

# SALIENCY-GUIDED HIDDEN ASSOCIATIVE REPLAY FOR CONTINUAL LEARNING

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Continual Learning (CL) is a burgeoning domain in next-generation AI, focusing on training neural networks over a sequence of tasks akin to human learning. While CL provides an edge over traditional supervised learning, its central challenge remains to counteract *catastrophic forgetting* and ensure the retention of prior tasks during subsequent learning. Amongst various strategies to tackle this, replay-based methods have emerged as preeminent, echoing biological memory mechanisms. However, these methods are memory-intensive, often preserving entire data samples—an approach inconsistent with humans’ selective memory retention of salient experiences. While some recent works have explored the storage of only significant portions of data in episodic memory, the inherent nature of partial data necessitates innovative retrieval mechanisms. Current solutions, like inpainting, approximate full data reconstruction from partial cues, a method that diverges from genuine human memory processes. Addressing these nuances, this paper presents the **S**aliency-Guided **H**idden **A**ssociative **R**eplay for **C**ontinual Learning (**SHARC**). This novel framework synergizes associative memory with replay-based strategies. SHARC primarily archives salient data segments via sparse memory encoding. Importantly, by harnessing associative memory paradigms, it introduces a content-focused memory retrieval mechanism, promising swift and near-perfect recall, bringing CL a step closer to authentic human memory processes. Extensive experimental results demonstrate the effectiveness of our proposed method for various continual learning tasks.

## 1 INTRODUCTION

Continual learning (CL) represents a vital advancement for next-generation AI, allowing neural networks to sequentially learn tasks like humans do Parisi et al. (2019). While traditional supervised learning is well-established, CL remains in its nascent stages. The main challenge is to prevent *Catastrophic Forgetting* McCloskey & Cohen (1989) as agents acquire new tasks, ensuring they retain earlier knowledge. In essence, CL strives to balance updating the model with retention across a series of tasks. In order to address this problem, researchers have put forward several strategies. *Replay-based methods* Rebuffi et al. (2017); Aljundi et al. (2019); Arani et al. (2022), which utilizes a small memory to store previous data and reuse them when learning new tasks, have emerged as a particularly effective solution, offering superior performance and drawing inspiration from biological systems Robins (1995). However, a potential bottleneck of this approach is its memory-intensive nature, as entire data samples are conserved. This mechanism contrasts starkly with the human brain’s approach to memory retention. Humans typically do not remember every detail but tend to recall fragments or the most salient features of experiences Rolls (2013). The vast storage requirements of replay-based methods and their divergence from natural memory processes necessitate exploration into more efficient and human-like strategies for continual learning.

While there are pioneering works Saha & Roy (2023); Bai et al. (2023) in replay-based CL that have begun to explore the idea of storing only the salient or partial aspects of data into episodic memory, challenges arise due to the inherent nature of partial data. Since these fragments are not directly usable as model input, an effective retrieval technique becomes indispensable. A straightforward solution is inpainting Elharrouss et al. (2020), which, through rule-based or generative models, attempts to recreate the full data from the available partial cue. This methodology, however, essentially approximates the entirety of the data by generating similar samples from a given distribution and may suffer from inaccurate retrieval under large noise or corruption (Col 3 in Figure 1). On the contrary,

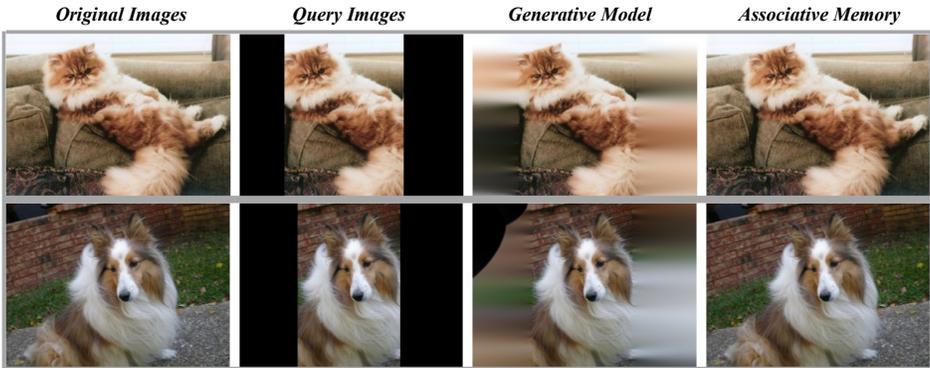


Figure 1: Content-based (associative memory) v.s. generative model retrieval. Pixels from non-salient areas are masked in query images. For a fair comparison, we train an autoencoder-based inpainting model Peng et al. (2021) and a Hopfield Network Ramsauer et al. (2020) with similar number of parameters. Associative memory achieves almost perfect recall even under large noise or corruption.

the human brain, especially the hippocampus, employs associative recall for *content-based* memory retrieval Hopfield (1982); Ramsauer et al. (2020), achieving a remarkable recall accuracy close to perfection (Col 4 in Figure 1). As such, for systems aiming to emulate human-like continual learning, there is an evident inspiration to design techniques that mirror the associative and content-based retrieval processes inherent in human cognition.

To address the aforementioned challenges, this paper introduces the **S**aliency-Guided **H**idden **A**ssociative **R**eplay for **C**ontinual Learning (**SHARC**), marking the inception of a Continual Learning framework that seamlessly integrates associative memory into replay-based techniques. As depicted in Figure 1, SHARC distinguishes itself from existing replay-based CL methodologies in two pivotal aspects: First, rather than archiving complete samples within episodic memory, SHARC conserves only the most salient segments through sparse memory encoding. More crucially, drawing inspiration from the principles of associative memory, we have crafted a content-centric memory retrieval module that boasts swift and impeccable recall capabilities.

Our contribution includes, 1). We develop a novel neural-inspired replay-based continual learning framework to handle catastrophic forgetting. 2). We propose to leverage associative memory for efficient memory storage and recovery. 3). We demonstrate our model’s efficacy and superiority with extensive experiments.

## 2 RELATED WORK

**Continual Learning (CL).** Catastrophic forgetting is a long-standing problem Robins (1995) in continual learning which has been recently tackled in a variety of visual tasks such as image classification Kirkpatrick et al. (2017); Rebuffi et al. (2017), object detection Zhou et al. (2020), etc.

Existing techniques in CL can be divided into three main categories Parisi et al. (2019): 1) *regularization-based approaches*, 2) *dynamic architectures* and 3) *replay-based approaches*. Regularization-based approaches alleviate catastrophic forgetting by either adding a regularization term to the objective function Kirkpatrick et al. (2017) or knowledge distillation over previous tasks Li & Hoiem (2017). Dynamic architecture approaches adaptively accommodate the network architecture (e.g., adding more neurons or layers) in response to new information during training. Dynamic architectures can be explicit if new network branches are grown, or implicit, if some network parameters are only available for certain tasks. Replay-based approaches alleviate the forgetting of deep neural networks by replaying stored samples from the previous history when learning new ones and have been shown to be the most effective method for mitigating catastrophic forgetting.

**Replay-based CL.** Replay-based methods mainly include three directions: rehearsal methods, constrained optimization, and pseudo rehearsal. *Rehearsal methods* directly retrieve previous samples from a limited size memory together with new samples for training Chaudhry et al. (2019); Hayes et al. (2020); Arani et al. (2022). While simple in nature, this approach is prone to overfitting the old samples from the memory. As an alternative, *constrained optimization* methods formulate

backward/forward transfer as constraints in the objective function. GEM Lopez-Paz & Ranzato (2017) constrains new task updates to not interfere with previous tasks by projecting the estimated gradient on the feasible region outlined by previous task gradients through first-order Taylor series approximation. A-GEM Chaudhry et al. (2018) further extended GEM and made the constraint computationally more efficient. Finally, *pseudo-rehearsal* methods typically utilize generative models such as GAN Goodfellow et al. (2020) or VAE Pu et al. (2016) to generate previous samples from random inputs and have shown the ability to generate high-quality images recently Robins (1995). Readers may refer to Parisi et al. (2019) for a more comprehensive survey on continual learning.

**Associative Memory (AM).** In general, the attractor-based mechanism Amit & Amit (1989) is typically used for the implementation of AMs, which are models that store and recall patterns. Pattern recall (associative recall) is a process whereby an associative memory, upon receiving a potentially corrupted memory query, retrieves the associated value from memory. One of the earliest and probably the most well-known associative memory are Hopfield Networks Hopfield (1982). Hopfield networks are a class of recurrent artificial neural networks that have gained prominence for their ability to model associative memory and pattern recognition. The modern Hopfield network refers to an updated version of the original Hopfield network Ramsauer et al. (2020); Krotov & Hopfield (2020). The modern Hopfield network incorporates enhancements and modifications to improve its performance and overcome some limitations of the original model. More recently, predictive coding networks Huang & Rao (2011) have provided a new perspective for the design of AM, and such works Salvatori et al. (2021); Yoo & Wood (2022) have shown strong performance on recall tasks.

### 3 PROBLEM FORMULATION

We consider supervised continual learning in this paper. Following the learning protocol in Chaudhry et al. (2018), we consider a training set  $\mathcal{D} = \{\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_T\}$  consisting of  $T$  tasks, where  $\mathcal{D}_t = \{(\mathbf{x}_i^{(t)}, \mathbf{y}_i^{(t)})\}_{i=1}^{n_t}$  contains  $n_t$  input-target pairs  $(\mathbf{x}_i^{(t)}, \mathbf{y}_i^{(t)}) \in \mathcal{X} \times \mathcal{Y}$ . While each learning task arrives sequentially, we make the assumption of *locally i.i.d.*, i.e.,  $\forall t, (\mathbf{x}_i^{(t)}, \mathbf{y}_i^{(t)}) \stackrel{iid}{\sim} P_t$ , where  $P_t$  denotes the data distribution for task  $t$  and *i.i.d.* for *independent and identically distributed*. Given such a stream of tasks, the goal is to train a learning agent  $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ , parameterized by  $\theta$ , which can be queried *at any time* to predict the target  $\mathbf{y}$  given associated unseen input  $\mathbf{x}$  and task id  $t$ . Moreover, we require that such a learning agent can only store a small amount of seen samples in an episodic memory  $\mathcal{M}$  with a fixed budget. Given predictor  $f_\theta$ , the loss on the episodic memory of task  $k$  is defined as

$$\ell(f_\theta, \mathcal{M}_k) := |\mathcal{M}_k|^{-1} \sum_{(\mathbf{x}_i, \mathbf{y}_i)} \phi(f_\theta(\mathbf{x}_i, k), \mathbf{y}_i), \quad \forall k < t, \quad (1)$$

where  $\phi$  can be *e.g.* cross-entropy or MSE. In general, a large body of *replay-based* continual learning methods seeks to optimize for the following loss function at  $t$ -th task

$$\min_{\theta} \mathcal{L}_{CL}(\theta), \quad \text{where } \mathcal{L}_{CL}(\theta) = \sum_{(\mathbf{x}, t, \mathbf{y})} \ell(f_\theta(\mathbf{x}, t), \mathbf{y}) + \sum_{k < t} \ell(f_\theta, \mathcal{M}_k), \quad (2)$$

which is an aggregation of the losses on the current task and replay data. After the training of task  $t$ , a subset of training samples will be stored in the episodic memory, i.e.,  $\mathcal{M} = \mathcal{M} \cup \{(\mathbf{x}_i^{(t)}, \mathbf{y}_i^{(t)})\}_{i=1}^{m_t}$ , where  $m_t$  is the memory buffer size for the current task.

## 4 PROPOSED METHOD

In this section, we introduce our proposed Saliency-Guided Hidden Associative Replay for Continual Learning. We innovatively utilize saliency methods to select the most important channels of feature maps for each image and only store those channels in the episodic memory, thus achieving controllable and better memory efficiency. During the training phase, we leverage pattern association techniques for memory completion, where each partially stored image will be restored by a brain-inspired associative memory. An overview of our framework is shown in Figure 2.

### 4.1 SALIENCY-GUIDED MEMORY ENCODING WITH STRUCTURED SPARSITY

In this section, we discuss the memory encoding process of our method. According to the *hippocampal indexing theory* Teyler & Rudy (2007), there are two major characteristics of how the human brain encodes its memories. First, the encoded representations in the human hippocampus are highly *sparse*, meaning that only a small subset of neurons in the hippocampus is activated for each specific

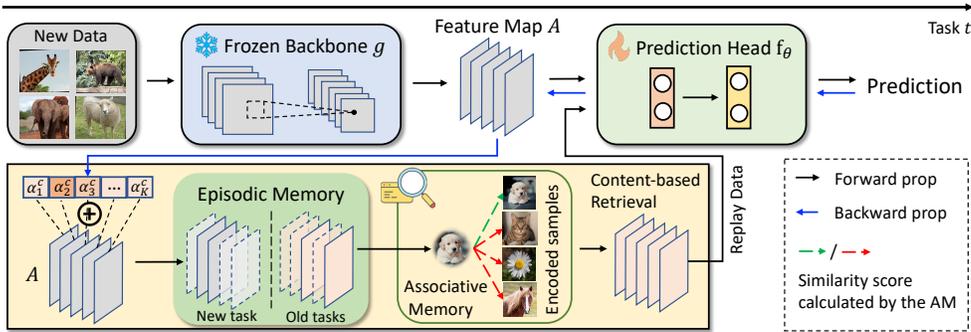


Figure 2: Overview of our proposed SHARC framework (best viewed in color). When new data comes, a pre-trained backbone is used to extract feature maps. Then, the saliency score of the saliency map is calculated via backpropagation and we drop channels with lower saliency, thus achieving structured sparsity and memory efficiency. During memory replay, previous feature maps are retrieved via an associative memory, which essentially picks the top-1 stored feature maps based on the similarity with the query feature map.

memory. Second, the stored representations for replay are not exact reproductions (e.g., raw pixels) Ji & Wilson (2007); instead, its visual inputs originate higher in the visual processing hierarchy rather than from the primary visual cortex or the retina Insausti et al. (2017). Motivated by this, our goal is to propose a computational model for memory encoding that satisfies both characteristics.

Different from many earlier works that store raw images for replay Lopez-Paz & Ranzato (2017); Ebrahimi et al. (2021); Arani et al. (2022), we consider first encoding raw data into high-level representations and store them. Formally, our model  $y = f_{\theta}(g(\mathbf{x}))$  is composed of a pre-trained backbone  $g$  and a trainable prediction head  $f_{\theta}$ . In this work, the pre-trained backbone is general and can be instantiated as various vision models such as VGG, ResNet, and ViT.

The output of  $g(\mathbf{x})$  is a tensor  $A \in \mathbb{R}^{H \times W \times K}$ , where  $H$ ,  $W$  are the dimension of the feature map and  $K$  is the number of channels. To achieve sparse representation, we consider saliency-based methods Selvaraju et al. (2017) that measure the importance of the neurons by their first-order gradients. Specifically, the saliency approach computes the gradient of the score for class  $c$ , i.e.,  $f_{\theta}^{(c)}(A)$ , with respect to feature map activations  $A$ , followed by a global-average-pooling over the width and height dimension to obtain the neuron importance weights  $\alpha_k^c$ :

$$\forall k, \alpha_k^c = \frac{1}{H} \sum_{i=1}^H \frac{1}{W} \sum_{j=1}^W \partial f_{\theta}^{(c)}(A) / \partial A_{ij}^k. \quad (3)$$

Intuitively speaking, feature maps with higher magnitudes in the saliency are more likely to be involved in the region of the target class object while those with lower magnitudes are more likely to be the non-target objects or background regions. In addition, existing work Sun et al. (2015) has proved that hidden representation learned by convolutional neural networks is highly sparse in the hidden space. In this work, we consider the saliency score as a measure and mask out feature maps with lower saliency scores. Formally, denote  $\alpha^c = [\alpha_1^c, \alpha_2^c, \dots, \alpha_K^c]$ , the masked feature map  $A'$  is

$$A' = \text{TR}_{H,W,K}(\mathbb{1}\{|\alpha^c| > Q_{\mu}\} \otimes \mathbf{J}_{H,W}) \odot A, \quad (4)$$

where  $\mathbb{1}\{\cdot\}$  is the indicator function,  $Q_{\mu}$  is the threshold for masking out the bottom  $\mu$  quantile of channels.  $\mathbf{J}$  denotes the all-one matrix and  $\text{TR}(\cdot)$  denotes tensor reshaping operation.

One advantage of our design is that the channel-wise operations result in **structured sparsity**, which is hardware-friendly and can lead to memory cost reduction instantly without further system-level efforts. To see this, we discard those channels that have a saliency score lower than the threshold, and the rest feature maps have a regular tensor shape. We only need to keep track of the channel index which is only a 1d vector and cheap to store.

#### 4.2 ASSOCIATIVE MEMORY RETRIEVAL FOR REPLAY

In this section, we discuss the memory retrieval process of our method. Associative memory plays an important role in human intelligence and its mechanisms have been linked to attention in machine learning Ramsauer et al. (2020). Recently, the machine learning community’s interest in associative

**Algorithm 1** SHARC Training

---

**Require:** Continual learning classifier  $f_\theta$ , associative memory  $\mathcal{A}(\cdot, \omega)$ , training continuum  $\mathcal{D}^{train}$ , dropping threshold  $\mu$ , optimizer OPT, forgetting frequency  $R$ , total number of tasks  $T$ .

- 1:  $\mathcal{M}_t \leftarrow \{\}, \forall t = 1, 2, \dots, T$  ▷ Initialize episodic memory
- 2: **for**  $t = 1$  **to**  $T$  **do**
- 3:      $\tilde{\mathcal{M}}_k \leftarrow \text{OPT}_{\mathbf{x}}(\mathbf{x}, \mathcal{M}_{k < t}, \omega), \forall k < t$  as Eq. 2 ▷ Associative memory *read*
- 4:     **for**  $\mathcal{B}_t \sim \mathcal{D}_t^{train}$  **do**
- 5:          $\theta \leftarrow \text{OPT}_\theta(\theta, \mathcal{B}_t, \tilde{\mathcal{M}}_{k < t})$  as Eq. 2 ▷ Train the classifier
- 6:         **for**  $(x, y) \in \mathcal{B}_t$  **do**
- 7:              $A = g(x)$
- 8:              $A' = \text{TR}_{H,W,K}(\mathbb{1}\{|\alpha^c| > Q_\mu\} \otimes \mathbf{J}_{H,W}) \odot A$  ▷ Channel-wise sparsity
- 9:              $\mathcal{M}_t \leftarrow \mathcal{M}_t \cup (A', y)$  ▷ Update episodic memory
- 10:         **end for**
- 11:     **end for**
- 12:      $\omega \leftarrow \text{OPT}_\omega(\omega, A_t)$  as Eq. 7 ▷ Associative memory *write*
- 13:     **if**  $t \% R == 0$  **then**
- 14:         Bayesian Training by  $\omega \leftarrow \text{OPT}_\omega(\omega)$  ▷ Associative memory *forgetting*
- 15:     **end if**
- 16: **end for**

---

memories has been rekindled, and several works have been proposed to achieve strong memory recall performance. However, we notice that how to leverage associative memory in the continual learning setting is under-explored.

In this work, our goal is to design a neuro-inspired continual learning method to improve memory efficiency and mitigate forgetting. To this end, associative memory becomes a natural fit for us due to its properties such as content-based retrieval, fast and efficient recall, and high noise tolerance. Content-based retrieval and noise tolerance allow us to increase the sparsity of the masked feature maps and achieve maximal memory saving. The fast recall reduces the computational overhead for memory retrieval which is critical for our method to be applied to various replay-based baselines.

Formally, an associative memory  $\mathcal{A}(\mathbf{x}, \omega)$  can be implemented as a recurrent or feed-forward neural network, where  $\mathbf{x}$  and  $\omega$  denote the input and model parameters of the associative memory, respectively. Corresponding to the “memorize” and “recall” in the human brain, associative memory has *read* and *write* operations which are implemented based on an *energy function*. For example, the predictive-coding-based energy function Rao & Ballard (1999) is the sum of prediction errors across all network layers, i.e.,

$$E(\mathbf{x}_{0:L}, \omega_{0:L}) = \|\mathbf{x}_L - \omega_L\|_2^2 + \lambda \sum_{\ell=0}^{L-1} \|\mathbf{x}_\ell - \mathcal{A}_\ell(\mathbf{x}_{\ell+1}, \omega_\ell)\|_2^2, \quad (5)$$

where  $\ell$  is the layer index and  $\lambda$  is a coefficient. During training, we write the ground-truth feature map  $A$  into associative memory, by minimizing Eq. 5 w.r.t. parameter  $\omega$  while keeping the input  $\mathbf{x} = A$ . During inference, given the masked feature map  $A'$  defined in Eq. 4, we retrieve the ground-truth  $A$  for memory replay, by minimizing Eq. 5 w.r.t. input  $\mathbf{x}$  initialized as  $A'$  while keeping the parameter  $\omega$  fixed.

### 4.3 TRAINING PIPELINE

Our proposed method involves training the classifier  $f_\theta$  and the associative memory  $\mathcal{A}$  while maintaining a small episodic memory  $\mathcal{M}$ . The overall training procedure is shown in Algorithm 1.

**Training classifier.** In each incremental phase, we update the model parameters of the continual learning classifier  $f_\theta$  by using the new coming data and the replay samples. The key here is that we use associative memory to retrieve the “complete” feature map and then feed it to the classifier for memory replay. Formally, the training objective of the classifier can be formulated as follows

$$\theta^* = \underset{\theta}{\operatorname{argmin}} \sum_{(\mathbf{x}, t, \mathbf{y})} \ell(f_\theta(g(\mathbf{x}), t), \mathbf{y}) + \sum_{k < t} \ell(f_\theta, \tilde{\mathcal{M}}_k) \quad (6)$$

where  $\tilde{\mathcal{M}}_k = \underset{\mathbf{x}}{\operatorname{argmin}} E(\mathbf{x}_{0:L}, \omega_{0:L})$  with  $\mathbf{x}_0$  initialized as  $A'_k \in \mathcal{M}_k$ ,

where the first row is the continual learning objective as defined in Eq. 2.  $\tilde{\mathcal{M}}_k$  denotes the retrieved episodic memory, i.e., the feature maps recalled by associative memory. As mentioned earlier, the retrieval by associative memory corresponds to solving an optimization problem as shown in the second row, where  $A'_k$  is the masked (before retrieval) feature map from task  $k$ .

**Training associative memory.** Given feature maps coming from new tasks in each incremental, we need to write those feature maps into the associative memory such that we can ask it to retrieve the complete feature map given a partial cue at a later timestamp. Formally, given feature maps  $A_t$  from new task  $t$ , writing them into associative memory corresponds to solving the following optimization problem

$$\omega = \operatorname{argmin}_{\omega} E(\mathbf{x}_{0:L}, \omega_{0:L}) \text{ with } \mathbf{x}_0 \text{ fixed as } A_t. \quad (7)$$

We also proposed a memory-forgetting mechanism for associative memory to avoid potential memory overload during continual learning. As more data are observed when new tasks keep coming, it is natural for the classifier and associative memory to be biased more towards new tasks' data rather than old ones. To this end, our forgetting mechanism *erases* old tasks' data more times than new ones thus satisfying the inductive bias we want to introduce.

**Update Episodic Memory.** For simplicity, we assume that the memory is populated with the last  $m$  examples from each task, although better memory update strategies could be employed (such as building a coreset per task, reservoir sampling, etc.)

## 5 EXPERIMENT

In this section, we evaluate our proposed method SHARC on Class-IL and Task-IL CL. Both performance tables and learning curves over entire tasks are provided. In addition, we demonstrate sensitivity analyses over the masking threshold and comparison of different associative memories. All experiments are conducted on a 64-bit machine with an NVIDIA T4 Tensor Core GPU which has 320 Turing Tensor cores, 2560 CUDA cores, 16GB memory, and Intel® Xeon® Platinum 8259CL CPU @ 2.50GHz. The anonymous code of our method can be found [here](#).

### 5.1 EXPERIMENT SETTING

**Dataset.** In our research, we conducted experiments on three datasets: Split CIFAR-10, Split CIFAR-100, and Split mini-ImageNet Chaudhry et al. (2019). CIFAR-10 consists of 50,000 RGB training images and 10,000 test images, categorized into 10 object classes. Similarly, CIFAR-100 extends this classification task by including 100 classes, with each class containing 600 images. ImageNet-50 comprises 50 classes with 1300 training images and 50 validation images per class. We divided the Split CIFAR-10 dataset into 5 tasks, each with 2 classes. For Split CIFAR-100 and Split mini-ImageNet, we expanded our investigation to 20 tasks, each with 5 classes.

**Comparison Methods.** We compare our method with several replay-based continual learning methods, including: **ER**, a rehearsal-based method that utilizes the average of parameter update gradients from the current task's samples alongside samples from episodic memory to update the learning agent. **MER**, a rehearsal-based model that harnesses the power of episodic memory Riemer et al. (2018). **GEM**, one ensures that valuable information from prior tasks is retained while accommodating new learning Lopez-Paz & Ranzato (2017). **A-GEM**, takes a step further than GEM by incorporating an adaptive mechanism that updates the model's parameters based on both the current task's gradient and the gradients of previous tasks stored in the episodic memory. **CLS-ER**, an innovative algorithm that utilizes a dual-memory learning mechanism to enhance performance in continual learning tasks Arani et al. (2022). **DER++**, a combination of rehearsal, knowledge distillation, and regularization techniques Buzzega et al. (2020).

**Evaluation Metrics.** We assess the classification performance using the ACC metric, which represents the average test classification accuracy across all tasks. We also measure backward transfer (BWT Lopez-Paz & Ranzato (2017)) to evaluate the impact of new learning on previous knowledge. Negative BWT indicates forgetting, so a higher value is preferable. Detailed experimental settings can be found in the appendix.

**Training Details.** We utilized a frozen pre-trained model ImageNet-1K, retaining only the MLP part for training. We directly employed the feature map as the input and output of the associative memory, storing it in a memory buffer. For further details, please refer to the appendix.

Table 1: **Performance comparison on image classification datasets (Task-IL)**. The mean and standard deviation are calculated based on five runs with varying seeds.  $^+$  denotes the corresponding method combined with our SHARC framework. In most cases, our proposed SHARC framework significantly improves the method.

Buffer	Model	S-CIFAR-10		S-CIFAR-100		S-Mini-ImgNet		
		ACC ( $\uparrow$ )	BWT ( $\uparrow$ )	ACC ( $\uparrow$ )	BWT ( $\uparrow$ )	ACC ( $\uparrow$ )	BWT ( $\uparrow$ )	
-	JOINT	93.49 $\pm$ 0.61	43.14 $\pm$ 2.07	87.57 $\pm$ 0.89	67.99 $\pm$ 1.53	74.95 $\pm$ 0.7	70.02 $\pm$ 0.81	
-	SGD	92.31 $\pm$ 0.54	-0.38 $\pm$ 0.82	85.83 $\pm$ 0.35	3.08 $\pm$ 2.14	76.2 $\pm$ 0.41	3.98 $\pm$ 0.75	
200	GEM	88.44 $\pm$ 1.11	-4.6 $\pm$ 2.24	82.82 $\pm$ 0.62	0.2 $\pm$ 1.69	72.23 $\pm$ 1.26	-1 $\pm$ 1.82	
	<b>GEM<math>^+</math></b>	91.01 $\pm$ 1.04	-1.5 $\pm$ 1.31	83.88 $\pm$ 0.52	-0.04 $\pm$ 1.2	76.13 $\pm$ 0.98	3.49 $\pm$ 1.47	
	A-GEM	90.52 $\pm$ 3.29	-1.07 $\pm$ 1.57	85.33 $\pm$ 0.58	2 $\pm$ 1	75.18 $\pm$ 1.11	2.33 $\pm$ 1.5	
	<b>A-GEM<math>^+</math></b>	91.72 $\pm$ 1.02	-0.35 $\pm$ 2.08	85.55 $\pm$ 0.88	1.25 $\pm$ 0.63	76.54 $\pm$ 0.97	4.14 $\pm$ 1.71	
	ER	86.36 $\pm$ 1.33	-5.04 $\pm$ 1.72	82.55 $\pm$ 0.47	-0.71 $\pm$ 1.59	71.66 $\pm$ 1.44	-1.53 $\pm$ 2.04	
	<b>ER<math>^+</math></b>	91.48 $\pm$ 1.18	-0.93 $\pm$ 1.54	84.55 $\pm$ 0.62	0.71 $\pm$ 0.61	73.68 $\pm$ 0.59	0.71 $\pm$ 0.97	
	MER	87.32 $\pm$ 1.39	-2.3 $\pm$ 3.83	82.04 $\pm$ 0.63	-0.83 $\pm$ 1.49	71.2 $\pm$ 1.43	-1.99 $\pm$ 1.53	
	<b>MER<math>^+</math></b>	91.14 $\pm$ 1.62	-0.9 $\pm$ 2.46	84.3 $\pm$ 0.92	0.41 $\pm$ 1.26	73.54 $\pm$ 0.58	0.7 $\pm$ 1.06	
	DER++	84.94 $\pm$ 1.95	-6.45 $\pm$ 1.91	83.27 $\pm$ 0.76	0.32 $\pm$ 1.47	72.92 $\pm$ 1.09	-0.13 $\pm$ 1.44	
	<b>DER++<math>^+</math></b>	89.89 $\pm$ 1.34	-2.72 $\pm$ 2.04	84.96 $\pm$ 0.97	0.62 $\pm$ 1.44	74.59 $\pm$ 0.87	2 $\pm$ 1.27	
	CLS-ER	80.97 $\pm$ 2.11	-12.6 $\pm$ 4.39	82.97 $\pm$ 0.32	-1.95 $\pm$ 1.5	73.67 $\pm$ 1.05	-1.75 $\pm$ 0.4	
	<b>CLS-ER<math>^+</math></b>	91.39 $\pm$ 0.7	-0.94 $\pm$ 1.18	85 $\pm$ 0.41	1.27 $\pm$ 0.68	77 $\pm$ 0.45	2.7 $\pm$ 0.97	
	500	GEM	88.02 $\pm$ 2.61	-3.84 $\pm$ 1.19	82.81 $\pm$ 0.66	0.06 $\pm$ 1.66	73.6 $\pm$ 1.13	0.23 $\pm$ 1.27
		<b>GEM<math>^+</math></b>	91.53 $\pm$ 1.17	-0.05 $\pm$ 1.58	84.37 $\pm$ 1.03	1.63 $\pm$ 0.79	75.56 $\pm$ 0.93	3.2 $\pm$ 1.61
A-GEM		90.81 $\pm$ 2.97	0.31 $\pm$ 2.81	85.44 $\pm$ 0.28	2.32 $\pm$ 0.84	75.59 $\pm$ 1.15	2.78 $\pm$ 1.61	
<b>A-GEM<math>^+</math></b>		92.32 $\pm$ 0.67	0.22 $\pm$ 0.7	85.89 $\pm$ 0.82	3.25 $\pm$ 1.01	75.65 $\pm$ 0.86	3.21 $\pm$ 1.56	
ER		88.05 $\pm$ 1.51	-1.88 $\pm$ 3.62	82.7 $\pm$ 0.63	-0.09 $\pm$ 0.5	71.83 $\pm$ 1.17	-1.36 $\pm$ 1.5	
<b>ER<math>^+</math></b>		91.64 $\pm$ 0.66	-0.82 $\pm$ 0.67	84.77 $\pm$ 1.57	1.85 $\pm$ 1.72	72.94 $\pm$ 0.63	0.41 $\pm$ 1.08	
MER		88.33 $\pm$ 1.87	-3.35 $\pm$ 1.6	82.11 $\pm$ 0.5	-0.36 $\pm$ 0.91	70.69 $\pm$ 1.07	-2.28 $\pm$ 1.41	
<b>MER<math>^+</math></b>		91.69 $\pm$ 0.73	-0.39 $\pm$ 1.56	84.06 $\pm$ 1.33	1.43 $\pm$ 1.46	72.63 $\pm$ 0.37	-0.28 $\pm$ 1.08	
DER++		86.73 $\pm$ 2.77	-4.79 $\pm$ 1.26	83.04 $\pm$ 0.58	0.76 $\pm$ 1.64	72.05 $\pm$ 0.87	-0.95 $\pm$ 1.52	
<b>DER++<math>^+</math></b>		90.46 $\pm$ 1.3	-2.71 $\pm$ 1.01	85.13 $\pm$ 1.57	1.81 $\pm$ 1.01	73.85 $\pm$ 0.79	1.5 $\pm$ 1.5	
CLS-ER		82.54 $\pm$ 3.06	-8.64 $\pm$ 5.52	81.34 $\pm$ 0.9	-2.27 $\pm$ 1.8	72.11 $\pm$ 0.38	-3.41 $\pm$ 0.88	
<b>CLS-ER<math>^+</math></b>		90.94 $\pm$ 1.49	-2.22 $\pm$ 1.51	85.36 $\pm$ 0.83	1.82 $\pm$ 1.66	76.27 $\pm$ 0.52	2.05 $\pm$ 0.82	

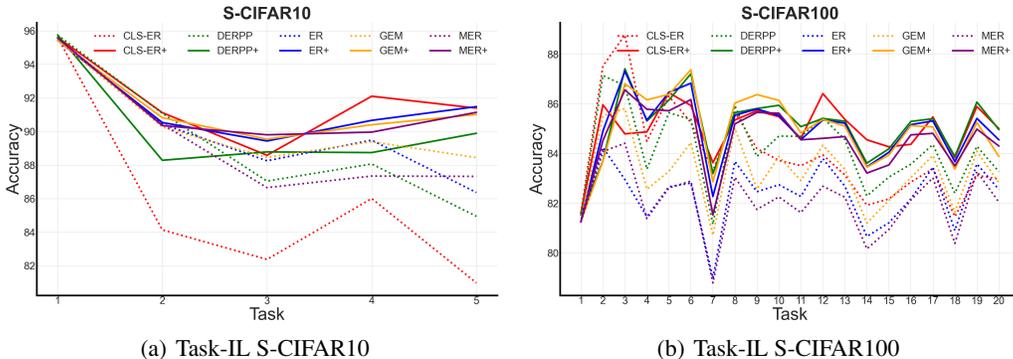


Figure 3: Learning curves of multiple models with/without SHARC on S-CIFAR10 and S-CIFAR100 in Task-IL scenario. Models with/without SHARC are shown in solid/dotted lines. The buffer size for all models is 200. Methods combined with our proposed framework SHARC significantly prevails.

## 5.2 PERFORMANCE COMPARISON

We demonstrate the impact of our proposed SHARC framework on several state-of-the-art replay-based approaches. Naive baselines such as SGD refer to standard training, while JOINT refers to joint training on all tasks, which provides an upper bound. The experimental results are shown in Table 1 and Table 2, which contain the results in the Task-IL scenario and Class-IL scenario, respectively. Buffer size controls the budget of episodic memory and is distributed evenly to all tasks. Each slot in the buffer contains feature maps of a sample, instead of an image.

Table 1 compares six replay-based methods before and after combining them with SHARC in the Task-IL scenario. Overall, in most cases, the methods used in conjunction with SHARC offer significant improvements. Such contrast exists in all settings (different datasets, models, and buffer sizes). In particular, CLS-ER equipped with SHARC achieves a 12.9% improvement in ACC on

Table 2: **Performance comparison on image classification datasets (Class-IL).** The mean and standard deviation are calculated based on five runs with varying seeds. <sup>+</sup> denotes the corresponding method combined with our SHARC framework. In most cases, our proposed SHARC framework significantly improves the method.

Buffer	Model	S-CIFAR-10		S-CIFAR-100		S-Mini-ImgNet	
		ACC ( $\uparrow$ )	BWT ( $\uparrow$ )	ACC ( $\uparrow$ )	BWT ( $\uparrow$ )	ACC ( $\uparrow$ )	BWT ( $\uparrow$ )
-	JOINT	72.85 $\pm$ 2.18	61.26 $\pm$ 8.55	45.87 $\pm$ 1.22	45.55 $\pm$ 1.45	47.08 $\pm$ 0.77	46.25 $\pm$ 0.98
-	SGD	20.47 $\pm$ 0.78	-90.16 $\pm$ 0.92	8.55 $\pm$ 1.39	-78.24 $\pm$ 0.93	12.21 $\pm$ 0.75	-67.11 $\pm$ 0.77
200	GEM	27.01 $\pm$ 6.16	-76.88 $\pm$ 6.93	16.38 $\pm$ 3.06	-67.72 $\pm$ 3.84	23.76 $\pm$ 2.65	-54.22 $\pm$ 3.47
	<b>GEM<sup>+</sup></b>	36.39 $\pm$ 4.04	-63.03 $\pm$ 7.86	20.77 $\pm$ 3.22	-63.08 $\pm$ 3.33	24.76 $\pm$ 1.86	-52.31 $\pm$ 2.62
	A-GEM	24.12 $\pm$ 6.92	-83.48 $\pm$ 3.5	12.75 $\pm$ 4.2	-73.96 $\pm$ 4.68	16.77 $\pm$ 2.05	-62.51 $\pm$ 2.71
	<b>A-GEM<sup>+</sup></b>	29.19 $\pm$ 2.31	-77.82 $\pm$ 3.52	17.55 $\pm$ 1.91	-69.61 $\pm$ 1.6	18.96 $\pm$ 1.53	-60.02 $\pm$ 2.11
	ER	29.35 $\pm$ 7.79	-65.76 $\pm$ 10.54	14.81 $\pm$ 2.75	-70.15 $\pm$ 3.41	21.71 $\pm$ 1.61	-56.46 $\pm$ 1.88
	<b>ER<sup>+</sup></b>	33.94 $\pm$ 6.24	-62 $\pm$ 10.91	20.58 $\pm$ 1.86	-63.09 $\pm$ 1.69	23.36 $\pm$ 1.76	-54.24 $\pm$ 2.46
	MER	30.02 $\pm$ 7.64	-61.33 $\pm$ 7.94	13.74 $\pm$ 3.33	-70.76 $\pm$ 3.48	21.46 $\pm$ 1.78	-56.59 $\pm$ 1.89
	<b>MER<sup>+</sup></b>	34.81 $\pm$ 5.44	-59.86 $\pm$ 8.47	19.46 $\pm$ 1.3	-64.09 $\pm$ 1.96	22.83 $\pm$ 1.88	-54.68 $\pm$ 2.46
	DER++	31.55 $\pm$ 4.61	-55.68 $\pm$ 10.28	14.44 $\pm$ 6.09	-69.28 $\pm$ 6.37	23.03 $\pm$ 1.64	-54.54 $\pm$ 1.74
	<b>DER++<sup>+</sup></b>	38.53 $\pm$ 4.84	-39.15 $\pm$ 7.8	21.01 $\pm$ 1.8	-62.08 $\pm$ 2.37	24.76 $\pm$ 1.3	-52.16 $\pm$ 1.58
	CLS-ER	27.3 $\pm$ 3.13	-56.75 $\pm$ 14.14	15.83 $\pm$ 1.84	-70.13 $\pm$ 2.49	21.77 $\pm$ 1.43	-58.55 $\pm$ 2.09
	<b>CLS-ER<sup>+</sup></b>	28.63 $\pm$ 5.43	-51.37 $\pm$ 8.46	18.77 $\pm$ 1.93	-64.2 $\pm$ 1.36	22.86 $\pm$ 1.81	-57.38 $\pm$ 2.31
500	GEM	30.12 $\pm$ 10.19	-63.83 $\pm$ 14.88	20.81 $\pm$ 5.66	-59.13 $\pm$ 6.09	30.88 $\pm$ 2.39	-45.43 $\pm$ 3.04
	<b>GEM<sup>+</sup></b>	33.52 $\pm$ 5.29	-51.17 $\pm$ 14.8	24.89 $\pm$ 1.19	-54.14 $\pm$ 3.14	29.8 $\pm$ 1.76	-46.46 $\pm$ 2.48
	A-GEM	25 $\pm$ 7.05	-80.94 $\pm$ 2.84	12.25 $\pm$ 3.87	-74.35 $\pm$ 4.45	17.04 $\pm$ 0.91	-62.21 $\pm$ 1.11
	<b>A-GEM<sup>+</sup></b>	29.21 $\pm$ 4.61	-77.99 $\pm$ 5.2	15.02 $\pm$ 2.02	-71.12 $\pm$ 2.31	17.33 $\pm$ 1.48	-61.45 $\pm$ 2.16
	ER	33.16 $\pm$ 8.18	-55.8 $\pm$ 12.65	19.37 $\pm$ 3.9	-61.39 $\pm$ 3.89	27.39 $\pm$ 1.52	-49.06 $\pm$ 2.18
	<b>ER<sup>+</sup></b>	36.18 $\pm$ 2.22	-53.67 $\pm$ 6.82	24.85 $\pm$ 1.41	-53.85 $\pm$ 1.38	26.69 $\pm$ 0.95	-49.48 $\pm$ 1.61
	MER	32.96 $\pm$ 8.35	-53.78 $\pm$ 14.59	20.51 $\pm$ 3.97	-58.51 $\pm$ 3.99	26.46 $\pm$ 1.83	-49.5 $\pm$ 2.45
	<b>MER<sup>+</sup></b>	36.04 $\pm$ 2.98	-51.5 $\pm$ 2.79	23.4 $\pm$ 1.52	-53.89 $\pm$ 1.8	26.1 $\pm$ 1.22	-50.08 $\pm$ 1.35
	DER++	34.21 $\pm$ 8.62	-43.95 $\pm$ 16.78	16.95 $\pm$ 5.86	-63.53 $\pm$ 5.67	26.48 $\pm$ 1.67	-49.82 $\pm$ 2.39
	<b>DER++<sup>+</sup></b>	30.57 $\pm$ 8.35	-45.5 $\pm$ 8	24.66 $\pm$ 1.83	-54.62 $\pm$ 2.09	25.54 $\pm$ 2.78	-50.72 $\pm$ 3.45
	CLS-ER	26.97 $\pm$ 3.32	-41.42 $\pm$ 5.8	22.1 $\pm$ 2.04	-58.79 $\pm$ 3.31	26.74 $\pm$ 1.49	-52.2 $\pm$ 1.73
	<b>CLS-ER<sup>+</sup></b>	29.64 $\pm$ 5.37	-77.81 $\pm$ 5.74	22.5 $\pm$ 2.52	-57.04 $\pm$ 3.65	26.31 $\pm$ 1.55	-52.68 $\pm$ 0.55

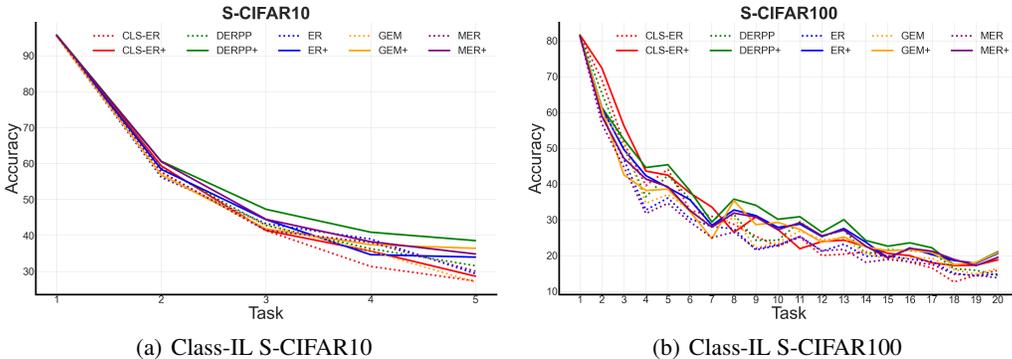


Figure 4: Learning curves of multiple models with/without SHARC on S-CIFAR10 and S-CIFAR100 in Class-IL scenario. Models with/without SHARC are shown in solid/dotted lines. The buffer size for all models is 200.

S-CIFAR10 with buffer size 200. From a methodological perspective, rehearsal-based methods (e.g., ER) offer greater improvements than constraint-based methods (e.g., GEM). As a typical example, the performance of A-GEM improves only slightly when used with SHARC on S-CIFAR100, which is reasonable since we keep the batch size of the retrieval process constant. Rehearsal-based methods can benefit more from masking because masking reduces the memory space for samples, allowing more previous samples to be reviewed. Furthermore, in most cases on S-CIFAR100 and S-MiniImgNet, the BWT increases or even becomes positive when using SHARC, indicating that SHARC is highly resistant to forgetting. As the buffer size decreases, the complexity of the task increases. Achieving good performance with smaller buffer sizes is the spirit of continual learning. Based on this consideration, we further investigate the learning curve for a minimum buffer size of 200. As shown in Figure 4, methods equipped with SHARC clearly prevail in the figure, indicating that they have been steadily improved during the learning process.

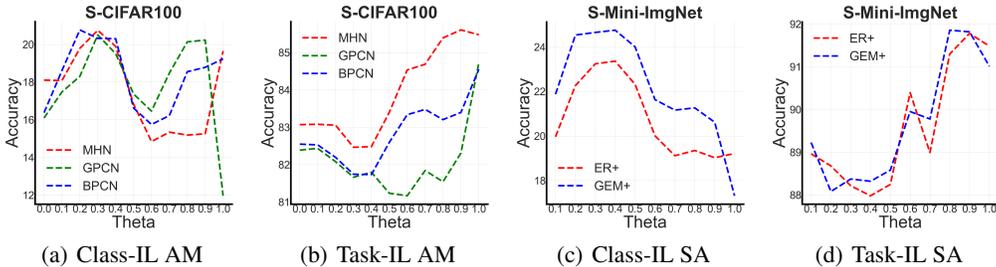


Figure 5: Average accuracy of GEM and ER using different associative memories for inpainting on S-CIFAR100. Figure (a) shows ACC of GEM using different associative memories in Class-IL scenario. Figure (b) shows ACC of ER using different associative memories in Task-IL scenario. The buffer size for all models is 200.

Since the corrupted feature maps after masking cannot be used for backpropagation, in order to keep the backbone network the same as the rest of the work, we froze the convolutional basis of the pre-trained ResNet18, leaving only the parameters of the fully connected layer available for training. A single-layer classifier may not be sufficient in Class-IL, causing all models to perform poorly. Still, this is enough to illustrate the effectiveness of our proposed SHARC framework. Table 2 compares six replay-based methods before and after combining them with SHARC in a Class-IL scenario. Overall, in most cases, the methods used in conjunction with SHARC offer significant improvements. In particular, DER++ equipped with SHARC achieves a 45.5% improvement in ACC on S-CIFAR100. The smaller the buffer, the more pronounced this contrast becomes. In particular, with a buffer size of 200, SHARC improves CLS-ER in ACC much greater compared to the buffer size of 500 on S-CIFAR100.

### 5.3 ASSOCIATIVE MEMORY COMPARISON

As shown in Figure 5(a) and Figure 5(b), we compare different configurations of the associative memory. We in general follow Yoo & Wood (2022) for the implementation. We found that Modern Hopfield Network Ramsauer et al. (2020) favors more towards the task-incremental setting while BayesPCN Yoo & Wood (2022) favors more towards class-incremental setting. This is potentially due to that, in BayesPCN a forgetting mechanism is implemented, which can help mitigate the memory overload of the associative memory when the samples to memorize are too many. Such design seems to play a more important role in the more challenging class-incremental setting.

### 5.4 MASKING THRESHOLD SENSITIVITY ANALYSIS

We conduct sensitivity analysis on the threshold  $Q_\mu$  in Eq 4. Masking threshold  $Q_\mu$  is defined as certain percentile value for 512 feature importance. Features with importance below the threshold will be masked.  $Q_\mu$  is controlled by a hyper-parameter Theta. Theta ranges between (0, 1], where 1 means that only the features with the highest importance are retained. While different methods may favor different optimal thresholds, a general sensitivity analysis is still helpful in determining optimal threshold settings. As shown in Figure 5(d), the optimal Theta in Task-IL is between 0.8 to 0.9. As shown in Figure 5(c), the optimal Theta in Class-IL is between 0.3 to 0.5. This indicates that the optimal thresholds show different trends in Task-IL and Class-IL. In particular, Task-IL requires a higher threshold to drop most features, while the opposite is true for Class-IL. This is reasonable because in Task-IL, the task ID is given as additional information. Whereas in Class-IL, the model can only get additional information from more features.

## 6 CONCLUSION

We propose SHARC, a novel framework that bridges the gap between current AI models and humans in continual learning. Combining associative memory and interpretive techniques, SHARC enables efficient, near-perfect recall of seen samples in a human-like manner. As a generic framework, SHARC can be seamlessly adapted to any replay-based approach, thus improving their performance in different continual learning scenarios. We demonstrate the effectiveness of our framework with abundant experimental results. Our proposed SHARC framework consistently improves several SOTA replay-based methods on multiple benchmark datasets.

## REFERENCES

- Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient based sample selection for online continual learning. *arXiv preprint arXiv:1903.08671*, 2019.
- Daniel J Amit and Daniel J Amit. *Modeling brain function: The world of attractor neural networks*. Cambridge university press, 1989.
- Elahe Arani, Fahad Sarfraz, and Bahram Zonooz. Learning fast, learning slow: A general continual learning method based on complementary learning system. *arXiv preprint arXiv:2201.12604*, 2022.
- Guangji Bai, Chen Ling, Yuyang Gao, and Liang Zhao. Saliency-augmented memory completion for continual learning. In *Proceedings of the 2023 SIAM International Conference on Data Mining (SDM)*, pp. 244–252. SIAM, 2023.
- Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: a strong, simple baseline. *Advances in neural information processing systems*, 33:15920–15930, 2020.
- Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a-gem. *arXiv preprint arXiv:1812.00420*, 2018.
- Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and Marc’Aurelio Ranzato. On tiny episodic memories in continual learning. *arXiv preprint arXiv:1902.10486*, 2019.
- Sayna Ebrahimi, Suzanne Petryk, Akash Gokul, William Gan, Joseph E Gonzalez, Marcus Rohrbach, and Trevor Darrell. Remembering for the right reasons: Explanations reduce catastrophic forgetting. *Applied AI letters*, 2(4):e44, 2021.
- Omar Elharrouss, Noor Almaadeed, Somaya Al-Maadeed, and Younes Akbari. Image inpainting: A review. *Neural Processing Letters*, 51:2007–2028, 2020.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.
- Tyler L Hayes, Kushal Kafle, Robik Shrestha, Manoj Acharya, and Christopher Kanan. Remind your neural network to prevent catastrophic forgetting. In *European Conference on Computer Vision*, pp. 466–483. Springer, 2020.
- Donald Hebb. *The organization of behavior*. New York, 1949.
- John J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558, 1982.
- Yanping Huang and Rajesh PN Rao. Predictive coding. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2(5):580–593, 2011.
- R Insausti, MP Marcos, A Mohedano-Moriano, MM Arroyo-Jiménez, M Córcoles-Parada, E Artacho-Pérula, MM Ubero-Martinez, and M Munoz-Lopez. The nonhuman primate hippocampus: neuroanatomy and patterns of cortical connectivity. *The hippocampus from cells to systems: Structure, connectivity, and functional contributions to memory and flexible cognition*, pp. 3–36, 2017.
- Daoyun Ji and Matthew A Wilson. Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nature neuroscience*, 10(1):100–107, 2007.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.
- Dmitry Krotov and John Hopfield. Large associative memory problem in neurobiology and machine learning. *arXiv preprint arXiv:2008.06996*, 2020.

- Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017.
- David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30:6467–6476, 2017.
- Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pp. 109–165. Elsevier, 1989.
- German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71, 2019.
- Jialun Peng, Dong Liu, Songcen Xu, and Houqiang Li. Generating diverse structure for image inpainting with hierarchical vq-vae. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10775–10784, 2021.
- Yunchen Pu, Zhe Gan, Ricardo Henao, Xin Yuan, Chunyuan Li, Andrew Stevens, and Lawrence Carin. Variational autoencoder for deep learning of images, labels and captions. *Advances in neural information processing systems*, 29:2352–2360, 2016.
- Hubert Ramsauer, Bernhard Schäfl, Johannes Lehner, Philipp Seidl, Michael Widrich, Thomas Adler, Lukas Gruber, Markus Holzleitner, Milena Pavlović, Geir Kjetil Sandve, et al. Hopfield networks is all you need. *arXiv preprint arXiv:2008.02217*, 2020.
- Rajesh PN Rao and Dana H Ballard. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience*, 2(1):79–87, 1999.
- Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp. 2001–2010, 2017.
- 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. *arXiv preprint arXiv:1810.11910*, 2018.
- Anthony Robins. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connection Science*, 7(2): 123–146, 1995.
- Edmund T Rolls. The mechanisms for pattern completion and pattern separation in the hippocampus. *Frontiers in systems neuroscience*, 7:74, 2013.
- Gobinda Saha and Kaushik Roy. Saliency guided experience packing for replay in continual learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 5273–5283, 2023.
- Tommaso Salvatori, Yuhang Song, Yujian Hong, Lei Sha, Simon Frieder, Zhenghua Xu, Rafal Bogacz, and Thomas Lukasiewicz. Associative memories via predictive coding. *Advances in Neural Information Processing Systems*, 34:3874–3886, 2021.
- Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pp. 618–626, 2017.
- Yi Sun, Xiaogang Wang, and Xiaoou Tang. Deeply learned face representations are sparse, selective, and robust. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2892–2900, 2015.
- Timothy J Teyler and Jerry W Rudy. The hippocampal indexing theory and episodic memory: updating the index. *Hippocampus*, 17(12):1158–1169, 2007.
- Jinsoo Yoo and Frank Wood. Bayespcn: A continually learnable predictive coding associative memory. *Advances in Neural Information Processing Systems*, 35:29903–29914, 2022.
- Wang Zhou, Shiyu Chang, Norma Sosa, Hendrik Hamann, and David Cox. Lifelong object detection. *arXiv preprint arXiv:2009.01129*, 2020.

## A APPENDIX

### A.1 EXPERIMENTAL DETAILS

#### A.1.1 DATASET DETAILS

We expand upon the datasets used for our experiments in this section. We highlighted the sentence that describes the domain drift within each dataset.

- **CIFAR-10:** The CIFAR-10 dataset is a comprehensive collection of 60,000  $32 \times 32$  color images divided into 10 distinct classes, with 6,000 images per class. This dataset is further split into 50,000 training images and 10,000 test images, allowing for effective model evaluation.
- **CIFAR-100:** The dataset is similar to CIFAR-10 and is composed of 100 classes, each containing 600 images. Specifically, within the 600 images, there are 500 images used for training and 100 images designated for testing purposes. It is important to note that the 100 classes are actually comprised of 20 classes, where each class is further divided into 5 sub-classes. Therefore, the total count of 100 classes is obtained by multiplying 5 and 20 ( $5 \times 20 = 100$ ).
- **Mini-ImageNet:** The Mini-ImageNet dataset contains 100 classes with a total of 60,000 color images. Each class has 600 samples, and the size of each image is  $84 \times 84$  pixels. Typically, the class distribution between the training and testing sets of this dataset is 80 : 20. Compared to the CIFAR-10 dataset, the Mini-ImageNet dataset is more complex but is better suited for prototype design and experimental research. Chaudhry et al. (2019)

#### A.1.2 DETAILS OF COMPARISON METHODS

In this paper, we compare our proposed SHARC with several SOTA replay-based methods as well as regularization-based methods. Specifically,

- **ER**, a rehearsal-based method that utilizes the average of parameter update gradients from the current task’s samples alongside samples from episodic memory to update the learning agent. This method, known as ER (Episodic Regularization), offers a computationally efficient alternative to GEM (Gradient Episodic Memory) and has demonstrated successful performance when dealing with a limited memory buffer.
- **MER**, a rehearsal-based model that harnesses the power of an episodic memory. MER employs a unique loss function that approximates the dot products of the gradients of current and previous tasks, thereby mitigating the issue of forgetting. To ensure a fair and comprehensive comparison with other methods, we adjust the experimental setting by setting the number of inner gradient steps to 1 for each outer meta-update, while maintaining a mini-batch size of 10. This adjustment allows us to establish a more consistent framework for evaluating the performance of MER alongside other approaches, specifically in terms of the number of stochastic gradient descent (SGD) updates. By presenting these findings, we aim to shed light on the effectiveness and practicality of MER as a rehearsal-based model in the context of meta-learning Riemer et al. (2018).
- **GEM**, who utilizes an episodic memory buffer to store past experiences and gradients. By incorporating both the current task’s gradient and the gradients of previous tasks from the episodic memory, GEM ensures that valuable information from prior tasks is retained while accommodating new learning. To prevent catastrophic forgetting, the algorithm employs a constrained optimization approach, projecting the current gradient onto a subspace that preserves knowledge from previous tasks Lopez-Paz & Ranzato (2017).
- **A-GEM**, takes a step further than GEM by incorporating an adaptive mechanism that updates the model’s parameters based on both the current task’s gradient and the gradients of previous tasks stored in the episodic memory. This allows AGEM to effectively preserve knowledge from prior tasks while adapting to new tasks.
- **CLS-ER**, an innovative algorithm that utilizes a dual-memory learning mechanism to enhance performance in continual learning tasks. In this approach, the episodic memory serves as a repository for storing samples encountered during the learning process. On the other hand, semantic memories play a crucial role in constructing short-term and long-term memories of the learned representations from the working model Arani et al. (2022).

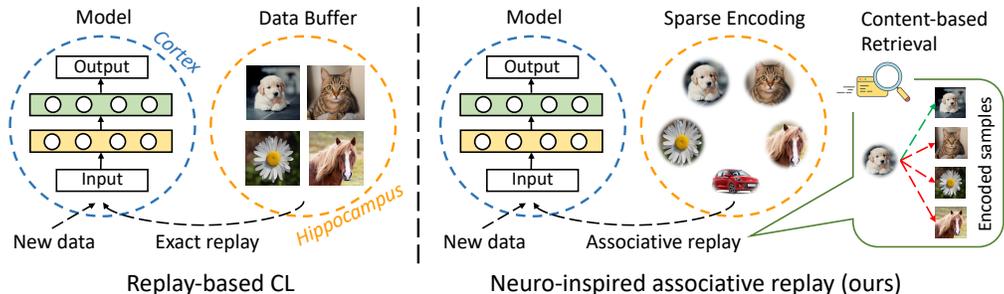


Figure 6: Motivation and overview of our proposed framework. **Left:** Typical replay-based approaches maintain and replay a small episodic memory of previous samples, which is inspired by the cortex and hippocampus in the human brain. **Right:** In FaR, our memory buffer is equipped with a forgetting mechanism to drop uninformative episodes, and a content-addressable associative memory is used to achieve fast and high-accuracy data retrieval.

- **DER++**, a combination of rehearsal, knowledge distillation, and regularization techniques. This approach leverages the network’s logits sampled at different stages of the optimization trajectory. This approach promotes consistency with the network’s past experiences Buzzega et al. (2020).

### A.1.3 HYPER-PARAMETER SETTING

All experiments are conducted on a 64-bit machine with an NVIDIA T4 Tensor Core GPU which has 320 Turing Tensor cores, 2560 CUDA cores, 16GB memory, and Intel® Xeon® Platinum 8259CL CPU @ 2.50GHz. The learning rate for all datasets is uniformly set to be 0.1. Down below we report the hyper-parameter unique to some models applied in our experiment.

- **ER:**  
'lr': 0.1
- **MER:**  
'lr': 0.1, 'gamma': 0.5, 'batch num': 1
- **GEM:**  
'lr': 0.1, 'gamma': 0.5
- **AGEM:**  
'lr': 0.1
- **DER++:**  
'lr': 0.1, 'alpha': 0.1, 'beta': 0.5
- **CLSER:**  
'lr': 0.1, 'reg weight': 0.15, 'stable model update freq': 0.1, 'stable model alpha': 0.999, 'plastic model update freq': 0.3, 'plastic model alpha': 0.999

## A.2 ADDITIONAL PRELIMINARIES

In this section, we provide more discussion about some preliminaries in this paper.

### A.3 MOTIVATION OF SHARC

#### A.3.1 ASSOCIATIVE MEMORY

When applied to Computer Science problems, associative memories come in two high-level forms: *auto-associative* and *hetero-associative* memories. While both are able to recall patterns given a set of inputs, auto-associative memories are primarily focused on recalling a pattern X when provided a partial or noisy variant of X. By contrast, hetero-associative memories are able to recall not only patterns of different sizes from their inputs but can be leveraged to map concepts between categories (hence “hetero-associative”). One common example from the literature is a hetero-associative memory that might recall the embedded animal concept of “monkey” given the embedded food concept of “banana”. Since all forms of AM are focused on the actual content being stored and retrieved, they are also commonly referred to as *content-addressable memories* (CAM) in the literature.

**Classical Hopfield Network.** One of the earliest and probably the most well-known auto-associative memory are Hopfield Networks Hopfield (1982). The original Hopfield Networks are discrete where they operate by storing binary-pattern inputs into the weights of a fully-connected neural network using a local update rule. For an input  $\mathbf{x} \in \{-1, 1\}^d$  containing  $d$  binary values, a Hopfield Network contains  $d^2$  real-valued connections, i.e.,  $\mathbf{W} \in \mathbb{R}^{d \times d}$ . The Hopfield learning algorithm specifies a write over  $n$  binary memories  $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$ , represented as column vectors, by accumulating their outer product

$$\mathbf{W} = \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^\top, \quad \mathbf{W}[p, q] = \sum_{i=1}^n \mathbf{x}_i[p] \mathbf{x}_i^\top[q]. \quad (8)$$

This specific write update is termed the *Hebbian update rule* as it follows the "fire together, wire together" principle proposed by psychologist Donald Hebb as a model of synaptic learning Hebb (1949). The Hopfield Network *read* operation involves minimizing an *energy function*

$$E(\mathbf{W}, \boldsymbol{\xi}) = -\frac{1}{2} \boldsymbol{\xi}^\top \mathbf{W} \boldsymbol{\xi}, \quad (9)$$

where  $\boldsymbol{\xi} \in \mathbb{R}^d$  is the *state* of the network, initialized as the initial query, and then optimized to a stable state known as the *attractor*.

**Modern Hopfield Network.** Hopfield Networks serve as an interesting weight-based method of memory storage. However, although they use optimization as a method of memory retrieval, their learning rule is not differentiable due to the use of discrete states. Modern Hopfield Network (MHN Ramsauer et al. (2020)) introduces a new energy function instead of that in Eq. 9. Specifically, MHN generalizes the energy function to continuous-valued patterns and adds a quadratic term, i.e.,

$$E(\mathbf{X}, \boldsymbol{\xi}, \beta) = -\text{LSE}(\beta, \mathbf{X}^\top \boldsymbol{\xi}) + \frac{1}{2} \boldsymbol{\xi}^\top \boldsymbol{\xi} + \beta^{-1} \log(N) + \frac{1}{2} M^2, \quad (10)$$

where  $\mathbf{X}$  is the matrix form of  $N$  *continuous* stored patterns  $\mathbf{x}_i$ ,  $i = 1, 2, \dots, N$ ,  $M$  is the largest norm of all stored patterns, LSE stands for the LogSumExp function with coefficient  $\beta$ .

**Predictive Coding Network.** The Predictive Coding Network is a computational model that aims to explain how the brain processes sensory information and makes predictions about future sensory inputs. It is based on the concept of predictive coding, which suggests that the brain constantly generates predictions about upcoming sensory inputs and updates these predictions based on the actual sensory feedback it receives. In machine learning, the predictive coding network is implemented as an energy-based associative memory model that has set the state-of-the-art on a number of image associative recall tasks.

#### A.4 ADDITIONAL EXPERIMENTAL RESULTS

Table 3: **Performance comparison on image classification datasets (Task-II).** The mean and standard deviation are calculated based on five runs with varying seeds. <sup>+</sup> denotes the corresponding method combined with our SHARC framework.

Buffer	Model	S-CIFAR-10		S-CIFAR-100		S-Mini-ImgNet	
		ACC ( $\uparrow$ )	BWT ( $\uparrow$ )	ACC ( $\uparrow$ )	BWT ( $\uparrow$ )	ACC ( $\uparrow$ )	BWT ( $\uparrow$ )
-	JOINT	72.85 $\pm$ 2.18	61.26 $\pm$ 8.55	45.87 $\pm$ 1.22	45.55 $\pm$ 1.45	47.08 $\pm$ 0.77	46.25 $\pm$ 0.98
-	SGD	20.47 $\pm$ 0.78	-90.16 $\pm$ 0.92	8.55 $\pm$ 1.39	-78.24 $\pm$ 0.93	12.21 $\pm$ 0.75	-67.11 $\pm$ 0.77
1000	GEM	89.07 $\pm$ 1.09	-3.28 $\pm$ 1.61	82.38 $\pm$ 0.93	2.59 $\pm$ 1.08	73.78 $\pm$ 1	0.38 $\pm$ 1.53
	GEM <sup>+</sup>	92.14 $\pm$ 0.39	-0.66 $\pm$ 1.93	85.56 $\pm$ 0.76	1.97 $\pm$ 1.8	75.11 $\pm$ 0.59	3.69 $\pm$ 1.79
	A-GEM	91.01 $\pm$ 2.86	-0.46 $\pm$ 1.09	85.7 $\pm$ 0.45	2.53 $\pm$ 1.43	75.86 $\pm$ 0.92	3.02 $\pm$ 1.38
	A-GEM <sup>+</sup>	92.13 $\pm$ 0.49	-0.44 $\pm$ 2.05	86.67 $\pm$ 0.87	3.46 $\pm$ 1.05	75.42 $\pm$ 0.88	4.02 $\pm$ 1.74
	ER	89.56 $\pm$ 1.09	-1.33 $\pm$ 3.47	83.58 $\pm$ 1.12	1.39 $\pm$ 1.17	71.45 $\pm$ 1.01	-1.71 $\pm$ 1.37
	ER <sup>+</sup>	91.81 $\pm$ 0.25	-0.83 $\pm$ 1.82	85.71 $\pm$ 0.63	2.34 $\pm$ 1.2	74.18 $\pm$ 0.82	2.53 $\pm$ 1.54
	MER	89.05 $\pm$ 2.37	-2.82 $\pm$ 3.1	82.5 $\pm$ 1.05	0.1 $\pm$ 2.23	70.17 $\pm$ 1.22	-3.16 $\pm$ 1.66
	MER <sup>+</sup>	91.7 $\pm$ 0.55	-0.78 $\pm$ 1.01	85.08 $\pm$ 1.19	1.53 $\pm$ 1.06	72.37 $\pm$ 0.87	0.29 $\pm$ 1.7
	DER++	87.56 $\pm$ 1.87	-2.14 $\pm$ 4.63	83.57 $\pm$ 0.64	1.62 $\pm$ 1.31	71.83 $\pm$ 1.12	-0.98 $\pm$ 1.46
	DER++ <sup>+</sup>	90.4 $\pm$ 0.63	-3.1 $\pm$ 0.77	85.87 $\pm$ 0.92	2.54 $\pm$ 1.22	74.09 $\pm$ 0.85	2.82 $\pm$ 2
	CLS-ER	83.02 $\pm$ 2.81	-8.88 $\pm$ 4.33	82.67 $\pm$ 0.73	-0.28 $\pm$ 1.5	71.97 $\pm$ 0.43	-3.69 $\pm$ 1.02
	CLS-ER <sup>+</sup>	89.81 $\pm$ 1.77	0.85 $\pm$ 3.75	85.17 $\pm$ 0.61	1.24 $\pm$ 2.02	76.35 $\pm$ 0.41	1.97 $\pm$ 0.52

In general, when combined with SHARC, the methods used show notable enhancements in most scenarios. Specifically, in the experiments conducted with a buffer size of 1000 and CIFAR-10, the maximum improvement in accuracy reaches approximately 7%.

Table 4: **Performance comparison on image classification datasets (Class-IL).** The mean and standard deviation are calculated based on five runs with varying seeds. <sup>+</sup> denotes the corresponding method combined with our SHARC framework.

Buffer	Model	S-CIFAR-10		S-CIFAR-100		S-Mini-ImgNet	
		ACC ( $\uparrow$ )	BWT ( $\uparrow$ )	ACC ( $\uparrow$ )	BWT ( $\uparrow$ )	ACC ( $\uparrow$ )	BWT ( $\uparrow$ )
-	JOINT	72.85 $\pm$ 2.18	61.26 $\pm$ 8.55	45.87 $\pm$ 1.22	45.55 $\pm$ 1.45	47.08 $\pm$ 0.77	46.25 $\pm$ 0.98
-	SGD	20.47 $\pm$ 0.78	-90.16 $\pm$ 0.92	8.55 $\pm$ 1.39	-78.24 $\pm$ 0.93	12.21 $\pm$ 0.75	-67.11 $\pm$ 0.77
1000	GEM	32.13 $\pm$ 6.79	-56.21 $\pm$ 13.61	23.39 $\pm$ 5.47	-48.35 $\pm$ 5.51	31.92 $\pm$ 5.29	-42.6 $\pm$ 6.71
	<b>GEM<sup>+</sup></b>	40.68 $\pm$ 3.18	-49.08 $\pm$ 11.47	27.49 $\pm$ 3.49	-44.72 $\pm$ 3.92	33.54 $\pm$ 2.24	-40.21 $\pm$ 4.06
	A-GEM	26.2 $\pm$ 8.53	-80.68 $\pm$ 7.31	13.38 $\pm$ 3.81	-73.22 $\pm$ 4.95	16.52 $\pm$ 1.54	-62.86 $\pm$ 2.1
	<b>A-GEM<sup>+</sup></b>	28.93 $\pm$ 3.67	-78.13 $\pm$ 5.81	16.04 $\pm$ 3.47	-70.52 $\pm$ 3.49	17.58 $\pm$ 1.71	-60.32 $\pm$ 1.7
	ER	33.82 $\pm$ 8.42	-53.69 $\pm$ 15.06	23.13 $\pm$ 5.06	-53.35 $\pm$ 5.2	30.77 $\pm$ 1.85	-43.9 $\pm$ 2.47
	<b>ER<sup>+</sup></b>	37.74 $\pm$ 2.45	-43.36 $\pm$ 14.88	26.75 $\pm$ 1.1	-49.72 $\pm$ 0.86	30.26 $\pm$ 1.95	-43.64 $\pm$ 1.07
	MER	32.87 $\pm$ 8.8	-50.42 $\pm$ 18.37	23.59 $\pm$ 3.26	-51.08 $\pm$ 4.59	29.73 $\pm$ 2.36	-44.68 $\pm$ 2.92
	<b>MER<sup>+</sup></b>	36.94 $\pm$ 5.38	-44.32 $\pm$ 14.72	25.22 $\pm$ 0.8	-49.68 $\pm$ 1.81	29.71 $\pm$ 1.37	-43.58 $\pm$ 1.47
	DER++	31.37 $\pm$ 7.22	-52.59 $\pm$ 16.29	19.43 $\pm$ 7.88	-59.2 $\pm$ 7.87	29.33 $\pm$ 1.83	-45.58 $\pm$ 2.29
	<b>DER++<sup>+</sup></b>	35.51 $\pm$ 7.44	-36.06 $\pm$ 11.08	26.35 $\pm$ 1.84	-52.05 $\pm$ 2.85	29.42 $\pm$ 2.02	-44.78 $\pm$ 2.61
	CLS-ER	24.13 $\pm$ 8.02	-33.43 $\pm$ 23.68	25.71 $\pm$ 2.91	-51.68 $\pm$ 2.84	30.3 $\pm$ 1.51	-47.22 $\pm$ 1.7
	<b>CLS-ER<sup>+</sup></b>	26.34 $\pm$ 5.38	-43.81 $\pm$ 7.09	23.99 $\pm$ 4.36	-53.42 $\pm$ 5.22	30.21 $\pm$ 1.25	-46.88 $\pm$ 1.73

Our method demonstrates strong performance on the image classification dataset task, with improvements observed across various metrics compared to the original method. Notably, even with a large buffer, we achieve an average improvement of approximately 3% in accuracy, providing compelling evidence of the effectiveness of our approach.

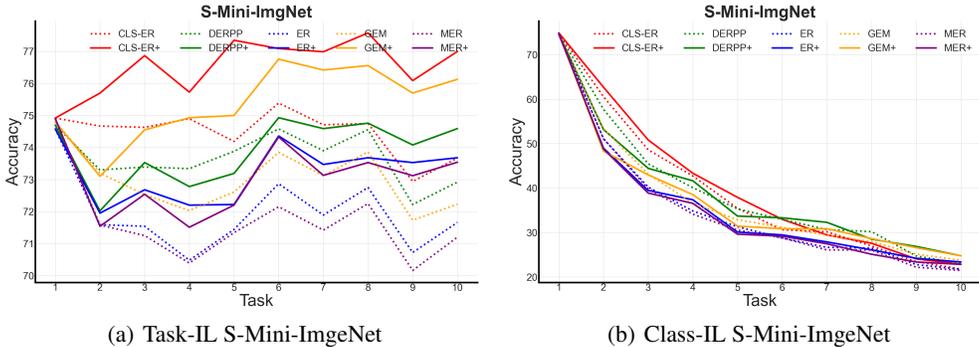


Figure 7: Learning curves of multiple models with / without SHARC on S-Mini-ImageNet. Models with / without SHARC are shown in solid / dotted lines. Buffer size for all models is 200.

It is evident that our approach significantly enhances the model across multiple dimensions. Notably, in terms of task accuracy, all algorithms show an improvement of approximately 2%, with this value consistently increasing as the number of tasks grows. Particularly noteworthy is the observation that while the accuracy of the original algorithms tends to decrease with larger tasks, our method continues to increase in accuracy, demonstrating its effectiveness in handling a large number of multi-tasks. These findings strongly indicate the superior performance of our method in scenarios involving numerous tasks.