

Attention Retention for Continual Learning with Vision Transformers

Yue Lu^{1†}, Xiangyu Zhou^{1†}, Shizhou Zhang^{1*}, Yinghui Xing¹, Guoqiang Liang¹, Wencong Zhang^{2*}

¹School of Computer Science, Northwestern Polytechnical University, China

²Hikrobot Co., Ltd., China

zgxld@mail.nwpu.edu.cn, 2021300119@mail.nwpu.edu.cn, szzhang@nwpu.edu.cn, xyh_7491@nwpu.edu.cn, gqliang@nwpu.edu.cn, zhangwencong@hikrobotics.com

Abstract

Continual learning (CL) empowers AI systems to progressively acquire knowledge from non-stationary data streams. However, *catastrophic forgetting* remains a critical challenge. In this work, we identify *attention drift* in Vision Transformers as a primary source of catastrophic forgetting, where the attention to previously learned visual concepts shifts significantly after learning new tasks. Inspired by neuroscientific insights into the selective attention in the human visual system, we propose a novel attention-retaining framework to mitigate forgetting in CL. Our method constrains attention drift by explicitly modifying gradients during backpropagation through a two-step process: 1) extracting attention maps of the previous task using a layer-wise rollout mechanism and generating instance-adaptive binary masks, and 2) when learning a new task, applying these masks to zero out gradients associated with previous attention regions, thereby preventing disruption of learned visual concepts. For compatibility with modern optimizers, the gradient masking process is further enhanced by scaling parameter updates proportionally to maintain their relative magnitudes. Experiments and visualizations demonstrate the effectiveness of our method in mitigating catastrophic forgetting and preserving visual concepts. It achieves state-of-the-art performance and exhibits robust generalizability across diverse CL scenarios.

Introduction

Humans are capable of continuously accumulating and integrating knowledge in dynamic environments. Similarly, for AI systems, the ability to learn continuously over time is essential for achieving long-term adaptation and intelligent evolution. *Continual learning* (CL), which aims to enable AI systems to learn adaptively in non-stationary data streams, akin to how humans acquire and develop knowledge.

Catastrophic forgetting remains the central challenge in continual learning for artificial neural networks. Extensive efforts have been devoted to addressing this issue. For instance, replay-based methods (Rebuffi et al. 2017; Wu et al. 2018) store and replay previously learned samples. Regularization-based approaches (Kirkpatrick et al. 2017;

*Shizhou Zhang and Wencong Zhang are co-corresponding authors.

† Yue Lu and Xiangyu Zhou are co-first authors.

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

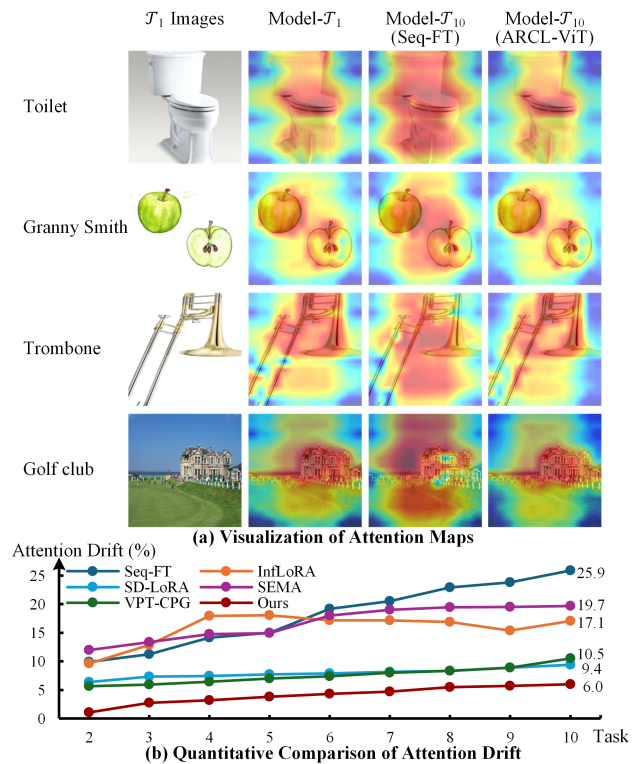


Figure 1: (a) Visualization of attention maps on the first task’s (\mathcal{T}_1) samples. Columns two to four show attention maps from ViT models that are (i) only fine-tuned on \mathcal{T}_1 , (ii) sequentially fine-tuned over 10 tasks (Model- \mathcal{T}_{10} Seq-FT), and (iii) trained with our method over 10 tasks. Seq-FT leads to severe attention drift, while our method effectively preserves the original attention and thus mitigates forgetting. (b) Quantitative comparison of attention drift relative to \mathcal{T}_1 .

Zeng et al. 2019) constrain changes in important parameters. Expansion-based methods (Yan, Xie, and He 2021; Smith et al. 2023) allocate independent network branches for new tasks. These strategies mimic various aspects of human learning, such as memory consolidation, synaptic plasticity regulation, and the formation of new neural connections. They offer valuable insights into building continual

learning systems that resemble biological intelligence.

In the visual system, limited neural resources force the brain to selectively process the sensory inputs. (Zhang et al. 2012) report that the primary visual cortex (also known as V1) generates a saliency map through neural activity to highlight attention-attracting regions in the visual field. This selective attention remains consistently anchored on the discriminative features of visual concepts, even as new concepts are acquired. Such inherent stability ensures that previously learned knowledge is preserved and not disrupted by subsequent learning, forming the biological foundation for non-forgetting in human visual cognition. Inspired by this neuroscientific insight, we argue that *preventing attention drift from previously learned concepts* is essential for CL systems to overcome catastrophic forgetting. However, Vision Transformers (Dosovitskiy et al. 2021) (ViTs) exhibit severe *attention drift* after learning new tasks, as illustrated in the third column of Figure 1(a). To address this issue and enhance the alignment between ViT attention mechanism and biological perception, we propose ARCL-ViT, an Attention-Retaining method that explicitly constrains attention drift on previous tasks during the learning of new tasks, thereby alleviating catastrophic forgetting.

However, directly optimizing the objective of preventing attention drift requires storing previous task data and is computationally infeasible. Instead, we approximately achieve this objective by *suppressing gradients contributing to shifting previous attention* during backpropagation, which eliminates the need for old task data. This is achieved through two steps: 1) Attention Mask Generation. After completing a task, we extract the attention maps corresponding to the learned categories and generate binary masks that identify their critical attention regions. In this step, we introduce a layer-wise rollout mechanism for attention map extraction and develop an instance-adaptive thresholding strategy to dynamically identify salient attention regions. 2) Gradient Masking. When learning the next task, we use the masks to zero out the gradients associated with these regions, so as to prevent parameter updates from disrupting the model’s original attention to previous visual concepts. To ensure compatibility with modern optimizers, we further scale the parameter updates proportionally to maintain the relative magnitude between the masked and original gradients. As a result, our method suppresses attention drift and mitigates forgetting, as shown in the last column of Figure 1(a).

Overall, we summarize our contributions as follows:

- We identify attention drift in Vision Transformers as a primary source of catastrophic forgetting in continual learning, and reveal the necessity of preserving selective attention to previous visual concepts.
- We propose a novel framework for retaining attention, which constrains attention drift via gradient masking and employs a layer-wise rollout and adaptive thresholding strategy to generate effective attention masks.
- Our method is extensively evaluated on challenging continual learning benchmarks, achieving state-of-the-art performance and robust generalization across diverse pre-training schemes and long-sequence settings.

Related Work

Replay-based methods mitigate catastrophic forgetting by emulating the hippocampal mechanism of memory reactivation, which involves periodically revisiting prior knowledge to consolidate learning. These methods can be broadly categorized into two types. 1) Memory-based replay retains and reuses samples from previous tasks, with the key challenge being the efficient selection of exemplars. For instance, (Rebuffi et al. 2017) use a herding strategy to select samples whose features closely approximate those of the entire training set. 2) Generative replay (Wu et al. 2018; Xiang et al. 2019), which synthesizes samples from past tasks using generative models to avoid storing real data. However, replay-based methods often raise concerns regarding data privacy and storage overhead. Consequently, recent CL approaches utilizing pre-trained ViTs have predominantly focused on the replay-free setting.

Expansion-based methods create new branches when learning new tasks, mimicking the formation of new synaptic connections in the brain. In CNN-based approaches, such as DER (Yan, Xie, and He 2021), expansion is implemented by adding new convolutional branches while freezing previously learned feature extractors. In ViT-based approaches, the expansion mechanism is implemented through dynamic extension of prompt pools (Wang et al. 2022b). Representative methods such as DualPrompt (Wang et al. 2022a) and NNPrompt (Lu et al. 2025a) learn a task-specific prompt for each task, which is then stored in a prompt pool. During inference, the model selects a relevant prompt from the pool for the input instance and incorporates the prompt into the prediction process. Therefore, each prompt essentially serves as a parameterized branch tailored to a specific task.

Regularization-based methods, inspired by biological synaptic plasticity that enables the brain to adapt to new environments, facilitate continual learning by adjusting parameter plasticity. Representative approaches include EWC (Kirkpatrick et al. 2017) and SI (Zenke, Poole, and Ganguli 2017), which estimate the importance of each parameter to previous tasks and penalize changes to critical weights when learning new tasks. Another line of regularization-based methods focuses on orthogonal projections, such as OWM (Zeng et al. 2019), NSCL (Wang et al. 2021; Lu et al. 2024; Kong et al. 2022) and GPM (Saha, Garg, and Roy 2021; Qiao et al. 2024). These methods aim to project model parameters onto a subspace that preserves the outputs from previous tasks while learning new ones, with differences in how they implement orthogonality.

Our approach can be categorized as a regularization-based method, as it directly constrains gradients during learning new tasks. By masking gradients associated with prior attention regions, it reduces the disruption of previously acquired attention and effectively mitigates forgetting.

Preliminaries

Continual Learning Problem Formulation

Under the continual learning protocol, a model is trained sequentially on a series of T tasks $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_T$. For the t -th task \mathcal{T}_t , the training set is denoted as $\mathcal{D}_t =$

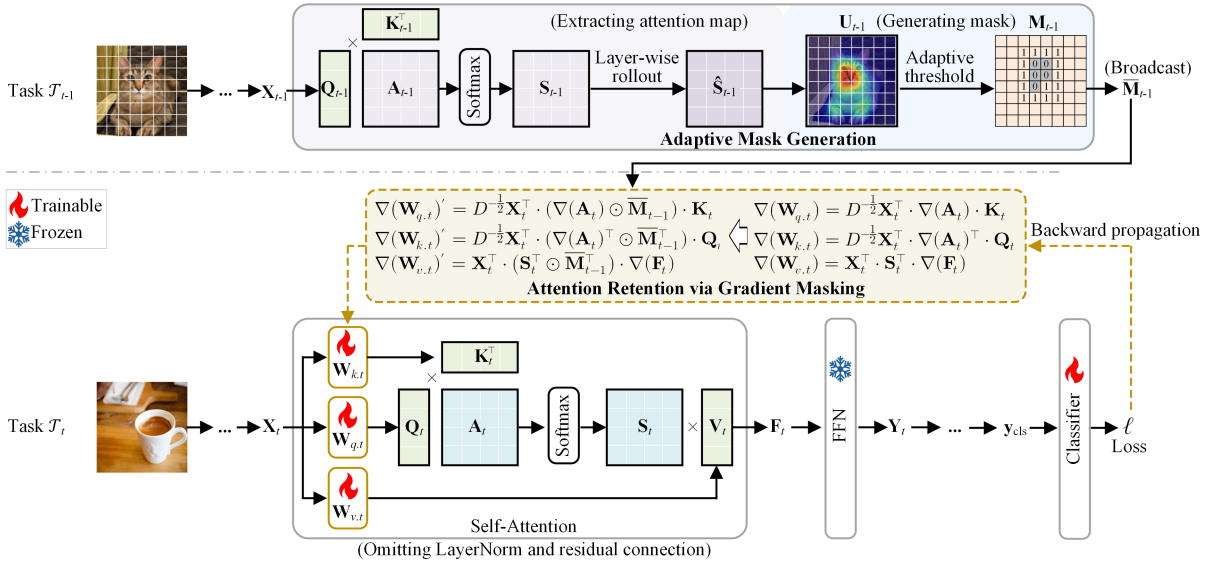


Figure 2: Illustration of our proposed attention-retaining continual learning (ARCL-ViT) framework. After finishing \mathcal{T}_{t-1} , the model extracts attention maps \mathbf{U}_{t-1} and generates the mask $\bar{\mathbf{M}}_{t-1}$ to identify attention regions. During subsequent training in \mathcal{T}_t , the mask $\bar{\mathbf{M}}_{t-1}$ is used to selectively zero out gradients in the corresponding attention regions. The masked gradients $\nabla(\mathbf{W}_{q,t})', \nabla(\mathbf{W}_{k,t})', \nabla(\mathbf{W}_{v,t})'$ are used to update the learnable weights $\mathbf{W}_{q,t}, \mathbf{W}_{k,t}, \mathbf{W}_{v,t}$.

$\{(\mathbf{x}_t^{(n)}, y_t^{(n)})\}_{n=1}^{|\mathcal{T}_t|}$, where $\mathbf{x}_t^{(n)}$ is the n -th input sample with label $y_t^{(n)}$, and $|\mathcal{T}_t|$ denotes the number of instances in this task. We focus on the class-incremental learning scenario, where the label space \mathcal{Y}_t of each task \mathcal{T}_t is disjoint from those of other tasks, *i.e.*, $\mathcal{Y}_t \cap \mathcal{Y}_{t'} = \emptyset$ for $t \neq t'$. After completing training on \mathcal{T}_t , the corresponding training set \mathcal{D}_t becomes unavailable for future tasks. The model is expected to accurately classify test samples from all previously learned classes, without access to task identity at inference time.

Vision Transformer

We adopt the Vision Transformer (Dosovitskiy et al. 2021) (ViT) as our base model. A given image is first divided into N non-overlapping patches and linearly projected into D -dimensional tokens. Then a class token is prepended to the sequence, followed by the addition of positional embeddings. Next, the resulting sequence is fed into L Transformer blocks. Within the l -th block $f^l(\cdot)$, the input tokens first undergo a LayerNorm, and the normalized tokens $\mathbf{X}^l \in \mathbb{R}^{(N+1) \times D}$ are transformed through three projections:

$$\mathbf{Q}^l = \mathbf{X}^l \mathbf{W}_q^l, \quad \mathbf{K}^l = \mathbf{X}^l \mathbf{W}_k^l, \quad \mathbf{V}^l = \mathbf{X}^l \mathbf{W}_v^l, \quad (1)$$

where $\mathbf{W}_q^l, \mathbf{W}_k^l, \mathbf{W}_v^l \in \mathbb{R}^{D \times D}$. The self-attention output is computed as: $\mathbf{F}^l = \text{softmax}\left(\frac{\mathbf{Q}^l (\mathbf{K}^l)^T}{\sqrt{D}}\right) \mathbf{V}^l$. As shown in the ‘‘Self-Attention’’ module in Figure 2, this operation can be decomposed into three steps:

$$\mathbf{A}^l = D^{-\frac{1}{2}} \mathbf{Q}^l (\mathbf{K}^l)^T, \quad \mathbf{S}^l = \text{softmax}(\mathbf{A}^l), \quad \mathbf{F}^l = \mathbf{S}^l \mathbf{V}^l. \quad (2)$$

\mathbf{F}^l is then passed through another LayerNorm followed by a feedforward network (FFN) to produce the output sequence \mathbf{Y}^l for this block. Finally, the class token from the last

layer’s output (denote as $\mathbf{y}_{\text{cls}}^L \in \mathbb{R}^D$) is extracted and fed into the classifier $g(\cdot)$ and loss function $\ell(\cdot)$ for optimization.

In our setting, we employ a ViT model that has been pre-trained on large-scale datasets. During continual learning, we freeze all pre-trained parameters except for the projection weights $\mathbf{W}_q^l, \mathbf{W}_k^l, \mathbf{W}_v^l$ in Transformer block $f^l(\cdot)$ and the classifier $g(\cdot)$. This design allows our model to adjust the attention weights to new tasks while preserving general representations from the pre-trained backbone.

Method

Overview of Our Approach

Our objective is to prevent attention drift from the previous task \mathcal{T}_{t-1} when optimizing parameters for the current task \mathcal{T}_t . Let $\psi^l(\cdot)$ denote the function extracting the attention map at the l -th layer. This goal is formulated as:

$$\begin{aligned} \min_{\mathbf{W}_{\theta,t}^l} \mathbb{E}_{(\mathbf{x}_t, y_t) \sim \mathcal{D}_t} [\ell(g(f^l(\mathbf{x}_t; \mathbf{W}_{\theta,t}^l), y_t))], \\ \text{s.t. } \|\psi^l(\mathbf{x}_{t-1}; \mathbf{W}_{\theta,t}^l) - \psi^l(\mathbf{x}_{t-1}; \mathbf{W}_{\theta,t-1}^l)\| = 0, \end{aligned} \quad (3)$$

where \mathbf{x}_{t-1} denotes samples from \mathcal{D}_{t-1} , and θ represents $\{q, k, v\}$. Directly optimizing this objective is infeasible due to: 1) the complexity of $\psi^l(\cdot)$ makes closed-form solutions intractable, and 2) the need to store \mathcal{D}_{t-1} incurring heavy storage overhead. Instead, we analyze how gradients impact attention regions and observe that updates along attention gradient directions directly alter attention values. By suppressing gradients in previously attended regions, we thus prevent such shifts and preserve prior attention patterns (see Appendix-‘‘Theoretical Justification’’), which enables us to approximately achieve the objective in a data-free manner.

As illustrated in Figure 2, our approach consists of two steps: *Adaptive Mask Generation* and *Gradient Masking*. The first step incorporates a layer-wise attention rollout and instance-adaptive thresholding to generate a mask $\bar{\mathbf{M}}_{t-1}$ for previous task, which is then applied in the second step as a gradient mask to modify gradients. In the following, we first assume a suitable attention mask $\bar{\mathbf{M}}_{t-1}$ is already available and introduce the proposed gradient masking for attention retention, followed by the procedure for mask generation.

Attention Retention via Gradient Masking

We first analyze the relationship between the gradients $\nabla(\cdot)$ w.r.t. the parameters $\mathbf{W}_{\theta,t}$ and the attention matrix \mathbf{A}_t for the current task \mathcal{T}_t . For simplicity, the layer superscript l is omitted. Based on Eq. (2) which defines \mathbf{A} , we have the following derivations: $\nabla(\mathbf{Q}_t) = D^{-\frac{1}{2}} \nabla(\mathbf{A}_t) \mathbf{K}_t$, $\nabla(\mathbf{K}_t) = D^{-\frac{1}{2}} \nabla(\mathbf{A}_t)^\top \mathbf{Q}_t$. Similarly, from Eq. (2) for \mathbf{F} , we obtain: $\nabla(\mathbf{V}_t) = \mathbf{S}_t^\top \nabla(\mathbf{F}_t)$. Building on the QKV computation in Eq. (1), we further derive the relationship between the gradients $\nabla(\mathbf{W}_{q,t})$, $\nabla(\mathbf{W}_{k,t})$, $\nabla(\mathbf{W}_{v,t})$ and \mathbf{A}_t :

$$\begin{cases} \nabla(\mathbf{W}_{q,t}) = D^{-\frac{1}{2}} \mathbf{X}_t^\top \cdot \nabla(\mathbf{A}_t) \cdot \mathbf{K}_t, \\ \nabla(\mathbf{W}_{k,t}) = D^{-\frac{1}{2}} \mathbf{X}_t^\top \cdot \nabla(\mathbf{A}_t)^\top \cdot \mathbf{Q}_t, \\ \nabla(\mathbf{W}_{v,t}) = \mathbf{X}_t^\top \cdot \mathbf{S}_t^\top \cdot \nabla(\mathbf{F}_t). \end{cases} \quad (4)$$

From the above equations, we observe that the attention matrix \mathbf{A}_t (and its activated counterpart \mathbf{S}_t) directly influences the gradients $\nabla(\mathbf{W}_{\theta,t})$. Under standard training, these gradients are applied via gradient descent to update the weights $\mathbf{W}_{\theta,t}$. Consequently, the attention computed by the updated weights shifts toward the *discriminative regions* (i.e., the most salient object features reflecting the visual concept, such as the head for a cat) of \mathcal{D}_t , thereby improving accuracy on the current task. However, the updated model may lose focus on \mathcal{D}_{t-1} , resulting in catastrophic forgetting and degraded performance.

To address this issue, we *remove gradients in the regions corresponding to the previous task's attention during back-propagation when training the current task*. Specifically, when completing the training of task \mathcal{T}_{t-1} , we extract the attention map \mathbf{U}_{t-1} and construct a binary mask $\bar{\mathbf{M}}_{t-1}$ by setting high-attention regions to 0 and low-attention regions to 1. During training of \mathcal{T}_t , we perform an element-wise multiplication (\odot) between $\bar{\mathbf{M}}_{t-1}$ and \mathbf{A}_t (or \mathbf{S}_t), thereby eliminating gradients corresponding to the previous task's attention regions in the updates of $\mathbf{W}_{\{q,t,k,t,v,t\}}$. Formally, the gradient masking is formulated as:

$$\begin{cases} \nabla(\mathbf{W}_{q,t})' = D^{-\frac{1}{2}} \mathbf{X}_t^\top \cdot (\nabla(\mathbf{A}_t) \odot \bar{\mathbf{M}}_{t-1}) \cdot \mathbf{K}_t, \\ \nabla(\mathbf{W}_{k,t})' = D^{-\frac{1}{2}} \mathbf{X}_t^\top \cdot (\nabla(\mathbf{A}_t)^\top \odot \bar{\mathbf{M}}_{t-1}^\top) \cdot \mathbf{Q}_t, \\ \nabla(\mathbf{W}_{v,t})' = \mathbf{X}_t^\top \cdot (\mathbf{S}_t^\top \odot \bar{\mathbf{M}}_{t-1}^\top) \cdot \nabla(\mathbf{F}_t), \end{cases} \quad (5)$$

where $\nabla(\cdot)'$ represents the masked gradients.

Nevertheless, the optimizer does not update parameters directly based on the raw gradients, i.e., $\Delta \mathbf{W}_{\theta,t} \neq \nabla(\mathbf{W}_{\theta,t})$, where $\Delta \mathbf{W}_{\theta,t}$ denotes the parameter update. For example, the Adam optimizer computes $\Delta \mathbf{W}_{\theta,t}$ based on the first and

second moments of past gradients. This can lead to excessively large or small updates when using the masked gradients, potentially hindering optimization. Therefore, we further adapt the proposed gradient masking to optimizers.

Specifically, we aim to scale the parameter updates so that the ratio between the masked and original updates matches the ratio of the corresponding gradients, formulated as:

$$\frac{\Delta \mathbf{W}'_{\theta,t}}{\Delta \mathbf{W}_{\theta,t}} = \frac{\nabla(\mathbf{W}_{\theta,t})'}{\nabla(\mathbf{W}_{\theta,t})}, \quad (6)$$

where $\Delta \mathbf{W}_{\theta,t}$ denotes the parameter update computed by the optimizer using the unmasked gradient $\nabla(\mathbf{W}_{\theta,t})$, while $\Delta \mathbf{W}'_{\theta,t}$ represents the expected parameter update. By rearranging the above equation, we derive:

$$\Delta \mathbf{W}'_{\theta,t} = \frac{\nabla(\mathbf{W}_{\theta,t})'}{\nabla(\mathbf{W}_{\theta,t})} \odot \Delta \mathbf{W}_{\theta,t}. \quad (7)$$

We then use $\Delta \mathbf{W}'_{\theta,t}$ to update the parameters:

$$\mathbf{W}_{\theta,t}^{<s>} = \mathbf{W}_{\theta,t}^{<s-1>} - \gamma \Delta \mathbf{W}'_{\theta,t}, \quad (8)$$

where $<s>$ denotes the s -th optimization step and γ is the learning rate. The same procedure can be similarly extended to all layers of the ViT architecture. Notably, this adaptation is optimizer-agnostic and can theoretically be integrated with any optimization algorithm. In our experiments, we employ the widely-used Adam optimizer, which empirically demonstrates satisfactory performance.

To implement Eq. (5), our method employs the mask $\bar{\mathbf{M}}_{t-1}$ constructed from the attention map \mathbf{U}_{t-1} of the previous task \mathcal{T}_{t-1} . The attention map should accurately capture the discriminative regions of the target object to ensure effective gradient modulation. To this end, we develop a tailored procedure for attention map extraction and mask generation, as detailed in the following subsection.

Adaptive Mask Generation

We propose a layer-wise rollout approach based on the rollout method (Abnar and Zuidema 2020) to extract attention maps, as well as an adaptive mask construction strategy. For the weight matrices $\mathbf{W}_{q,t}^l$, $\mathbf{W}_{k,t}^l$ and $\mathbf{W}_{v,t}^l$ in the l -th block ($l \in \{1, 2, \dots, L\}$), the computation of its associated mask $\bar{\mathbf{M}}_{t-1}^l$ is as follows.

We first consider a single sample from \mathcal{D}_{t-1} (omitting the sample identifier for brevity), which is fed into the model trained on \mathcal{T}_{t-1} for forward propagation. We extract the activated attention matrices from all layers, denoted as $\{\mathbf{S}_{t-1}^1, \mathbf{S}_{t-1}^2, \dots, \mathbf{S}_{t-1}^L\}$. The layer-wise rollout attention matrix for the l -th layer is computed as:

$$\hat{\mathbf{S}}_{t-1}^l = \tilde{\mathbf{S}}_{t-1}^1 \cdot \tilde{\mathbf{S}}_{t-1}^2 \cdot \dots \cdot \tilde{\mathbf{S}}_{t-1}^l, \quad (9)$$

where $\tilde{\mathbf{S}}_{t-1}^l = \mathbf{I} + \mathbf{S}_{t-1}^l$, and \mathbf{I} is the identity matrix to prevent self-suppression across layers (Abnar and Zuidema 2020). Based on $\hat{\mathbf{S}}_{t-1}^l \in \mathbb{R}^{(N+1)^2}$, we extract its class attention map $\mathbf{U}_{t-1}^l \in \mathbb{R}^{\sqrt{N^2}}$ by selecting the attention weights between the class token (query) and all image tokens (keys) and reshaping it into a matrix. Specifically, \mathbf{U}_{t-1}^l corresponds to the first row of $\hat{\mathbf{S}}_{t-1}^l$ excluding the first element,

Method	Venue	10S-ImageNet-R		20S-ImageNet-R		10S-CIFAR-100		10S-DomainNet	
		Acc.↑	Forgetting↓	Acc.↑	Forgetting↓	Acc.↑	Forgetting↓	Acc.↑	Forgetting↓
L2P	CVPR'22	61.57 \pm 0.66	9.73 \pm 0.47	59.38 \pm 0.50	5.89 \pm 0.36	83.83 \pm 0.04	7.63 \pm 0.30	81.17 \pm 0.83 [†]	8.98 \pm 1.25
DualPrompt	ECCV'22	68.13 \pm 0.49	4.68 \pm 0.20	63.21 \pm 0.49	5.28 \pm 0.45	86.51 \pm 0.33	5.16 \pm 0.09	81.70 \pm 0.78 [†]	8.04 \pm 0.31
CODA-Prompt	CVPR'23	75.45 \pm 0.56	1.64 \pm 0.10	72.37 \pm 1.19	0.96 \pm 0.15	86.25 \pm 0.74	1.67 \pm 0.26	80.04 \pm 0.79 [†]	10.16 \pm 0.35
APG	ICCV'23	73.27	8.59	71.22	7.39	89.35	6.01	-	-
ESN	AAAI'23	71.07 \pm 0.29	4.99 \pm 0.49	64.77 \pm 0.71	6.65 \pm 1.24	86.34 \pm 0.52	4.76 \pm 0.14	79.22 \pm 2.04 [‡]	10.62 \pm 2.12
OS-Prompt++	ECCV'24	75.67 \pm 0.40	1.27 \pm 0.10	73.77 \pm 0.19	0.79 \pm 0.07	86.68 \pm 0.67	1.18 \pm 0.21	-	-
EvoPrompt	AAAI'24	76.83 \pm 0.08	2.78 \pm 0.06	74.41 \pm 0.23	2.56 \pm 0.22	87.97 \pm 0.30	2.60 \pm 0.42	79.50 \pm 0.29 [‡]	3.81 \pm 0.36
PGP	ICLR'24	69.34 \pm 0.05	4.53 \pm 0.04	-	-	86.92 \pm 0.05	5.35 \pm 0.19	80.41 \pm 0.25 [‡]	8.39 \pm 0.18
OVOR-Deep	ICLR'24	76.11 \pm 0.21	7.16 \pm 0.34	73.85 \pm 0.29	6.80 \pm 0.65	85.99 \pm 0.89	6.42 \pm 0.03	79.61 \pm 0.86 [‡]	4.77 \pm 0.94
ConvPrompt	CVPR'24	77.86 \pm 0.25	4.33 \pm 0.24	75.10 \pm 0.39	4.10 \pm 0.29	88.87 \pm 0.33	4.75 \pm 0.15	79.47 \pm 0.35 [‡]	6.49 \pm 0.43
InfLoRA	CVPR'24	75.65 \pm 0.14	5.73 \pm 0.44	71.01 \pm 0.45	6.83 \pm 0.44	87.06 \pm 0.25	6.22 \pm 0.39	81.45 \pm 0.68 [‡]	5.35 \pm 0.52
EASE	CVPR'24	76.17	7.82	73.27	8.51	87.76	5.94	78.89 [‡]	7.89
CPrompt	CVPR'24	77.14 \pm 0.11	5.97 \pm 0.68	74.79 \pm 0.28	7.34 \pm 0.65	87.82 \pm 0.21	5.06 \pm 0.50	82.97 \pm 0.34	7.45 \pm 0.93
VPT-CPG	AAAI'25	<u>78.63</u> \pm 0.52	7.18 \pm 0.62	75.33 \pm 0.49 [*]	7.08 \pm 0.52	<u>90.63</u> \pm 0.44	3.98 \pm 0.65	<u>83.21</u> \pm 0.67	7.09 \pm 0.82
CAPrompt	AAAI'25	71.67 \pm 0.45 [*]	5.24 \pm 0.48	71.24 \pm 0.80 [*]	6.45 \pm 0.78	87.27 \pm 0.39 [*]	4.90 \pm 0.22	80.97 \pm 0.98 [*]	4.21 \pm 0.34
SD-LoRA	ICLR'25	77.34 \pm 0.35	-	75.26 \pm 0.37	-	88.01 \pm 0.31	-	81.01 \pm 0.42 [*]	-
BiLoRA	CVPR'25	77.95 \pm 0.14	-	72.41 \pm 0.76	-	87.46 \pm 0.76	-	76.56 \pm 0.41	-
LoRA-DRS	CVPR'25	75.94 \pm 0.46	-	73.78 \pm 0.44	-	89.14 \pm 0.23	-	71.19 \pm 0.55 [*]	-
CPrompt-KAC	CVPR'25	78.07	-	<u>75.73</u>	-	87.19	-	-	-
SEMA	CVPR'25	78.00	-	74.53	-	86.98	-	80.92 [*]	-
Seq-FT	Baseline	49.43 \pm 1.72	44.31 \pm 0.94	39.63 \pm 2.43	53.98 \pm 2.23	52.85 \pm 1.89	49.24 \pm 2.11	44.06 \pm 0.87	59.43 \pm 1.15
ARCL-ViT	This work	81.17 \pm 0.59	4.84 \pm 1.45	79.40 \pm 0.56	6.59 \pm 1.01	90.87 \pm 0.26	3.92 \pm 0.34	83.94 \pm 0.47	7.04 \pm 1.21

Table 1: Comparison with existing methods. The standard deviation (*i.e.*, the value after \pm) is reported when available. Missing results in the corresponding papers are denoted by “-”. Accuracies marked with \dagger , \ddagger and $*$ are reproduced by (Gao, Cen, and Chang 2024), (Lu et al. 2025b) and us, respectively. ‘Seq-FT’ represents the sequential fine-tuning baseline using Eq. (4). The best accuracy among all methods is highlighted in **bold**, and the best accuracy among existing methods is underlined.

and can be formulated as $\mathbf{U}_{t-1}^l = \hat{\mathbf{S}}_{t-1}^l[1, 2:]$. We then construct the mask based on \mathbf{U}_{t-1}^l .

Since the size of discriminative regions varies across images, using a fixed threshold or proportion for mask generation is suboptimal. To address this, we propose an adaptive approach to distinguish between attention and background regions. Due to the sharpening effect of the softmax function, the activated values in the attention regions are typically much higher than those in the background. As a result, sorting the elements of \mathbf{U}_{t-1}^l in ascending order yields a curve that resembles a sigmoid function. The optimal threshold index k^* is determined by locating the point with the lowest second-order derivative in the sorted attention curve:

$$k^* = \arg \min_{2 \leq k \leq N-1} (u_{k+1} - 2u_k + u_{k-1}), \quad (10)$$

where u_1, u_2, \dots, u_N denotes the sorted elements of \mathbf{U}_{t-1}^l , and the corresponding adaptive threshold is $\tau_{t-1}^l = u_{k^*}$. This point corresponds to the inflection point where the curve rises sharply from low values to the high-value plateau region. Next, all elements in \mathbf{U}_{t-1}^l greater than or equal to τ_{t-1}^l are classified as attention regions, and a binary mask \mathbf{M}_{t-1}^l is generated accordingly:

$$\mathbf{M}_{t-1,i}^l = \begin{cases} 0, & \text{if } \mathbf{U}_{t-1,i}^l \geq \tau_{t-1}^l \quad \forall i \in \{1, \dots, N\} \\ 1, & \text{otherwise} \end{cases} \quad (11)$$

To apply $\mathbf{M}_{t-1}^l \in \mathbb{R}^{\sqrt{N^2}}$ to $\nabla(\mathbf{A}_t^l) \in \mathbb{R}^{(N+1)^2}$ for gradient masking, we extend and align the mask as follows. First, we flatten \mathbf{M}_{t-1}^l into a row vector denoted as \mathbf{m}_{t-1}^l , and prepend a 0 to its left which aims to maintain the stability of the class token, resulting in $\bar{\mathbf{m}}_{t-1}^l \in \mathbb{R}^{N+1}$. Next, to enforce a consistent attention retention rule for all queries, we broadcast $\bar{\mathbf{m}}_{t-1}^l$ to all $N+1$ rows, yielding the mask matrix $\bar{\mathbf{M}}_{t-1}^l$ of shape $(N+1)^2$. This mask is used in Eq. (5) for element-wise multiplication with $\nabla(\mathbf{A}_t^l)$ and \mathbf{S}_t^l during gradient masking.

In practice, we further extend \mathbf{M}_{t-1}^l from each instance to all samples of a class and previous tasks by averaging, resulting in continuous masks within $[0, 1]$. For Transformers with multi-head attention, Eq. (5)~(11) can be applied independently to each head. More details and the full algorithm are provided in the Appendix.

Experiments

Experimental Setups

We evaluate our method on four class-incremental learning (CIL) benchmarks, including 10-split and 20-split ImageNet-R, 10-split CIFAR-100 and 10-split DomainNet. Note that the 10-split DomainNet benchmark was organized by (Wang et al. 2023) for challenging cross-domain CIL, consisting of the 200 classes with the most images from the

Index	Attention			Mask		10S-ImageNet-R		10S-DomainNet	
	<i>a.</i>	<i>b.</i>	<i>c.</i>	<i>d.</i>	<i>e.</i>	Acc.	Forg.	Acc.	Forg.
1	×	×	×	×	×	49.43	44.31	44.06	59.43
2	√	×	×	√	×	77.63	12.84	76.28	20.86
3	√	×	×	×	√	80.33	6.31	82.17	10.33
4	×	√	×	×	√	80.43	7.42	82.28	11.10
5	×	×	√	√	×	76.72	13.86	75.60	21.61
6	×	×	√	×	√	81.17	4.84	83.94	7.04

Table 2: Ablation study on different attention map extraction (*a.* raw attention, *b.* naive rollout, *c.* layer-wise rollout) and mask thresholding (*d.* fixed, *e.* adaptive) strategies.

Masked regions	10S-ImageNet-R		10S-DomainNet	
	Acc.	Forgetting	Acc.	Forgetting
Random	73.27	4.22	79.52	6.70
Non-attention	69.55	7.56	76.08	9.47
Attention	81.17	4.84	83.94	7.04

Table 3: Ablation study on different masked regions for attention maps.

original DomainNet (Peng et al. 2019).

We use ViT-B/16 pre-trained on ImageNet-21k (Rusakovsky et al. 2015) as the backbone by default in our experiments. Each task employs an independent classifier. During training, the $\mathbf{W}_{\{q,k,v\}}$ layers in all 12 ViT blocks as well as the classifier are fine-tuned. We employ the Adam optimizer with a learning rate of 0.0001 for the $\mathbf{W}_{\{q,k,v\}}$ layers and 0.01 for the classifier. During inference, all classifiers are concatenated, and the backbone from the final task is used for prediction. We report the final average accuracy and final average forgetting, each computed over three runs with different random seeds. Additional experimental details are provided in the Appendix. Our code is available at <https://github.com/zugexiaodui/AttentionRetentionCL>.

Comparison with Existing Methods

We compare the proposed ARCL-ViT with existing methods on the four benchmarks in Table 1. The competitors include recent state-of-the-art works such as VPT-CPG (Lu et al. 2025b), SD-LoRA (Wu et al. 2025), BiLoRA (Zhu et al. 2025) and KAC (Hu et al. 2025). The Seq-FT baseline refers to sequentially fine-tuning the $\mathbf{W}_{\{q,k,v\}}$ layers and the classifier without any anti-forgetting technique. The only difference between Seq-FT and our method is that Seq-FT updates parameters using Eq. (4) while our method applies Eq. (5)~(8). The direct comparison to Seq-FT validates the core effectiveness of our approach, which achieves 37% higher accuracy and 46% less forgetting by removing attention drift from previous tasks. Furthermore, our method surpasses 15 recent approaches from 2024 and 2025 and achieves state-of-the-art performance, with up to 3.7% and an average of 1.8% higher accuracy across the benchmarks.

Weights	Method	10S-ImageNet-R		10S-DomainNet	
		Acc.	Forgetting	Acc.	Forgetting
DINO-1k	L2P	56.71	4.60	50.72	16.89
	CODA-P	64.02	7.70	66.74	8.91
	EASE	65.97	9.24	70.23	9.64
	CPrompt	64.66	9.26	75.31	9.93
	InfLoRA	68.31	8.81	75.05	7.70
	VPT-CPG	<u>68.32</u>	2.78	<u>77.46</u>	10.18
	Seq-FT	39.93	49.27	35.67	67.98
ARCL-ViT	72.37	4.91	79.39	7.90	
iBOT-1k	L2P	60.80	6.65	53.32	17.12
	CODA-P	68.02	6.23	65.80	12.35
	EASE	69.91	9.22	72.39	10.25
	CPrompt	69.14	8.82	77.45	9.28
	InfLoRA	71.84	9.05	76.60	8.51
	VPT-CPG	<u>72.16</u>	5.80	<u>77.79</u>	13.02
	Seq-FT	44.48	47.96	37.85	66.06
ARCL-ViT	76.48	7.68	80.88	9.85	

Table 4: Results of using DINO and iBOT weights pre-trained on ImageNet-1k (*i.e.*, DINO-1k and iBOT-1k).

Ablation Study

Construction of Attention Masks We investigate two factors in the construction of attention masks: 1) the method for extracting attention maps, and 2) the strategy for determining the mask threshold. Specifically, three attention map extraction methods are considered: *a*) Raw attention in each layer, *i.e.*, replacing Eq. (9) with $\hat{\mathbf{S}}_{t-1}^l = \mathbf{S}_{t-1}^l$. *b*) Naive rollout, which simply accumulates the attention maps across all layers, *i.e.*, $\hat{\mathbf{S}}_{t-1}^l = \tilde{\mathbf{S}}_{t-1}^1 \cdot \dots \cdot \tilde{\mathbf{S}}_{t-1}^L$. *c*) Layer-wise rollout using Eq. (9). The thresholding strategy include two alternatives: *d*) Fixed threshold, where the top 20% (determined by hyper-parameter search) positions in attention maps are set to 0. *e*) Adaptive threshold described in Eq. (10).

The results of different configurations are presented in Table 2. Several key observations can be drawn: 1) Experiment 2, serving as a simplified variant of our approach, demonstrates the effectiveness of our core attention preservation concept for CL. By employing basic raw-attention with fixed thresholding, this configuration achieves a substantial 30% accuracy improvement and 35% reduction in forgetting compared to the Seq-FT baseline (Index 1). 2) Layer-wise rollout attention exhibits superior compatibility with our framework. Comparing experiments 3, 4, and 6, layer-wise rollout achieves an average accuracy improvement of 1.3% over both raw attention and naive rollout methods. 3) Comparison between experiments 2 vs. 3 and 5 vs. 6 shows that the adaptive thresholding strategy yields an average accuracy gain of 5.4% over the fixed threshold. This result highlights the importance of adopting adaptive thresholding, which removes the need for manual tuning and enables more accurate identification of discriminative regions. When layer-wise rollout is combined with adaptive thresholding, our method achieves the best performance.

Method	Venue	50S-ImageNet-R		50S-DomainNet		100S-ImageNet-R		100S-DomainNet	
		Acc.	Forgetting	Acc.	Forgetting	Acc.	Forgetting	Acc.	Forgetting
L2P	CVPR'22	51.38	12.34	63.13	11.19	41.51	14.48	54.83	14.95
OVOR-Deep	ICLR'24	63.25	5.23	68.29	6.85	43.02	7.30	52.09	8.66
ConvPrompt	CVPR'24	64.61	7.12	71.76	6.37	44.32	8.97	56.21	6.27
InfLoRA	CVPR'24	62.81	10.37	71.87	9.20	42.23	14.17	48.06	15.43
EASE	CVPR'24	70.27	6.73	65.34	9.64	51.56	7.56	37.26	29.12
CPrompt	CVPR'24	70.75	7.44	70.74	9.01	59.90	9.52	57.60	9.42
VPT-CPG	AAAI'25	<u>73.08</u>	12.12	<u>72.27</u>	19.93	<u>64.63</u>	12.18	<u>60.81</u>	27.43
Seq-FT	Baseline	22.57	70.01	9.89	90.34	10.65	76.04	8.87	88.65
ARCL-ViT	This work	76.45	8.20	74.32	15.95	70.72	12.90	62.35	25.63

Table 5: Performance on long-sequence continual learning experiments using 50 and 100 splits of ImageNet-R and DomainNet.

Different Masked Regions To validate the necessity of masking *attention* regions specifically (*i.e.*, Eq. (11)), we conduct two comparative experiments: 1) Random masking. We randomly select positions to be set to zero in the mask, with the number of masked pixels matching that of the attention pixels identified by Eq. (11). 2) Non-attention Masking. We set non-attention regions to 0 by modifying the condition in Eq. (11) to $\mathbf{U}_{t-1,i}^l < \tau_{t-1}^l$. As shown in Table 3, masking the non-attention regions yields the poorest performance and leads to severe forgetting. Interestingly, it still performs better than Seq-FT, which is because the non-attention regions for one class may overlap with the attention regions of another class, thereby indirectly providing partial retention to relevant attention areas. Our proposed approach, which removes the attention regions of previous tasks, substantially outperforms the random masking baseline. This systematic comparison confirms the necessity of masking attention regions to effectively reduce attention drift for CL.

Generalization Experiments

Different Pre-training Weights To validate the generalizability of our method across different pre-training weights, we follow InfLoRA (Liang and Li 2024) and conduct experiments using ViT-B/16 initialized with self-supervised DINO (Caron et al. 2021) and iBOT (Zhou et al. 2022) weights pre-trained on ImageNet-1k (denoted as DINO-1k and iBOT-1k, respectively). For comparison, we also reproduce 6 representative methods when official results are unavailable. As presented in Table 4, ARCL-ViT exhibits significant improvements over the Seq-FT baseline, achieving an average accuracy gain of 37.8% while reducing forgetting by 50.2%. It surpasses existing approaches by 2.9% in accuracy with DINO-1k weights and by 3.7% with iBOT-1k weights. These results verify that our approach effectively generalizes to diverse pre-training weights. Our method also generalizes well across different ViT architectures, with experimental results provided in the Appendix.

Long-Sequence Continual Learning We further evaluate the generalizability of our method in long-sequence CL scenarios. In this setting, the ImageNet-R and DomainNet datasets, each containing 200 classes, are partitioned into 50 and 100 tasks. As shown in Table 5, which compares our

results with seven competitors (implementations from VPT-CPG (Lu et al. 2025b)), ARCL-ViT outperforms the previous state-of-the-art approach VPT-CPG by 3.3% in accuracy. Although our method exhibits comparable forgetting, this result indicates that it achieves an effective trade-off between plasticity and stability. This enables our model to continually acquire new knowledge through controlled forgetting. Consequently, our approach attains superior overall accuracy in the challenging long-sequence CL scenarios.

Analysis of Attention Retention

We visualize the attention maps for samples from \mathcal{T}_1 of 10-split ImageNet-R and DoaminNet. As shown in Figure 1(a), our method effectively preserves the original attention on key visual concepts. We further make a quantitative comparison of attention drift across different methods, as illustrated in Figure 1(b). The attention drift value at the t -th task ($t > 2$) is calculated as follows: For each sample in \mathcal{D}_1 , attention maps are extracted using both the \mathcal{T}_1 and \mathcal{T}_t models. The Frobenius norm of their difference is calculated and then normalized by the Frobenius norm of the \mathcal{T}_1 attention map to serve as the attention drift value. Our method achieves the lowest attention drift at the end of training (6.0%), which is significantly lower than that of Seq-FT (25.9%) and other competitors. These results provide quantitative evidence that our method retains attention to previously learned visual concepts, thereby mitigating catastrophic forgetting.

Conclusion

In this work, we introduce a novel attention-retaining framework to address the challenge of catastrophic forgetting in continual learning for Vision Transformers. Concretely, we propose a gradient masking method to preserve selective attention to previously learned visual concepts. Our method employs a layer-wise rollout and adaptive thresholding strategy for attention mask generation, and incorporates optimizer-compatible updates to ensure effective training. Extensive experimental results on multiple continual learning scenarios demonstrate both the superiority and generalizability of our approach. In future work, we will extend this framework to other Transformer-based architectures and explore the integration of additional biological principles.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (NSFC) under Grant 62576282, 62476223, 62376218; in part by the National Key Research and Development Program of China under Grant 2024YFF1306501; in part by Innovation Capability Support Program of Shaanxi (Program No. 2024ZC-KJXX-043); in part by Natural Science Basic Research Program of Shaanxi Province (2024JC-DXWT-07).

References

- Abnar, S.; and Zuidema, W. 2020. Quantifying Attention Flow in Transformers. In *Annual Meeting of the Association for Computational Linguistics*, 4190–4197.
- Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging Properties in Self-Supervised Vision Transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9630–9640.
- Chaudhry, A.; Khan, N.; Dokania, P. K.; and Torr, P. H. S. 2020. Continual Learning in Low-rank Orthogonal Subspaces. In *Proceedings of the Conference on Neural Information Processing Systems*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houshy, N. 2021. An Image Is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *Proceedings of the International Conference on Learning Representations*.
- Gao, Z.; Cen, J.; and Chang, X. 2024. Consistent Prompting for Rehearsal-Free Continual Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 28463–28473.
- Hu, Y.; Liang, Z.; Yang, F.; Hou, Q.; Liu, X.; and Cheng, M.-M. 2025. KAC: Kolmogorov-Arnold Classifier for Continual Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N.; Veness, J.; Desjardins, G.; Rusu, A. A.; Milan, K.; Quan, J.; Ramalho, T.; Grabska-Barwinska, A.; Hassabis, D.; Clopath, C.; Kumaran, D.; and Hadsell, R. 2017. Overcoming Catastrophic Forgetting in Neural Networks. *Proceedings of the National Academy of Sciences*, 114(13): 3521–3526.
- Kong, Y.; Liu, L.; Wang, Z.; and Tao, D. 2022. Balancing Stability and Plasticity Through Advanced Null Space in Continual Learning. In *Proceedings of the European Conference on Computer Vision*, volume 13686, 219–236.
- Liang, Y.-S.; and Li, W.-J. 2024. InfLoRA: Interference-Free Low-Rank Adaptation for Continual Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23638–23647.
- Lu, Y.; Tan, J.; Zhang, S.; Xing, Y.; Liang, G.; and Zhang, Y. 2025a. Nearest-Neighbor Class Prototype Prompt and Simulated Logits for Continual Learning. *Pattern Recognition*, 169: 111933.
- Lu, Y.; Zhang, S.; Cheng, D.; Liang, G.; Xing, Y.; Wang, N.; and Zhang, Y. 2025b. Training Consistent Mixture-of-Experts-Based Prompt Generator for Continual Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Lu, Y.; Zhang, S.; Cheng, D.; Xing, Y.; Wang, N.; Wang, P.; and Zhang, Y. 2024. Visual Prompt Tuning in Null Space for Continual Learning. In *Proceedings of the Conference on Neural Information Processing Systems*.
- Peng, X.; Bai, Q.; Xia, X.; Huang, Z.; Saenko, K.; and Wang, B. 2019. Moment Matching for Multi-Source Domain Adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1406–1415.
- Qiao, J.; Zhang, Z.; Tan, X.; Chen, C.; Qu, Y.; Peng, Y.; and Xie, Y. 2024. Prompt Gradient Projection for Continual Learning. In *Proceedings of the International Conference on Learning Representations*.
- Rebuffi, S.-A.; Kolesnikov, A.; Sperl, G.; and Lampert, C. H. 2017. iCaRL: Incremental Classifier and Representation Learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2001–2010.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M. S.; Berg, A. C.; and Fei-Fei, L. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, 115(3): 211–252.
- Saha, G.; Garg, I.; and Roy, K. 2021. Gradient Projection Memory for Continual Learning. In *Proceedings of the International Conference on Learning Representations*.
- 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.
- Wang, S.; Li, X.; Sun, J.; and Xu, Z. 2021. Training Networks in Null Space of Feature Covariance for Continual Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 184–193.
- Wang, Y.; Ma, Z.; Huang, Z.; Wang, Y.; Su, Z.; and Hong, X. 2023. Isolation and Impartial Aggregation: A Paradigm of Incremental Learning without Interference. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 10209–10217.
- Wang, Z.; Zhang, Z.; Ebrahimi, S.; Sun, R.; Zhang, H.; Lee, C.-Y.; Ren, X.; Su, G.; Perot, V.; Dy, J. G.; and Pfister, T. 2022a. DualPrompt: Complementary Prompting for Rehearsal-Free Continual Learning. In *Proceedings of the European Conference on Computer Vision*, volume 13686, 631–648.
- 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.

- Wu, C.; Herranz, L.; Liu, X.; Van De Weijer, J.; Raducanu, B.; et al. 2018. Memory Replay GANs: Learning to Generate New Categories without Forgetting. *Proceedings of the Conference on Neural Information Processing Systems*, 31.
- Wu, Y.; Piao, H.; Huang, L.-K.; Wang, R.; Li, W.; Pfister, H.; Meng, D.; Ma, K.; and Wei, Y. 2025. SD-LoRA: Scalable Decoupled Low-Rank Adaptation for Class Incremental Learning. *Proceedings of the International Conference on Learning Representations*.
- Xiang, Y.; Fu, Y.; Ji, P.; and Huang, H. 2019. Incremental Learning Using Conditional Adversarial Networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6618–6627.
- Yan, S.; Xie, J.; and He, X. 2021. DER: Dynamically Expandable Representation for Class Incremental Learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3014–3023.
- Zeng, G.; Chen, Y.; Cui, B.; and Yu, S. 2019. Continual Learning of Context-Dependent Processing in Neural Networks. *Nature Machine Intelligence*, 1(8): 364–372.
- Zenke, F.; Poole, B.; and Ganguli, S. 2017. Continual Learning Through Synaptic Intelligence. In *Proceedings of the International Conference on Machine Learning*, volume 70, 3987–3995.
- Zhang, X.; Zhaoping, L.; Zhou, T.; and Fang, F. 2012. Neural Activities in V1 Create a Bottom-Up Saliency Map. *Neuron*, 73(1): 183–192.
- Zhou, J.; Wei, C.; Wang, H.; Shen, W.; Xie, C.; Yuille, A. L.; and Kong, T. 2022. Image BERT Pre-training with Online Tokenizer. In *Proceedings of the International Conference on Learning Representations*.
- Zhu, H.; Zhang, Y.; Dong, J.; and Koniusz, P. 2025. BiLoRA: Almost-Orthogonal Parameter Spaces for Continual Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.