

PEARL: Input-Agnostic Prompt Enhancement with Negative Feedback Regulation for Class-Incremental Learning

Yongchun Qin^{1,2}, Pengfei Fang^{1,2*}, Hui Xue^{1,2*}

¹School of Computer Science and Engineering, Southeast University, Nanjing 210096, China
²Key Laboratory of New Generation Artificial Intelligence Technology and Its Interdisciplinary Applications (Southeast University), Ministry of Education, China
 {ycqin, fangpengfei, hxue}@seu.edu.cn

Abstract

Class-incremental learning (CIL) aims to continuously introduce novel categories into a classification system without forgetting previously learned ones, thus adapting to evolving data distributions. Researchers are currently focusing on leveraging the rich semantic information of pre-trained models (PTMs) in CIL tasks. Prompt learning has been adopted in CIL for its ability to adjust data distribution to better align with pre-trained knowledge. This paper critically examines the limitations of existing methods from the perspective of prompt learning, which heavily rely on input information. To address this issue, we propose a novel PTM-based CIL method called Input-Agnostic Prompt Enhancement with Negative Feedback Regulation (PEARL). In PEARL, we implement an input-agnostic global prompt coupled with an adaptive momentum update strategy to reduce the model’s dependency on data distribution, thereby effectively mitigating catastrophic forgetting. Guided by negative feedback regulation, this adaptive momentum update addresses the parameter sensitivity inherent in fixed-weight momentum updates. Furthermore, it fosters the continuous enhancement of the prompt for new tasks by harnessing correlations between different tasks in CIL. Experiments on six benchmarks demonstrate that our method achieves state-of-the-art performance.

Code — <https://github.com/qinyongchun/PEARL>

Introduction

In the fields of computer vision and machine learning, Class-Incremental Learning (CIL) (Rebuffi et al. 2017a; Castro et al. 2018a; Hou et al. 2019a) has become a pivotal paradigm, designed to enable models to acquire new tasks over time without forgetting previously learned information. This approach differs from traditional batch learning, which processes the entire dataset in one go; instead, CIL gradually introduces new categories, allowing the model to adjust to changing data distributions and real-world conditions where new categories may appear dynamically. Integrating Pre-Trained Models (PTMs) into CIL capitalizes on their extensive feature representations, which can speed up convergence and enhance initial performance (Wang et al. 2022b,a). However, in the face of the dynamic challenges

*Corresponding author.
 Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

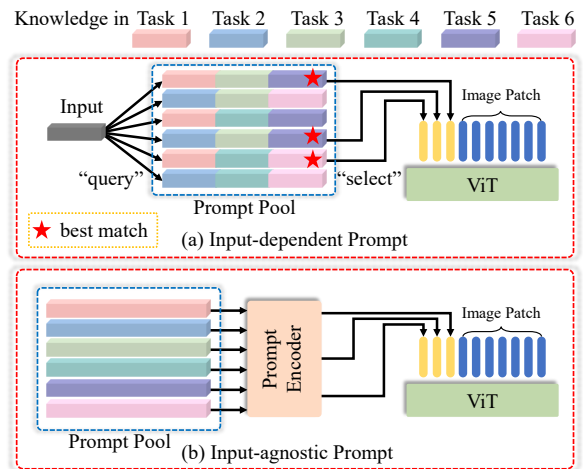


Figure 1: The comparison of (a) input-dependent prompt and (b) input-independent prompt. L2P and DualPrompt following the paradigm in (a), select the best matched prompts, while CODA-Prompt assemble the prompt pool with learnable components during the “select” phase.

posed by real-world environments, merely fine-tuning PTMs proves inadequate. Fine-tuning on new classes may lead to catastrophic forgetting, a phenomenon where a model’s performance on old classes deteriorates significantly as new information overwrites existing representations (McCloskey and Cohen 1989).

To address the challenges inherent in incremental learning, prompt learning has been introduced as a groundbreaking approach (Wang et al. 2022b,a; Wang, Huang, and Hong 2022; Smith et al. 2023). Prompts serve as task-specific instructions or contextual guides that help the model process input data (Li and Liang 2021; Lester, Al-Rfou, and Constant 2021). In the realm of Class-Incremental Learning (CIL), prompts are dynamically evolved to accommodate new tasks. To prevent the erasure of previous knowledge by subsequent tasks, current strategies typically involve selective sampling from the prompt pool based on input data (Wang et al. 2022b,a; Smith et al. 2023; Gao, Cen, and Chang 2024). This approach ensures that only a subset of the prompt pool is updated at any time, thus reducing the

risk of catastrophic forgetting. This is described as “prompt-selection” by Gao *et al.* (Gao, Cen, and Chang 2024), which we further summarize as a “query-select” mechanism. For example, L2P (Wang et al. 2022b) identifies optimal prompts by evaluating their similarity to the input and the prompt pool. DualPrompt (Wang et al. 2022a) categorizes prompts into “expert” and “general” types and applies them using prefix tuning. Meanwhile, Coda-Prompt (Smith et al. 2023) overcomes the non-differentiability challenges seen in L2P and DualPrompt by linearly combining the entire prompt pool during the selection phase to generate a consistent-length prompt.

A recent review in (Zhou et al. 2024a) reveals that prompt-based methods generally underperform compared to other approaches. Conversely, Jia *et al.* demonstrate that prompt learning is effective on comparable datasets within a supervised learning context (Jia et al. 2022). This indicates that although fixed-length prompts possess adequate expressive potential, the “query-select” mechanism fails to fully harness this capability. In this paper, we propose a novel theory to elucidate the observed phenomenon and offer insights into surmounting the limitations of existing prompt-based methods. We introduce the concept of a “knowledge container” to detail the shortcomings of these approaches. Each prompt in the pool serves as a knowledge container, accumulating insights from incoming tasks. As illustrated in Fig. 1, the “query-select” mechanism frequently amalgamates knowledge from various tasks in a disorganized fashion. This haphazard mixing results in inconsistent knowledge representations and challenges in preserving coherent, task-specific information. The lack of a systematic method to manage and safeguard task-specific knowledge significantly detracts from the effectiveness of existing approaches.

To overcome the limitations of the “query-select” mechanism, we propose the creation of an input-agnostic prompt suitable for all instances within a single session. Considering the temporal dynamics of incremental learning, we conceptualize the prompt adaptation process as a sequential problem, facilitating steady incremental learning by progressively capturing task correlations. We refer to this approach as **Sequential Prompt Adaptation (SPA)**. During each session, a session-sharing prompt encoder processes the prompt pool to generate a global prompt. Following (Smith et al. 2023), we create independent knowledge containers for each task by freezing specific parameters in the prompt pool, while ensuring uniform knowledge representation through the prompt encoder. Further, we introduce a segmented positional encoding to maintain consistency within each segment of the prompt pool.

By eliminating the “query-select” mechanism and introducing a session-sharing prompt encoder, our model risks overfitting to the current task. To counteract this, we introduce the **Negative-feedback Knowledge Accumulation (NKA)** mechanism. This approach updates the prompt using a momentum-based method, where the momentum weight is influenced by the model’s output, which in turn is affected by the prompt itself. The weight is adjusted based on the divergence between current and previous outputs. Low divergence, indicating well-retained old knowledge, allows

the prompt to integrate more new knowledge (*i.e.*, reduce weight). Conversely, high divergence, showing poor retention of old knowledge, requires increased focus on preserving previous parameters (*i.e.*, increase weight). Essentially, this mechanism ensures knowledge retention as a prerequisite for learning new tasks. Additionally, the NKA mechanism helps to reveal potential correlations between tasks, thereby enhancing knowledge accumulation.

Our contributions include:

- We propose a novel CIL framework called **PEARL**, where a prompt encoder generates uniform prompts infused with global knowledge and accumulates knowledge through a momentum-based update strategy driven by negative feedback regulation.
- We introduce the **SPA** module, which enables a global prompt to simultaneously encapsulate knowledge from different tasks, overcoming the shortcomings of current “query-select” mechanism.
- The proposed **NKA** mechanism effectively implements an adaptive momentum update, achieving efficient knowledge accumulation by leveraging inherent data correlations.
- Through extensive experiments, we demonstrate that our method achieves state-of-the-art performance, surpassing the second-best results by an average of 2.23% in accuracy across six benchmarks.

Related Work

Class-Incremental Learning

CIL is one of the research hotspots in machine learning. Its main challenge is “catastrophic forgetting”, which occurs when the model overfits the current task and loses knowledge from previous tasks. According to (Masana et al. 2022), there are three main technical approaches for existing CIL researches. Rehearsal-based approaches reduce forgetting by either retaining a limited set of representative samples or generating pseudo-samples (Rebuffi et al. 2017b; Shin et al. 2017; Xiang et al. 2019; Ostapenko et al. 2019). Methods based on regularization consider to impose constraints on the representation or weight of the model (Jung et al. 2016; Kirkpatrick et al. 2017; Li and Hoiem 2017; Aljundi et al. 2018; Chaudhry et al. 2018), and usually use knowledge distillation technology (Hinton, Vinyals, and Dean 2015) to enhance the memory ability of the model. The bias-correction approach aims to solve the domain shift problem by aligning feature distribution between different tasks to alleviate overfitting when the model is faced with new tasks (Castro et al. 2018b; Hou et al. 2019b; Wu et al. 2019).

In recent years, with the rise of PTMs, many researches focus on PTM-based CIL. According to the recent review research (Zhou et al. 2024a), the existing methods can be divided into three categories. The prompt-based methods focus on prompt learning in CIL (Wang et al. 2022b,a; Smith et al. 2023; Huang, Chen, and Hsu 2024; Kim et al. 2025). Leveraging the robust representational capabilities of PTMs, researchers have demonstrated that effective continuous learning on downstream tasks can be achieved by

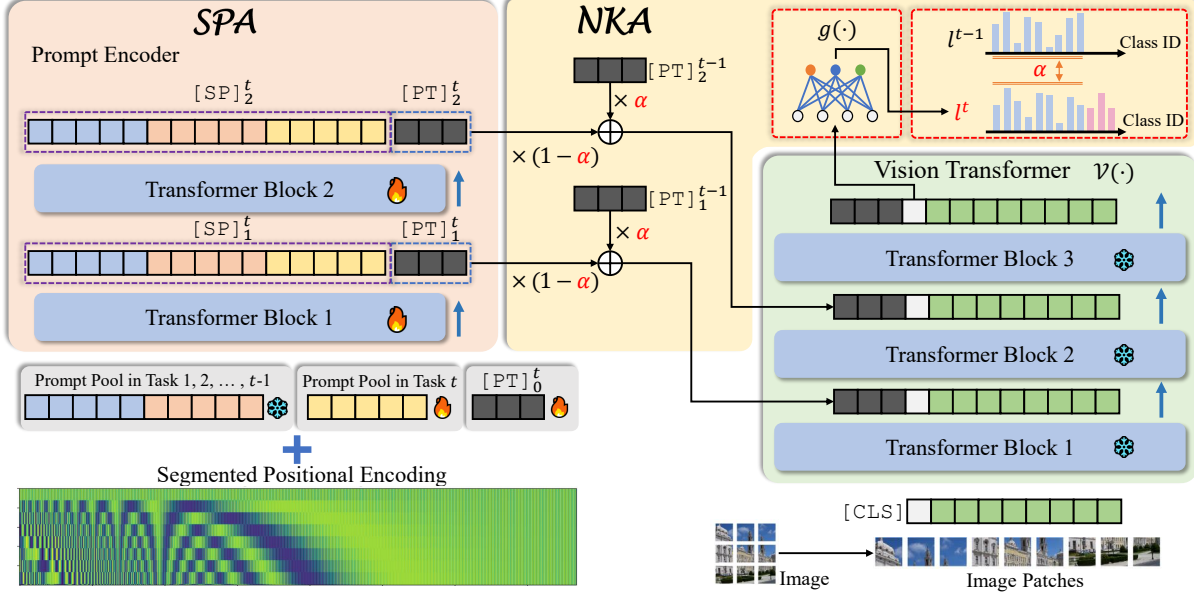


Figure 2: The illustration of the proposed PEARL. The ViT consists of 12 blocks, with a L -layer prompt encoder added to the last few blocks of the ViT. The ViT is frozen during training, while the prompt encoder and prompt pool remain learnable.

fine-tuning a prompt. Similarly, the representation-based approach focuses on rich representations of PTMs and improve their generalization in CIL through regularization and metric learning (Zhou et al. 2023; McDonnell et al. 2024; Zhou et al. 2024b). In addition, the approach based on model-mixture has also received attention (Wang et al. 2024, 2023). The key of model-mixture is to assign corresponding expert models to different tasks through ensemble learning, so as to improve the overall capability of the model. Zhou *et al.* (Zhou et al. 2024a), through comprehensive comparison, find that the prompt-based methods underperform compared to the other two mainstream approaches. Through in-depth analysis, we propose a “knowledge container” theory to explain this phenomenon, and designs a novel prompt method based on this theory, which reaches the state-of-the-art result.

Prompt Learning for Pre-Trained Models

With the success of PTMs in the field of natural language processing (NLP), similar techniques have been introduced to computer vision (CV) tasks. Prompt Tuning (Lester, Al-Rfou, and Constant 2021; Zhang et al. 2024) is a method of tuning models by adding learnable prompt tokens before input data. Prompt Tuning helps models perform better by aligning distribution of downstream tasks and pre-trained data. Prefix Tuning (Li and Liang 2021) establishes a flexible attention mechanism, where learnable prompts are appended to the attention parameters. Additionally, researchers have also developed several other fine-tuning methods (*e.g.* BitFit (Zaken, Ravfogel, and Goldberg 2021) and LoRA (Hu et al. 2021)) to further improve the efficiency of tuning.

Methodology

Problem Formulation

In CIL, training data appears in the form of data stream and each session in the stream contains a task. A data stream with \mathcal{N} sessions can be referred to as: $\{\mathcal{D}^1, \dots, \mathcal{D}^{\mathcal{N}}\}$, where $\mathcal{D}^t = \{(x_i, y_i)\}_{0 < t \leq \mathcal{N}}$ is the training set for the t -th session and y_i belongs to the class set \mathcal{C}^t . Each task contains the same number of categories, *i.e.* $|\mathcal{C}^t| = K$, for $t = 1, 2, \dots, \mathcal{N}$. The settings of CIL require that datasets at different sessions cannot share class labels, *i.e.* $\mathcal{C}^i \cap \mathcal{C}^j = \emptyset$, for any $i \neq j$. At the same time, the model is required to retain the memory of all previous tasks during testing, so the test set at the t -th session needs to contain all previous labels, *i.e.* $\mathcal{C}_{\text{test}}^t = \mathcal{C}^0 \cup \mathcal{C}^1 \dots \cup \mathcal{C}^t$.

Overview on PEARL

Recognizing the limitations of existing “query-select” methods, we propose the **PEARL** to build an input-agnostic prompt. PEARL consists of two components: the **SPA** module and the **NKA** mechanism. In SPA, we model the prompt learning process as a sequential problem, leveraging the inherent temporal dynamics of incremental learning. Following CODA-Prompt (Smith et al. 2023), only a proportion of the prompt pool is updated in the corresponding session to facilitate explicit knowledge management. The prompt encoder \mathbf{E} and the prompt pool \mathcal{P} are defined as follows:

$$\begin{aligned} \mathbf{E} &= \{\mathbf{B}_1, \mathbf{B}_2, \dots, \mathbf{B}_L\}, \\ \mathcal{P} &= \{\mathbf{P}_1, \mathbf{P}_2, \dots, \mathbf{P}_{\mathcal{M}}\}, \end{aligned} \quad (1)$$

where, \mathbf{B}_i represent the i -th block in the prompt encoder, L represents the depth of the prompt encoder and \mathcal{P} contains learnable prompts with the number of \mathcal{M} .

We choose ViT (Dosovitskiy et al. 2020) as the backbone model $\mathcal{V}(\cdot) \in \mathbb{R}^d$ and implement a classification head $g(\cdot)$ which contains no trainable parameters and updated following the RanPAC manner (McDonnell et al. 2024). Inspired by the [CLS] token in ViT, we design a prompt token [PT] to obtain a fixed-length prompt after aggregating the prompt pool. During session t , the model can derive a prompt token $[\text{PT}]_i^t \in \mathbb{R}^{H \times d}$ with length H from the i -th block of $\mathbf{E}(\cdot)$, generating diverse representations and forming a set of prompts: $\{[\text{PT}]_i^t\}_{1 \leq i \leq L} \in \mathbb{R}^{L \times H \times d}$. For a certain instance \mathbf{x} , The prediction logit l^t is computed by:

$$l^t = g \circ \mathcal{V}(\mathbf{x}, \{[\text{PT}]_i^{\text{mem}}\}), \quad (2)$$

where $[\text{PT}]_i^{\text{mem}}$ will be defined by Eq. (3).

From the perspective of knowledge containers, $[\text{PT}]_i^t$ aggregates previous knowledge from the prompt pool. However, this knowledge serves primarily as a good initialization and is prone to overfitting. The proposed NKA mechanism address this problem by introducing a momentum update strategy:

$$[\text{PT}]_i^{\text{mem}} = \alpha^\tau \cdot [\text{PT}]_i^{t-1} + (1 - \alpha^\tau) \cdot [\text{PT}]_i^t, \quad (3)$$

where the momentum prompt $[\text{PT}]_i^{\text{mem}}$ will be sent into the backbone and α^τ represents the momentum weight which is obtained through the negative feedback regulation. By mixing knowledge from different sessions, the NKA mechanism ensures the stability of old knowledge while also acquiring new knowledge.

The backbone $\mathcal{V}(\cdot)$ remains frozen during incremental learning, and the primary objective is to identify the optimal prompt encoder and prompt pool:

$$\mathbf{E}^*, \mathcal{P}^* = \arg \max_{\mathbf{E}, \mathcal{P}} \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}^t} \mathbb{I}(y \neq l^t). \quad (4)$$

The pseudo code is provided in supplementary materials ¹.

Sequential Prompt Adaptation

The primary challenge in constructing the input-agnostic prompt is establishing cross-task information interaction. This is because CIL needs to be backward-compatible: new tasks must build upon previous knowledge rather than requiring a complete rebuild.

Existing prompt-based methods (Wang et al. 2022b,a; Smith et al. 2023) achieve this interaction through a ‘‘prompt-input-prompt’’ link. However, the input-agnostic prompt cannot establish such a link. We propose to omit the intermediate link and enable direct interaction between prompts in the form of a sequence. We utilize the Transformer architecture (Vaswani et al. 2017) to capture the sequential relationship. The prompt encoder, defined in Eq. (1), consists of L blocks, with the output function after the i -th block denoted as:

$$\mathcal{E}_i(\cdot) = \mathbf{B}_1 \circ \mathbf{B}_2 \circ \dots \circ \mathbf{B}_i(\cdot), \quad (5)$$

and the subset of the prompt pool is denoted as:

$$\mathcal{P}[1 : k] = \{\mathbf{P}_1, \mathbf{P}_2, \dots, \mathbf{P}_k\}. \quad (6)$$

During the t -th session, only $\mathcal{P}[1 + \frac{\mathcal{M}}{N} \times (t-1) : \frac{\mathcal{M}}{N} \times t]$ are learnable while other prompts are frozen to keep previous knowledge. The process of prompt encoding can be formulated as below:

$$\begin{aligned} [\text{PT}]_1^t, [\text{SP}]_1^t &= \mathcal{E}_1(\text{ConCat}([\text{PT}]_0^t, \mathcal{P}[1 : \frac{\mathcal{M}}{N} \times t])), \\ [\text{PT}]_2^t, [\text{SP}]_2^t &= \mathcal{E}_2(\text{ConCat}([\text{PT}]_1^t, [\text{SP}]_1^t)), \\ &\vdots \\ [\text{PT}]_i^t, [\text{SP}]_i^t &= \mathcal{E}_i(\text{ConCat}([\text{PT}]_{i-1}^t, [\text{SP}]_{i-1}^t)), \end{aligned} \quad (7)$$

where $[\text{SP}]_i^t$ is short for sequential prompts after the i -th block and is the intermediate variable during encoding. The data flow can be seen in Fig. 2. We adopt the prefix-tuning manner (Li and Liang 2021), and the $[\text{PT}]_i^t$ can be further embedded as learnable prefixes: $\mathbf{p}_K, \mathbf{p}_V \in \mathbb{R}^{L \times H \times d}$. The learnable prefixes are attached in the Multi-head Self-attention (MSA):

$$f_{\text{prefix}} = \text{MSA}(\mathbf{h}_Q, \text{ConCat}(\mathbf{p}_K, \mathbf{h}_K), \text{ConCat}(\mathbf{p}_V, \mathbf{h}_V)), \quad (8)$$

where $\mathbf{h}_Q, \mathbf{h}_K, \mathbf{h}_V$ are attention parameters.

Instead of encoding each position individually, SPA encodes the input sequence according to the task number. To achieve this, we introduce a segmented positional encoding (SPE):

$$\begin{aligned} \text{SPE}_{(\text{pos}, 2j)} &= \sin\left(\frac{\lfloor \text{pos} / \frac{\mathcal{M}}{N} \rfloor}{10000^{2j/d}}\right), \\ \text{SPE}_{(\text{pos}, 2j+1)} &= \cos\left(\frac{\lfloor \text{pos} / \frac{\mathcal{M}}{N} \rfloor}{10000^{2j/d}}\right), \end{aligned} \quad (9)$$

where pos, j indicates the location and d denotes the dimension of feature vectors. For $\text{pos} \in \{H+1, H + \frac{\mathcal{M}}{N}\}$, the $[\text{SP}]_i^t$ share the same positional encoding because they represent knowledge from the same task. SPE incorporates session information into the sequential prompts, thereby enhancing the model’s ability to learn and retain task-specific knowledge. Additional details and visualization of SPE is available in the supplementary material.

Negative-feedback Knowledge Accumulation

Inspired by the negative feedback regulation, we propose a knowledge accumulation mechanism. Specifically, we assess knowledge retention by computing the divergence between the current logits and those from the previous task. This divergence serves as a feedback signal used to dynamically adjust the weights of the prior prompt token (*i.e.* $\{[\text{PT}]_i^{t-1}\}$) and the current prompt token (*i.e.* $\{[\text{PT}]_i^t\}$). The flow of the NKA mechanism is shown in Fig. 3. Given two logits l^t and l^{t-1} , the divergence is computed as Mean Absolute Error (MAE):

$$\text{mae} = \text{MAE}(l^t[0 : K(t-1)] \cdot \lambda, l^{t-1} \cdot \lambda), \quad (10)$$

¹<https://github.com/qinyongchun/PEARL>




Figure 3: The illustration of the proposed NKA mechanism. A low mae indicates good knowledge retention, enabling the model to focus more on the new task.

where λ serves as a scale factor. The computation involves only the first $K(t-1)$ terms of the current logits, as the decision space is expanding and the last K terms represent new knowledge rather than previous knowledge. The momentum weight α^τ is further computed by:

$$\alpha^\tau = \gamma \cdot \alpha^{\tau-1} + (1 - \gamma) \cdot \sigma(mae), \quad (11)$$

where τ is the iteration number and $\sigma(\cdot)$ is the sigmoid-like activation function with an upper bound θ_{\max} and a lower bound θ_{\min} . We employ a momentum update for α^τ to ensure numerical stability and prevent fluctuations that could lead to drastic changes in $\{[PT]_i^{t-1}\}$. When updating $\{[PT]_i^{\text{mem}}\}$, the momentum update is denoted by Eq. (3).

Eq. (3) and Eq. (11) defines the feedback process of negative-feedback regulation. Further, the forward process of NKA is denoted as follows:

$$\begin{aligned} l^{t-1} &= g \circ \mathcal{V}(\mathbf{x}, \{[PT]_i^{t-1}\}), \\ l^t &= g \circ \mathcal{V}(\mathbf{x}, \{[PT]_i^{\text{mem}}\}). \end{aligned} \quad (12)$$

The current momentum token $\{[PT]_i^{\text{mem}}\}$ is stored in memory and serves as $\{[PT]_i^t\}$ during the $(t+1)$ -th session. Since PEARL is an input-agnostic method, it requires only additional memory space of size $L \times H \times d$, avoiding inference costs for computing $\{[PT]_i^{t-1}\}$.

In optimization, the Cross-Entropy loss is used during the t -th session:

$$\mathcal{L}_{\text{cls}} = \mathbb{E}_{(\mathbf{x}, y) \sim \mathbf{D}^t} \text{CE}(l^t, y). \quad (13)$$

Our method does not require any knowledge retention loss (e.g. Knowledge Distillation), because knowledge retention is accounted for by the adaptive weighting factor α^τ . Our model begins to fit new tasks only if there is no forgetting of old tasks; otherwise, the prompt token will continuously backtrack. Consequently, PEARL elevates prompt-based methods to state-of-the-art performance, challenging existing beliefs about their efficacy for PTM-based CIL.

Experiments

Datasets and Implementation Details

Datasets. In order to comprehensively examine the model performance, we follow (Zhou et al. 2024c) and conduct ex-

periments on six datasets including CIFAR00 (Krizhevsky, Hinton et al. 2009), CUB200 (Wah et al. 2011), ImageNet-R (Hendrycks et al. 2021a), ImageNet-A (Hendrycks et al. 2021b), Omnibenchmark (Zhang et al. 2022) and VTAB (Zhai et al. 2019). There are 50 classes in VTAB, 100 classes in CIFAR-100, 200 classes in CUB200, ImageNet-R and ImageNet-A and 300 classes in Omnibenchmark.

Evaluation metrics. For an incremental learning task with \mathcal{N} sessions in total, the classification accuracy of the model on the t -th session is denoted as \mathcal{A}^t . Followed by (Zhou et al. 2024c), we adopt two evaluation metrics: the average accuracy $\bar{\mathcal{A}} = \frac{1}{\mathcal{N}} \sum_{t=1}^{\mathcal{N}} \mathcal{A}^t$ and the final accuracy $\mathcal{A}^{\mathcal{N}}$.

Implementation details. Following the experiment settings of (Wang et al. 2022b, 2023; Zhou et al. 2024b,c), We choose ViT as the backbone initialized with **ViT-B/16-IN21K** and **ViT-B/16-IN1K** parameters. L and H equals 2 and 4, respectively and the length of prompt pool is 100. In NKA mechanism, the initial α^0 is set as 0.99 and λ equals 12500 during the training process. The upper and lower bounds of $\sigma(\cdot)$ are set as 0.999 and 0.7, respectively. We train the model with SGD optimizer and cosine annealing with epoch as 10 and batchsize as 32. Our results are the average of three random runs and conducted with PyTorch(Paszke et al. 2019) and PILOT(Sun et al. 2023). All experiments are conducted on one RTX 4090.

Comparison with State-of-the-art Methods

The compared methods include prompt-based (e.g. L2P (Wang et al. 2022b), DualPrompt (Wang et al. 2022a), CODA-Prompt (Smith et al. 2023)), representation-based (e.g. SimpleCIL (Zhou et al. 2023), ADAM (Zhou et al. 2023), RanPAC (McDonnell et al. 2024), EASE (Zhou et al. 2024b)) and model mixture-based (e.g. HiDe-Prompt (Wang et al. 2024), ESN (Wang et al. 2023)).

As reported in Table 1, PEARL achieves the best performance among all six benchmarks. The experiments span various sequential lengths, and our method performs well in a variety of settings, demonstrating its superiority. Compared to other prompt-based methods (i.e. L2P, DualPrompt and CODA-Prompt), our method demonstrates a significant advantage, with average improvements of 13.58% and 15.02% on $\bar{\mathcal{A}}$ and $\mathcal{A}^{\mathcal{N}}$. This improvement is attributed to the integration of the proposed SPA module and NKA mechanism, which will be further analyzed in the ablation study. Compared to RanPAC, the second-best method, PEARL achieves an average improvement of 2.24% and 1.65% on $\bar{\mathcal{A}}$ and $\mathcal{A}^{\mathcal{N}}$, respectively. In RanPAC, the model updates only at $t = 1$ and adjusts the classification head based on features from later sessions, limiting its ability to effectively learn from subsequent tasks. In contrast, PEARL performs continuous updates across all sequential tasks, ensuring that the latest knowledge is consistently learned.

Ablation Study

Effect of SPA module. To verify the generality of the proposed SPA module, we add the prompt encoder to other prompt-based methods and make comparisons. As reported in Table 2, the model benefits only when both the prompt

Method	CIFAR 20-tasks		CUB 20-tasks		IN-R 40-tasks		IN-A 10-tasks		Omni 10-tasks		VTAB 5-tasks		Average	
	\bar{A}	$\mathcal{A}^{\mathcal{N}}$	\bar{A}	$\mathcal{A}^{\mathcal{N}}$	\bar{A}	$\mathcal{A}^{\mathcal{N}}$	\bar{A}	$\mathcal{A}^{\mathcal{N}}$	\bar{A}	$\mathcal{A}^{\mathcal{N}}$	\bar{A}	$\mathcal{A}^{\mathcal{N}}$	\bar{A}	$\mathcal{A}^{\mathcal{N}}$
L2P	85.94	79.93	67.05	56.25	66.53	59.22	49.39	41.71	73.36	64.49	77.11	77.10	69.90	63.12
DualPrompt	87.87	81.15	77.47	66.54	63.31	55.22	53.71	41.67	73.92	65.52	83.36	81.23	73.27	65.22
CODA-Prompt	89.11	81.96	84.00	73.37	64.42	55.08	53.54	42.73	77.03	68.09	83.90	83.02	75.33	67.38
SimpleCIL	87.57	81.26	92.20	86.73	62.58	54.55	59.77	48.91	79.34	73.15	85.99	84.38	77.91	71.50
ADAM + VPT-D	88.46	82.17	91.02	84.99	68.79	60.48	58.48	48.52	81.05	74.47	86.59	83.06	79.07	72.28
ADAM + SSF	87.78	81.98	91.72	86.13	68.94	60.60	61.30	50.03	80.53	74.00	85.66	81.92	79.32	72.44
ADAM + Adapter	90.65	85.15	92.21	86.73	72.35	64.33	60.47	49.37	80.75	74.37	85.95	84.35	80.40	74.05
RanPAC	<u>93.51</u>	<u>89.30</u>	<u>93.13</u>	89.40	75.74	68.75	64.16	52.86	<u>85.95</u>	<u>79.55</u>	92.56	91.83	<u>84.18</u>	78.62
HiDe-Prompt	91.22	89.92	89.75	<u>89.46</u>	76.20	74.56	61.41	49.27	76.60	77.01	91.24	92.78	81.07	<u>78.83</u>
ESN	87.15	80.37	65.69	63.10	60.69	55.13	44.06	31.07	75.32	66.57	81.52	62.15	69.07	59.73
EASE	91.51	85.80	92.23	86.81	<u>78.31</u>	70.58	<u>65.34</u>	<u>55.04</u>	81.66	74.85	<u>93.61</u>	93.55	83.78	77.77
PEARL (Ours)	93.64	89.02	94.48	89.65	79.54	<u>72.33</u>	67.41	57.87	86.87	79.68	96.52	<u>93.02</u>	86.41	80.26

Table 1: Comparison results on six benchmarks with ViT-B/16-IN21K as the backbone. Experiments are labeled as ‘‘Dataset- \mathcal{N} -tasks’’ where \mathcal{N} represents the length of the data stream. ‘‘IN-R’’ is short for ImageNet-R, ‘‘IN-A’’ is short for ImageNet-A, and ‘‘Omni’’ is short for Omnibenchmark. Bold texts: the best results, underline texts: the second-best results.

Method	PE	Mom	IN-A 10-tasks		VTAB 5-tasks	
			\bar{A}	$\mathcal{A}^{\mathcal{N}}$	\bar{A}	$\mathcal{A}^{\mathcal{N}}$
L2P	-	-	53.36	43.45	80.84	61.40
	✓	-	3.00	1.18	5.36	2.76
	✓	✓	58.16	49.18	88.51	69.74
DualPrompt	-	-	57.05	46.61	83.03	66.32
	✓	-	56.71	45.69	80.22	63.51
	✓	✓	60.50	50.69	87.83	77.27
CODA-Prompt	-	-	59.67	47.33	81.79	84.75
	✓	-	15.93	3.55	42.46	32.05
	✓	✓	61.83	51.68	84.82	86.60

Table 2: Ablation study on SPA with ViT-B/16-IN1K as the backbone. ‘‘PE’’ is short for prompt encoder and ‘‘Mom’’ is short for momentum update with weight equals 0.9.

encoder and momentum update are combined; using the prompt encoder alone leads to significant drawbacks. This is because the prompt encoder mixes task-specific knowledge, making the ‘‘query-select’’ mechanism ineffective. DualPrompt partially maintained its performance due to its unique ‘‘general-expert’’ prompt design, whereas L2P and CODA-Prompt experienced significant degradation. However, a simple fixed-weight momentum update addresses this issue by enabling smooth knowledge accumulation. We conclude that implementing an input-agnostic prompt effectively requires both a global encoder and a momentum update strategy; relying on either alone is insufficient. Table 3 reports the impact of different positional encodings. We compare the continual positional encoding (Vaswani et al. 2017) with the proposed segmented positional encoding. The experiments demonstrate that segmented positional

Positional Encoding	IN-A 10-tasks		VTAB 5-tasks	
	\bar{A}	$\mathcal{A}^{\mathcal{N}}$	\bar{A}	$\mathcal{A}^{\mathcal{N}}$
Continual	66.91	56.35	96.47	93.00
Segmented	67.65	57.14	96.59	93.07

Table 3: Ablation study on the segmented positional encoding with ViT-B/16-IN1K as the backbone.

outperforms on both ImageNet-A and VTAB. Further details are provided in the supplementary material.

Effect of NKA mechanism. We compare the NKA update with a fixed-weight momentum update. As shown in Table 4, when α falls below 0.9, performance deteriorates rapidly, resulting in the complete failure of the fixed-weight momentum update. This suggests that a fixed α makes the model highly sensitive to the initial value of α^0 . Meanwhile, when α is updated based on the NKA mechanism, the model achieves better results across different initial conditions, with \bar{A} and $\mathcal{A}^{\mathcal{N}}$ improved by an average of 14.37% and 18.09%, respectively. Additionally, we perform further analysis of the update process of α^T , which will be discussed in the following part.

Further Analysis

We visualize the NKA update process across six settings as reported in Table 4. Fig. 4 shows the update curves of mae and α under various initial conditions. Both mae and α consistently converge to a fixed value, suggesting this value represents the inherent correlation coefficient between tasks. As depicted in Fig. 5, although this coefficient varies across tasks, the model reliably converges to it regardless of initial conditions. This demonstrates that the NKA mechanism effectively reveals the inherent correlation in the data stream, promoting stable knowledge accumulation across different

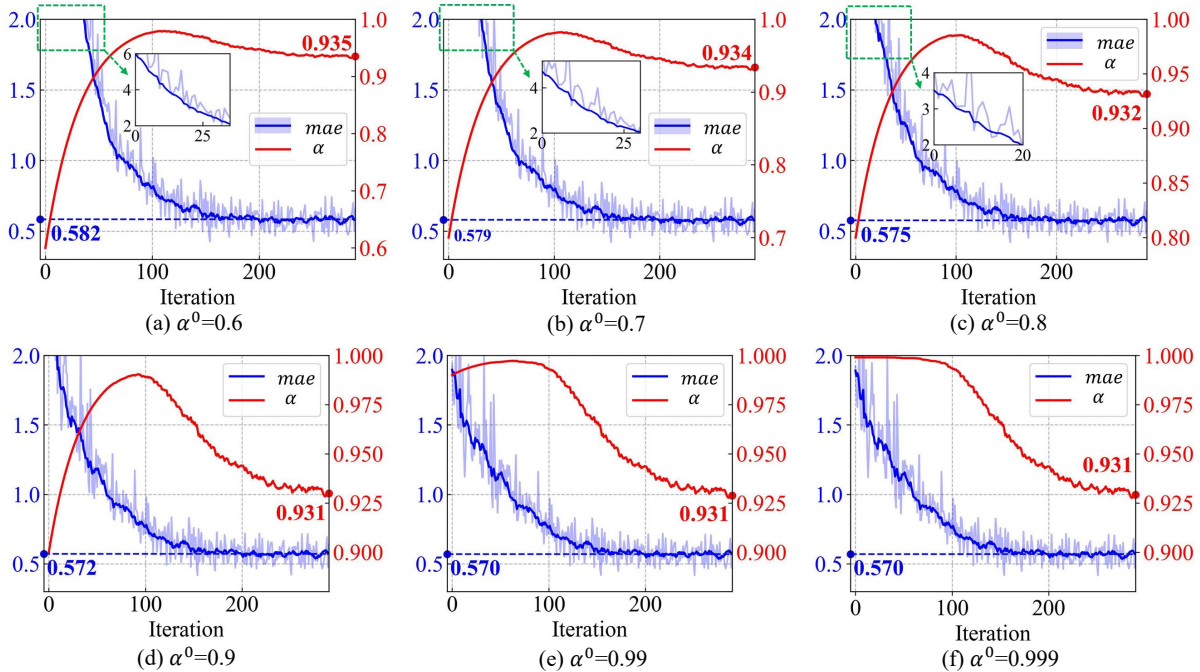


Figure 4: The curves of mae and α , across different initial value of α^0 . Results are derived from the second session of CUB.

α^0	Fixed α		NKA α	
	$\bar{\mathcal{A}}$	$\mathcal{A}^{\mathcal{N}}$	$\bar{\mathcal{A}}$	$\mathcal{A}^{\mathcal{N}}$
0.60	60.04	44.83	81.56	73.88
0.70	61.41	47.96	84.47	76.89
0.80	68.79	63.74	87.28	79.90
0.90	86.32	82.15	89.73	83.40
0.99	82.21	70.31	91.26	84.73
0.999	80.64	66.16	91.33	84.86

Table 4: Ablation study on the NKA mechanism with ViT-B/16-IN1K as the backbone. The results are obtained on CUB-200, and \mathcal{N} equals 10.

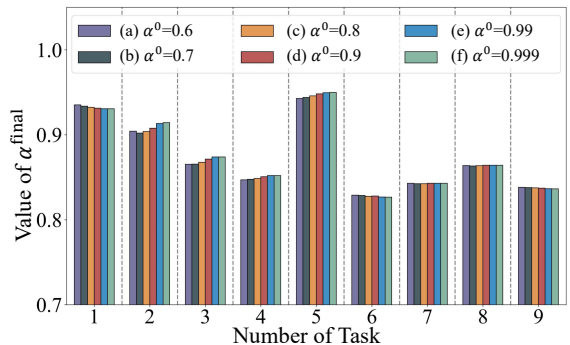


Figure 5: The value of α^{final} across different settings.

tasks. This observation explains the significant performance improvement of PEARL, which is likely due to the NKA mechanism’s ability to address the distribution divergence between the downstream and pre-trained datasets by uncovering potential correlations.

The mae metric indicates that a higher value reflects poorer knowledge preservation ability by the model. As shown in the inset figures of Fig. 4, the initial value of mae is negatively correlated with the choice of α^0 . This suggests that the weighted mixing method for prompts effectively regulates knowledge mixing, thereby ensuring the NKA mechanism’s effectiveness. The update curve of α exhibits an overshoot, consistent with a typical negative feedback response. This confirms the successful application of negative feedback regulation in our method. In the future, we will conduct further research on applying negative feedback regulation in deep learning.

Conclusion

This paper presents PEARL, an input-agnostic prompt method designed to address the issue of knowledge interference caused by the “query-select” mechanism of existing input-dependent prompt methods. Our method systematically manages and integrates task-specific knowledge through a global prompt, which helps mitigate catastrophic forgetting across incremental learning. Additionally, the proposed negative feedback based momentum update mechanism reveals potential correlations within the dataset, facilitating smooth and efficient knowledge accumulation. We hope our work can offer good insights into the field of CIL and provide some inspiration to other researchers.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (No. 62476056, 62076062, and 62306070) and the Social Development Science and Technology Project of Jiangsu Province (No. BE2022811). Furthermore, the work was also supported by the Big Data Computing Center of Southeast University.

References

- Aljundi, R.; Babiloni, F.; Elhoseiny, M.; Rohrbach, M.; and Tuytelaars, T. 2018. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European conference on computer vision (ECCV)*, 139–154.
- Castro, F. M.; Marín-Jiménez, M. J.; Guil, N.; Schmid, C.; and Alahari, K. 2018a. End-to-end incremental learning. In *Proceedings of the European Conference on Computer Vision*, 233–248.
- Castro, F. M.; Marín-Jiménez, M. J.; Guil, N.; Schmid, C.; and Alahari, K. 2018b. End-to-end incremental learning. In *Proceedings of the European conference on computer vision (ECCV)*, 233–248.
- Chaudhry, A.; Dokania, P. K.; Ajanthan, T.; and Torr, P. H. 2018. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *Proceedings of the European conference on computer vision (ECCV)*, 532–547.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- 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.
- 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.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Hou, S.; Pan, X.; Loy, C. C.; Wang, Z.; and Lin, D. 2019a. Learning a unified classifier incrementally via rebalancing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 831–839.
- Hou, S.; Pan, X.; Loy, C. C.; Wang, Z.; and Lin, D. 2019b. Learning a unified classifier incrementally via rebalancing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 831–839.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Huang, W.-C.; Chen, C.-F.; and Hsu, H. 2024. OVOR: OnePrompt with Virtual Outlier Regularization for Rehearsal-Free Class-Incremental Learning. *arXiv preprint arXiv:2402.04129*.
- 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, H.; Ju, J.; Jung, M.; and Kim, J. 2016. Less-forgetting learning in deep neural networks. *arXiv preprint arXiv:1607.00122*.
- Kim, Y.; Fang, J.; Zhang, Q.; Cai, Z.; Shen, Y.; Duggal, R.; S Raychaudhuri, D.; Tu, Z.; Xing, Y.; and Dabeer, O. 2025. Open-World Dynamic Prompt and Continual Visual Representation Learning. In *European Conference on Computer Vision*, 357–374. Springer.
- 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.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Lester, B.; Al-Rfou, R.; and Constant, N. 2021. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*.
- Li, X. L.; and Liang, P. 2021. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*.
- Li, Z.; and Hoiem, D. 2017. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12): 2935–2947.
- 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.
- McCloskey, M.; and Cohen, N. J. 1989. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, 109–165. Elsevier.
- 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.
- Ostapenko, O.; Puscas, M.; Klein, T.; Jahnichen, P.; and Nabi, M. 2019. Learning to remember: A synaptic plasticity driven framework for continual learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11321–11329.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance

- deep learning library. *Advances in neural information processing systems*, 32.
- Rebuffi, S.-A.; Kolesnikov, A.; Sperl, G.; and Lampert, C. H. 2017a. iCaRL: Incremental classifier and representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2001–2010.
- Rebuffi, S.-A.; Kolesnikov, A.; Sperl, G.; and Lampert, C. H. 2017b. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2001–2010.
- Shin, H.; Lee, J. K.; Kim, J.; and Kim, J. 2017. Continual learning with deep generative replay. *Advances in neural information processing systems*, 30.
- Smith, J. S.; Karlinsky, L.; Gutta, V.; Cascante-Bonilla, P.; Kim, D.; Arbelles, 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.
- Sun, H.-L.; Zhou, D.-W.; Ye, H.-J.; and Zhan, D.-C. 2023. PILOT: A Pre-Trained Model-Based Continual Learning Toolbox. *arXiv preprint arXiv:2309.07117*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wah, C.; Branson, S.; Welinder, P.; Perona, P.; and Belongie, S. 2011. The caltech-ucsd birds-200-2011 dataset.
- Wang, L.; Xie, J.; Zhang, X.; Huang, M.; Su, H.; and Zhu, J. 2024. Hierarchical decomposition of prompt-based continual learning: Rethinking obscured sub-optimality. *Advances in Neural Information Processing Systems*, 36.
- Wang, Y.; Huang, Z.; and Hong, X. 2022. S-prompts learning with pre-trained transformers: An occam’s razor for domain incremental learning. *Advances in Neural Information Processing Systems*, 35: 5682–5695.
- 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*, volume 37, 10209–10217.
- 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.
- Wu, Y.; Chen, Y.; Wang, L.; Ye, Y.; Liu, Z.; Guo, Y.; and Fu, Y. 2019. Large scale incremental learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 374–382.
- 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*, 6619–6628.
- Zaken, E. B.; Ravfogel, S.; and Goldberg, Y. 2021. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. *arXiv preprint arXiv:2106.10199*.
- Zhai, X.; Puigcerver, J.; Kolesnikov, A.; Ruyssen, P.; Riquelme, C.; Lucic, M.; Djolonga, J.; Pinto, A. S.; Neumann, M.; Dosovitskiy, A.; et al. 2019. A large-scale study of representation learning with the visual task adaptation benchmark. *arXiv preprint arXiv:1910.04867*.
- Zhang, C.; Zhang, L.; Wu, J.; Zhou, D.; and He, Y. 2024. Causal prompting: Debiasing large language model prompting based on front-door adjustment. *arXiv preprint arXiv:2403.02738*.
- Zhang, Y.; Yin, Z.; Shao, J.; and Liu, Z. 2022. Benchmarking omni-vision representation through the lens of visual realms. In *European Conference on Computer Vision*, 594–611. Springer.
- Zhou, D.-W.; Sun, H.-L.; Ning, J.; Ye, H.-J.; and Zhan, D.-C. 2024a. Continual learning with pre-trained models: A survey. *arXiv preprint arXiv:2401.16386*.
- Zhou, D.-W.; Sun, H.-L.; Ye, H.-J.; and Zhan, D.-C. 2024b. Expandable subspace ensemble for pre-trained model-based class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23554–23564.
- Zhou, D.-W.; Wang, Q.-W.; Qi, Z.-H.; Ye, H.-J.; Zhan, D.-C.; and Liu, Z. 2024c. Class-Incremental Learning: A Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Zhou, D.-W.; Ye, H.-J.; Zhan, D.-C.; and Liu, Z. 2023. Revisiting class-incremental learning with pre-trained models: Generalizability and adaptivity are all you need. *arXiv preprint arXiv:2303.07338*.