning the final task; Overall CL performance: the average accuracy across all tasks after learning the final task.

stability. As illustrated in Fig. 1, while existing CL methods using frozen PTMs exhibit strong stability (with approximately 10% forgetting), their overall CL performance remains limited by weak plasticity. Integrating ACL not only significantly enhances plasticity but also preserves stability, yielding superior CL performance.

The contributions of this study are outlined as follows: (i) We identify that prevailing PTM-based CL methods achieve suboptimal CL performance due to inherent limitations in plasticity, underscoring the necessity for an effective adaptation mechanism. (ii) We demonstrate theoretically the core objectives of adaptation, *i.e.*, enhancing plasticity while preserving stability, can be effectively achieved by encouraging embeddings to converge toward their original class prototypes and diverge from others. (iii) We propose ACL, a novel, plug-and-play CL framework designed to achieve a improved stability-plasticity trade-off. Extensive experiments across diverse benchmarks and established CL methods validate the effectiveness and broad applicability of ACL.

Related Works

Continual Learning (CL)

CL aims to enable neural networks to sequentially acquire knowledge from a series of tasks without forgetting previously learned concepts (Masana et al. 2022; Van de Ven, Tuytelaars, and Toliás 2022). Traditional CL methods can be broadly categorized into three types. *Replay-based* methods retain a subset of previous data information in a memory buffer, which is subsequently utilized to recover old data distributions (Aljundi et al. 2019; Liu et al. 2020; Iscen et al. 2020; Zhao et al. 2021). *Regularization-based* methods incorporate penalty terms that constrain model updates during the learning of new tasks (Kirkpatrick et al. 2017; Zenke, Poole, and Ganguli 2017; Li and Hoiem 2017; Feng, Wang, and Yuan 2022). *Architecture-based* methods allocate task-specific parameter spaces within the network for each new task, thereby mitigating forgetting (Kang et al. 2022; Konishi et al. 2023; Yan, Xie, and He 2021; Zhou et al. 2023).

CL with PTMs. With the growing prevalence of PTMs (Dosovitskiy 2020; Radford et al. 2021; Gan and Wei 2024), PTM-based CL has recently garnered significant attention. Given that PTMs have been equipped with generalizable knowledge, these methods often freeze the pre-trained backbones and utilize additional trainable modules to learn task-specific knowledge (Zhou et al. 2024b; Luo et al. 2025). Early research primarily focuses on applying visual prompt tuning (Jia et al. 2022) to CL, enabling models to learn new tasks without modifying the pre-trained weights (Smith et al. 2023; Jung et al. 2023), *e.g.*, L2P (Wang et al. 2022b) and DualPrompt (Wang et al. 2022a). Recently, some studies have demonstrated that adapter-based tuning outperforms prompt-based methods in PTM-based CL (Tan et al. 2024; Gao et al. 2024), *e.g.*, SSIAT (Tan et al. 2024) and MOS (Sun et al. 2024). In addition to developing additional trainable modules, several studies have focused on optimizing the classification head to enhance CL performance, *e.g.*, FeCAM (Goswami et al. 2024) and RanPAC (McDonnell et al. 2024).

Prototypical Networks

Prototypical networks (Snell, Swersky, and Zemel 2017) involve learning an embedding space where samples are classified by minimizing their distance to the mean embedding (prototype) of their respective class (Li et al. 2020; Zhang, Song, and Tao 2022). In our research, we demonstrate, both theoretically and empirically, that the integration of this core principle into the adaptation phase of PTMs achieves a desirable balance between plasticity and stability.

Contrastive Learning

Contrastive learning (Oord, Li, and Vinyals 2018) has emerged as a powerful framework in self-supervised learning and supervised learning (Khosla et al. 2020), which brings similar examples closer together in the feature space while pushing dissimilar examples apart. In the context of CL, several studies have leveraged contrastive learning to enhance stability (Nagata and Hotta 2023; Wen et al. 2024), *e.g.*, Co²L (Cha, Lee, and Shin 2021) and PCL (Lin et al.

2023). These approaches typically contrast embeddings from current task data against replayed samples from previous tasks (exemplar replay) to preserve learned representations. Unlike these methods that primarily utilize contrastive learning with replayed data to bolster stability, our work focuses on applying contrastive principles exclusively to the current task’s data. Furthermore, our primary objective through this application is to enhance plasticity, with stability being an emergent benefit of our formulation.

ACL: Adapt Before Continual Learning

To enhance plasticity in CL with PTMs, we propose a novel CL framework, ACL, which introduces a novel adaptation phase before learning each incremental task. This phase aims to adapt the PTM’s weights to the characteristics of the new data, thereby improving feature discriminability for the current task, while preserving the existing knowledge. In the subsequent sections, we first outline the overall procedure of ACL. We then detail the specific adaptation loss function employed and provide a theoretical analysis demonstrating how it addresses the stability-plasticity trade-off.

Preliminaries

For clarity, we decompose the CL model into two primary components: $f(x) = \mathcal{C}(\phi(x))$. Here, $\phi(\cdot) : \mathbb{R}^D \rightarrow \mathbb{R}^d$ represents the PTM backbone, which maps input samples x into feature embeddings. The classification head, $\mathcal{C}(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}^{|\mathcal{Y}_k|}$, takes these embeddings and produces classification outputs. Existing PTM-based CL approaches often freeze the backbone $\phi(\cdot)$ throughout the learning process to preserve this general knowledge. To acquire knowledge from new tasks, these methods typically introduce additional trainable modules, denoted by Θ (e.g., prompts and adapters). The model then becomes $f(x) = \mathcal{C}(\phi(x), \Theta)$, and only Θ and the classification head \mathcal{C} are updated during CL.

Overall Procedure of ACL

The ACL framework operates in two distinct phases for each new task, as illustrated in Fig. 2.

Phase 1: Adaptation. At the beginning of learning the k -th task, given the data \mathcal{D}_k , we adapt the parameters of the PTM backbone ϕ_{k-1} and the lightweight modules Θ_{k-1} . This adaptation aims to make the features generated by the model more discriminative for the current task. Formally, this process is defined as:

$$\phi_{k-1}^*, \Theta_{k-1}^* = \mathcal{A}(\phi_{k-1}, \Theta_{k-1}, \mathcal{D}_k), \quad (1)$$

where \mathcal{A} denotes the adaptation algorithm. The output consists of the adapted backbone ϕ_{k-1}^* and adapted lightweight modules Θ_{k-1}^* . Critically, the design of \mathcal{A} ensures that the separability of features for \mathcal{D}_k is enhanced while attempting to maintain the existing knowledge encoded in ϕ_{k-1} and Θ_{k-1} . The specific design of \mathcal{A} is detailed in the next section.

Phase 2: Core Learning. Following adaptation, the adapted backbone ϕ_{k-1}^* is frozen. Subsequently, the classification head \mathcal{C}_{k-1} and the adapted lightweight modules Θ_{k-1}^* are fine-tuned to learn the classification task for \mathcal{D}_k .

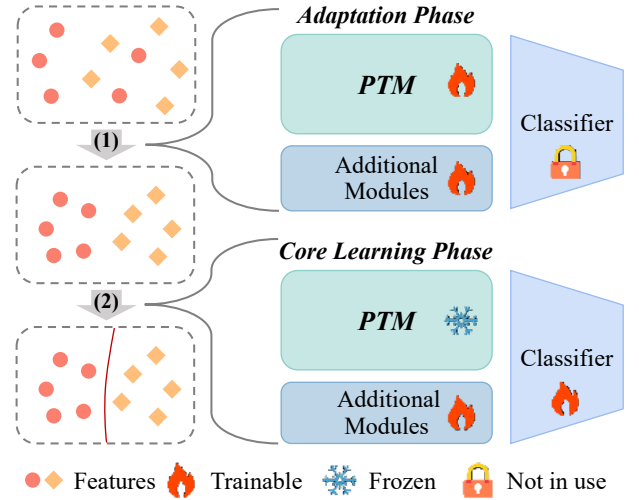


Figure 2: Illustration of ACL. ACL comprises two phases per task: (1) Adapting the PTM weights to enhance feature discriminability for the current task, and (2) Learning classification using the frozen adapted PTM and trainable modules.

Algorithm 1: Procedure of ACL

- 1: **Input:** PTM backbone ϕ_0 , Additional lightweight module Θ_0 , Classification Head \mathcal{C}_0 , Incremental datasets $\{\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_K\}$;
- 2: **Output:** Updated PTM;
- 3: **for** task $k = 1, 2, \dots, K$ **do**
- 4: Get the training set of the incremental dataset \mathcal{D}_k ;
- 5: Optimize ϕ_{k-1} and Θ_{k-1} to obtain ϕ_{k-1}^* and Θ_{k-1}^* via Eq. 1;
- 6: Save ϕ_{k-1}^* , i.e., $\phi_k = \phi_{k-1}^*$;
- 7: Optimize Θ_{k-1}^* and \mathcal{C}_{k-1} to obtain Θ_k and \mathcal{C}_k via Eq. 2 with frozen ϕ_k ;
- 8: **end for**

This step leverages the adapted PTM’s knowledge for the new task. The process is formalized as:

$$\begin{aligned} \phi_k &= \phi_{k-1}^*, \\ \mathcal{C}_k, \Theta_k &= \mathcal{F}(\mathcal{C}_{k-1}, \Theta_{k-1}^*, \phi_k, \mathcal{D}_k), \end{aligned} \quad (2)$$

where \mathcal{F} represents the CL method integrated into the ACL framework, e.g., L2P (Wang et al. 2022b).

Algorithm 1 presents the pseudo-code for ACL. It iteratively applies the adaptation (Eq. 1) and core learning (Eq. 2) phases for each incremental task from $k = 1$ to K .

Adaptation Algorithm Within ACL

This section introduces the adaptation loss function used within the \mathcal{A} phase of ACL, termed the ACL loss. This loss is specifically designed to navigate the stability-plasticity trade-off inherent in CL. We provide a theoretical justification showing that the ACL loss simultaneously (1) promotes *plasticity* by minimizing an upper bound on the current task’s classification error and (2) maintains *stability* by implicitly regularizing feature deviation.

For analytical simplicity, we consider a cosine classifier, a common choice in PTM-based CL (Zhou et al. 2024a; Sun et al. 2024). This classifier assigns a sample x to the class c whose prototype p_c has the highest cosine similarity with the feature embedding. We note that, although the theoretical analysis is based on the cosine classifier, the empirical results demonstrate the applicability of ACL to CL methods using other classifiers, e.g., the linear classifier.

ACL Loss Let ϕ be the PTM backbone prior to adaptation and ϕ^* the backbone after adaptation. For the current task \mathcal{D}_k , the adaptation process computes class prototypes $p_c = \mathbb{E}_{(x,y) \in \mathcal{D}_k, y=c}[\phi(x)]$ for every class c . Note that all feature embeddings $\phi(x)$, $\phi^*(x)$ and prototypes p_c are ℓ_2 -normalized to the unit hypersphere \mathbb{S}^{d-1} . The ACL loss encourages high cosine similarity between an adapted embedding and its corresponding class prototype, while penalizing similarities to incorrect prototypes:

$$\mathcal{L}_{\text{ACL}}(x_i, y_i) = -\log \frac{\exp(\cos(\phi^*(x_i), p_{y_i})/\tau)}{\sum_j \exp(\cos(\phi^*(x_i), p_j)/\tau)}, \quad (3)$$

where τ is a temperature parameter.

Plasticity Analysis We first establish that minimizing the ACL loss promotes plasticity by directly reducing an upper bound on the classification error for the current task.

Proposition 1. *For the adapted model $\phi^*(\cdot)$, the probability of misclassifying a sample x_i is upper-bounded by the expected ACL loss:*

$$P(\text{misclassify}) \leq \frac{\mathbb{E}[\mathcal{L}_{\text{ACL}}]}{\log 2}. \quad (4)$$

Proof. By rewriting the ACL loss in terms of log-sum-exp, we obtain:

$$\begin{aligned} \mathcal{L}_{\text{ACL}}(x_i, y_i) &= -\log \frac{\exp(S_{y_i})}{\exp(S_{y_i}) + \sum_{j \neq y_i} \exp(S_j)} \\ &= \log \frac{\exp(S_{y_i}) + \sum_{j \neq y_i} \exp(S_j)}{\exp(S_{y_i})} \\ &= \log(1 + \sum_{j \neq y_i} \exp(S_j - S_{y_i})). \end{aligned} \quad (5)$$

where $S_j = \cos(\phi^*(x_i), p_j)/\tau$; $S_{y_i} = \cos(\phi^*(x_i), p_{y_i})/\tau$.

By definition, for the adapted model ϕ^* using a cosine classifier, a sample (x_i, y_i) is misclassified if and only if there exists an incorrect class $k \neq y_i$ such that $\cos(\phi(x_i), p_k) \geq \cos(\phi(x_i), p_{y_i})$. Therefore, in the event of misclassification, there exists at least one k such that $S_k - S_{y_i} \geq 0$, making $\exp(S_k - S_{y_i}) \geq 1$. Thus, the sum $\sum_{j \neq y_i} \exp(S_j - S_{y_i}) \geq 1$, and $\mathcal{L}_{\text{ACL}}(x_i, y_i) \geq \log(1 + 1) = \log 2$. That is, $\text{misclassify} \implies \mathcal{L}_{\text{ACL}}(x_i, y_i) \geq \log(1 + 1) = \log 2$.

Therefore, $P(\text{misclassify}) \leq P(\mathcal{L}_{\text{ACL}} \geq \log 2)$. Applying Markov's inequality to \mathcal{L}_{ACL} yields the final bound:

$$P(\text{misclassify}) \leq P(\mathcal{L}_{\text{ACL}} \geq \log 2) \leq \frac{\mathbb{E}[\mathcal{L}_{\text{ACL}}]}{\log 2}. \quad (6)$$

□

Conclusion on plasticity. Minimizing the expected ACL loss directly minimizes a probabilistic upper bound on the current-task classification error, thereby enhancing plasticity.

Stability Analysis We now demonstrate that the ACL loss maintains stability through an implicit regularization mechanism that constrains feature deviation.

Lemma 1. *For any two ℓ_2 -normalized vectors $a, b \in \mathbb{S}^{d-1}$, the squared Euclidean distance is directly related to their cosine similarity:*

$$\|a - b\|_2^2 = 2(1 - \cos(a, b)). \quad (7)$$

Lemma 2. *For a give class y_i , the prototype is the unique point that minimizes the expected squared Euclidean distance to all features.*

$$p_{y_i} = \arg \min_z \mathbb{E}_{(x,y) \in \mathcal{D}_k, y=y_i} [\|\phi(x) - z\|_2^2]. \quad (8)$$

Proposition 2. *The ACL loss implicitly regularizes the feature deviation. Specifically, it enforces a tight upper bound on the expected feature deviation $\mathbb{E}[\|\phi^*(x) - \phi(x)\|_2^2]$.*

Proof. Consider the expected squared feature change over the current task distribution \mathcal{D}_k :

$$\mathbb{E}_{(x,y) \sim \mathcal{D}_k} [\|\phi^*(x) - \phi(x)\|_2^2].$$

For any sample (x_i, y_i) from \mathcal{D}_k , we can bound the squared feature change $\|\phi^*(x_i) - \phi(x_i)\|_2^2$ by using the class prototype p_{y_i} as an intermediate point and applying a fundamental inequality for squared norms:

$$\|\phi^*(x_i) - \phi(x_i)\|_2^2 \leq 2(\|\phi^*(x_i) - p_{y_i}\|_2^2 + \|\phi(x_i) - p_{y_i}\|_2^2). \quad (9)$$

Taking expectations over $x \sim \mathcal{D}_k$ on both sides preserves the inequality due to the linearity of expectation:

$$\begin{aligned} \mathbb{E}_{(x,y) \sim \mathcal{D}_k} [\|\phi^*(x) - \phi(x)\|_2^2] &\leq \\ 2(\mathbb{E}_{(x,y) \sim \mathcal{D}_k} [\|\phi^*(x) - p_y\|_2^2] &+ \mathbb{E}_{(x,y) \sim \mathcal{D}_k} [\|\phi(x) - p_y\|_2^2]). \end{aligned} \quad (10)$$

Let us now analyze each expected term:

- $\mathbb{E}_{(x,y) \sim \mathcal{D}_k} [\|\phi^*(x) - p_y\|_2^2]$. By the Lemma 1, this equals $2 \mathbb{E}_{(x,y) \sim \mathcal{D}_k} [1 - \cos(\phi^*(x), p_y)]$. ACL loss directly minimizes this term by maximizing the expected cosine similarity $\mathbb{E}_{(x,y) \sim \mathcal{D}_k} [\cos(\phi^*(x), p_y)]$.
- $\mathbb{E}_{(x,y) \sim \mathcal{D}_k} [\|\phi(x) - p_y\|_2^2]$. Since ϕ is fixed during adaptation and the prototypes p_y are computed from ϕ , this term is constant with respect to the adaptation process. Moreover, by Lemma 2, the use of class prototypes ensures that this term is minimized over all possible anchors, thereby tightening the overall bound.

As a result, the ACL loss serves as an implicit regularization mechanism that constrains feature deviation. □

Conclusion on stability. Minimizing the ACL loss implicitly constrains the deviation between original and adapted features, thereby maintaining stability. Intuitively, this is achieved by anchoring the adapted features to the most representative points in the original feature space, i.e., prototypes.

Method	ImageNet-R-Inc20		ImageNet-R-Inc10		ImageNet-A-Inc20		ImageNet-A-Inc10	
	LA	AIA	LA	AIA	LA	AIA	LA	AIA
L2P	71.91±0.27	76.76±0.45	69.24±0.78	74.61±0.61	42.58±0.39	50.42±1.12	34.93±0.96	44.24±1.25
w/ Ours	75.47±0.53	80.09±0.40	73.07±0.57	78.49±0.56	48.65±0.55	55.01±1.24	41.92±1.61	49.34±2.06
Improvement	+3.56	+3.33	+3.83	+3.88	+6.07	+4.59	+6.99	+5.10
DualPrompt	69.43±0.51	74.85±0.18	65.71±0.24	71.89±0.34	45.35±1.04	54.72±1.64	39.04±1.83	49.46±2.26
w/ Ours	74.97±0.25	79.88±0.44	72.07±0.32	77.52±0.40	53.22±0.70	60.02±1.91	46.65±1.19	54.97±2.01
Improvement	+5.54	+5.03	+6.36	+5.63	+7.87	+5.30	+7.61	+5.51
FeCAM	60.39±1.30	66.15±1.24	55.60±0.26	61.98±0.42	33.43±0.18	41.89±0.95	33.79±0.10	42.96±0.65
w/ Ours	65.82±0.80	70.33±0.86	63.05±1.48	67.96±1.25	41.28±0.61	46.67±1.89	38.62±0.44	45.56±0.98
Improvement	+5.43	+4.18	+7.45	+5.98	+7.85	+4.78	+4.83	+2.60
RanPAC	76.07±0.85	81.18±0.94	72.84±0.23	78.47±0.55	58.16±0.46	66.73±1.47	57.33±1.26	65.79±1.55
w/ Ours	79.14±0.21	83.29±0.50	78.20±0.25	82.37±0.34	64.45±0.37	70.59±1.93	61.57±1.75	66.22±3.48
Improvement	+3.07	+2.11	+5.36	+3.90	+6.29	+3.86	+4.24	+0.43
SSIAT	78.76±0.24	81.64±0.34	77.18±0.15	80.04±0.34	59.57±0.32	66.54±1.36	56.34±0.70	65.62±1.43
w/ Ours	79.13±0.22	82.80±0.33	77.93±0.28	81.68±0.27	63.91±0.39	69.84±1.42	59.97±0.57	68.14±1.28
Improvement	+0.37	+1.16	+0.75	+1.64	+4.34	+3.30	+3.63	+2.52
MOS	74.07±0.36	78.84±0.43	71.50±0.20	76.92±0.22	57.71±0.55	65.84±1.00	56.06±0.08	65.71±1.02
w/ Ours	77.03±0.43	81.68±0.51	76.54±0.37	80.97±0.26	62.87±0.82	68.91±1.68	61.54±0.31	68.50±1.01
Improvement	+2.96	+2.84	+5.04	+4.05	+5.16	+3.07	+5.48	+2.79

Table 1: Performance (%) of six state-of-the-art CL methods with/without ACL. ‘Improvement’ represents the boost of ACL.

Experiments

Experiment Setup

Dataset. Given that PTMs are typically trained on ImageNet series datasets (Ridnik et al. 2021), evaluation on the standard ImageNet benchmark is meaningless due to the overlapping data distribution (Zhou et al. 2024a). Hence, we evaluate ACL on two datasets that exhibit a significant domain gap (Zhou et al. 2024a) with ImageNet, *i.e.*, ImageNet-R (Hendrycks et al. 2021a) and ImageNet-A (Hendrycks et al. 2021b). To simulate a CL scenario, both datasets are equally divided into multiple tasks without overlapping data. Specifically, we create two task configurations: (1) 20 tasks with **10** classes each (Inc-10) and (2) 10 tasks with **20** classes each (Inc-20).

Baselines. We compare our proposed method against six state-of-the-art PTM-based CL methods: L2P (Wang et al. 2022b), DualPrompt (Wang et al. 2022a), RanPAC (McDonnell et al. 2024), FeCAM (Goswami et al. 2024), SSIAT (Tan et al. 2024), and MOS (Sun et al. 2024). Since our framework is designed as a plug-and-play component, we integrate it into these baseline methods to systematically assess its effectiveness. We note that L2P and DualPrompt use a linear classifier, and other methods use a cosine classifier or its variants. Furthermore, we include a comparison with Aper (Zhou et al. 2024a), a method that fine-tunes the PTM on the first task using standard classification loss and freezes the resulting model for future tasks. Aper is specialized through various adaptation algorithms, including full **Finetune**, Visual Prompt Tuning (**VPT**), Scale and Shift (**SSF**), and **Adapter-**

based tuning. In particular, Aper with VPT has two variants: **VPT-Deep**, which prepends the prompts at every attention layer, and **VPT-Shallow**, which only prepends the prompts at the first layer (Zhou et al. 2024a). Following these baselines, our validation focuses on a general and realistic CL scenario *i.e.*, class incremental learning (Wang et al. 2023).

Implementation Details. We select a representative PTM, denoted as ViT-B/16-IN1K, for our experiments. This PTM is initially pre-trained on ImageNet21K (Ridnik et al. 2021) and subsequently finetuned on ImageNet1K (Deng et al. 2009). To ensure consistency and reproducibility, we adhere to the hyperparameter configurations provided by the open-source library PILOT (Sun et al. 2023) for all CL methods. For each incremental task, we limit the adaptation phase to 1 training epoch to minimize computational overhead and set $\tau = 0.1$ for ACL loss. For all results, we report mean \pm std of 5 runs with different task orders.

Evaluation Metrics. In line with established conventions (Zhou et al. 2024b; Wang et al. 2023), the CL performance is evaluated using two key metrics: *Last Accuracy* (LA) and *Average Incremental Accuracy* (AIA). LA measures the model’s performance across all classes after completing the final task, while AIA quantifies the average performance of the model after learning each incremental task (Lu et al. 2024). Formally, let K denote the total number of tasks, and let A_b represent the classification accuracy evaluated on the test set encompassing all classes learned up to and including the b -th task. These metrics are defined as $LA = A_K$ and $AIA = \frac{1}{K} \sum_{b=1}^K A_b$. Higher values for LA and AIA indicate

Method	ImageNet-R-Inc20		ImageNet-R-Inc10		ImageNet-A-Inc20		ImageNet-A-Inc10	
	LA	AIA	LA	AIA	LA	AIA	LA	AIA
SimpleCIL	61.35±0.00	66.97±0.46	61.35±0.00	67.58±0.47	49.24±0.00	58.35±1.16	49.24±0.00	59.33±1.01
Aper w/ Finetune	63.60±1.16	71.77±0.91	64.19±1.11	71.54±1.02	51.74±1.91	60.65±1.94	50.44±2.34	60.71±2.09
Aper w/ VPT-Deep	68.70±5.76	75.08±6.13	68.00±1.05	74.71±1.34	46.11±3.25	56.03±3.22	42.15±4.09	53.01±4.91
Aper w/ VPT-Shallow	64.50±0.72	70.21±0.91	64.83±0.38	71.20±0.71	46.90±1.46	56.42±0.83	45.61±1.84	56.55±2.36
Aper w/ SSF	<u>70.07±0.37</u>	<u>76.29±0.80</u>	67.84±0.06	74.31±0.44	50.24±1.47	59.65±0.94	47.93±1.57	58.59±1.16
Aper w/ Adapter	67.25±1.21	73.13±1.54	62.47±0.17	68.70±0.66	49.22±0.03	58.37±1.17	49.23±0.07	59.34±1.03
ACL (Ours)	73.93±0.38	77.90±0.57	72.26±0.43	76.33±0.46	56.88±0.31	63.50±1.85	55.18±0.34	62.64±1.53
Improvement	+3.86	+1.61	+4.26	+1.62	+5.14	+2.85	+4.74	+1.93

Table 2: Performance comparison (%) between Aper and ACL. **Bolded** indicates the best performance, underline denotes the second best. ‘Improvement’ represents the boost of ACL towards the best variants of Aper.

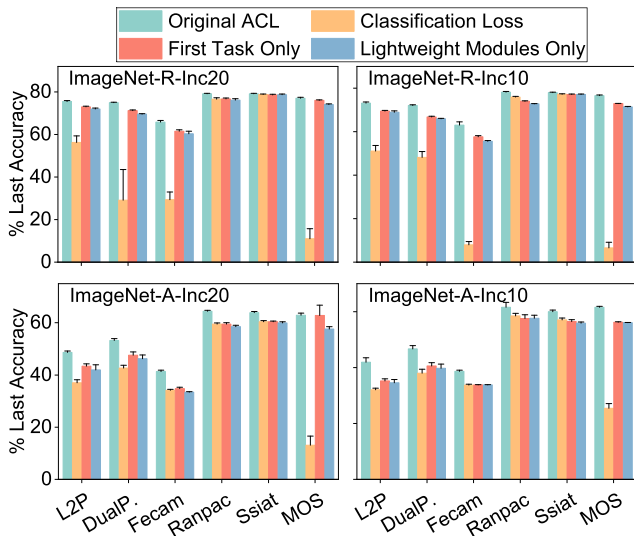


Figure 3: Performance of original ACL and its ablation variants, including (1) using standard **classification loss** for adaptation, (2) adapting for the **first task only**, and (3) adapting **lightweight modules only** with frozen backbone.

superior CL performance, reflecting a better balance between stability and plasticity.

Main Results

Integration with Existing Methods. We begin by assessing ACL’s effectiveness when incorporated into six state-of-the-art PTM-based CL methods. As shown in Tab. 1, ACL consistently improves the performance of these methods across diverse datasets and incremental steps. Notably, ACL achieves gains of up to 7.87% in LA and 5.98% in AIA compared to the original methods. These findings highlight the effectiveness of ACL in achieving a better stability-plasticity trade-off, thereby enhancing the performance of existing CL methods.

Comparison with Aper. We further compare ACL with Aper, a method that adapts the PTM using a standard classification loss on the first task only. To ensure a fair comparison,

we integrate ACL with SimpleCIL (Zhou et al. 2024a), which uses the same cosine classifier as Aper and involves no training after adaptation. As shown in Tab. 2, ACL consistently surpasses Aper across all datasets and incremental steps. Notably, ACL achieves performance gains of up to 5.14% in LA and 2.85% in AIA compared to the best-performing variants of Aper. These results demonstrate ACL’s superiority in leading to a better stability-plasticity balance than Aper.

Ablation Study

Key Components for Adaptation. We first conduct an ablation study to systematically investigate the impact of various strategies employed during the adaptation phase, focusing on three key aspects: loss design, adaptation steps, and adapted network components. As illustrated in Fig. 3, all ablation variants exhibit inferior performance compared to the original ACL framework across all integrated CL methods. These results underscore the critical importance of *continually adapting the entire PTM backbone using the proposed ACL loss for all incremental tasks*, which facilitates more effective adaptation and knowledge retention.

Full PTM Adaptation vs. Multi-Epoch Adaptation. To further show the advantages of adapting the entire PTM versus solely adapting lightweight modules, we extended our comparison by considering the impact of multiple adaptation epochs. The results, presented in Fig. 4, demonstrate two key findings: (1) adapting the entire PTM consistently outperforms adapting only lightweight modules; (2) increasing adaptation epochs beyond two yields marginal or negligible performance gains for either strategy. These observations indicate that the performance benefits derived from full PTM adaptation cannot be replicated merely by increasing the adaptation epochs. This reinforces our claim that freezing pre-trained weights results in a suboptimal balance between stability and plasticity.

Visualization

We employ t-SNE (Van der Maaten and Hinton 2008) to visualize the feature representations extracted by the final PTMs, comparing those obtained without and with the proposed ACL framework. For simplicity, we adopt Simple-

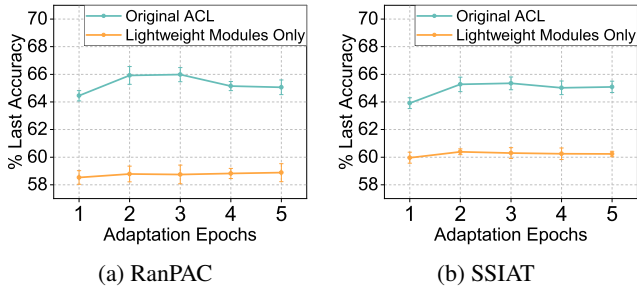


Figure 4: Performance with different adaptation epochs.

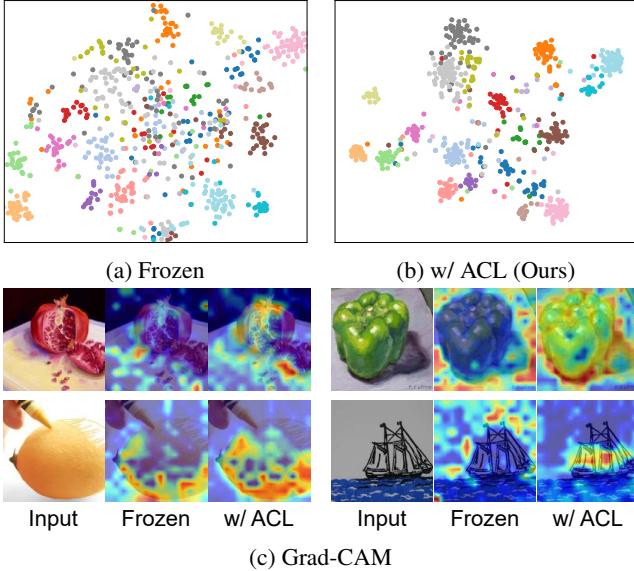


Figure 5: (a-b) Visualization of 2D feature representations using t-SNE. (c) Grad-CAM visualization, where important regions are highlighted with warm colors.

CIL (Zhou et al. 2024a) as the baseline method and evaluate on ImageNet-R-Inc20, selecting two classes per incremental task for clearer visualization. The results, shown in Fig. 5(a) and (b), demonstrate that the PTM adapted with ACL generates more discriminative feature representations than the frozen model, even for classes in previously learned tasks. This indicates that ACL effectively enhances feature discriminability across all incremental tasks, achieving a better stability-plasticity trade-off.

To further validate our approach, we visualize Grad-CAM (Selvaraju et al. 2017) results on samples with a large domain gap (Hendrycks et al. 2021a) relative to the pre-training data, which highlight critical image regions for concept prediction. As depicted in Fig. 5(c), the frozen PTM often attends to irrelevant background regions. In contrast, the PTM adapted via ACL focuses more accurately on class-specific features. These findings underscore the necessity of adapting PTMs to incremental data, especially when the target distribution significantly diverges from the pre-training domain.

Methods	Original	w/ ACL (Ours)	Improv.
L2P	39.83±1.15	45.49±0.56	+5.66
DualPrompt	43.02±1.29	45.57±1.56	+2.55
FeCAM	45.37±0.30	47.81±1.22	+2.44
RanPAC	54.59±0.84	58.97±0.21	+4.38
SSIAT	56.84±0.49	59.80±0.33	+2.96
MOS	54.17±0.45	58.55±0.24	+4.38

Table 3: Performance (%LA) using ViT-B/16-IN21K. ‘Improv.’ represents the boost of ACL towards original methods.

Method	LA	AIA
Continual CLIP	71.70±0.01	78.73±0.66
w/ Ours	74.98±0.25 (+3.28)	80.95±0.51 (+2.22)

Table 4: Performance (%) using CLIP with/without ACL.

Validation on Other Backbones

Validation on ViT-B/16-IN21K. To further validate the effectiveness of ACL, we conduct experiments on ViT-B/16-IN21K, a model pre-trained on ImageNet21K only, with ImageNet-A-Inc20 as the benchmark. As shown in Tab. 3, ACL consistently enhances the CL performance across various CL methods. These findings underscore the versatility and generalizability of our framework.

Validation on CLIP. While our study primarily focuses on visual models, the insights presented in our paper are potentially applicable to visual-language models, such as CLIP (Radford et al. 2021). To demonstrate this, we employ Continual CLIP (Thengane et al. 2022) as the baseline and evaluate ACL on the ImageNet-R-inc20 benchmark. Since the text labels for the same class are consistent, we only adapt the visual encoder using ACL. The experimental results, summarized in Tab. 4, indicate that ACL significantly enhances the CL performance of CLIP. These findings demonstrate the potential of ACL to improve CL in the context of visual-language models.

Conclusion

In this paper, we revisit CL with PTMs and argue that existing PTM-based CL methods overly prioritize stability at the expense of plasticity. To address this limitation, we propose ACL, a framework that can be orthogonally integrated with existing PTM-based CL methods to enhance plasticity while simultaneously maintaining stability. Extensive experiments demonstrate the effectiveness of ACL in enhancing plasticity and achieving a more balanced stability-plasticity trade-off. Future work will focus on exploring more effective or efficient adaptation algorithms within the ACL framework to further improve its performance and applicability.

Limitations. Adapting the entire PTM introduces additional GPU memory consumption, specifically, approximately 7GB with the experiment settings in Tab. 1.

Acknowledgements

This work was supported by National Natural Science Foundation of China under Grant 62276175 and Innovative Research Group Program of Natural Science Foundation of Sichuan Province under Grant 2024NSFTD0035.

References

- Aljundi, R.; Lin, M.; Goujaud, B.; and Bengio, Y. 2019. Gradient based sample selection for online continual learning. *Advances in neural information processing systems*, 32.
- Bian, A.; Li, W.; Yuan, H.; Wang, M.; Zhao, Z.; Lu, A.; Ji, P.; Feng, T.; et al. 2024. Make continual learning stronger via c-flat. *Advances in Neural Information Processing Systems*, 37: 7608–7630.
- Cha, H.; Lee, J.; and Shin, J. 2021. Co2l: Contrastive continual learning. In *Proceedings of the IEEE/CVF International conference on computer vision*, 9516–9525.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Dosovitskiy, A. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Feng, T.; Li, W.; Zhu, D.; Yuan, H.; Zheng, W.; Zhang, D.; and Tang, J. 2025. Zeroflow: Overcoming catastrophic forgetting is easier than you think. *arXiv preprint arXiv:2501.01045*.
- Feng, T.; Wang, M.; and Yuan, H. 2022. Overcoming catastrophic forgetting in incremental object detection via elastic response distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9427–9436.
- Gan, K.; and Wei, T. 2024. Erasing the bias: fine-tuning foundation models for semi-supervised learning. In *Proceedings of the 41st International Conference on Machine Learning*, 14453–14470.
- Gao, X.; Dong, S.; He, Y.; Wang, Q.; and Gong, Y. 2024. Beyond prompt learning: Continual adapter for efficient rehearsal-free continual learning. In *European Conference on Computer Vision*, 89–106. Springer.
- Goodfellow, I. J.; Mirza, M.; Xiao, D.; Courville, A.; and Bengio, Y. 2013. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*.
- Goswami, D.; Liu, Y.; Twardowski, B.; and van de Weijer, J. 2024. Fecam: Exploiting the heterogeneity of class distributions in exemplar-free continual learning. *Advances in Neural Information Processing Systems*, 36.
- Grossberg, S. 2013. Adaptive Resonance Theory: How a brain learns to consciously attend, learn, and recognize a changing world. *Neural networks*, 37: 1–47.
- Hendrycks, D.; Basart, S.; Mu, N.; Kadavath, S.; Wang, F.; Dorundo, E.; Desai, R.; Zhu, T.; Parajuli, S.; Guo, M.; et al. 2021a. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF international conference on computer vision*, 8340–8349.
- Hendrycks, D.; Zhao, K.; Basart, S.; Steinhardt, J.; and Song, D. 2021b. Natural adversarial examples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 15262–15271.
- Iscen, A.; Zhang, J.; Lazebnik, S.; and Schmid, C. 2020. Memory-efficient incremental learning through feature adaptation. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVI 16*, 699–715. Springer.
- Jia, M.; Tang, L.; Chen, B.-C.; Cardie, C.; Belongie, S.; Hariharan, B.; and Lim, S.-N. 2022. Visual prompt tuning. In *European Conference on Computer Vision*, 709–727. Springer.
- Jung, D.; Han, D.; Bang, J.; and Song, H. 2023. Generating instance-level prompts for rehearsal-free continual learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11847–11857.
- Kang, H.; Mina, R. J. L.; Madjid, S. R. H.; Yoon, J.; Hasegawa-Johnson, M.; Hwang, S. J.; and Yoo, C. D. 2022. Forget-free continual learning with winning subnetworks. In *International Conference on Machine Learning*, 10734–10750. PMLR.
- Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; and Krishnan, D. 2020. Supervised contrastive learning. *Advances in neural information processing systems*, 33: 18661–18673.
- Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N.; Veness, J.; Desjardins, G.; Rusu, A. A.; Milan, K.; Quan, J.; Ramalho, T.; Grabska-Barwinska, A.; et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13): 3521–3526.
- Konishi, T.; Kurokawa, M.; Ono, C.; Ke, Z.; Kim, G.; and Liu, B. 2023. Parameter-level soft-masking for continual learning. In *International Conference on Machine Learning*, 17492–17505. PMLR.
- Kumar, A.; Raghunathan, A.; Jones, R. M.; Ma, T.; and Liang, P. 2022. Fine-Tuning can Distort Pretrained Features and Underperform Out-of-Distribution. In *International Conference on Learning Representations*.
- Li, J.; Zhou, P.; Xiong, C.; and Hoi, S. C. 2020. Prototypical contrastive learning of unsupervised representations. *arXiv preprint arXiv:2005.04966*.
- Li, Z.; and Hoiem, D. 2017. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12): 2935–2947.
- Lin, H.; Zhang, B.; Feng, S.; Li, X.; and Ye, Y. 2023. Pcr: Proxy-based contrastive replay for online class-incremental continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24246–24255.
- Liu, Y.; Su, Y.; Liu, A.-A.; Schiele, B.; and Sun, Q. 2020. Mnemonics training: Multi-class incremental learning without forgetting. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, 12245–12254.

- Lu, A.; Feng, T.; Yuan, H.; Song, X.; and Sun, Y. 2024. Revisiting Neural Networks for Continual Learning: An Architectural Perspective. In *IJCAI*, 4651–4659.
- Lu, A.; Yuan, H.; Feng, T.; and Sun, Y. 2025. Rethinking the stability-plasticity trade-off in continual learning from an architectural perspective. *ICML*.
- Luo, M.-L.; Zhou, Z.-H.; Wei, T.; and Zhang, M.-L. 2025. LADA: Scalable Label-Specific CLIP Adapter for Continual Learning. *arXiv preprint arXiv:2505.23271*.
- Masana, M.; Liu, X.; Twardowski, B.; Menta, M.; Bagdanov, A. D.; and Van De Weijer, J. 2022. Class-incremental learning: survey and performance evaluation on image classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(5): 5513–5533.
- McDonnell, M. D.; Gong, D.; Parvaneh, A.; Abbasnejad, E.; and van den Hengel, A. 2024. Ranpac: Random projections and pre-trained models for continual learning. *Advances in Neural Information Processing Systems*, 36.
- Nagata, K.; and Hotta, K. 2023. Margin Contrastive Learning with Learnable-Vector for Continual Learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3570–3576.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Ridnik, T.; Ben-Baruch, E.; Noy, A.; and Zelnik-Manor, L. 2021. Imagenet-21k pretraining for the masses. *arXiv preprint arXiv:2104.10972*.
- Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, 618–626.
- Smith, J. S.; Karlinsky, L.; Gutta, V.; Cascante-Bonilla, P.; Kim, D.; Arbelle, A.; Panda, R.; Feris, R.; and Kira, Z. 2023. 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*, 11909–11919.
- Snell, J.; Swersky, K.; and Zemel, R. 2017. Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30.
- Sun, H.-L.; Zhou, D.-W.; Ye, H.-J.; and Zhan, D.-C. 2023. PI-LOT: A Pre-Trained Model-Based Continual Learning Toolbox. *arXiv preprint arXiv:2309.07117*.
- Sun, H.-L.; Zhou, D.-W.; Zhao, H.; Gan, L.; Zhan, D.-C.; and Ye, H.-J. 2024. MOS: Model Surgery for Pre-Trained Model-Based Class-Incremental Learning. *arXiv preprint arXiv:2412.09441*.
- Tan, Y.; Zhou, Q.; Xiang, X.; Wang, K.; Wu, Y.; and Li, Y. 2024. Semantically-Shifted Incremental Adapter-Tuning is A Continual ViTransformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23252–23262.
- Thengane, V.; Khan, S.; Hayat, M.; and Khan, F. 2022. CLIP model is an Efficient Continual Learner. *arXiv:2210.03114*.
- Van de Ven, G. M.; Tuytelaars, T.; and Tolias, A. S. 2022. Three types of incremental learning. *Nature Machine Intelligence*, 4(12): 1185–1197.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).
- Wang, L.; Zhang, X.; Su, H.; and Zhu, J. 2023. A comprehensive survey of continual learning: Theory, method and application. *arXiv preprint arXiv:2302.00487*.
- Wang, Z.; Zhang, Z.; Ebrahimi, S.; Sun, R.; Zhang, H.; Lee, C.-Y.; Ren, X.; Su, G.; Perot, V.; Dy, J.; et al. 2022a. Dual-prompt: Complementary prompting for rehearsal-free continual learning. In *European Conference on Computer Vision*, 631–648. Springer.
- Wang, Z.; Zhang, Z.; Lee, C.-Y.; Zhang, H.; Sun, R.; Ren, X.; Su, G.; Perot, V.; Dy, J.; and Pfister, T. 2022b. Learning to prompt for continual learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 139–149.
- Wen, Y.; Tan, Z.; Zheng, K.; Xie, C.; and Huang, W. 2024. Provable Contrastive Continual Learning. *arXiv preprint arXiv:2405.18756*.
- Yan, S.; Xie, J.; and He, X. 2021. Der: Dynamically expandable representation for class incremental learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3014–3023.
- Zenke, F.; Poole, B.; and Ganguli, S. 2017. Continual learning through synaptic intelligence. In *International conference on machine learning*, 3987–3995. PMLR.
- Zhang, X.; Song, D.; and Tao, D. 2022. Hierarchical prototype networks for continual graph representation learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4): 4622–4636.
- Zhao, H.; Wang, H.; Fu, Y.; Wu, F.; and Li, X. 2021. Memory-efficient class-incremental learning for image classification. *IEEE Transactions on Neural Networks and Learning Systems*, 33(10): 5966–5977.
- Zhou, D.-W.; Cai, Z.-W.; Ye, H.-J.; Zhan, D.-C.; and Liu, Z. 2024a. Revisiting class-incremental learning with pre-trained models: Generalizability and adaptivity are all you need. *International Journal of Computer Vision*, 1–21.
- Zhou, D.-W.; Sun, H.-L.; Ning, J.; Ye, H.-J.; and Zhan, D.-C. 2024b. Continual learning with pre-trained models: A survey. *arXiv preprint arXiv:2401.16386*.
- Zhou, D.-W.; Wang, Q.-W.; Ye, H.-J.; and Zhan, D.-C. 2023. A Model or 603 Exemplars: Towards Memory-Efficient Class-Incremental Learning. In *ICLR*.
- Zhou, K.; Liu, Z.; Qiao, Y.; Xiang, T.; and Loy, C. C. 2022. Domain generalization: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 45(4): 4396–4415.