

# DualCP: Rehearsal-Free Domain-Incremental Learning via Dual-Level Concept Prototype

Qiang Wang<sup>1</sup>, Yuhang He<sup>1\*</sup>, Songlin Dong<sup>1</sup>, Xiang Song<sup>1</sup>,  
Jizhou Han<sup>1</sup>, Haoyu Luo<sup>1</sup>, Yihong Gong<sup>1,2\*</sup>

<sup>1</sup>Xi'an Jiaotong University

<sup>2</sup>Shenzhen University of Advanced Technology  
qwang@stu.xjtu.edu.cn, heyuhang@xjtu.edu.cn,

{dsl972731417, songxiang, jizhou.han, luohaoyu}@stu.xjtu.edu.cn, ygong@mail.xjtu.edu.cn

## Abstract

Domain-Incremental Learning (DIL) enables vision models to adapt to changing conditions in real-world environments while maintaining the knowledge acquired from previous domains. Given privacy concerns and training time, Rehearsal-Free DIL (RFDIL) is more practical. Inspired by the incremental cognitive process of the human brain, we design Dual-level Concept Prototypes (DualCP) for each class to address the conflict between learning new knowledge and retaining old knowledge in RFDIL. To construct DualCP, we propose a Concept Prototype Generator (CPG) that generates both coarse-grained and fine-grained prototypes for each class. Additionally, we introduce a Coarse-to-Fine calibrator (C2F) to align image features with DualCP. Finally, we propose a Dual Dot-Regression (DDR) loss function to optimize our C2F module. Extensive experiments on the DomainNet, CDDB, and CORe50 datasets demonstrate the effectiveness of our method.

## Introduction

Domain-Incremental Learning (DIL) aims to train a unified model incrementally across continuously encountered domains. It has drawn remarkable attention in recent years (Zhang and Mueller 2022; Verwimp et al. 2023) and has a wide range of applications. For example, in a visual recognition model on an autonomous vehicle, the model is expected to incrementally learn and adapt to new and dynamic environments (Wang et al. 2024) such as forests, deserts, cities, *etc.* As the number of domains increases, the model may forget previously acquired knowledge. Many studies (Isele and Cosgun 2018; Zhao et al. 2021) address the retention of knowledge by preserving and retraining old samples, known as *rehearsal*. Rehearsal-based methods, however, require longer training time, and storing old domain images may raise privacy concerns (Wan et al. 2024a). Therefore, Rehearsal-Free DIL (RFDIL) is more practical for real-world scenarios compared to rehearsal-based approaches.

The primary challenge of RFDIL is the conflict between learning new domains and preventing forgetting of old ones.

\*Corresponding Authors

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

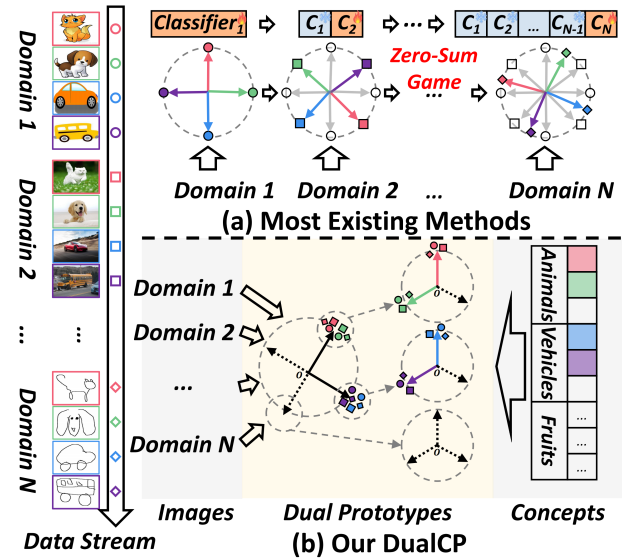


Figure 1: (a) Illustration of most existing methods, where new domain features and old domain features compete within the feature space. (b) Illustration of our method based on dual-level concept prototypes. Different colors (red, green, blue, and purple) represent different classes, while different shapes ( $\circ$ ,  $\square$ , and  $\diamond$ ) represent different domains. Best viewed in color.

On the one hand, learning new domain knowledge may overwrite parameters related to old domains, leading to *catastrophic forgetting* (McCloskey and Cohen 1989). On the other hand, adding artificial constraints to prevent forgetting old domains will hinder learning new ones. To mitigate this dilemma, the method (Zhu et al. 2021) proposes generating samples for the old domains as a supplement to retain knowledge; the methods (Rebuffi et al. 2017; Hou et al. 2019) suggest employing knowledge distillation strategies to slow down the forgetting of old knowledge; the methods (Douillard et al. 2022; Wang et al. 2022b; Smith et al. 2023; Wang, Huang, and Hong 2022) propose to train a set of learnable prompt tokens for knowledge retention. However, most methods aim to balance new-domain learning and old-domain forgetting, resulting in a zero-sum game. As il-

lustrated in Fig. 1 (a), the increase in the number of domains causes more features to crowd within the same feature space, escalating the competition between new and old domains. This raises a question: *are different domains necessarily in opposition?*

To address this question, we draw inspiration from the incremental cognitive process of humans. The cognitive science research (Bar 2003; Fenske et al. 2006) indicate that the early visual areas of the brain extract both low and high spatial frequency signals (LF & HF) from the visual stimuli after encountering a new object. The LF signals are projected to the prefrontal cortex, forming coarse-grained features within approximately 140 ms (Barcelo, Suwazono, and Knight 2000). For example, when seeing a photograph or a cartoon image of a cat, the brain initially forms a general impression of an animal, without determining the cat’s color or whether the cat is real. This process occurs so quickly that we are often unaware of it in everyday life. The HF signals are processed by the ventral stream to generate fine-grained features, such as the texture of the image (identifying the domain it belongs to) or the precise category of the object (*e.g.*, cat or dog within the animal category). From this research, we derive two key insights: **(a) The learning of new knowledge is hierarchical.** The brain first forms a general concept of a new object before focusing on its details. In contrast, existing deep classification models treat all given categories equally, without considering their inherent semantic relationships. **(b) New and old domains are not completely distinct.** Although domain shifts cause visual differences among different domains, the coarse-grained representations of the same concept are similar from a human perspective. *This answers the above question.*

Based on the above insights, we design a framework based on two principles. **(a) We propose treating images of the same class across multiple domains as a single concept.** This concept remains consistent regardless of the domain. For instance, the concept of a “cat” does not change whether it is represented by photos, cartoons, or simple sketches of cats. During the initial domain training in RFDIL, we employ neural networks to mine the commonalities of the same concept across different domains. Specifically, we use the semantic features of the concept to construct a concept prototype. To maximize the separability between different prototypes, we represent the concept prototype using a semantically guided simplex Equiangular Tight Frame (ETF) through the neural collapse theory. The ETF is a mathematical structure that maximizes the distances between feature pairs, with further details provided in the preliminary section. **(b) We introduce Dual-level Concept Prototypes (DualCP) to model the human cognitive process of recognizing new objects.** we first group all classes based on their superordinate, as shown in Fig. 1 (b). For example, cats and dogs belong to the animal group, while cars and buses belong to the vehicle group. First-level concept prototypes are designed for each superordinate concept to simulate the coarse-grained features formed by LF signals in the prefrontal cortex. Within each group, we design second-level concept prototypes to simulate the process by which HF signals in the ventral stream recognize specific classes.

Finally, we introduce a coarse-to-fine calibrator (C2F) to align image features with their corresponding concept prototypes. A Dual-level Dot Regression (DDR) loss function is employed to optimize the training of the C2F. We conduct extensive experiments on three benchmark datasets, including DomainNet, CORe50, and CDDDB, comparing our DualCP with state-of-the-art methods.

Our contributions can be summarized as follows:

- Inspired by the incremental cognition process of humans, we propose treating cross-domain images of the same class as a single concept. We construct concept prototypes for all domains to prevent a zero-sum game between new and old domains.
- We introduce DualCP, an RFDIL method based on dual-level concept prototypes, which further enhances the separability between classes by performing fine-grained classification on similar categories.
- We propose the C2F module and the DDR loss function to align objects with their corresponding coarse-grained and fine-grained concept prototypes.
- Extensive experiments demonstrate that the proposed DualCP method outperforms other RFDIL methods on three benchmark datasets, with a margin of up to 7.6% on the CDDDB dataset.

## Related Work

### Domain-Incremental Learning

Multiple studies focus on addressing *catastrophic forgetting* (McCloskey and Cohen 1989) in Domain-Incremental Learning (DIL). For instance, (Kirkpatrick et al. 2017; Akyürek et al. 2021; Shi et al. 2023) restrict the model’s plasticity to balance the learning of new domains with the forgetting of old ones, resulting in a zero-sum game between new and old domains. **Rehearsal-based** methods (Rebuffi et al. 2017; Chaudhry et al. 2019) mitigate the forgetting of old domains by replaying a subset of representative samples from the old domains, which raises privacy concerns. **Rehearsal-free** methods (Van De Ven, Li, and Tolia 2021; Wang et al. 2023; Wan et al. 2024b) use synthetic data replay from old domains to reduce forgetting. (Wang et al. 2022b; Liu, Peng, and Zhou 2024; Gao et al. 2024; Wang et al. 2025) dynamically expand the model to incrementally store knowledge from new domains, aiming to separate the learning of new and old domain knowledge to avoid conflicts.

### Neural Collapse

The Neural Collapse (NC) theory (Papayan, Han, and Donoho 2020) posits that features of the same class in the final layer collapse to a single point on a hypersphere in a well-trained classification model. Besides, the distance between features of any two different classes is maximized, and these features of all classes collectively can be defined by a simplex Equiangular Tight Frame (ETF). This indicates that a “training endpoint” can be easily constructed for the model. We can align image features to the pre-designed endpoint, which is independent of the model’s initial parameters. The papers (Mixon, Parshall, and Pi 2020; Ji et al. 2021;

Zhou et al. 2022) have shown that the accuracy of models trained using the simplex ETF is comparable to those trained with conventional cross-entropy loss.

### Applications of Neural Collapse in Other Tasks

Recent studies have utilized the above characteristics of NC to address various tasks, such as imbalanced learning (Yang et al. 2022), few-shot class-incremental learning (Yang et al. 2023), and federated learning (Huang et al. 2023). Compared with existing research, the proposed DualCP is **not a simple combination of NC and RFDIL, but is a novel method inspired by the cognitive science research, i.e.,** the learning of new knowledge is hierarchical and the new and old domains are not completely distinct. We build concept prototypes based on the two insights to solve the RFDIL problem, while the ETF serves as a tool to realize our idea.

For the classification of a large number of classes, (Jiang et al. 2023) proposed maximizing the minimum one-vs-rest margins to achieve generalized neural collapse. However, this method has overlooked the inherent similarities among classes. Our proposed DualCP instead groups the numerous classes using text features, then performs coarse-grained classification between groups and fine-grained classification within groups. Our proposed dual-level concept prototypes are not only easy to implement but also accommodate up to  $d^2$  classes when the feature dimension is  $d$ , which is sufficient for most real-world scenarios.

## Preliminary

### Problem Formulation

Rehearsal-Free Domain-Incremental Learning (RFDIL) aims to train a unified model progressively on the data from  $T$  domains  $\mathcal{D} = \{\mathcal{D}_t\}_{t=1}^T$ , where  $T$  is the total number of domains. The data in each domain is split into the training set  $\mathcal{X}_t$  and the test set  $\mathcal{Z}_t$ , denoted as  $\mathcal{D}_t = (\mathcal{X}_t, \mathcal{Z}_t)$ . At the  $t$ -th stage of training, the model is only allowed to train on  $\mathcal{X}_t$  and does not have access to  $\mathcal{X}_{1 \sim t-1} = \bigcup_{\tau=1}^{t-1} \mathcal{X}_\tau$ . After the  $t$ -th training, the model is tested on  $\mathcal{Z}_{1 \sim t} = \bigcup_{\tau=1}^t \mathcal{Z}_\tau$ . The training set  $\mathcal{X}_t$  is composed of  $N_t$  tuples, denoted as  $\mathcal{X}_t = \{(x_{t,i}, y_{t,i})\}_{i=1}^{N_t}$ , where  $x_{t,i}$  signifies the  $i$ -th image from the  $t$ -th domain, and  $y_{t,i}$  denotes the label of  $x_{t,i}$ . The test set  $\mathcal{Z}_t$  follows a similar structure. Moreover, define  $\mathcal{C}_t$  as the set of labels for the  $t$ -th domain, i.e.,  $\forall i \in [1, N_t], y_{t,i} \in \mathcal{C}_t$ . In other words,  $|\mathcal{C}_t|$ , the number of classes in each domain, remains constant in the same dataset.

### Definition of Simplex ETF

For a well-trained classification model with  $K$  classes, the within-class means correspond to  $K$  prototypes, denoted as  $\mathbf{m}_i \in \mathbb{R}^d, i = 1, 2, \dots, K$ , where  $K \leq d+1$ . The collection of these prototypes, i.e.,  $\mathbf{M} = [\mathbf{m}_1, \dots, \mathbf{m}_K]$  in  $\mathbb{R}^{d \times K}$ , is called a simplex ETF, which means:

$$\mathbf{M} = \sqrt{\frac{K}{K-1}} \mathbf{U} (\mathbf{I}_K - \frac{1}{K} \mathbf{1}_K \mathbf{1}_K^T), \quad (1)$$

where  $\mathbf{U}$  satisfies  $\mathbf{U}^T \mathbf{U} = \mathbf{I}_K$ , and  $\mathbf{1}_K$  is a  $K$ -dimensional all-ones vector. All prototypes  $\mathbf{m}_i$  have the same  $l_2$  norm,

i.e.,  $\|\mathbf{m}_i\| = 1, i = 1, 2, \dots, K$ , and the same pair-wise angle, i.e.,

$$\mathbf{m}_i^T \mathbf{m}_j = \frac{K}{K-1} \delta_{i,j} - \frac{1}{K-1}, \forall i, j \in [1, K], \quad (2)$$

where  $\delta_{i,j}$  equals to 1 when  $i = j$  and 0 otherwise. The pairwise angle  $-\frac{1}{K-1}$  is the maximal equiangular separation of  $K$  prototypes in the  $d$ -dimension feature space.

Based on the definition of simplex ETF, the NC phenomenon can be summarized as:

**(NC1)** Feature collapse. The last-layer features of the same class will collapse to their within-class mean, i.e.,  $\sum_W \rightarrow 0$ , where  $\sum_W = \text{Avg}\{(\mu_{k,i} - \mu_k)(\mu_{k,i} - \mu_k)^T\}$ ,  $\mu_{k,i}$  is the feature of the  $i$ -th sample of the  $k$ -th class and  $\mu_k$  is the within-class mean of the  $k$ -th class.

**(NC2)** Convergence to simplex ETF. The within-class means of all  $K$  classes is centered by the global mean  $\mu_G = \text{Avg}_{i,k}(\mu_{k,i})$ . These means  $\mu_k$  will converge to the  $K$  prototypes of a simplex ETF, i.e.,  $\hat{\mu}_k = (\mu_k - \mu_G) / \|\mu_k - \mu_G\|, 1 \leq k \leq K$ , where  $\hat{\mathbf{M}} = [\hat{\mu}_1, \dots, \hat{\mu}_K]$  satisfies Eq. (1), and  $\hat{\mu}_k$  satisfies Eq. (2).

**(NC3)** Self-duality. The within-class feature means will be aligned with their corresponding classifier weights  $\mathbf{w}_k$ , i.e.,  $\hat{\mu}_k = \mathbf{w}_k / \|\mathbf{w}_k\|$ .

**(NC4)** Based on (NC1)-(NC3), the model prediction using the classifier can be simplified to select the nearest class center, i.e.,  $\text{argmax}_k \langle \mu, \mathbf{w}_k \rangle = \text{argmin}_k \|\mu - \mu_k\|$ , where  $\langle \cdot \rangle$  is the inner product operator, and  $\mu$  is the last-layer feature of a sample for prediction.

## Method

### Overall Framework

Fig. 2 illustrates the framework of our proposed DualCP. Our method contains three main components: (a) a pre-trained image feature extractor  $f_i(\cdot; \theta_i)$ , (b) the Concept Prototype Generator (CPG) with a pre-trained text feature extractor  $f_t(\cdot; \theta_t)$ , and (c) the Coarse-to-Fine calibrator (C2F), which includes a coarse-grained layer  $g_C(\cdot; \varphi_C)$  and multiple fine-grained layers  $g_{F_k}(\cdot; \varphi_{F_k})$ .

For simplicity, we abbreviate  $x_{t,i}$  as  $x$  and  $y_{t,i}$  as  $y$ . Given a set of images  $x$  and their class names  $y$ , we first extract the image features by  $\mathbf{x} = f_i(x; \theta_i)$ . Then, the class names are fed into the CPG to extract text features and construct dual-level concept prototypes, which include coarse-grained prototypes  $\mathbf{E}_C$  and fine-grained prototypes  $\mathbf{E}_{F_i}$ . Finally, we train the C2F to align the image features  $\mathbf{x}$  with the corresponding concept prototypes by minimizing the proposed Dual Dot-Regression (DDR) loss function. For further details on the training and inference of our DualCP framework, please refer to the appendix.

### Concept Prototype Generator

**Preparation for generating prototypes.** First, we collect the names of all classes, denoted as  $\mathcal{C}_t = \{y_1, y_2, \dots, y_K\}$ , e.g.,  $\mathcal{C}_t = \{\text{airplane, bike, cat, ..., zebra}\}$ .  $K = |\mathcal{C}_t|$  is the

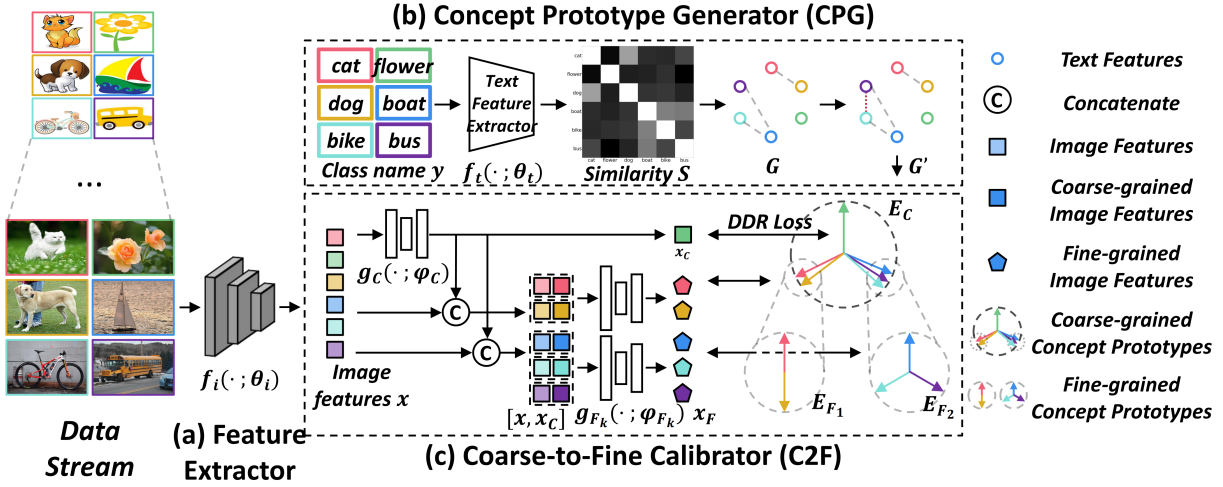


Figure 2: **The framework of the proposed DualCP.** DualCP comprises three main components: (a) a feature extractor to get the image features, (b) the CPG module to construct the dual-level concept prototype based on the text features of the class names, and (c) the C2F module to align the image features with the corresponding prototypes. We selected six common classes from the DomainNet dataset, *i.e.*, cat, flower, dog, boat, bike, and bus, to further illustrate our method. Similar classes were grouped, such as cats and dogs. We constructed coarse-grained prototypes between groups and fine-grained prototypes within groups. This coarse-to-fine classification approach helps the model better distinguish similar categories. Best viewed in color.

number of classes. Text features are extracted by:

$$y_i = \frac{f_t(y_i; \theta_t)}{\|f_t(y_i; \theta_t)\|}, \quad (3)$$

$$\mathbf{Y} = [y_1, y_2, \dots, y_K], \quad (4)$$

where  $f_t(\cdot; \theta_t)$  denotes a pre-trained text feature extractor, such as the CLIP text encoder (Radford et al. 2021). The symbol  $[\cdot]$  represents the concatenation operation, the matrix  $\mathbf{Y} \in \mathbb{R}^{d \times K}$  is the collection of text features, with  $d$  being the feature dimension. Note that text features have been normalized, *i.e.*,  $\forall i \in [1, K], \|y_i\| = 1$ .

Second, we introduce two strategies of the CPG, namely, vanilla concept prototype and dual-level concept prototype. We will elucidate our approaches successively.

**Vanilla Concept Prototype.** We constructed a simplex ETF based on text features to maximize the separability between each pair of generated concept prototypes. Specifically, we perform a QR decomposition on  $\mathbf{Y}$  to obtain the orthogonal basis of text features by:

$$\mathbf{Y} = \mathbf{Q}\mathbf{R}, \quad (5)$$

where  $\mathbf{Q} = [\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_K] \in \mathbb{R}^{d \times K}$  is an orthogonal matrix,  $\mathbf{q}_i^T \mathbf{q}_j = 1$ , when  $i = j$ , and 0 otherwise.  $\mathbf{R} \in \mathbb{R}^{K \times K}$  is an upper triangular matrix. Then we can compute the vanilla concept prototypes following Eq. (1) by:

$$\mathbf{E} = \sqrt{\frac{K}{K-1}} \mathbf{Q} (\mathbf{I}_K - \frac{1}{K} \mathbf{1}_K \mathbf{1}_K^T), \quad (6)$$

where  $\mathbf{I}_K$  is a  $K \times K$  identity matrix, and  $\mathbf{1}_K$  represents a  $K$ -dimensional all-one vector. Besides,  $\mathbf{E} = [\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_K] \in \mathbb{R}^{d \times K}$  encompasses the concept prototypes for all classes, where  $\mathbf{e}_i \in \mathbb{R}^d$  denotes the prototype

for the  $i$ -th class. The similarity between pairwise prototypes can be expressed as:

$$\mathbf{e}_i^T \mathbf{e}_j = \frac{K}{K-1} \mathbf{q}_i^T \mathbf{q}_j - \frac{1}{K-1}, \forall i, j \in [1, K], \quad (7)$$

**Dual-level Concept Prototype.** The construction of vanilla concept prototypes treats each class equally, without considering the similarity between classes. To enhance the distinction between similar categories, a direct idea is to perform more fine-grained differentiation for similar classes. Specifically, we set the similar classes as a group and construct dual-level concept prototypes, comprising coarse-grained prototypes for all groups and fine-grained prototypes for all classes within each group. We commence by computing the similarity matrix  $\mathbf{S}$  between all classes as  $\mathbf{S} = \mathbf{Y}^T \mathbf{Y}$ . We use the matrix  $\mathbf{S}$  to obtain the adjacency matrix  $\mathbf{A}$  with the hyperparameter  $p$  by:

$$\mathbf{A}_{ij} = \{\mathbf{S}_{ij} > p\}, \forall i, j \in [1, K], \quad (8)$$

where  $\{\cdot\}$  denotes an Iverson bracket, a mathematical notation used to represent a logical value based on a condition. It equals 1 if the condition is true and 0 if false.

Subsequently, we construct a connectivity graph  $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ , representing the relationships between all classes.  $\mathcal{V}$  denotes the set of nodes (classes), and  $|\mathcal{V}| = K$ .  $\mathcal{E}$  denotes the set of edges (whether the two classes are similar), and  $|\mathcal{E}| = \sum_{i=1}^K \sum_{j=1}^K \mathbf{A}_{ij}$ . To group similar nodes based on the adjacency matrix  $\mathbf{A}$ , we conduct a connectivity analysis on the graph  $\mathcal{G}$  and get another graph  $\mathcal{G}'$ . Connectivity analysis refers to the process where if a path exists between two classes, they should be placed in the same group. The details can be found in the appendix. Please refer to the illustration in Fig. 2 (b) for a simple example.

Based on the above algorithms, we can group the text features  $\mathbf{Y} = [\mathbf{Y}_1, \mathbf{Y}_2, \dots, \mathbf{Y}_{N_g}]$ , where  $N_g$  is the number of groups.  $\mathbf{Y}_i \in \mathbb{R}^{|g_k| \times d}$  represents the text features for the  $i$ -th group, and  $|g_k|$  is the number of concepts in the  $k$ -th group. Then we compute the average text feature for each group by:

$$\bar{\mathbf{Y}} = \left[ \frac{1}{|g_1|} \sum_{i=1}^{|g_1|} \mathbf{Y}_{1,i}, \frac{1}{|g_2|} \sum_{i=1}^{|g_2|} \mathbf{Y}_{2,i}, \dots, \frac{1}{|g_{N_g}|} \sum_{i=1}^{|g_{N_g}|} \mathbf{Y}_{N_g,i} \right]. \quad (9)$$

Based on Eqs. (5) to (7) and (9), we can calculate the coarse-grained concept prototypes (CCP) through  $\bar{\mathbf{Y}}$ , denoted as:

$$\mathbf{E}_C = [\mathbf{e}_{C,1}, \mathbf{e}_{C,2}, \dots, \mathbf{e}_{C,N_g}] \in \mathbb{R}^{d \times N_g}. \quad (10)$$

Additionally, we can calculate the fine-grained concept prototypes (FCP) using  $\mathbf{Y}_i$ , represented as:

$$\mathbf{E}_{F_k} = [\mathbf{e}_{F_k,1}, \mathbf{e}_{F_k,2}, \dots, \mathbf{e}_{F_k,|g_k|}] \in \mathbb{R}^{d \times |g_k|}. \quad (11)$$

Now we construct the dual-level concept prototypes for all classes. In other words, our CPG generates a coarse-grained prototype  $\mathbf{e}_C$  and a fine-grained prototype  $\mathbf{e}_F$  for each class.

### Coarse-to-Fine Calibrator

The coarse-to-fine calibrator (C2F) is proposed to align the image features with the corresponding prototypes. Given an input image  $x$  and its class name  $y$ , we extract features of the image  $x$  as  $\mathbf{x} = f_i(x; \theta_i) \in \mathbb{R}^d$ , where  $f_i(\cdot; \theta_i)$  is a pre-trained model, such as ViT (Dosovitskiy et al. 2020) or CLIP image encoder (Radford et al. 2021). Assuming  $y$  is the  $j$ -th class from the  $i$ -th group, its concept prototypes are referred to as  $\mathbf{e}_{C,i}$  and  $\mathbf{e}_{F_i,j}$ . The C2F module consists of the coarse-grained layer  $g_C(\cdot; \varphi_C)$  and the fine-grained layers  $g_{F_k}(\cdot; \varphi_{F_k})$ . The coarse-grained feature  $\mathbf{x}_C$  and fine-grained features  $\mathbf{x}_F$  can be computed by:

$$\begin{aligned} \mathbf{x}_C &= g_C(\mathbf{x}; \varphi_C) \in \mathbb{R}^d, \\ \mathbf{x}_F &= g_{F_i}([\mathbf{x}, \mathbf{x}_C]; \varphi_F) \in \mathbb{R}^d. \end{aligned} \quad (12)$$

**Dual Dot-Regression Loss.** To train the C2F module, we propose the Dual Dot-Regression (DDR) loss function, denoted as  $\mathcal{L}$ . We train the parameter sets  $\varphi_C$  and  $\varphi_F$  by minimize the DDR loss with a hyperparameter  $\alpha$ :

$$\begin{aligned} \min_{\varphi_C, \varphi_F} \mathcal{L}(\mathbf{x}_C, \mathbf{x}_F, \mathbf{e}_{C,i}, \mathbf{e}_{F_i,j}) = \\ \alpha(\mathbf{x}_C^T \mathbf{e}_{C,i} - 1)^2 + (1 - \alpha)(\mathbf{x}_F^T \mathbf{e}_{F_i,j} - 1)^2, \end{aligned} \quad (13)$$

### Theoretical Analysis

**Theorem 1.** *The angle between any pair of CCPs or FCPs is larger than or equal to the angle between any pair of vanilla concept prototypes:*

$$\left. \begin{aligned} \langle \mathbf{e}_{C,m}, \mathbf{e}_{C,n} \rangle &\geq \langle \mathbf{e}_i, \mathbf{e}_j \rangle, \\ \langle \mathbf{e}_{F_k,m}, \mathbf{e}_{F_k,n} \rangle &\geq \langle \mathbf{e}_i, \mathbf{e}_j \rangle, \forall k, \end{aligned} \right\} \forall i \neq j, m \neq n. \quad (14)$$

This theorem demonstrates that our proposed dual-level concept prototype has larger inter-class angles, indicating better classification capability than the vanilla concept prototype. The proof of the theorem is provided in the appendix.

Method	Buffer(↓)	$A_T$ (↑)	$F_T$ (↑)
DyTox (Douillard et al. 2022)		62.94	-
DARE (Jee., Ara., and Zon. 2024)	50/class	32.32*	-22.98
DARE++ (Jee., Ara., and Zon. 2024)		40.51*	-
EWC (Kirkpatrick et al. 2017)		47.62	-
LwF (Li and Hoiem 2017)		49.19	-5.01
SimCLR (Chen et al. 2020)		44.20	-
BYOL (Grill et al. 2020)		49.70	-
Barlow Twins (Zbontar et al. 2021)		48.90	-
SupCon (Khosla et al. 2020)		50.90	-
L2P (Wang et al. 2022b)	0/class	40.15†	-2.25
DualPrompt (Wang et al. 2022a)		43.79†	<u>-2.03</u>
S-iP (Wang, Huang, and Hong 2022)		50.62†	-2.85
CODA-P (Smith et al. 2023)		47.42†	-3.46
C-Prompt (Liu, Peng, and Zhou 2024)		<u>58.68†</u>	-
DualCP (ours)		<b>60.13†</b>	<b>-1.96</b>

Table 1: **Experimental results on the DomainNet dataset.** † denotes that the method is based on the pre-trained ViT-B/16 model. \* denotes that DARE is based on ResNet-18. The best result within rehearsal-free methods is indicated by **bold**, and the second is marked by underline.

## Experiments

### Experimental Settings

**Datasets.** We conducted experiments on three multi-domain datasets, include DomainNet (Peng et al. 2019), CDDDB (Li et al. 2023), and CORE50 (Lomonaco and Maltoni 2017). **DomainNet** emerges as a large-scale dataset for DIL and domain adaptation, whose images are sourced from six domains marked by prominent inter-domain variations, with each domain including 345 categories. The training set of DomainNet consists of 409,832 images, while the test set comprises 176,743 images. **CORE50** is an object recognition dataset involving 11 distinct domains (50 classes per domain), with 8 domains for training and 3 for testing. **CDDDB** is specifically crafted for deepfake detection, encompassing 12 distinct deepfake methodologies and 3 different evaluation scenarios. We opt for the most challenging HARD track as suggested by S-Prompts (Wang, Huang, and Hong 2022).

**Evaluation Metrics.** There are three commonly used evaluation metrics for DIL: (1) the average accuracy ( $A_T$ ) at the end of training on all  $T$  domains; (2) the forgetting degree ( $F_T$ ) following (Li et al. 2023), and the formulas for calculating  $A_T$  and  $F_T$  are detailed in the appendix; (3) the ‘‘Buffer’’ represents additional data stored by the model for incremental learning. This data may include images from old domains used for rehearsal.

**Implementation Details.** We employ an SGD optimizer with an initial learning rate of 0.1 and a cosine decay schedule. Additionally, we apply the weight decay of  $2e^{-4}$  for regularization to mitigate overfitting. The training consists of 20 epochs on all datasets except DomainNet, which extends to 30 epochs. The mini-batch size is set to 128. The hyperparameters  $p$  and  $\alpha$  are set to 0.85 and 0.5, respectively. The ablations of  $p$  and  $\alpha$  are provided in the appendix.

Method	Buffer( $\downarrow$ )	$A_T$ ( $\uparrow$ )	$F_T$ ( $\uparrow$ )
LRCIL (Pellegrini et al. 2020)		76.39	-4.39
iCaRL (Marra et al. 2019)	100/class	79.76	-8.73
LUCIR (Hou et al. 2019)		82.53	-5.34
LRCIL (Pellegrini et al. 2020)		74.01	-8.62
iCaRL (Marra et al. 2019)	50/class	73.98	-14.50
LUCIR (Hou et al. 2019)		80.77	-7.85
DyTox (Douillard et al. 2022)		86.21	-1.55
EWC (Kirkpatrick et al. 2017)		50.59	-42.62
LwF (Li and Hoiem 2017)		60.94	-13.53
DyTox (Douillard et al. 2022)		51.27	-45.85
L2P (Wang et al. 2022b)		61.28 $\dagger$	-9.23
DualPrompt (Wang et al. 2022a)	0/class	64.80 $\dagger$	-8.74
S-iP (Wang, Huang, and Hong 2022)		<u>74.51<math>\dagger</math></u>	<u>-1.30</u>
CODA-P (Smith et al. 2023)		70.54 $\dagger$	-5.53
DualCP (ours)		<b>82.16<math>\dagger</math></b>	<b>-0.73</b>

Table 2: **Experimental results on the Hard track of CDDB.**  $\dagger$  denotes that the method is based on the pre-trained ViT-B/16 model. The best result within rehearsal-free methods is indicated by **bold**, and the second is marked by underline.

## Main Results

**Baselines.** The current DIL methods can be broadly categorized into rehearsal-based and rehearsal-free methods. **Rehearsal-based** methods select and retrain a subset of images as exemplars of the domain when training. Representative methods include ER, LRCIL, iCaRL, and LUCIR. **Rehearsal-free** methods do not require saving images from the old domain. Representative methods include EWC, LwF, DyTox, L2P, S-Prompts, *etc.* Note that rehearsal-based methods often require storing thousands of images, ranging in size from 100MB to 3GB. In contrast, rehearsal-free methods may require only a small amount of learnable parameters, occupying 1-50MB of space. Therefore, rehearsal-free methods significantly outperform rehearsal-based methods in terms of storage space requirements. The proposed DualCP belongs to the rehearsal-free setting, so we prioritize comparing it with similar methods. Additionally, we also list the state-of-the-art rehearsal-based methods for reference.

**Comparison with State-of-the-arts.** We compare our approach with other state-of-the-art (SOTA) methods on three DIL benchmark datasets. The methods are grouped based on the number of images per class to be retained, with “0/class” indicating a rehearsal-free approach.

Table 1 illustrates the comparison results on the DomainNet dataset. Our DualCP surpasses the SOTA method C-Prompt (60.13% vs. 58.68%), even coming close to the performance of the DyTox method that utilizes rehearsal. Besides, our method achieved the best performance in preventing forgetting among all methods (-1.96%). Table 2 showcases results on CDDB, where our DualCP surpasses the best rehearsal-free method by a large margin (82.46% vs. 74.51%), and is comparable to LUCIR, which utilizes a substantial buffer (“100/class”). It falls just behind the rehearsal-based DyTox method. Additionally, our method

Method	Buffer( $\downarrow$ )	$A_T$ ( $\uparrow$ )
ER (Chaudhry et al. 2019)		80.10
GDumb (Prabhu, Torr, and Dokania 2020)		74.92
BiC (Wu et al. 2019)		79.28
DER++ (Buzzega et al. 2020)	50/class	79.70
Co <sup>2</sup> L (Cha, Lee, and Shin 2021)		79.75
DyTox (Douillard et al. 2022)		79.21
L2P (Wang et al. 2022b)		81.07
EWC (Kirkpatrick et al. 2017)		74.82
LwF (Li and Hoiem 2017)		75.45
L2P (Wang et al. 2022b)		78.33 $\dagger$
DualPrompt (Wang et al. 2022a)		80.25 $\dagger$
S-iP (Wang, Huang, and Hong 2022)	0/class	83.13 $\dagger$
CODA-P (Smith et al. 2023)		<u>85.68<math>\dagger</math></u>
C-Prompt (Liu, Peng, and Zhou 2024)		85.31 $\dagger$
DualCP (ours)		<b>88.10<math>\dagger</math></b>

Table 3: **Experimental results on the CORE50 dataset.** Note that  $F_T$  is not applicable to CORE50 because the training and test domains do not overlap.  $\dagger$  denotes that the method is based on the pre-trained ViT-B/16 model. The best result within rehearsal-free methods is indicated by **bold**, and the second is marked by underline.

Methods \ Datasets	DomainNet (345 classes)		CDDB (2 classes)		CORE50 (50 classes)
	$A_T$ ( $\uparrow$ )	$F_T$ ( $\uparrow$ )	$A_T$ ( $\uparrow$ )	$F_T$ ( $\uparrow$ )	$A_T$ ( $\uparrow$ )
VanillaCP	56.22	-2.80	<b>82.16</b>	<b>-0.73</b>	86.27
DualCP (proposed)	<b>60.13</b>	<b>-1.96</b>	<b>82.16*</b>	<b>-0.73*</b>	<b>88.10</b>

Table 4: **Comparison of different concept prototype designs in our framework.** \* represents that CDDB only contains two classes: real and fake, thus it cannot be divided into more groups and cannot apply to the DualCP.

outperforms the SOTA method in mitigating forgetting (-0.73% vs. -1.30%). Table 3 presents the comparisons on CORE50, revealing that our method achieves the best results in both rehearsal-based (88.10% vs. 81.07%) and rehearsal-free (88.10% vs. 85.68%) tracks.

## Ablation Study

**Effectiveness of Dual Concept Prototype.** We introduce two solutions to generate the concept prototypes: a single-level concept prototype (VanillaCP) and a dual-level concept prototype (DualCP). Table 4 demonstrates a consistent improvement of DualCP over VanillaCP, especially on datasets containing a large number of classes. This indicates that our DualCP contributes to distinguishing similar classes.

**Comparison on Different Backbones.** To validate the generality of our approach, we conducted experiments on both ViT and CLIP, a multimodal model based on contrastive learning with an image encoder and a text encoder. Table 5 presents the accuracy of our DualCP under different settings, as well as that of S-Prompts and MoP-CLIP.

**Image Feature Extractor.** We extract image features using two optional setting, *i.e.*, ViT-B/16 or CLIP image en-

Methods	Image Feature Extractor		CPG Guidance		DomainNet		CDDB		CORe50
	ViT-B/16	CLIP-Image	ViT-B/16-BD	CLIP-Text	$A_T(\uparrow)$	$F_T(\uparrow)$	$A_T(\uparrow)$	$F_T(\uparrow)$	$A_T(\uparrow)$
S-Prompts	✓				50.62	-2.85	74.51	-1.30	83.13
MoP-CLIP		✓		✓*	67.78	-1.64	88.65	-0.69	89.06
DualCP (ours)	✓		✓		60.13	-1.96	82.16	-0.73	88.10
	✓			✓	62.73	-1.81	83.05	-0.76	88.55
		✓	✓		69.31	-1.49	91.86	-0.35	89.98
		✓		✓	<b>72.46</b>	<b>-1.26</b>	<b>92.34</b>	<b>-0.32</b>	<b>90.59</b>

Table 5: **Comparison on different backbones.** ViT-B/16-BD means that image features extracted from the base domain by ViT-B/16 are used to guide the CPG module. \* indicates that S-Prompts and MoP-CLIP utilize CLIP text encoder but are unrelated to the CPG module. The proposed CPG module is only used for our DualCP.




Figure 3: **Ablation study of the C2F module on DomainNet dataset.**  $A_T$ ,  $F_T$  denotes the average accuracy and the forgetting degree, respectively. The hidden dimensions are set as multiples of the image feature dimensions of 768 in the ViT-B/16 backbone. “0.5x, 1x, 2x, 4x” correspond to hidden dimensions of 384, 768, 1536, and 3072, respectively.

coder, as the feature extractor ( $f_i(\cdot; \theta_i)$ ).

**CPG Guidance.** We proposed two methods to extract the concept prototype in the CPG module. One is based on the CLIP text encoder, as shown in Eq. (3). When CLIP is unavailable, we introduce an alternative approach using ViT. We use ViT-B/16 to extract image features from the base domain, and then calculate the mean features  $\mathbf{y}'_i$  for each class. We use  $\mathbf{Y}' = \{\mathbf{y}'_1, \mathbf{y}'_2, \dots, \mathbf{y}'_k\}$  to guide the construction of the concept prototype, as shown in Eq. (5).

**Ablations of C2F Design.** Our C2F comprises a coarse-grained layer  $g_C$  and multiple fine-grained layers  $g_{F_i}$ , which are implemented as a multi-layer perceptron (MLP). As depicted in Fig. 3, we conduct ablation experiments on the number of layers and the hidden dimension of our C2F to assess the influence of MLP on model performance.

### Visualization

We utilize t-SNE visualization to demonstrate the effectiveness of our DualCP. We choose six common classes (cat, flower, dog, boat, bicycle, and bus) from DomainNet for presentation. As shown in Fig. 4, (a) represents the image features extracted by the pre-trained ViT model. (b-d) depict the image features extracted by our DualCP. Our model categorizes similar classes into the same group based on semantics, such as cat and dog. (b) represents the coarse-grained features extracted by DualCP. (c) and (d) represent

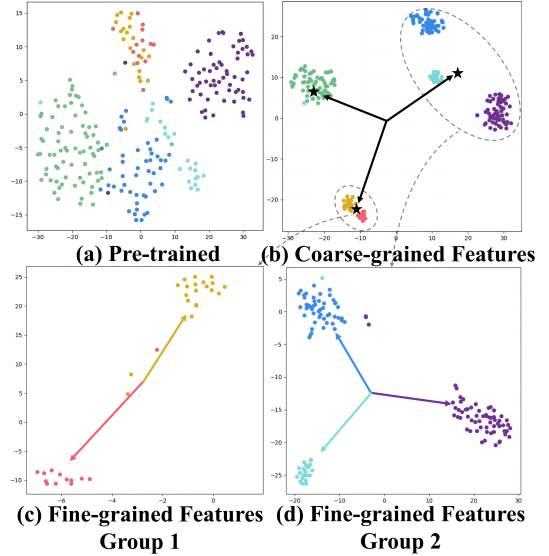


Figure 4: **t-SNE visualization** of feature space for common classes in DomainNet, with flower represented in green, cat and dog in (c), and boat, bicycle, and bus in (d). The pentagrams represent the average image features of a group.

fine-grained features of different groups. As illustrated in Fig. 4 (c)(d), our dual-level concept prototypes effectively distinguish similar classes.

### Conclusion and Future Works

This paper introduces a novel approach to RFDIL inspired by humans’ incremental cognitive processes. We proposed constructing dual-level concept prototypes (DualCP) for each class across domains to address the zero-sum game between learning new domains and preserving old ones. By aligning features from different domains to the same feature space, we avoid compromising the feature space of old domains while accommodating new domain features. Extensive experiments on three datasets with different backbones consistently show that our DualCP outperforms existing SOTA methods. Furthermore, our method is expected to be extended to other applications, such as domain-incremental object detection. It may also provide a reference for the development of generalized neural collapse.

## Acknowledgments

This work was funded by the National Natural Science Foundation of China under Grant No.U21B2048 and No.62302382, Shenzhen Key Technical Projects under Grant CJGJZD2022051714160501, China Postdoctoral Science Foundation No.2024M752584, and Natural Science Foundation of Shaanxi Province No.2024JC-YBQN-0637.

## References

- Akyürek, A. F.; Akyürek, E.; Wijaya, D. T.; and Andreas, J. 2021. Subspace regularizers for few-shot class incremental learning. *arXiv preprint arXiv:2110.07059*.
- Bar, M. 2003. A cortical mechanism for triggering top-down facilitation in visual object recognition. *Journal of cognitive neuroscience*, 15(4): 600–609.
- Barcelo, F.; Suwazono, S.; and Knight, R. T. 2000. Prefrontal modulation of visual processing in humans. *Nature neuroscience*, 3(4): 399–403.
- Buzzega, P.; Boschini, M.; Porrello, A.; Abati, D.; and Calderara, S. 2020. Dark experience for general continual learning: a strong, simple baseline. *Advances in neural information processing systems*, 33: 15920–15930.
- Cha, H.; Lee, J.; and Shin, J. 2021. Co2l: Contrastive continual learning. In *Proceedings of the IEEE/CVF International conference on computer vision*, 9516–9525.
- Chaudhry, A.; Rohrbach, M.; Elhoseiny, M.; Ajanthan, T.; Dokania, P. K.; Torr, P. H.; and Ranzato, M. 2019. On tiny episodic memories in continual learning. *arXiv preprint arXiv:1902.10486*.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- 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*.
- Douillard, A.; Ramé, A.; Couairon, G.; and Cord, M. 2022. Dytox: Transformers for continual learning with dynamic token expansion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9285–9295.
- Fenske, M. J.; Aminoff, E.; Gronau, N.; and Bar, M. 2006. Top-down facilitation of visual object recognition: object-based and context-based contributions. *Progress in brain research*, 155: 3–21.
- Gao, X.; Dong, S.; He, Y.; Wang, Q.; and Gong, Y. 2024. Beyond Prompt Learning: Continual Adapter for Efficient Rehearsal-Free Continual Learning. *arXiv preprint arXiv:2407.10281*.
- Grill, J.-B.; Strub, F.; Altché, F.; Tallec, C.; Richemond, P.; Buchatskaya, E.; Doersch, C.; Avila Pires, B.; Guo, Z.; Gheshlaghi Azar, M.; et al. 2020. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33: 21271–21284.
- Hou, S.; Pan, X.; Loy, C. C.; Wang, Z.; and Lin, D. 2019. Learning a unified classifier incrementally via rebalancing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 831–839.
- Huang, C.; Xie, L.; Yang, Y.; Wang, W.; Lin, B.; and Cai, D. 2023. Neural Collapse Inspired Federated Learning with Non-iid Data. *arXiv preprint arXiv:2303.16066*.
- Isele, D.; and Cosgun, A. 2018. Selective experience replay for lifelong learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Jee, K.; Ara., E.; and Zon., B. 2024. Gradual Divergence for Seamless Adaptation: A Novel Domain Incremental Learning Method. *arXiv preprint arXiv:2406.16231*.
- Ji, W.; Lu, Y.; Zhang, Y.; Deng, Z.; and Su, W. J. 2021. An unconstrained layer-peeled perspective on neural collapse. *arXiv preprint arXiv:2110.02796*.
- Jiang, J.; Zhou, J.; Wang, P.; Qu, Q.; Mixon, D.; You, C.; and Zhu, Z. 2023. Generalized neural collapse for a large number of classes. *arXiv preprint arXiv:2310.05351*.
- Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; and Krishnan, D. 2020. Supervised contrastive learning. *Advances in neural information processing systems*, 33: 18661–18673.
- Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N.; Veness, J.; Desjardins, G.; Rusu, A. A.; Milan, K.; Quan, J.; Ramalho, T.; Grabska-Barwinska, A.; et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13): 3521–3526.
- Li, C.; Huang, Z.; Paudel, D. P.; Wang, Y.; Shahbazi, M.; Hong, X.; and Van Gool, L. 2023. A continual deepfake detection benchmark: Dataset, methods, and essentials. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 1339–1349.
- Li, Z.; and Hoiem, D. 2017. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12): 2935–2947.
- Liu, Z.; Peng, Y.; and Zhou, J. 2024. Compositional Prompting for Anti-Forgetting in Domain Incremental Learning. *International Journal of Computer Vision*, 1–18.
- Lomonaco, V.; and Maltoni, D. 2017. Core50: a new dataset and benchmark for continuous object recognition. In *Conference on robot learning*, 17–26. PMLR.
- Marra, F.; Saltori, C.; Boato, G.; and Verdoliva, L. 2019. Incremental learning for the detection and classification of gan-generated images. In *2019 IEEE international workshop on information forensics and security (WIFS)*, 1–6. IEEE.
- 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.
- Mixon, D. G.; Parshall, H.; and Pi, J. 2020. Neural collapse with unconstrained features. *arXiv preprint arXiv:2011.11619*.
- Papayan, V.; Han, X.; and Donoho, D. L. 2020. Prevalence of neural collapse during the terminal phase of deep learning

- training. *Proceedings of the National Academy of Sciences*, 117(40): 24652–24663.
- Pellegrini, L.; Graffieti, G.; Lomonaco, V.; and Maltoni, D. 2020. Latent replay for real-time continual learning. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 10203–10209. IEEE.
- 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.
- Prabhu, A.; Torr, P. H.; and Dokania, P. K. 2020. Gdumb: A simple approach that questions our progress in continual learning. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, 524–540. Springer.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Rebuffi, S.-A.; Kolesnikov, A.; Sperl, G.; and Lampert, C. H. 2017. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2001–2010.
- Shi, Y.; Shi, D.; Qiao, Z.; Wang, Z.; Zhang, Y.; Yang, S.; and Qiu, C. 2023. Multi-granularity knowledge distillation and prototype consistency regularization for class-incremental learning. *Neural Networks*, 164: 617–630.
- 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.
- Van De Ven, G. M.; Li, Z.; and Tolias, A. S. 2021. Class-incremental learning with generative classifiers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3611–3620.
- Verwimp, E.; Yang, K.; Parisot, S.; Hong, L.; McDonagh, S.; Pérez-Pellitero, E.; De Lange, M.; and Tuytelaars, T. 2023. Clad: A realistic continual learning benchmark for autonomous driving. *Neural Networks*, 161: 659–669.
- Wan, C.; He, Y.; Song, X.; and Gong, Y. 2024a. Prompt-Agnostic Adversarial Perturbation for Customized Diffusion Models. *arXiv preprint arXiv:2408.10571*.
- Wan, C.; Luo, X.; Cai, Z.; Song, Y.; Zhao, Y.; Bai, Y.; He, Y.; and Gong, Y. 2024b. GRID: Visual Layout Generation. *arXiv preprint arXiv:2412.10718*.
- Wang, Q.; He, Y.; Dong, S.; Gao, X.; Wang, S.; and Gong, Y. 2025. Non-exemplar Domain Incremental Learning via Cross-Domain Concept Integration. In *European Conference on Computer Vision*, 144–162. Springer.
- Wang, S.; Shi, W.; He, Y.; Yu, Y.; and Gong, Y. 2023. Non-Exemplar Class-Incremental Learning via Adaptive Old Class Reconstruction. In *Proceedings of the 31st ACM International Conference on Multimedia*, 4524–4534.
- Wang, S.; Yu, Y.; He, Y.; and Gong, Y. 2024. Enhancing Pre-trained ViTs for Downstream Task Adaptation: A Locality-Aware Prompt Learning Method. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 797–806.
- 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, 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.
- Yang, Y.; Chen, S.; Li, X.; Xie, L.; Lin, Z.; and Tao, D. 2022. Inducing Neural Collapse in Imbalanced Learning: Do We Really Need a Learnable Classifier at the End of Deep Neural Network? *Advances in Neural Information Processing Systems*, 35: 37991–38002.
- Yang, Y.; Yuan, H.; Li, X.; Lin, Z.; Torr, P.; and Tao, D. 2023. Neural collapse inspired feature-classifier alignment for few-shot class incremental learning. *arXiv preprint arXiv:2302.03004*.
- Zbontar, J.; Jing, L.; Misra, I.; LeCun, Y.; and Deny, S. 2021. Barlow twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning*, 12310–12320. PMLR.
- Zhang, H.; and Mueller, F. 2022. CLAIRE: Enabling Continual Learning for Real-time Autonomous Driving with a Dual-head Architecture. In *2022 IEEE 25th International Symposium On Real-Time Distributed Computing (ISORC)*, 1–10. IEEE.
- Zhao, H.; Wang, H.; Fu, Y.; Wu, F.; and Li, X. 2021. Memory-efficient class-incremental learning for image classification. *IEEE Transactions on Neural Networks and Learning Systems*, 33(10): 5966–5977.
- Zhou, J.; Li, X.; Ding, T.; You, C.; Qu, Q.; and Zhu, Z. 2022. On the optimization landscape of neural collapse under mse loss: Global optimality with unconstrained features. In *International Conference on Machine Learning*, 27179–27202. PMLR.
- Zhu, F.; Zhang, X.-Y.; Wang, C.; Yin, F.; and Liu, C.-L. 2021. Prototype augmentation and self-supervision for incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5871–5880.