



**HAL**  
open science

# Memory-efficient Continual Learning with Neural Collapse Contrastive

Trung-Anh Dang, Vincent Nguyen, Ngoc-Son Vu, Christel Vrain

► **To cite this version:**

Trung-Anh Dang, Vincent Nguyen, Ngoc-Son Vu, Christel Vrain. Memory-efficient Continual Learning with Neural Collapse Contrastive. Winter Conference on Applications of Computer Vision (WACV), IEEE; CVF, Feb 2025, Tucson, United States. <hal-05075376>

**HAL Id: hal-05075376**

**<https://hal.science/hal-05075376v1>**

Submitted on 20 May 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

# Memory-efficient Continual Learning with Neural Collapse Contrastive

Trung-Anh Dang<sup>1</sup>, Vincent Nguyen<sup>1</sup>, Ngoc-Son Vu<sup>2</sup>, Christel Vrain<sup>1</sup>,

<sup>1</sup>Université d’Orléans, INSA CVL, LIFO UR 4022, Orléans, France

<sup>2</sup>ETIS - CY Cergy Paris University, ENSEA, CNRS, France

{trung-anh.dang, vincent.nguyen, christel.vrain}@univ-orleans.fr, son.vu@ensea.fr

## Abstract

*Contrastive learning has significantly improved representation quality, enhancing knowledge transfer across tasks in continual learning (CL). However, catastrophic forgetting remains a key challenge, as contrastive based methods primarily focus on “soft relationships” or “softness” between samples, which shift with changing data distributions and lead to representation overlap across tasks. Recently, the newly identified Neural Collapse phenomenon has shown promise in CL by focusing on “hard relationships” or “hardness” between samples and fixed prototypes. However, this approach overlooks “softness”, crucial for capturing intra-class variability, and this rigid focus can also pull old class representations toward current ones, increasing forgetting. Building on these insights, we propose Focal Neural Collapse Contrastive (FNC<sup>2</sup>), a novel representation learning loss that effectively balances both soft and hard relationships. Additionally, we introduce the Hardness-Softness Distillation (HSD) loss to progressively preserve the knowledge gained from these relationships across tasks. Our method outperforms state-of-the-art approaches, particularly in minimizing memory reliance. Remarkably, even without the use of memory, our approach rivals rehearsal-based methods, offering a compelling solution for data privacy concerns.*

## 1. Introduction

Unlike human learning, which is adaptive and ongoing, current deep neural networks (DNNs) are typically trained in batches, using all available data at once. To emulate human learning, these models need to support continual learning (CL), which involves acquiring new knowledge while retaining previously learned information. However, this process is hindered by the phenomenon known as “catastrophic forgetting” (CF) [36, 44], where learning new tasks with fresh data causes the model to overwrite its prior knowledge, leading to a drastic decline in performance on older tasks. To tackle this challenge, CL has gained sig-

nificant attention in recent years [5, 27, 37, 41, 45, 48, 53]. CL aims to develop methods that enable models to learn from a continuous stream of data by balancing the retention of prior knowledge with the ability to adapt to new information. Achieving this balance, known as the stability-plasticity trade-off, is crucial for preventing performance loss when integrating new tasks.

Current CL methods are predominantly based on supervised strategies, which have proven effective in preserving knowledge [3, 50, 57]. Notably, approaches that decouple representation learning from classifier training have shown greater resistance to forgetting compared to joint training methods [5, 12, 34, 53]. Building on this decoupling, several supervised contrastive representation learning methods have delivered strong results [5, 53]. However, ① most of these methods depend on replay buffers for storing past samples, limiting their use when data privacy is a concern. ② Another limitation is their reliance on inter-sample relationships, which can lead to **representation drift and overlap** with new tasks—one of the main causes of forgetting in CL. Indeed, this issue has been highlighted in both DNNs [4, 58] and neuroscience [10].

Neural collapse (NC), a recently discovered phenomenon characterized by highly structured and aligned neural network features, has attracted considerable attention in the deep learning community [11, 15, 20, 33, 49, 63]. It shows particular promise for CL by reducing representation overlap, enhancing class separation, and mitigating CF. NC achieves this by aligning feature representations with fixed prototypes, which act as optimal class-specific reference points and remain constant throughout training, significantly enhancing class separation. Additionally, prototypes can serve as class representatives, and their integration into contrastive learning reduces reliance on memory buffers. Leveraging these advantages, several CL methods [46, 55, 56] have incorporated NC. However, focusing only on sample-prototype relationships can reduce diversity, disrupt within-class data distribution, and lead to forgetting as older representations shift toward current task prototypes. To address this, we propose a supervised method using

fixed, evenly spaced prototypes to enhance representation quality, minimize task overlap, and reduce memory dependence while preserving intra-class data distribution.

To summarize, the main contributions of our work are:

- We reveal that “softness” (inter-sample relationships) and “hardness” (relationships with fixed prototypes)<sup>1</sup> are critical in NC-based contrastive CL for effectively acquiring new tasks while retaining prior knowledge.
- In both plasticity and stability phases of learning, we introduce new loss functions that address both hardness and softness. For plasticity, we propose Focal Neural Collapse Contrastive (FNC<sup>2</sup>), a loss function that combines hard and soft semantics to enhance representation learning by focusing on challenging samples. For stability, we introduce the Hardness-Softness Distillation (HSD) loss function, which preserves knowledge from both hard and soft relationships, significantly reducing forgetting.
- Our model surpasses state-of-the-art (SoTA) results in both replay-based and memory-free scenarios, especially excelling in settings with no stored exemplars, making it ideal for applications with strict data privacy requirements.

## 2. Related Work

### 2.1. Continual Learning

CL approaches can be broadly classified into three main categories. Rehearsal-based approaches [3, 5, 34, 53] store a small amount of data in a memory buffer and replay them to prevent forgetting. Regularization-based methods [12, 21, 39, 48, 60] penalize changes in network parameters of the current task with respect to the previous task. Meanwhile, instead of using shared parameters, the architectural-based approaches [28, 45, 57] construct task-specific parameters and allow network expansion during CL. This work focuses on regularization-based methods by devising a specific regularization loss to align the current model with the previous one. Additionally, our approach is capable of performing well in both replay-based and memory-free scenarios.

Most of the current CL methods target primarily to achieve a balance between acquiring new tasks (**plasticity**) and retaining knowledge of previous tasks (**stability**) [5, 12, 47, 53]. To attain this balance, regularization-based approaches typically employ knowledge distillation (KD)

<sup>1</sup>In unsupervised contrastive-based learning, *alignment* ensures similar features are mapped to similar samples, while *uniformity* ensures a feature distribution that maximizes information and ideally forms a uniform hypersphere. In a supervised setting, we use the term “softness” to represent alignment, and introduce “hardness” to describe the relationships between samples and NC-based prototypes.

[18], which aims to transfer knowledge from the previous trained model (teacher) to the current one (student). Recent works based on KD directly minimize the divergence between their intermediate representations [19] or final outputs [6, 12]. Additionally, several methods have explored relational KD, which enhances knowledge retention by leveraging the sample-to-sample relationships, such as Instance-wise Relation Distillation (IRD) [5, 53]. Other approaches, such as [2, 26], focus on relationships between *learnable* class prototypes and individual samples, rather than between samples themselves. In this work, we utilize KD through IRD, alongside a prototype-based distillation method. Our approach differs from [2, 26] in both the use of *fixed* rather than learnable prototypes, and in how the sample-prototype relationships are applied.

Beyond balancing plasticity and stability, recent research emphasizes the importance of **cross-task** consolidation for improving representation and reducing forgetting [29, 61]. We *implicitly* address this by designing a plasticity loss function crafted to enhance cross-task separability.

### 2.2. Contrastive Learning

Contrastive learning has emerged as a prominent representation learning approach, demonstrating its SoTA for different downstream tasks [8, 22, 51]. Numerous contrastive methods have been proposed and widely applied in both unsupervised [8, 9, 16, 51] and supervised settings [22].

In the context of CL, many studies [5, 12, 26, 35, 53] have shown that contrastive learning is highly effective in acquiring task-invariant representations, which significantly mitigates the primary cause of forgetting — data imbalances between previous and current tasks. Among these, Co<sup>2</sup>L [5] is the first method to apply supervised contrastive learning in CL. Subsequently, CILA [53] improved Co<sup>2</sup>L by analyzing the importance of coefficients for the distillation loss. Recently, CCLIS [26] emerged as a SoTA method by preserving knowledge through importance sampling to recover previous data distributions. However, unlike Co<sup>2</sup>L and CILA, CCLIS cannot operate without a memory buffer, which limits its efficiency in many real-world applications where data privacy is a concern.

Contrastive learning generates augmented views of each sample, bringing positive pairs closer and pushing negative pairs apart, promoting representation invariance to augmentations. Most contrastive methods focus on learning representations through relationships between samples, referred to as soft relationships or “softness”. In CL, relying solely on softness preserves class diversity and data distribution but can cause task representation overlap, as shown in Fig. 1a. In this paper, we propose a novel supervised contrastive loss for CL that reduces memory dependence and resolves this overlap while maintaining class distribution.

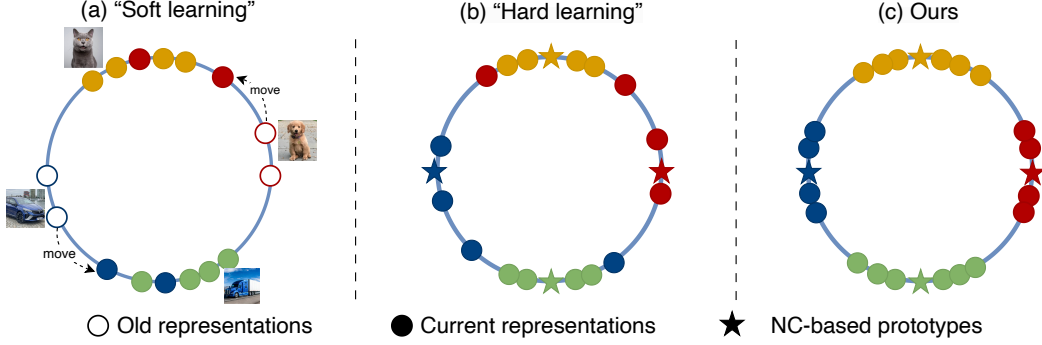


Figure 1. Our method mitigates the drawbacks of both “soft” and “hard” learning by using fixed, equidistant prototypes to minimize cluster overlap. (a) Overlap between new and old classes happens due to non-fixed class clusters and representation drift. (b) Strong alignment with prototypes can overly cluster current task representations and pull older ones into new clusters, especially in CL with few old samples. Moreover, mixed-feature samples should be placed between classes, not tightly aligned with prototypes. (c) Our method considers both inter-sample and sample-prototype relationships, maintaining distinct cluster representations and preserving distribution within each cluster.

### 2.3. Neural Collapse (NC)

A recent study [38] identified neural collapse (NC) phenomenon, where, at the end of training on a balanced dataset, class features collapse to their class means, aligned with a simplex equiangular tight frame (ETF). This finding has led to further research showing that NC represents global optimality in balanced training with cross-entropy [11, 14, 20, 33, 64] and mean squared error [15, 40, 49, 63] loss functions. Inspired by NC, studies such as [24] have used fixed simplex ETF points with modified contrastive loss to achieve NC, while [54] induced NC under imbalanced data conditions by fixing the classifier.

**Inducing NC for CL.** Building on [13], which showed that NC persists when transferring models to new samples or classes, several CL studies have leveraged NC to mitigate forgetting [55, 56]. These works pre-assign a group of classifier prototypes as a simplex ETF for all tasks and then align sample representations to their corresponding prototypes. The relationship between samples and prototypes, which this learning approach focuses on, is referred to as hard relationships, or “hardness”. For instance, [55, 56] employed the dot-regression (DR) loss proposed in [54] for NC-based CL, which is considered a hard learning method. Since the prototype group remains consistent, these methods prevent task overlap but risk reducing diversity and disrupting within-class distribution, potentially leading to forgetting. Our method addresses this issue by integrating NC directly into the loss function, combining both softness and hardness to preserve data distribution and avoid task overlap in CL.

The concept of NC can be presented as follows.

**Definition 1.** A simplex Equiangular Tight Frame (ETF) is a collection of  $K$  vectors:  $\mathbf{Q} = \{\mathbf{q}_k\}_{k=1}^K$ , each vector  $\mathbf{q}_k \in \mathbb{R}^d$ ,  $K \leq d + 1$ , which satisfies:

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

where  $\mathbf{U} \in \mathbb{R}^{d \times K}$  is an orthogonal basis and  $\mathbf{U}^T \mathbf{U} = \mathbf{I}_K$ ,  $\mathbf{I}_K$  is an identity matrix and  $\mathbf{1}_K$  is an all-ones vector.

Each vector  $\mathbf{q}_k$  has the same  $\ell_2$  norm, and any two distinct vectors consistently produce an inner product of  $-\frac{1}{K-1}$ , which is the lowest possible cosine similarity for  $K$  equiangular vectors in  $\mathbb{R}^d$ . This geometric relationship can be described as

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

where  $\delta_{i,j} = 1$  in case of  $i = j$ , and 0 otherwise.

After that, the NC phenomenon can be formally characterized by the following four attributes [38]:

NC1: Features from the last layer within the same class converge to their intra-class mean, such that the covariance  $\Sigma_V^k \rightarrow \mathbf{0}$ . Here,  $\Sigma_V^k = \text{Avg}_i \{(\mathbf{v}_{k,i} - \boldsymbol{\mu}_k)(\mathbf{v}_{k,i} - \boldsymbol{\mu}_k)^T\}$ , where  $\mathbf{v}_{k,i}$  is the feature of sample  $i$  in class  $k$ , and  $\boldsymbol{\mu}_k$  is the intra-class feature mean.

NC2: After centering by the global mean, intra-class means align with simplex ETF vertices, i.e.,  $\{\tilde{\boldsymbol{\mu}}_k\}$ ,  $1 \leq k \leq K$  satisfy Eq. (2), where  $\tilde{\boldsymbol{\mu}}_k = (\boldsymbol{\mu}_k - \boldsymbol{\mu}_G) / \|\boldsymbol{\mu}_k - \boldsymbol{\mu}_G\|$  and global mean  $\boldsymbol{\mu}_G = \frac{1}{K} \sum_{k=1}^K \boldsymbol{\mu}_k$ ;

NC3: Intra-class means centered by the global mean align with their classifier weights, leading to the same simplex ETF, i.e.,  $\tilde{\boldsymbol{\mu}}_k = \mathbf{w}_k / \|\mathbf{w}_k\|$ , where  $1 \leq k \leq K$  and  $\mathbf{w}_k$  is the classifier weight of class  $k$ ;

NC4: When NC1-NC3 hold, model predictions simplify to selecting the nearest class center, represented as  $\text{argmax}_k \langle \mathbf{z}, \mathbf{w}_k \rangle = \text{argmin}_k \|\mathbf{z} - \boldsymbol{\mu}_k\|$ , where  $\langle \cdot, \cdot \rangle$  denotes the inner product operator, and  $\mathbf{z}$  is the output of the model.

### 3. Preliminaries

#### 3.1. Problem Setup

In the general supervised CL scenario, we have a sequence of training datasets, which is drawn from non-stationary data distributions for each task. Namely, let  $t$  be the task index, where  $t \in \{1, \dots, T\}$ , and  $T$  represents the maximum number of tasks. The dataset for the  $t$ -th task, denoted by  $\mathcal{D}_t$ , consists of  $N_t$  supervised pairs:  $\mathcal{D}_t = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N_t}$ , along with the set of classes  $\mathcal{C}_t$ . CL comprises a variety of scenarios; however, in this work, we focus specifically on two popular settings: class-incremental learning (Class-IL) and task-incremental learning (Task-IL). In both settings, there is no overlap in class labels across tasks, ensuring  $\mathcal{C}_t \cap \mathcal{C}_{t'} = \emptyset$  for any two distinct tasks  $t \neq t'$ . For Task-IL, the learned model additionally has access to the task label during the testing phase.

#### 3.2. Supervised Contrastive Learning

This section details the SupCon algorithm [22], which inspires the “softness” component of our approach. Suppose that in each batch  $B$  of  $N$  samples, SupCon firstly creates two randomly augmented versions of each sample in the batch, making each batch now contain  $2N$  views:  $|B| = 2N$ . After that, given the feature extractor  $f$ , each view  $\mathbf{x}_i$  in the batch is mapped into a unit d-dimensional Euclidean sphere through a linear projector  $g$  as  $\mathbf{z}_i = h(\mathbf{x}_i)$ , where  $h = g \circ f$ . Consequently, generic presentations are learned through the minimization of the following loss:

$$\mathcal{L}_{SupCon} = \sum_{i=1}^{2N} \frac{-1}{|P(i)|} \sum_{j \in P(i)} \log\left(\frac{e^{\langle \mathbf{z}_i, \mathbf{z}_j \rangle / \tau}}{\sum_{k \in A(i)} e^{\langle \mathbf{z}_i, \mathbf{z}_k \rangle / \tau}}\right) \quad (3)$$

where  $\langle \cdot \rangle$  is the cosine similarity,  $\tau > 0$  is the temperature factor,  $A(i) = \{1..2N\} \setminus \{i\}$ , and  $P(i)$  is the index set of positive views with the anchor  $\mathbf{x}_i$ , denoted as:

$$P(i) = \{p \in \{1..2N\} | y_p = y_i, p \neq i\} \quad (4)$$

**Focal contrastive learning.** Despite their advantages, contrastive learning methods often struggle to reduce intra-class feature dispersion. They rely heavily on positive/negative pairs, but most samples are easy to contrast, resulting in minimal loss gradients and scattered intra-class samples. To address this issue, inspired by focal loss [30], [62] introduced the focal contrastive loss. This approach emphasizes hard positive views—those with low cosine similarity to the anchor and thus lower prediction probability. Since hard positives are more influential in contrastive loss, they lead to clearer class clustering. Building on these insights, we propose focal contrastive loss in the context of NC for CL, as detailed in Sec. 4.3.

### 4. Methodology

#### 4.1. Motivation

Our method uses a two-stage learning process, as in [5]. First, we learn a representation, which is then used to train the classifier. The main objective of CL is to balance two goals: learning new tasks (plasticity) and preserving the knowledge from previous tasks (stability). The overall loss can be described as:

$$\mathcal{L}_{overall} = \mathcal{L}_{plasticity} + \mathcal{L}_{stability} \quad (5)$$

**Remark.** While plasticity and stability are central to most CL algorithms, some very recent works have highlighted the importance of cross-task consolidation [29, 61]. Although our approach does not explicitly incorporate a cross-task consolidation term, we will show that it is implicitly accounted for in our  $\mathcal{L}_{plasticity}$ .

**a. Plasticity.** Previous contrastive learning approaches, like SupCon [22], use sample relationships but allow class representations to shift, as shown in Fig. 1a, causing overlap with current clusters and leading to forgetting. We term these methods “soft plasticity”. To address this issue, inspired by Neural Collapse, recent works [55, 56] propose using fixed, equidistant prototypes as the optimal class means at the end of training, as shown in Fig. 1b. These methods, which we term “hard plasticity”, focus on aligning sample representations with their assigned prototypes.

While these “hard” methods tightly align representations with prototypes, they have drawbacks. First, they neglect sample relationships, leading to representations clustered only around prototypes, which is not ideal since some samples share characteristics with multiple classes and should lie between them. Second, tightly aligning samples with current prototypes can pull old class representations towards the current task’s prototypes, causing forgetting, as shown in Fig. 1b. To harness the benefits of NC while avoiding the pitfalls of hard learning, we integrate both soft and hard learning. This approach maintains cluster distribution within each class and preserves separation from clusters of both current and past classes, as shown in Fig. 1c.

**b. Stability.** Co<sup>2</sup>L, the pioneering work in continual contrastive learning, employs IRD as  $\mathcal{L}_{stability}$ . The “soft stability” IRD is designed to preserve the relationships between samples in the old and new feature spaces. However, as training progresses, its effectiveness diminishes, as demonstrated in Fig. 2:

1. At the start of training task  $t$ , the representations of all current samples match those from task  $t - 1$  because no training has occurred yet. Initially, sample representations are consistent across tasks, but as training progresses, they gradually shift towards their respective class clusters.

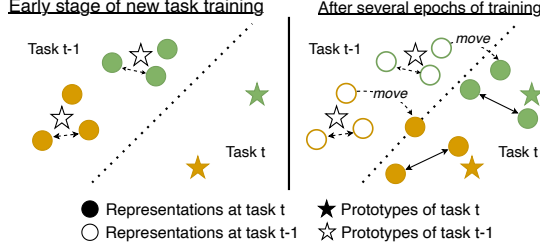


Figure 2. Shifts in representations of current samples at task  $t$  between the beginning of training and after several epochs.

- Over many epochs, the increasing discrepancy between the current and past feature spaces leads to a rise in IRD loss. This shift occurs because representations approach optimal positions, reducing the significance of  $\mathcal{L}_{stability}$  based on IRD to the overall loss  $\mathcal{L}_{overall}$  over time.

To address this limitation, we introduce a new distillation loss, termed “hard stability”, inspired by NC phenomenon, along with a novel strategy to enhance IRD’s effectiveness throughout the training process. Both of these will be detailed in Sec. 4.4. As presented in Fig. 1, both soft and hard learning methods suffer from representation drift, especially in settings with limited or no memory. Our approach effectively mitigates this issue. However, as the memory capacity increases, the impact of representation drift in soft and hard learning methods diminishes. Consequently, our method demonstrates a significant advantage in low-memory regimes, where the problem of representation drift is more pronounced.

## 4.2. Overview of the proposed method

Fig. 3 provides an overview of our method (see Sec. 3.1 for setting details). In our approach, memory is optional, and when used, we employ the Reservoir sampling strategy [52] to fill a fixed-size buffer  $\mathcal{M}$ . After each task beyond the first, current samples are combined with buffered samples, and each sample is drawn independently with equal probability for the mini-batch. We first predefine a set of fixed equidistant prototypes as the vertices of an ETF. We denote this prototype set as  $\mathbf{P} = \{\mathbf{p}_i\}_{i=1}^K$ ,  $K$  is the number of prototypes, corresponding to the number of classes. These prototypes are used as equidistant optimal points in the feature space. We utilize these prototype points in both learning new tasks and distilling old tasks through their direct use in the corresponding loss functions.

The overall objective of our method is:

$$\mathcal{L} = \mathcal{L}_{FNC^2} + \mathcal{L}_{HSD} \quad (6)$$

with details on each new loss function provided in the following sections.

## 4.3. Hardness-Softness Plasticity

Inspired by NC phenomenon [38] and the concept of focal loss for addressing batch imbalance during training, we introduce a novel loss called Focal Neural Collapse Contrastive (FNC<sup>2</sup>), defined as:

$$\mathcal{L}_{FNC^2} = - \sum_{i=1}^{2N} \frac{1}{|P(i)| + 1} \left( \sum_{\mathbf{z}_j \in P(i)} (1 - c_{ij})^\gamma \log(c_{ij}) + (1 - r_i)^\gamma \log(r_i) \right) \quad (7)$$

where

$$c_{ij} = \frac{e^{\langle \mathbf{z}_i, \mathbf{z}_j \rangle / \tau}}{\sum_{k \neq i} e^{\langle \mathbf{z}_i, \mathbf{z}_k \rangle / \tau} + \underbrace{\sum_{\mathbf{p}_l \in \mathbf{P}_{1:t-1}} e^{\langle \mathbf{z}_i, \mathbf{p}_l \rangle / \tau}}_{\text{cross-task}}} \quad (8)$$

$$r_i = \frac{e^{\langle \mathbf{z}_i, \mathbf{p}_{z_i} \rangle / \tau}}{\sum_{k \neq i} e^{\langle \mathbf{z}_i, \mathbf{z}_k \rangle / \tau} + \sum_{\mathbf{p}_l \in \mathbf{P}_{1:t-1}} e^{\langle \mathbf{z}_i, \mathbf{p}_l \rangle / \tau}} \quad (9)$$

Here,  $P(i)$  is the set of positive indexes for each anchor, as defined in Eq. (4),  $|P(i)|$  represents its cardinality, and  $\gamma$  is the focusing hyperparameter. Additionally,  $\mathbf{p}_{z_i}$  is the prototype (specifically, the ETF vertex corresponding to label  $y_i$ ), and  $\mathbf{P}_{1:t-1}$  represents the set of prototypes used in all previous tasks.

Intuitively, for each sample in the current task, we pull positive samples closer to the anchor and push negative samples away, forming clusters that are then pulled towards their optimal prototypes, as shown in Eq. (8) and Eq. (9). This approach helps the model learn both hardness and softness information. By incorporating  $(1 - c_{ij})^\gamma$  and  $(1 - r_i)^\gamma$ ,  $\mathcal{L}_{FNC^2}$  emphasizes hard samples - those that are positive but far from the anchor or far from their prototype. These hard samples are crucial in contrastive learning, as they significantly affect intra-class sample distribution, unlike easy-to-contrast samples. A larger  $\gamma$  further focuses the model on training from samples that are distant from their positive views and prototypes. To reduce reliance on memory, we use prototypes from previous tasks as representative points for past samples and include their cosine similarity with the anchor in  $c_{ij}$  and  $r_i$ . This allows our method to use old prototypes as negative points, similar to **cross-task consolidation** as described in [29, 61], aiding in distinguishing between current and old classes. This approach effectively mimics stored samples and facilitates robust learning even without past task samples, serving as a form of **pseudo-replay** where old prototypes are effectively reintroduced. We will present performance in both memory-based and memory-free settings in Sec. 5.2.

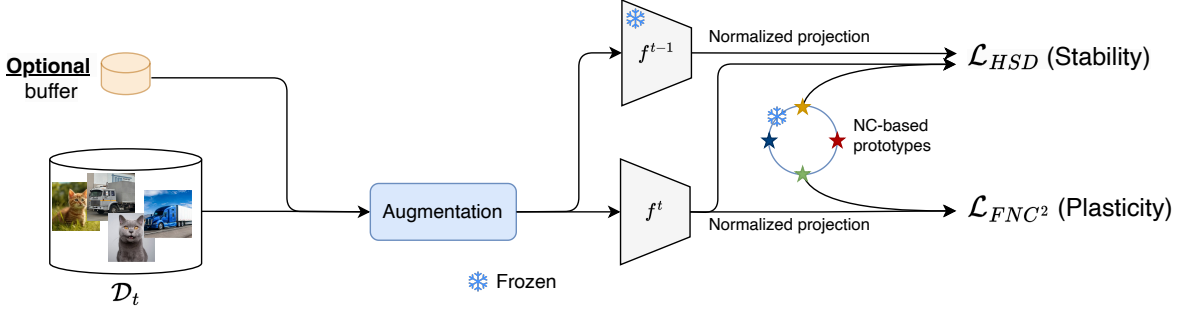


Figure 3. **Overall architecture of our method.** Augmented samples from each batch are fed into the current model  $f^t$  to learn new knowledge via  $\mathcal{L}_{FNC^2}$  and the frozen previous model  $f^{t-1}$  for distillation using  $\mathcal{L}_{HSD}$ . The buffer is optional, and NC-based prototypes are directly involved in both loss functions during training.

When using memory, we employ an asymmetric version of this loss, where only current samples serve as anchors and buffer samples are used solely as negative points.

#### 4.4. Hardness-Softness Distillation

Our new distillation loss, which capitalizes on both hardness and softness plasticity during training, is described as:

$$\begin{aligned} \mathcal{L}_{HSD} = & (1 - \alpha) \underbrace{\sum_{i=1}^{2N} -\mathbf{o}^{t-1}(\mathbf{x}_i) \cdot \log(\mathbf{o}^t(\mathbf{x}_i))}_{\mathcal{L}_{IRD}} \\ & + \alpha \underbrace{\sum_{i=1}^{2N} -\mathbf{q}^{t-1}(\mathbf{x}_i; \mathbf{P}_{1:t}) \cdot \log(\mathbf{q}^t(\mathbf{x}_i; \mathbf{P}_{1:t}))}_{\mathcal{L}_{S-PRD}} \end{aligned} \quad (10)$$

with  $\alpha = \max(0, \frac{e-e_0}{E})$ ,  $e$  is the epoch index,  $e_0$  is the number of epochs used for the warm-up period, and  $E$  is the number of epochs (details are provided in the Appendix). Here,  $\mathcal{L}_{IRD}$  (used in [5]) represents ‘‘soft stability’’ while the proposed Sample-Prototype Relation Distillation loss  $\mathcal{L}_{S-PRD}$  embodies ‘‘hard stability’’.

In  $\mathcal{L}_{IRD}$ ,  $\mathbf{o}^t(\mathbf{x}_i) = [o_{i,1}^t, \dots, o_{i,i-1}^t, o_i^t, o_{i,i+1}^t, \dots, o_{i,2N}^t]$ , each  $o_{i,j}^t$  is:  $o_{i,j}^t = \frac{e^{\langle \mathbf{z}_i^t, \mathbf{z}_j^t \rangle / \kappa}}{\sum_{k \neq i} e^{\langle \mathbf{z}_i^t, \mathbf{z}_k^t \rangle / \kappa}}$ , where  $t \geq 1$  is the task index and  $\kappa$  is the temperature hyperparameter. In  $\mathcal{L}_{S-PRD}$ ,  $\mathbf{P}_{1:t} = \{\mathbf{p}_s\}_{s=1}^S$  is the set of prototypes used from task 1 to task  $t$ . Besides, we have  $\mathbf{q}^t(\mathbf{x}_i; \mathbf{P}_{1:t}) = [q_{i,1}^t, q_{i,2}^t, \dots, q_{i,S}^t]$ , and each  $q_{i,j}^t$  is computed as:  $q_{i,j}^t = \frac{e^{\langle \mathbf{z}_i^t, \mathbf{p}_j \rangle / \zeta}}{\sum_{s=1}^S e^{\langle \mathbf{z}_i^t, \mathbf{p}_s \rangle / \zeta}}$ , given the temperature factor  $\zeta$ .

$\mathcal{L}_{S-PRD}$  preserves hardness by maintaining sample-prototype relationships, while  $\mathcal{L}_{IRD}$  focuses on softness, capturing the evolving relationships between samples during learning. While  $\mathcal{L}_{S-PRD}$  strongly preserves old knowledge by regularizing sample-prototype relationships, it may

hinder learning if used exclusively, as it restricts the adaptation of representations. Early in training,  $\mathcal{L}_{IRD}$  is crucial as it supports learning plasticity, while  $\mathcal{L}_{S-PRD}$  is less effective. As training progresses and representations approach their prototypes,  $\mathcal{L}_{S-PRD}$  becomes more beneficial, enhancing distillation. That is the motivation behind the idea of combining  $\mathcal{L}_{IRD}$  and  $\mathcal{L}_{S-PRD}$  in  $\mathcal{L}_{HSD}$  (Hardness-Softness Distillation) with weights that shift over time. After an initial warm-up period ( $e_0$  epochs), we gradually decrease  $\mathcal{L}_{IRD}$  and increase  $\mathcal{L}_{S-PRD}$  to balance learning and information preservation effectively.

## 5. Experiments

### 5.1. Datasets and Implementation details

**Datasets.** As in other continual contrastive learning works, we use three datasets: Seq-Cifar-10, Seq-Cifar-100, and Seq-Tiny-ImageNet for all experiments. We consider the class-IL and task-IL settings, described in Sec. 3.1. Seq-Cifar-10 is created from Cifar-10 [23] and divided into five tasks with two classes for each task, Seq-Cifar-100 contains five tasks (20 classes/task) built from Cifar-100 [23], and Seq-Tiny-ImageNet has 10 tasks (20 classes/task) built from Tiny-ImageNet [25].

**Implementation details.** We train our method using ResNet-18 [17] as the backbone, and we remove the last layer of the backbone as [5, 12]. Similar to previous works [5, 8, 59], we add a two-layer projection MLP on top of the backbone, followed by ReLU activation functions to map the output of the backbone into  $d$ -dimension embedding space, where  $d = 128$  with Seq-Cifar-10 and Seq-Cifar-100,  $d = 256$  with Seq-Tiny-ImageNet. With a batch size of 512, we train the backbone with 500 epochs for the initial task and 100 epochs for other tasks as in [5] for all datasets. Besides, we run with buffer sizes 0, 200 and 500 to evaluate the performance of the model with different memory settings. Regarding the value of  $\gamma$  in  $\mathcal{L}_{FNC^2}$ , we choose  $\gamma = 1$  for Seq-Cifar-10, and  $\gamma = 4$  for all other datasets.

Details about the selection of  $\gamma$  and other hyperparameters are provided in the Appendix. Since Co<sup>2</sup>L [5] does not report results on Seq-Cifar-100 and results in the memory-free setting on Seq-Tiny-ImageNet, we run Co<sup>2</sup>L on these cases for comparison. For evaluation, we train a classifier on top of the learned backbone using last-task samples and buffered samples with 100 epochs.

**Evaluation metrics.** Similar to [5, 53], we evaluate the quality of the learned encoder  $f^t$  by training a classifier  $h^t$  on top of the frozen encoder  $f^t$ , using only the current training dataset  $\mathcal{D}_t$  and the samples from the memory  $\mathcal{M}$ . As defined in [7] as well as in other CL methods [5, 12, 53], we compute the Average Accuracy (AA) on the test dataset and all accuracies  $A_{T,k}$  of each task  $k$  after learning the final task  $T$ . The equation for AA is defined as:

$$AA = \frac{1}{T} \sum_{k=1}^T A_{T,k} \quad (11)$$

**Considered methods.** To focus on learning ability with limited or no memory, we compare our method only with recent approaches that can operate in both settings. Since CCLIS [26] and LODE [29] cannot operate in memory-free scenarios, their results are excluded in Tab. 2. As shown in Tab. 1, only Co<sup>2</sup>L [5], the recent method CILA [53], and ours can run without a memory buffer. Additionally, we compare our results with other well-known supervised methods, including ER [43], iCaRL [42], GEM [31], GSS [1], DER [3], Co<sup>2</sup>L [5], GCR [50], and CILA [53].

## 5.2. Main Results

Tab. 2 reports results for buffer sizes 0 and 200 to highlight memory-free and limited memory scenarios. Results for a buffer size of 500 are provided in the Appendix. Although CILA [53] can operate without memory, results for this setting are not available due to the lack of publicly accessible code, so we cannot include them.

As shown in Tab. 2, our model significantly outperforms recent replay-based methods across various CL settings, datasets, and buffer sizes. Notably, without using a buffer, our method outperforms Co<sup>2</sup>L in all settings and datasets. For Seq-Cifar-10, our method exceeds Co<sup>2</sup>L by approximately 10% in Class-IL and 7.76% in Task-IL. Similarly, for Seq-Cifar-100, our method outperforms Co<sup>2</sup>L by 5.68% for Class-IL and by 5.96% for Task-IL. On Seq-Tiny-ImageNet, our method surpasses Co<sup>2</sup>L in both Class-IL and Task-IL by 1.11% and 3.6%, respectively. This improvement is less pronounced than on the CIFAR-10/100 datasets because the current evaluation protocol trains the classifier with samples from the current task and only a limited number from each old class. Consequently, the less significant improvements may stem from the classifier rather than the representation.

Type	Method
Memory-based	ER [43], iCaRL [42], GEM [31], GSS [1], DER [3], Co <sup>2</sup> L [5], GCR [50], LODE [29], CILA [53], CCLIS [26], Ours
Memory-free ability	Co <sup>2</sup> L [5], CILA [53], Ours

Table 1. Among supervised CL, ours is one of the few that operates without memory.

When compared to other methods with a memory buffer size of 200, our method outperforms all except for Seq-Cifar-100 in the Task-IL setting, where it lags behind GCR [50]. Notably, in the memory-free scenario, our method surpasses all replay methods with a buffer size of 200 on Seq-Cifar-10 and nearly matches the top result on Seq-Tiny-ImageNet, achieved by CILA. Although memory-free results for CILA are not available, our method outperforms their buffer size 200 results by 2.2% on Seq-Cifar-10, highlighting its robustness in both limited memory and memory-free settings. With SoTA performance in memory-free scenarios, our model is well-suited for real-world applications where data storage is a concern.

Additionally, the Appendix includes average forgetting results to assess how well our method retains knowledge of previous tasks compared to other baselines.

## 5.3. Ablation Studies

**Effectiveness of FNC<sup>2</sup>.** To evaluate the efficiency of learning plasticity via  $\mathcal{L}_{FNC^2}$ , we compare it with the asymmetric version of SupCon loss ( $\mathcal{L}_{SupCon}^{asym}$ ) [5], in combination with different using distillation methods scenario in the Class-IL setting. Note that with buffer size 0, the asymmetric loss  $\mathcal{L}_{SupCon}^{asym}$  become the original SupCon loss ( $\mathcal{L}_{SupCon}$ ). The results in Tab. 3 show that when stability loss is absent, the performance with both  $\mathcal{L}_{FNC^2}$  and  $\mathcal{L}_{SupCon}^{asym}$  drops significantly, with both yielding approximately the same results. When using distillation methods,  $\mathcal{L}_{FNC^2}$  outperforms  $\mathcal{L}_{SupCon}^{asym}$  in all cases, especially in the case of no memory or small memory size (200).

**Effectiveness of HSD.** In assessing the ability of  $\mathcal{L}_{HSD}$  to preserve prior knowledge, we compare cases without and with different distillation methods, incorporating the plasticity loss  $\mathcal{L}_{FNC^2}$ . The results in Tab. 4 demonstrate that using either  $\mathcal{L}_{IRD}$  [5] or  $\mathcal{L}_{S-PRD}$  individually yields minimal changes in performance. However, when both methods are combined in  $\mathcal{L}_{HSD}$ , performance consistently improves across all buffer sizes.

**Effectiveness of pseudo-replay prototypes.** We conduct ablation experiments using the class-IL setup on the Seq-Cifar-10 and Seq-Cifar-100 datasets with buffer size 0 and 200 to evaluate the role of pseudo-replay prototypes in our method. As shown in Tab. 5, this approach consistently improves performance across all datasets and buffer sizes. No-

Buffer	Dataset Scenario	Seq-Cifar-10		Seq-Cifar-100		Seq-Tiny-ImageNet	
		Class-IL	Task-IL	Class-IL	Task-IL	Class-IL	Task-IL
0	Co <sup>2</sup> L [5]	58.89±2.61	86.65±1.05	26.89±0.78	51.91±0.63	13.43±0.57	40.21±0.68
	Ours	<b>69.26±0.32</b>	<b>94.41±0.43</b>	<b>32.57±0.55</b>	<b>57.87±0.62</b>	<b>14.54±0.52</b>	<b>43.81±0.47</b>
200	ER [43]	44.79±1.86	91.19±0.94	21.78±0.48	60.19±1.01	8.49±0.16	38.17±2.00
	iCaRL [42]	49.02±3.20	88.99±2.13	28.00±0.91	51.43±1.47	7.53±0.79	28.19±1.47
	GEM [31]	25.54±0.76	90.44±0.94	20.75±0.66	58.84±1.00	-	-
	GSS [1]	39.07±5.59	88.80±2.89	19.42±0.29	55.38±1.34	-	-
	DER [3]	61.93±1.79	91.40±0.92	31.23±1.38	63.09±1.09	11.87±0.78	40.22±0.67
	Co <sup>2</sup> L [5]	65.57±1.37	93.43±0.78	27.38±0.85	53.94±0.76	13.88±0.40	42.37±0.74
	GCR [50]	64.84±1.63	90.80±1.05	33.69±1.40	<b>64.24±0.83</b>	13.05±0.91	42.11±1.01
	Ours	<b>72.63±0.78</b>	<b>95.31±0.32</b>	<b>34.04±0.42</b>	59.46±0.65	<b>15.52±0.53</b>	<b>44.59±0.72</b>

Table 2. Results for our method compared with other supervised baselines, with memory sizes 0 and 200, are averaged over five trials (best results in each column are highlighted in bold).

Plasticity	Stability	Buffer size		
		0	200	500
$\mathcal{L}_{FNC^2}$	$\times$	53.59±0.63	53.62±0.81	58.71±0.93
$\mathcal{L}_{SupCon}^{asym}$	$\times$	53.25±1.70	53.57±1.03	58.56±0.85
$\mathcal{L}_{SupCon}^{asym}$	$\mathcal{L}_{IRD}$	58.89±2.61	65.57±1.37	74.26±0.77
$\mathcal{L}_{FNC^2}$	$\mathcal{L}_{IRD}$	63.65±0.55	70.54±0.95	74.81±1.12
$\mathcal{L}_{FNC^2}$	$\mathcal{L}_{HSD}$	<b>69.26±0.32</b>	<b>72.63±0.78</b>	<b>75.51±0.52</b>

Table 3. Performance comparison of  $\mathcal{L}_{FNC^2}$  with  $\mathcal{L}_{SupCon}^{asym}$  in Class-IL setting on Seq-Cifar-10. The test  $\mathcal{L}_{SupCon}^{asym}$  with  $\mathcal{L}_{S-PRD}$  is omitted due to incompatibility, and in case of buffer size 0,  $\mathcal{L}_{SupCon}^{asym}$  become  $\mathcal{L}_{SupCon}$ .

Description	Buffer size	Accuracy (%)
w/o distillation	0	53.59±0.63
w/ $\mathcal{L}_{IRD}$	0	63.65±0.55
w/ $\mathcal{L}_{S-PRD}$	0	64.17±0.41
w/ $\mathcal{L}_{HSD}$	0	<b>69.26±0.32</b>
w/o distillation	200	53.62±0.81
w/ $\mathcal{L}_{IRD}$	200	70.54±0.95
w/ $\mathcal{L}_{S-PRD}$	200	69.20±0.58
w/ $\mathcal{L}_{HSD}$	200	<b>72.63±0.78</b>

Table 4. Ablation study on the effectiveness of  $\mathcal{L}_{HSD}$  in class-IL. We run all tests on the Seq-Cifar-10 dataset with the plasticity loss  $\mathcal{L}_{FNC^2}$  and different distillation methods.

tably, in the memory-free scenario, accuracy with pseudo-replay prototypes is close to that achieved with a buffer size of 200 without them. These outcomes validate our hypothesis that prior prototypes can act as effective representatives of past samples, enhancing cross-task consolidation when used as negative points in current tasks.

## 6. Conclusion

We explore the roles of hard (“hardness”) and soft (“softness”) relationships in NC-based CL for both *plasticity* and

Buffer	Pseudo-replay	Dataset	
		Seq-Cifar-10	Seq-Cifar-100
0	$\times$	67.79±0.43	30.80±0.49
	$\checkmark$	<b>69.26±0.32</b>	<b>32.57±0.55</b>
200	$\times$	70.71±0.89	33.12±0.81
	$\checkmark$	<b>72.63±0.78</b>	<b>34.04±0.42</b>

Table 5. Influence of pseudo-replay prototypes.

*stability*. To address these, we propose two loss functions:  $\mathcal{L}_{FNC^2}$  for plasticity, which uses fixed prototypes to guide representations towards optimal points while emphasizing hard samples and implicitly incorporating *cross-task* consolidation through pseudo-replay of old prototypes, and  $\mathcal{L}_{HSD}$  for distilling both hardness and softness over time. Our approach achieves SoTA in memory-free settings across various datasets and remains competitive with limited buffer sizes in memory-based scenarios.

**Limitations and future work.** Like other NC-inducing methods in CL, our approach is limited by the need to pre-define prototypes, which is impractical when the number of prototypes is unknown. In future work, inspired by recent advancements [32], we will tackle this issue by pre-defining a maximum number of prototypes and learning their distributions, allowing samples to be associated with multiple prototypes and different weights. Additionally, we plan to explore alternative memory-free evaluation methods, as current approaches including Co<sup>2</sup>L [5], CILA [53] rely on buffers, which are less effective with limited samples for old classes. We also aim to identify easily forgotten samples and focus on distilling only the core knowledge.

## Acknowledgements

This work is partially supported by the ANR-21-ASRO-0003 ROV-Chasseur project.

## References

- [1] Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient based sample selection for online continual learning. In *NeurIPS*, 2019. 7, 8, 12
- [2] Nader Asadi, Mohammad Davar, Sudhir P. Mudur, Rahaf Aljundi, and Eugene Belilovsky. Prototype-sample relation distillation: Towards replay-free continual learning. In *ICML*, 2023. 2
- [3] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: a strong, simple baseline. In *NeurIPS*, 2020. 1, 2, 7, 8, 12
- [4] Lucas Caccia, Rahaf Aljundi, Nader Asadi, Tinne Tuytelaars, Joelle Pineau, and Eugene Belilovsky. New insights on reducing abrupt representation change in online continual learning. In *ICLR*, 2021. 1
- [5] Hyuntak Cha, Jaeho Lee, and Jinwoo Shin. Co<sup>2</sup>S: Contrastive Continual Learning. In *ICCV*, 2021. 1, 2, 4, 6, 7, 8, 11, 12
- [6] Sungmin Cha, Kyunghyun Cho, and Taesup Moon. Regularizing with pseudo-negatives for continual self-supervised learning. In *ICML*, 2024. 2
- [7] Arslan Chaudhry, Puneet K. Dokania, Thalaiyasingam Ajanthan, and Philip H. S. Torr. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *ECCV*, 2018. 7, 11
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 2, 6
- [9] Xinlei Chen, Haoqi Fan, Ross B. Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint 2003.04297*, 2020. 2
- [10] Laura N. Driscoll, Lea Duncker, and Christopher D. Harvey. Representational drift: Emerging theories for continual learning and experimental future directions. *Current Opinion in Neurobiology*, page 102609, 2022. 1
- [11] Cong Fang, Hangfeng He, Qi Long, and Weijie J Su. Exploring deep neural networks via layer-peeled model: Minority collapse in imbalanced training. *PNAS*, 2021. 1, 3
- [12] Enrico Fini, Victor G. Turrise da Costa, Xavier Alameda-Pineda, Elisa Ricci, Karteek Alahari, and Julien Mairal. Self-supervised models are continual learners. In *CVPR*, 2022. 1, 2, 6, 7
- [13] Tomer Galanti, Andr'as Gyorgy, and Marcus Hutter. On the role of neural collapse in transfer learning. In *ICLR*, 2022. 3
- [14] Florian Graf, Christoph D. Hofer, Marc Niethammer, and Roland Kwitt. Dissecting supervised contrastive learning. In *ICML*, 2021. 3
- [15] X. Y. Han, Vardan Papyan, and David L. Donoho. Neural collapse under mse loss: Proximity to and dynamics on the central path. In *ICLR*, 2022. 1, 3
- [16] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 2
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 6
- [18] Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In *NIPS Workshops*, 2015. 2
- [19] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *CVPR*, 2019. 2
- [20] Wenlong Ji, Yiping Lu, Yiliang Zhang, Zhun Deng, and Weijie J. Su. An unconstrained layer-peeled perspective on neural collapse. In *ICLR*, 2022. 1, 3
- [21] Sangwon Jung, Hongjoon Ahn, Sungmin Cha, and Taesup Moon. Continual learning with node-importance based adaptive group sparse regularization. In *NeurIPS*, 2020. 2
- [22] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. In *NeurIPS*, 2020. 2, 4
- [23] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical report, University of Toronto, Toronto, 2009. 6
- [24] Nicolas Larue, Ngoc-Son Vu, Vitomir Struc, Peter Peer, and Vassilis Christophides. Seeable: Soft discrepancies and bounded contrastive learning for exposing deepfakes. In *ICCV*, 2022. 3
- [25] Ya Le and Xuan S. Yang. Tiny imagenet visual recognition challenge, 2015. 6
- [26] Jiyong Li, Dilshod Azizov, Yang Li, and Shangsong Liang. Contrastive continual learning with importance sampling and prototype-instance relation distillation. In *AAAI*, 2024. 2, 7
- [27] Shiyao Li, Xuefei Ning, Shanghang Zhang, Lidong Guo, Tianchen Zhao, Huazhong Yang, and Yu Wang. Tcp: Triplet contrastive-relationship preserving for class-incremental learning. In *WACV*, 2024. 1
- [28] Xilai Li, Yingbo Zhou, Tianfu Wu, Richard Socher, and Caiming Xiong. Learn to grow: A continual structure learning framework for overcoming catastrophic forgetting. *arXiv preprint 1904.00310*, 2019. 2
- [29] Yan-Shuo Liang and Wu-Jun Li. Loss decoupling for task-agnostic continual learning. In *NeurIPS*, 2023. 2, 4, 5, 7
- [30] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, 2017. 4
- [31] David Lopez-Paz and Marc'Aurelio Ranzato. Gradient episodic memory for continual learning. In *NeurIPS*, 2017. 7, 8, 12
- [32] Haodong Lu, Dong Gong, Shuo Wang, Jason Xue, Lina Yao, and Kristen Moore. Learning with mixture of prototypes for out-of-distribution detection. In *ICLR*, 2024. 8
- [33] Jianfeng Lu and Stefan Steinerberger. Neural collapse with cross-entropy loss. *arXiv preprint 2012.08465*, 2020. 1, 3
- [34] Divyam Madaan, Jaehong Yoon, Yuanchun Li, Yunxin Liu, and Sung Ju Hwang. Representational continuity for unsupervised continual learning. In *ICLR*, 2022. 1, 2
- [35] Z. Mai, R. Li, H. Kim, and S. Sanner. Supervised contrastive replay: Revisiting the nearest class mean classifier in online class-incremental continual learning. In *CVPRW*, 2021. 2
- [36] Michael McCloskey and Neal J. Cohen. Catastrophic Interference in Connectionist Networks: The Sequential Learning Problem. *Psychology of Learning and Motivation*, 1989. 1

- [37] Antoine Montmaur, Nicolas Larue, and Ngoc-Son Vu. Neural collapse inspired contrastive continual learning. In *BMVC*, 2024. 1
- [38] Vardan Papyan, X. Y. Han, and David L. Donoho. Prevalence of neural collapse during the terminal phase of deep learning training. *PNAS*, 2020. 3, 5
- [39] Dongmin Park, Seokil Hong, Bohyung Han, and Kyoung Mu Lee. Continual learning by asymmetric loss approximation with single-side overestimation. In *ICCV*, 2019. 2
- [40] Tomaso Poggio and Qianli Liao. Explicit regularization and implicit bias in deep network classifiers trained with the square loss. *arXiv preprint 2101.00072*, 2020. 3
- [41] Julien Pourcel, Ngoc-Son Vu, and Robert M. French. Online task-free continual learning with dynamic sparse distributed memory. In *ECCV*, 2022. 1
- [42] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H. Lampert. icarl: Incremental classifier and representation learning. In *CVPR*, 2017. 7, 8, 12
- [43] Matthew Riemer, Ignacio Cases, Robert Ajemian, Miao Liu, Irina Rish, Yuhai Tu, and Gerald Tesauro. Learning to learn without forgetting by maximizing transfer and minimizing interference. In *ICLR*, 2019. 7, 8, 12
- [44] Anthony V. Robins. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connect. Sci.*, pages 123–146, 1995. 1
- [45] Andrei A. Rusu, Neil C. Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint 1606.04671*, 2016. 1, 2
- [46] Minhyuk Seo, Hyun woo Koh, Wonje Jeung, Minjae Lee, San Kim, Hankook Lee, Sungjun Cho, Sungik Choi, Hyunwoo Kim, and Jonghyun Choi. Learning equi-angular representations for online continual learning. In *CVPR*, 2024. 1
- [47] F. Szatkowski, M. Pyla, M. Przewiezlikowski, S. Cygert, B. Twardowski, and T. Trzcinski. Adapt your teacher: Improving knowledge distillation for exemplar-free continual learning. In *ICCVW*, 2023. 2
- [48] C. Tang, L. Qendro, D. Spathis, F. Kawsar, C. Mascolo, and A. Mathur. Kaizen: Practical self-supervised continual learning with continual fine-tuning. In *WACV*, 2024. 1, 2
- [49] Tom Tirer and Joan Bruna. Extended unconstrained features model for exploring deep neural collapse. In *ICML*, 2022. 1, 3
- [50] Rishabh Tiwari, Krishnateja Killamsetty, Rishabh Iyer, and Pradeep Shenoy. Gcr: Gradient coreset based replay buffer selection for continual learning. In *CVPR*, 2022. 1, 7, 8, 11, 12
- [51] Aäron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint 1807.03748*, 2018. 2
- [52] Jeffrey S. Vitter. Random sampling with a reservoir. *ACM Transactions on Mathematical Software*, pages 37–57, 1985. 5
- [53] Yichen Wen, Zhiqian Tan, Kaipeng Zheng, Chuanlong Xie, and Weiran Huang. Provable contrastive continual learning. In *ICML*, 2024. 1, 2, 7, 8, 12
- [54] Yibo Yang, Shixiang Chen, Xiangtai Li, Liangru Xie, Zhouchen Lin, and Dacheng Tao. Inducing neural collapse in imbalanced learning: Do we really need a learnable classifier at the end of deep neural network? In *NeurIPS*, 2022. 3
- [55] Yibo Yang, Haobo Yuan, Xiangtai Li, Zhouchen Lin, Philip Torr, and Dacheng Tao. Neural collapse inspired feature-classifier alignment for few-shot class-incremental learning. In *ICLR*, 2023. 1, 3, 4
- [56] Yibo Yang, Haobo Yuan, Xiangtai Li, Jianlong Wu, Lefei Zhang, Zhouchen Lin, Philip Torr, Dacheng Tao, and Bernard Ghanem. Neural collapse terminus: A unified solution for class incremental learning and its variants, 2023. 1, 3, 4
- [57] Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong learning with dynamically expandable networks. *arXiv preprint 1708.01547*, 2018. 1, 2
- [58] L. Yu, B. Twardowski, X. Liu, L. Herranz, K. Wang, Y. Cheng, S. Jui, and J. van de Weijer. Semantic drift compensation for class-incremental learning. In *CVPR*, 2020. 1
- [59] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *ICML*, 2021. 6
- [60] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *ICML*, 2017. 2
- [61] Yipeng Zhang, Laurent Charlin, Richard Zemel, and Mengye Ren. Integrating present and past in unsupervised continual learning. In *CoLLAs*, 2024. 2, 4, 5
- [62] Yifan Zhang, Bryan Hooi, Dapeng Hu, Jian Liang, and Jiashi Feng. Unleashing the power of contrastive self-supervised visual models via contrast-regularized fine-tuning. In *NeurIPS*, 2021. 4
- [63] Jinxin Zhou, Xiao Li, Tianyu Ding, Chong You, Qing Qu, and Zhihui Zhu. On the optimization landscape of neural collapse under mse loss: Global optimality with unconstrained features. In *ICML*, 2022. 1, 3
- [64] Zhihui Zhu, Tianyu Ding, Jinxin Zhou, Xiao Li, Chong You, Jeremias Sulam, and Qing Qu. A geometric analysis of neural collapse with unconstrained features. In *NeurIPS*, 2021. 3

## APPENDIX

### A. Hyperparameter selection

To select hyperparameters, we employ a grid search strategy, using a randomly drawn 10% of the training data as the validation set. The considered hyperparameters are:

- Learning rate ( $\eta$ )
- Batch size ( $bsz$ )
- Number of start epochs ( $E_1$ )
- Number of epochs of  $t$ -th task ( $E_{t \geq 2}$ )
- Temperature for plasticity loss ( $\tau$ ): We use the same  $\tau$  for both Focal Neural Collapse Contrastive (FNC<sup>2</sup>) and Asymmetric SupCon loss [5]
- Focusing hyperparameters ( $\gamma$ ) for FNC<sup>2</sup> loss
- Temperature for instance-wise relation distillation loss ( $\mathcal{L}_{IRD}$ ): As in [5], we use different temperature hyperparameters for the past ( $\kappa_{past}$ ) and current ( $\kappa_{current}$ ) similarity vectors
- Temperature for sample-prototype relation distillation loss ( $\mathcal{L}_{S-PRD}$ ): We utilize  $\zeta_{past}$  for the past and  $\zeta_{current}$  for the current similarity vectors
- Number of warm-up epochs in hardness-softness distillation loss ( $\mathcal{L}_{HSD}$ ) ( $e_0$ )

The corresponding search space of these hyperparameters are provided in Tab. 6. The selections of these hyperparameters are based on the average test accuracy over five independent trials, and the final chosen values are detailed in Tab. 7. For the sake of conciseness and to maintain focus, we omit those hyperparameters previously discovered in the literature.

**Focusing hyperparameter ( $\gamma$ ).** In the FNC<sup>2</sup> loss function,  $\gamma$  plays a crucial role in determining the level of focus on hard samples (i.e., positive samples that are far from

Hyperparameter	Values
$\eta$	{0.1, 0.5, 1.0}
$bsz$	{256, 512}
$E_1$	{500}
$E_{t \geq 2}$	{50, 100}
$\tau$	{0.1, 0.5, 1.0}
$\gamma$	{0, 1, 2, 4, 7, 10}
$\kappa_{past}$	{0.01, 0.05, 0.1}
$\kappa_{current}$	{0.1, 0.2}
$\zeta_{past}$	{0.01, 0.05, 0.1}
$\zeta_{current}$	{0.1, 0.2}
$e_0$	{10, 20, 30}

Table 6. Search spaces of hyperparameters.

Method	Buffer size	Dataset	Hyperparameter
Our	0, 200, 500	Seq-Cifar-10	$\eta$ : 0.5, $\gamma$ : 1, $bsz$ : 512, $E_1$ : 500, $E_{t \geq 2}$ : 100, $\tau$ : 0.5, $e_0$ : 30, $\kappa_{past}$ : 0.01, $\kappa_{current}$ : 0.2, $\zeta_{past}$ : 0.01, $\zeta_{current}$ : 0.2
	0, 200, 500	Seq-Cifar-100	$\eta$ : 0.5, $\gamma$ : 4, $bsz$ : 512, $E_1$ : 500, $E_{t \geq 2}$ : 100, $\tau$ : 0.5, $e_0$ : 30, $\kappa_{past}$ : 0.01, $\kappa_{current}$ : 0.2, $\zeta_{past}$ : 0.1, $\zeta_{current}$ : 0.2
	0, 200, 500	Seq-Tiny-ImageNet	$\eta$ : 0.1, $\gamma$ : 4, $bsz$ : 512, $E_1$ : 500, $E_{t \geq 2}$ : 50, $\tau$ : 0.5, $e_0$ : 20, $\kappa_{past}$ : 0.1, $\kappa_{current}$ : 0.1, $\zeta_{past}$ : 0.1, $\zeta_{current}$ : 0.2
Co <sup>2</sup> L	0, 200, 500	Seq-Cifar-100	$\eta$ : 0.5, $bsz$ : 512, $E_1$ : 500, $E_{t \geq 2}$ : 100, $\tau$ : 0.5, $\kappa_{past}$ : 0.01, $\kappa_{current}$ : 0.2
	0	Seq-Tiny-ImageNet	$\eta$ : 0.1, $bsz$ : 512, $E_1$ : 500, $E_{t \geq 2}$ : 50, $\tau$ : 0.5, $\kappa_{past}$ : 0.1, $\kappa_{current}$ : 0.1

Table 7. Selected hyperparameters in our experiments.

the anchor or their prototypes). To explore this role and select the most suitable  $\gamma$  for each dataset, we conduct experiments across different datasets to observe how the performance of our method changes as  $\gamma$  varies. The test accuracy results in Fig. 4 show that our method performs best at different values of  $\gamma$  for each dataset. Specifically, as reported in Tab. 7, the chosen  $\gamma$  for the Seq-Cifar-100 and Seq-Tiny-ImageNet datasets ( $\gamma = 4$  for both) are larger than that for the Seq-Cifar-10 dataset ( $\gamma = 1$ ). This difference arises because Seq-Cifar-100 and Seq-Tiny-ImageNet have a large number of classes per task (both have 20 classes/task), which increases the likelihood of samples being close to the prototypes of other class clusters. In contrast, the Seq-Cifar-10 dataset has only 2 classes each task, making it less complex and not requiring a large  $\gamma$ .

### B. Additional Experiments

#### B.1. Average accuracy results with buffer size 500

In addition to the results with small buffer sizes (0 and 200), we run experiments with a buffer size of 500 across different datasets to further assess the effectiveness of our method with a larger buffer. As shown in Tab. 8, although our method does not surpass state-of-the-art methods, it achieves results close to them on Seq-Cifar-10 and Seq-Tiny-ImageNet, underperforming only on Seq-Cifar-100 compared to GCR [50]. This further demonstrates that our method, aside from excelling in memory-free and small buffer settings, remains effective with larger buffers.

#### B.2. Average forgetting results

We utilize the Average Forgetting metric as defined in [7] to quantify how much information the model has forgotten about previous tasks, which as

$$F = \frac{1}{T-1} \sum_{i=1}^{T-1} \max_{t \in \{1, \dots, T-1\}} (A_{t,i} - A_{T,i}) \quad (12)$$

Tab. 9 report the average forgetting results of our method compared to all other baselines. The results show that our method can effectively mitigate forgetting, especially even without using additional buffers.

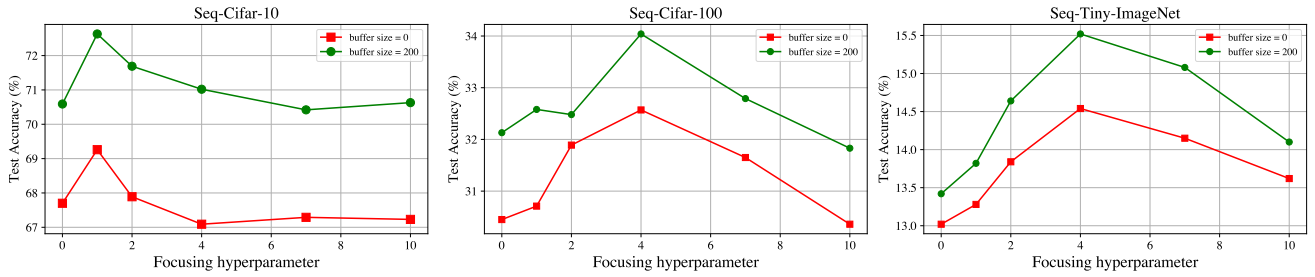


Figure 4. Test accuracy over different values of  $\gamma$ .

Buffer	Dataset Scenario	Seq-Cifar-10		Seq-Cifar-100		Seq-Tiny-ImageNet	
		Class-IL	Task-IL	Class-IL	Task-IL	Class-IL	Task-IL
500	ER [43]	57.74±0.27	93.61±0.27	27.66±0.61	66.23±1.52	9.99±0.29	48.64±0.46
	iCaRL [42]	47.55±3.95	88.22±2.62	33.25±1.25	58.16±1.76	9.38±1.53	31.55±3.27
	GEM [31]	26.20±1.26	92.16±0.64	25.54±0.65	66.31±0.86	-	-
	GSS [1]	49.73±4.78	91.02±1.57	21.92±0.34	60.28±1.18	-	-
	DER [3]	70.51±1.67	93.40±0.39	41.36±1.76	71.73±0.74	17.75±1.14	51.78±0.88
	Co <sup>2</sup> L [5]	74.26±0.77	95.90±0.26	37.02±0.76	62.44±0.36	20.12±0.42	53.04±0.69
	GCR [50]	74.69±0.85	94.44±0.32	<b>45.91±1.30</b>	<b>71.64±2.10</b>	19.66±0.68	52.99±0.89
	CILA [53]	<b>76.03±0.79</b>	<b>96.40±0.21</b>	-	-	<b>20.64±0.59</b>	<b>54.13±0.72</b>
Ours	75.51±0.52	96.14±0.25	40.25±0.58	65.85±0.44	20.31±0.34	53.46±0.59	

Table 8. Additional results with buffer size 500 (best results in each column are bold).

Buffer	Dataset Scenario	Seq-Cifar-10		Seq-Cifar-100		Seq-Tiny-ImageNet	
		Class-IL	Task-IL	Class-IL	Task-IL	Class-IL	Task-IL
0	Co <sup>2</sup> L [5]	35.81±1.08	14.33±0.87	66.51±0.28	39.63±0.62	62.80±0.77	39.54±1.08
	Ours	<b>23.85±0.30</b>	<b>4.72±0.28</b>	<b>52.03±0.63</b>	<b>36.20±0.48</b>	<b>53.97±0.63</b>	<b>37.57±0.88</b>
200	ER [43]	59.30±2.48	6.07±1.09	75.06±0.63	27.38±1.46	76.53±0.51	40.47±1.54
	GEM [31]	80.36±5.25	9.57±2.05	77.40±1.09	29.59±1.66	-	-
	GSS [1]	72.48±4.45	8.49±2.05	77.62±0.76	32.81±1.75	-	-
	iCaRL [42]	<b>23.52±1.27</b>	25.34±1.64	<b>47.20±1.23</b>	36.20±1.85	<b>31.06±1.91</b>	42.47±2.47
	DER [3]	35.79±2.59	6.08±0.70	62.72±2.69	25.98±1.55	64.83±1.48	40.43±1.05
	Co <sup>2</sup> L [5]	36.35±1.16	6.71±0.35	67.82±0.41	38.22±0.34	73.25±0.21	47.11±1.04
	GCR [50]	32.75±2.67	7.38±1.02	57.65±2.48	<b>24.12±1.17</b>	65.29±1.73	40.36±1.08
	CILA [53]	-	-	-	-	-	-
Ours	25.24±0.69	<b>4.28±0.32</b>	52.40±0.83	33.66±0.24	52.07±0.46	<b>33.76±0.58</b>	
500	ER [43]	43.22±2.10	3.50±0.53	67.96±0.78	17.37±1.06	75.21±0.54	30.73±0.62
	GEM [31]	78.93±6.53	5.60±0.96	71.34±0.78	20.44±1.13	-	-
	GSS [1]	59.18±4.00	6.37±1.55	74.12±0.42	26.57±1.34	-	-
	iCaRL [42]	28.20±2.41	22.61±3.97	40.99±1.02	27.90±1.37	<b>37.30±1.42</b>	39.44±0.84
	DER [3]	24.02±1.63	3.72±0.55	49.07±2.54	25.98±1.55	59.95±2.31	28.21±0.97
	Co <sup>2</sup> L [5]	25.33±0.99	3.41±0.80	51.23±0.65	26.30±0.57	65.15±0.26	39.22±0.69
	GCR [50]	<b>19.27±1.48</b>	<b>3.14±0.36</b>	<b>39.20±2.84</b>	<b>15.07±1.88</b>	56.40±1.08	27.88±1.19
	CILA [53]	-	-	-	-	-	-
Ours	22.59±1.02	3.21±0.25	41.66±0.78	24.84±0.91	46.08±0.56	<b>26.45±0.79</b>	

Table 9. Average forgetting (lower is better) across five independent trials: Comparison of our method with all baselines in continual learning.