

# Exemplar Masking for Multimodal Incremental Learning

Yi-Lun Lee<sup>1,3</sup>   Chen-Yu Lee<sup>2</sup>   Wei-Chen Chiu<sup>1</sup>   Yi-Hsuan Tsai<sup>3</sup>

<sup>1</sup>National Yang Ming Chiao Tung University   <sup>2</sup>Google   <sup>3</sup>Atmanity

{yllee10727, walon}@cs.nycu.edu.tw, chenylulee@google.com, yhtsai@atmanity.io

## Abstract

Multimodal incremental learning needs to digest the information from multiple modalities while concurrently learning new knowledge without forgetting the previously learned information. There are numerous challenges for this task, mainly including the larger storage size of multimodal data in exemplar-based methods and the computational requirement of finetuning on huge multimodal models. In this paper, we leverage the parameter-efficient tuning scheme to reduce the burden of fine-tuning and propose the exemplar masking framework to efficiently replay old knowledge. Specifically, the non-important tokens are masked based on the attention weights and the correlation across different modalities, significantly reducing the storage size of an exemplar and consequently saving more exemplars under the same memory buffer. Moreover, we design a multimodal data augmentation technique to diversify exemplars for replaying prior knowledge. In experiments, we not only evaluate our method in existing multimodal datasets but also extend the ImageNet-R dataset to a multimodal dataset as a real-world application, where captions are generated by querying multimodal large language models (e.g., InstructBLIP). Extensive experiments show that our exemplar masking framework is more efficient and robust to catastrophic forgetting under the same limited memory buffer.

## 1. Introduction

Human perception of the real-world environment acquires knowledge from multiple senses in a continual and sequential manner. Recently, multimodal transformers [10, 11, 14, 22, 23, 34] pre-trained on large-scale datasets demonstrate promising performance across a range of multimodal tasks, including visual recognition [14, 22], object detection [10], multimodal sentiment analysis [23, 34], etc. However, when facing the scenario that the model keeps updated with the new data coming in sequentially, such strong models suffer from the issue of catastrophic forgetting which may lead to severe performance drop of the old knowledge.

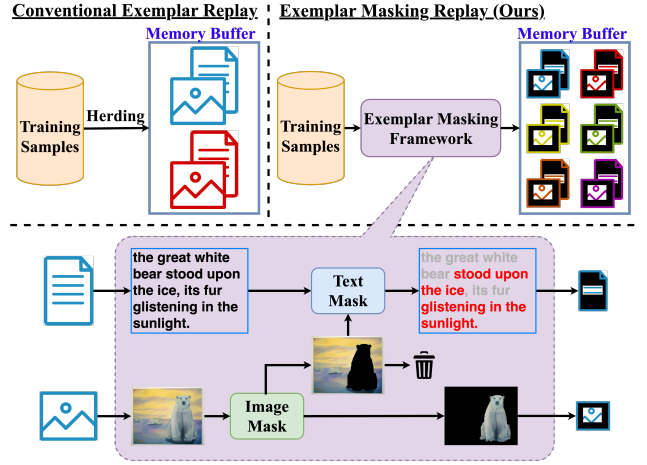


Figure 1. Illustration of exemplar replay for a new class “ice bear”. In the conventional exemplar replay framework, only very few data samples can be stored in the limited memory buffer due to the high storage demand. In contrast, our exemplar masking framework preserves the important regions of the image and discards the non-important ones to reduce the storage space. Moreover, we propose to preserve the information of discarded regions via another modality (i.e., text) to retain as much information as possible. Under the same memory buffer, our framework can store more samples, contributing to more effective knowledge replay.

Moreover, with the growth in the size of multimodal models, finetuning the entire model with multimodal data becomes increasingly impractical under the limited computation resources. To this end, we delve into the realm of multimodal class-incremental learning (MCIL), which is a more practical learning scenario for AI agents.

In this paper, we tackle two practical challenges: 1) heavy multimodal model fine-tuning, and 2) catastrophic forgetting when learning with new data. First, nowadays it is common to use multimodal models with billions of parameters [2, 25, 26] as the pre-trained models and transfer it on the downstream tasks with finetuning. However, as the multimodal model size explosively increases, finetuning the entire models when new data comes inevitably results in a heavy computation cost, which is even no longer applicable under the limited computation resources.

To this end, we adopt parameter-efficient tuning (PET) [8, 15, 16, 18, 40] in our MCIL framework to achieve the model training efficiency while reducing the negative fine-tuning effect in retaining the old knowledge.

Second, to further alleviate catastrophic forgetting, exemplar replay proposed by iCaRL [27] has emerged as a widely-used technique in incremental learning. The core concept is to retain a limited number of past samples per class as exemplars, facilitating the retention of previous knowledge when learning new classes. For our multimodal setup, as the volume of multimodal data increases, the capacity of storing complete exemplars in the memory buffer decreases, which makes it more difficult to keep the old knowledge. To address this challenge and select more representative exemplars within the same buffer size, a few data-efficient methods [21, 32] are proposed. However, these approaches still retain redundant information, e.g., non-essential regions in images, resulting in substantial storage overhead. In addition, for each type of data in multimodal tasks, the method to reduce redundant information may vary accordingly, which is not addressed in the prior work.

In this paper, we design an MCIL framework by accounting the property of each data type, thereby optimizing the space to store exemplars dynamically for each data with different modalities. Specifically, we propose an exemplar masking method to select important tokens of exemplars according to their attention weights, so only a portion of information in representative exemplars for each class is needed to store (see Figure 1). Our motivation is that, the discriminative tokens with higher attention weights may be more valuable to be preserved for replaying old knowledge, while the masked tokens containing non-discriminative information can be removed. Moreover, to replay the masked exemplars more effectively, we further interchange the data of different modalities in the same class as a data augmentation strategy to improve data diversity.

In experiments, we not only validate our proposed method in existing multimodal datasets but also create a multimodal dataset extended from the existing ImageNet-R [5] dataset in a practical manner for the MCIL task. We generate captions for the images with multimodal large language models (MLLMs) [4]. We conduct extensive experiments to explore different masking and selection strategies to demonstrate that the proposed method is able to achieve efficient and effective MCIL against other baseline approaches. Our contributions are summarized as follows:

- We explore multimodal incremental learning in both data-efficient and memory-efficient manners, which is practical for AI agents.
- We propose an exemplar masking method for the exemplar-based incremental learning framework, which highly reduces the storage space of multimodal samples via adaptive masking, and thus more samples

are able to be saved in the same memory buffer size.

- To replay the multimodal masked exemplars more efficiently, we further propose multimodal data augmentation to enrich the old exemplars, encouraging models to replay the old knowledge effectively.
- We extend the image classification dataset to a multimodal one by generating rich captions via MLLMs.

## 2. Related Work

### 2.1. Incremental Learning

The ability to learn new concepts without forgetting the previously acquired knowledge is essential for AI agents. There are three groups of methods for alleviating the forgetting issue. **Parameter regularization methods** [12, 17] estimate the discrepancy between new and old models, and then adopt corresponding penalization terms to the objectives. **Model-based methods** [1, 20, 31, 39] preserve the model parameters for learning new classes to prevent overwriting the learned weights for previous classes. **Replay-based methods**, as a long-lasting and widely-used method, assume there exists a limited memory buffer to store a few old samples (named as exemplars) for replaying previous knowledge. iCaRL [27] first introduces this paradigm for class-incremental learning, motivating various works to improve the performance with a limited replay buffer. [6, 38, 42] aim at mitigating the issue of biased classifiers, which arises due to the substantial imbalance between the number of old exemplars and newly acquired samples. In addition, some approaches share the same high-level idea as exemplars but with different storing types, including topology-based [30], feature-based [9], GAN-based [29, 37], and prompt-based [35, 36] methods.

Recently, [13, 19, 21, 24, 32] have been proposed to improve the memory efficiency of the replay-based methods. Mnemonics [19] distill the current new training samples into exemplars via a bi-level optimization to store more representative samples with the same quantity. On the other hand, numerous works [13, 24] focus on augmenting the exemplars with augmentation methods like mixup [41] to diversify the exemplars and avoid the over-memorizing issue. Furthermore, considering the trade-off between quality and quantity, MRDC [32] adopts JPEG compression codec to compress the samples into more compact data, while CIM [21] proposes class-incremental masking based on the class activation map (CAM) to select the discriminative regions and downsample other regions. However, these downsampled and non-discriminative regions sometimes do not provide useful information for replaying old knowledge, leading to waste the storage space. Also, these image compression and downsampling methods cannot be generalized to different modality data for the multimodal incremental learning scenario. Recently, as a pioneer in tackling

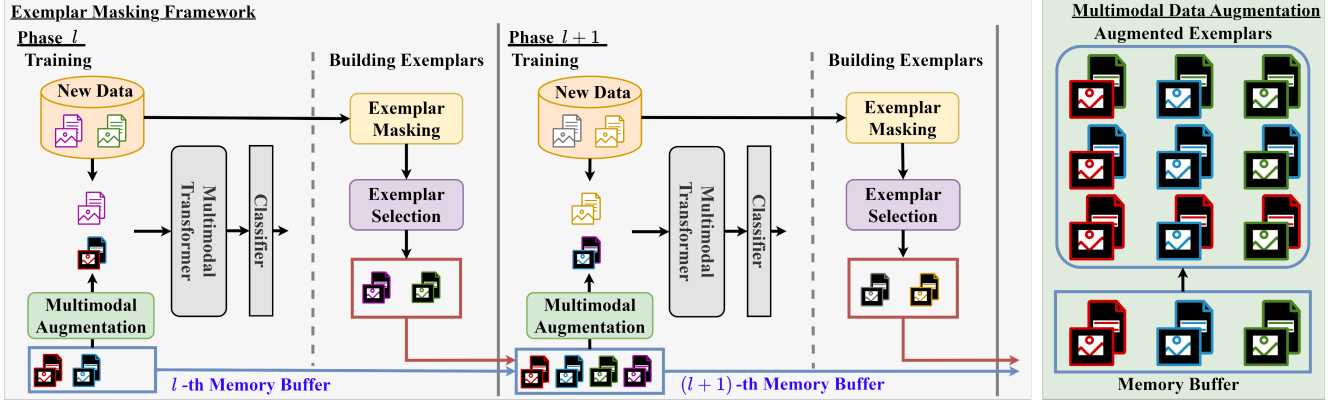


Figure 2. Overview of the proposed exemplar masking framework for multimodal class-incremental learning, including exemplar masking, exemplar selection, and multimodal data augmentation. In the  $l$ -th incremental phase, we first apply multimodal data augmentation on the  $l$ -th memory buffer and train the model with both new data and augmented exemplars. After training, we generate the masked exemplars from new data via the proposed exemplar masking and exemplar selection methods, then combine them with the memory buffer.

MCIL, GMM [3] considers it as a generative task and leverages the strong generative capacity of LLMs to predict new classes. In our work, we adopt the multimodal transformer as the backbone where the multimodal data is embedded in tokens, and we propose an exemplar masking framework to discard the non-important tokens according to the attention map.

## 2.2. Parameter-efficient tuning

As the explosion of the parameters in huge pre-trained multimodal models, finetuning the entire model on downstream tasks is not available under the limited computation resources. To alleviate the burden of finetuning entire models without hurting the performance, there are various parameter-efficient tuning (PET) methods to alleviate this issue. As the pioneer, [7] proposes to transfer the pre-trained models on downstream tasks with lightweight adapter modules. Prompt tuning [15] and prefix tuning [16] propose to insert the learnable tokens (named prompts) into the input sequence in order to instruct the pre-trained models performing the downstream tasks. Bitfix [40] adapts pre-trained models on small datasets with only training the bias-terms in the layers. SSF [18] modifies the features via a linear transformation with learnable scale and shift features for transferring to new data. LoRA [8] injects trainable rank decomposition matrices into each transformer block to learn the adaptation on downstream tasks. Compared to finetuning entire models, these PET methods reach a competitive or even better performance with very few learnable parameters (i.e.,  $< 1\%$  total model parameters). In our work, we adopt the SSF [18] as our PET method for efficient multimodal incremental learning in the practical scenario.

## 3. Proposed Method

In this paper, we focus on multimodal class-incremental learning based on the exemplar replay framework and pro-

pose an exemplar masking framework for it, as shown in Figure 2. Without loss of generality, we consider that the multimodal data is composed of two modalities: text  $T$  and image  $I$ . As the typical scenario of incremental learning, the data of different classes arrive sequentially in each phase, in which there are in total  $L$  incremental phases. Thus the entire multimodal incremental dataset is denoted as  $D = \{D^1, D^2, \dots, D^L\}$  where  $D^l$  indicates the data samples arriving in  $l$ -th phase. Basically, model in the  $l$ -th incremental phase is expected to learn the new classes in  $D^{\text{new}} = D^l$  while retaining the previous knowledge of the old classes in  $D^{\text{old}} = \{D^1, D^2, \dots, D^{l-1}\}$ . Due to limited size of the memory buffer, we cannot access the entire old training samples  $D^{\text{old}}$ , and thus we only preserve a few representative samples  $D^{\text{exp}}$  (named as *exemplars*) from previous data (i.e.,  $D^{\text{exp}} \subset D^{\text{old}}$ ) via the herding algorithm [27]. Overall, the model is trained with available training samples in  $D^{\text{new}} \cup D^{\text{exp}}$  during the  $l$ -th incremental phase.

We adopt the multimodal transformer ViLT [11] as our backbone model, which is pre-trained on multiple large image-text datasets. Instead of fine-tuning the entire model, we adopt parameter-efficient tuning (PET) methods to mitigate computational demands, where we utilize the SSF [18] as our main PET scheme in this paper.

### 3.1. Exemplar Masking for Replay

The typical manner of keeping several original/complete exemplars of old classes could lead to fewer stored samples when dealing with multimodal data due to limited memory buffer. Therefore, it is crucial to maximize the storage efficiency of the memory buffer such that it only preserves useful information of old data for replaying purposes. To this end, for each exemplar (which is composed of text  $x_T$ , image  $x_I$ , and the ground truth class label  $y$ ) to be kept in the buffer, we propose to only store its essential tokens while

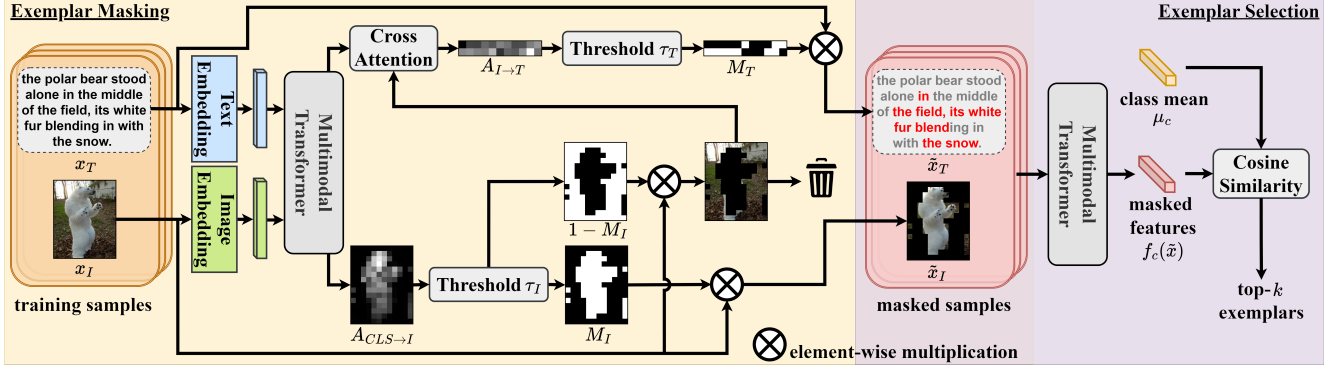


Figure 3. The overview of our proposed exemplar masking and exemplar selection methods. Given a training sample  $(x_T, x_I)$  of class  $c$ , we first calculate the attention of the class token for the image modality  $A_{CLS \rightarrow I}$  and obtain the image mask  $M_I$  according to the threshold  $\tau_I$ . Then the masked image regions are preserved in the masked image  $\tilde{x}_I$  while the others are discarded. To preserve contextual information from the discarded image regions, we calculate the cross attention  $A_{I \rightarrow T}$  between discarded image tokens and text tokens to obtain the text mask  $M_T$  via the threshold  $\tau_T$ . Hence the masked text  $\tilde{x}_T$  is produced via applying text mask  $M_T$  on  $x_T$ . Finally, we compute the cosine similarity between the feature  $f_c(\tilde{x})$  of the masked sample  $(\tilde{x}_T, \tilde{x}_I)$  and the mean  $\mu_c$  of the class  $c$ , in which the samples with the top- $k$  highest similarity are selected as the exemplars of class  $c$  and preserved without exceeding the memory size.

discarding the non-important ones according to the attention weight of the class token  $x_{CLS}$ <sup>1</sup>, as shown in Figure 3. To be detailed, given the attention map  $A_{CLS}$  of the class token  $x_{CLS}$ , we separate it into  $A_{CLS \rightarrow T}$  and  $A_{CLS \rightarrow I}$  for the text and image modalities respectively, followed by using them to create the masks for masking the text  $x_T$  and image  $x_I$ . As the image modality requires a larger storage space, we first consider masking the image tokens.

**Masking image exemplars.** Before moving to the phase  $l+1$ , the model has learned to recognize samples of the new classes  $C^{new(l)}$  added in the phase  $l$ , in which some selected exemplars of classes  $C^{new(l)}$  will be replayed in the next phase for retaining the knowledge. Therefore, the attention maps produced by the model are expected to focus more on the image content related to classes  $C^{new(l)}$ , meaning that the image tokens with higher attention weights involve the most information about  $C^{new(l)}$ . To effectively mask out the non-discriminative tokens and preserve the informative tokens, we find that the mean of the image attention map  $A_{CLS \rightarrow I}$  serves as a suitable threshold  $\tau_I$  for producing the mask of image modality  $M_I = \{m_I(i) \mid i = 0, 1, \dots, N_I\}$ , where  $N_I$  is the number of image tokens:

$$\tau_I = \frac{1}{N_I} \sum_{i=0}^{N_I} A_{CLS \rightarrow I}(i), \quad (1)$$

$$m_I(i) = \begin{cases} 1, & A_{CLS \rightarrow I}(i) \geq \tau_I, \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

Setting the threshold for masking to the mean of attention weights rather than a fixed value allows for dynamic adjustment based on input content. This adaptive threshold enhances flexibility of our method across various datasets.

<sup>1</sup>We note that our backbone is ViLT, where the input exemplar is tokenized through the transformer blocks to have the resultant class token.

**Masking text exemplars.** Although some image regions are discarded due to low attention weights (i.e., being less important in the image domain), they may still provide *contextual information* which is beneficial for recognizing the target classes (e.g., the background with an “ice lake” could be helpful to recognize the class “ice bear”). Hence, we would still like to maintain the contextual information hidden behind masked image tokens in our replay buffer.

To this end, we instead store such information via a memory-efficient manner by leveraging the corresponding text tokens. As the second modality, text tokens are encouraged to preserve not only the text-related information but also the discarded information from the image modality according to the attention weights across two modalities. To be specific, we first calculate cross-attention weights  $A_{I \rightarrow T}(i)$  of text tokens with respect to image tokens (i.e., the attention from image tokens to text tokens), especially on the attention of masked image tokens (i.e., with  $m_I(i) = 0$ ). Then we compute the threshold  $\tau_T$  as follows for extracting the text tokens that are more relevant to the masked image tokens, i.e., producing the mask of text modality  $M_T = \{m_T(i) \mid i = 0, 1, \dots, N_T\}$ :

$$\tau_T = \frac{1}{N_T} \sum_{i=0}^{N_T} \frac{1}{N_{m_I=0}} \sum_{j=0}^{N_I} \mathbb{I}(m_I(j) = 0) A_{I \rightarrow T}(j, i), \quad (3)$$

$$m_T(i) = \begin{cases} 1, & A_{I \rightarrow T}(i) \geq \tau_T, \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

where  $\mathbb{I}(\cdot)$  is an indicator function and  $A_{I \rightarrow T}(j, i)$  is the cross-attention value of text token  $i$  on the image token  $j$ , while  $N_T$  and  $N_{m_I=0}$  are the number of the text tokens and the masked image tokens respectively.

Eventually, the multimodal exemplar after applying our



masking strategy as masked exemplars becomes:

$$\tilde{x}_T = M_T \otimes x_T, \quad \tilde{x}_I = M_I \otimes x_I, \quad (5)$$

where  $\otimes$  is the element-wise multiplication. Please note that we leverage the attention map from the backbone model instead of applying additional models to identify the important regions to be preserved.

**Selection of masked exemplars.** In the typical setting of replay-based incremental-learning methods, the herding selection strategy is adopted to select  $k$  exemplars in which they are the most representative samples of the target class according to the class center/mean. Specifically, the first few samples being closest to the class center are selected based on the feature space extracted by the model. In the proposed method, we adopt the herding strategy used in iCaRL [27] as well, but we use the cosine similarity for measuring the distance between the class mean and masked training samples (i.e., the samples after applying our masking operation), in which the masked training samples with the top- $k$  highest similarity to their respective class mean are selected as exemplars (see the right side of Figure 3).

### 3.2. Multimodal Data Augmentation

Although our exemplar masking method helps store more exemplars for the old/learned classes in the replay buffer of fixed size, there still exists the imbalance issue between the old and new classes. That is, the amount of exemplars used to replay the knowledge of old classes is much less than the training samples of new classes. Moreover, using the same exemplars for replaying old classes in different incremental phases would cause the over-memorizing issue (i.e., the model only memorizes the exemplars and does not generalize to other data samples of the same class), as described in [13]. To this end, we propose a simple yet effective data augmentation technique upon the exemplars to enrich the training samples of replaying old classes, thus alleviating both class-imbalance and over-memorizing issues.

As mentioned previously, for a multimodal data sample, the masked image mainly represents the object of the corresponding class, while the masked text provides complementary and auxiliary information. Hence, our multimodal data augmentation interchanges either the images or text descriptions between an arbitrary pair of exemplars from the same class (see the right side of Figure 2). This creates numerous diverse training samples to replay the old knowledge.

### 3.3. Training Objective

In the incremental phase  $l$ , we would have the masked exemplars  $\tilde{D}^{\text{exp}}$  of old classes  $C^{\text{old}(l)}$  that are seen from previous phases, as well as the training samples  $D^{\text{new}}$  of new classes  $C^{\text{new}(l)}$  arriving at phase  $l$ . Our objective function for multimodal class-incremental learning is based on the cross-entropy function, where its optimization in the phase

$l$  is driven by the available training data  $\tilde{D}^{\text{exp}} \cup D^{\text{new}}$ . Here, each training sample is composed of the multimodal data  $\mathbf{x}$  (i.e., image and text) and the corresponding class label  $y$ . The  $\mathcal{L}_{\text{CE}}$  loss for every training sample is defined as:

$$\mathcal{L}_{\text{CE}}(\mathbf{x}, y) = \sum_{c=0}^{C^l} -\delta_{c=y} \log(p_c(\mathbf{x}), c), \quad (6)$$

where  $C^l = |C^{\text{old}(l)}| + |C^{\text{new}(l)}|$  denotes the total number of classes in the current incremental phase  $l$  and  $p_c(\mathbf{x})$  is the output prediction on the class  $c$ . Please note that, in order to avoid the catastrophic forgetting caused by overriding the previously learned classifier weights, we adopt the masked logit trick [35, 36]. This operation masks the logits of old classes for training samples of new classes, while masking the logits of new classes for exemplars of old classes.

## 4. Experimental Results

**Dataset.** The evaluation of multimodal class-incremental learning is based on the well-known multimodal classification dataset UPMC Food-101 [33] and a newly proposed dataset proposed (denoted as **MM-ImageNet-R**), which is stemmed from the ImageNet-R [5] but has the carefully-designed multimodal extension (i.e., extended from only having image modality to having both text and image modalities). Specifically, given an image from ImageNet-R, we derive its text description automatically by querying the multimodal large language model, InstructBLIP[4]. with the following prompts: “What’s in this image? Please use 100 words to describe the image content in detail.”

**Incremental setting.** We follow the common incremental setting in [27], where the dataset is equally split into  $L$  subsets without overlapping classes for  $L$  incremental phases. The default memory buffer size is set to support 5 raw multimodal samples per class. We adopt the average incremental learning accuracy  $\bar{A} = \sum_{l=1}^L A^l$  as our evaluation metric, where  $A^l$  is the accuracy of all the seen classes at the end of  $l$ -th incremental phase. We will make the datasets, source codes, and models available to the public.

### 4.1. Quantitative Results

In our experiments, we mainly focus on two critical properties: 1) the data efficiency of the stored exemplars, and 2) the parameter-efficient methods for multimodal incremental learning. The compared baselines include: 1) finetuning the entire model (denoted as **FT**) in each incremental phase, and 2) the parameter-efficient tuning (PET) method (i.e., **SSF** [18] is adopted in experiments) that only learns the additional few parameters in each incremental phase. Both of them are based on the conventional exemplar-replay framework. Moreover, in order to compare with other data-efficient exemplar-replay methods on the scenario of multimodal incremental learning, a baseline is built by following

the concept of **CIM** [21] which selects the important image regions according to the CAM map and downsamples the non-discriminative image regions by 4 times (i.e., half on both width and length) to compress the size of exemplars.

We summarize quantitative results of our proposed methods on MM-ImageNet-R and UPMC Food-101 in Table 1. We draw several observations:

1) SSF reaches a comparable performance with FT and the required parameters for training are less than 0.3% of total model parameters, showing that SSF (i.e., a parameter-efficient tuning method) is more applicable with limited computation resources.

2) In comparison to baselines, our exemplar masking framework without multimodal data augmentation (denoted as **Ours w/o MDA**) improves performance as the replaying is benefited from having more exemplars thus becomes more effective. In addition, with the multimodal data augmentation (denoted as **Ours**), the performance is further boosted to achieve the best across two training schemes (i.e., by FT or SST).

3) When the number of incremental phases  $L$  increases, the model variants adopting our multimodal data augmentation improve more, e.g., in Table 1, “FT + Ours” improves “FT + Ours (w/o MDA)” by 1.96% for  $L=20$  on MM-ImageNet-R, showing that our model is less sensitive to the forgetting issue under the long-term incremental case, with the help from more diverse exemplars generated by our multimodal data augmentation.

4) Although applying CIM to baselines also improves the performance, the gain is not optimal when compared with ours. The reason is that the non-important regions are preserved in low-resolution and still require additional storage space, leading to inefficient memory usage. In contrast, our exemplar masking framework discards non-important regions but still preserves the contextual information of discarded regions in another modality (i.e., text), which makes storing exemplars more efficient and replaying exemplars more effective.

## 4.2. Qualitative Results

In Figure 4, we visualize the masked exemplars and the corresponding attention map. For the masked texts, we highlight the words related to the target objects (with the yellow color box) and the words related to the contextual information from the discarded regions (with the melon color box). As shown in Figure 4, the attention map  $A_{CLS \rightarrow I}$  focuses on the region of the target objects, and thus the important regions related to objects are selected by our proposed exemplar masking method and preserved as the masked images (e.g., faces and bodies of the target animals). Moreover, the masked texts driven by our exemplar masking method indeed preserve words related to the class object and descriptions related to the discarded regions, both of which are ben-

eficial for the recognition of the corresponding classes. For example, as shown in the first row of Figure 4, the word “polar bear” and “animal” are related to the ice bear while the word “snowy ice landscape”, “girl”, and “water” are related to the discarded image regions that indicate the environment where the bear is.

## 4.3. Comparison with Other CIL Methods

To better validate the effectiveness of our proposed method, we then provide additional comparisons with: 1) another data-efficient exemplar-based method MRDC [32], which adopts JPEG to compress images in order to reduce the storage size, and 2) three state-of-the-art unimodal incremental learning approaches, L2P [36], DualPrompt [35], and EASE [43], which also adopt pre-trained models and perform parameter-efficient fine-tuning. Note that none of the above methods are originally designed for multimodal scenarios, so we adapt their implementations to our setting and tune the training hyperparameters if needed for fair comparisons. As results shown in Table 2, we show that simply adapting unimodal methods (L2P, DualPrompt, and EASE) has suboptimal performance even compared to our SSF baseline, demonstrating that the multimodal incremental learning is a more challenging task and requires more careful designs. Moreover, our exemplar masking framework consistently outperforms other data-efficient algorithms (MRDC, CIL), showing that our proposed framework can efficiently store more exemplars in the limited memory buffer and thus replay exemplars more effectively.

## 4.4. Ablation Studies

To validate our designs for the proposed exemplar masking framework, we conduct several ablation studies, including masking references, masking thresholds, and the usage of memory. We also discuss different masking methods in the supplementary materials. We define the “preserved ratio” to represent the proportion of the preserved tokens after masking, and “# of exemplars” indicates the actual number of masked exemplars stored in the memory buffer with the capacity of 5 raw multimodal samples per class. All the ablations are conducted under the default setting, which is evaluated on the MM-ImageNet-R dataset with the number of incremental phases  $L = 10$ .

**Masking references.** In this paper, exemplar masking preserves the important tokens and discards the non-important ones according to the masking reference. In addition to adopting the attention map as masking reference, we further experiment with several design choices, such as 1) **Entropy**, 2) class activation map (CAM), 3) **GradCAM** [28], and 4) **Random** masking. To be specific: 1) The entropy of attention weights for each token is calculated, where the high entropy indicates that the token attends to others more equally, while the low entropy implies that the token only

Table 1. Average incremental accuracy on our MM-ImageNet-R and UPMC Food-101 datasets with different numbers of incremental phases  $L=5, 10, 20$ . The memory size is 5 raw multimodal samples/class. **Ours** stands for our overall exemplar masking framework, while **MDA** stands for multimodal data augmentation.

Methods	# of Parameters	MM-ImageNet-R			UPMC Food-101		
		$L=5$	$L=10$	$L=20$	$L=5$	$L=10$	$L=20$
FT	112M	79.72	77.31	75.22	86.43	81.61	76.97
FT + CIM [21]		80.33	78.41	76.46	87.15	82.68	78.98
FT + Ours (w/o MDA)		82.06	79.39	76.23	88.03	84.54	79.49
FT + Ours		<b>82.97</b>	<b>80.72</b>	<b>78.19</b>	<b>88.08</b>	<b>85.91</b>	<b>80.53</b>
SSF	206K	80.58	77.03	75.26	86.66	81.62	76.91
SSF + CIM [21]		80.97	78.21	77.23	86.39	83.48	79.25
SSF + Ours (w/o MDA)		81.80	79.05	77.14	<b>87.20</b>	83.03	79.79
SSF + Ours		<b>82.76</b>	<b>80.55</b>	<b>78.64</b>	86.73	<b>84.51</b>	<b>80.68</b>

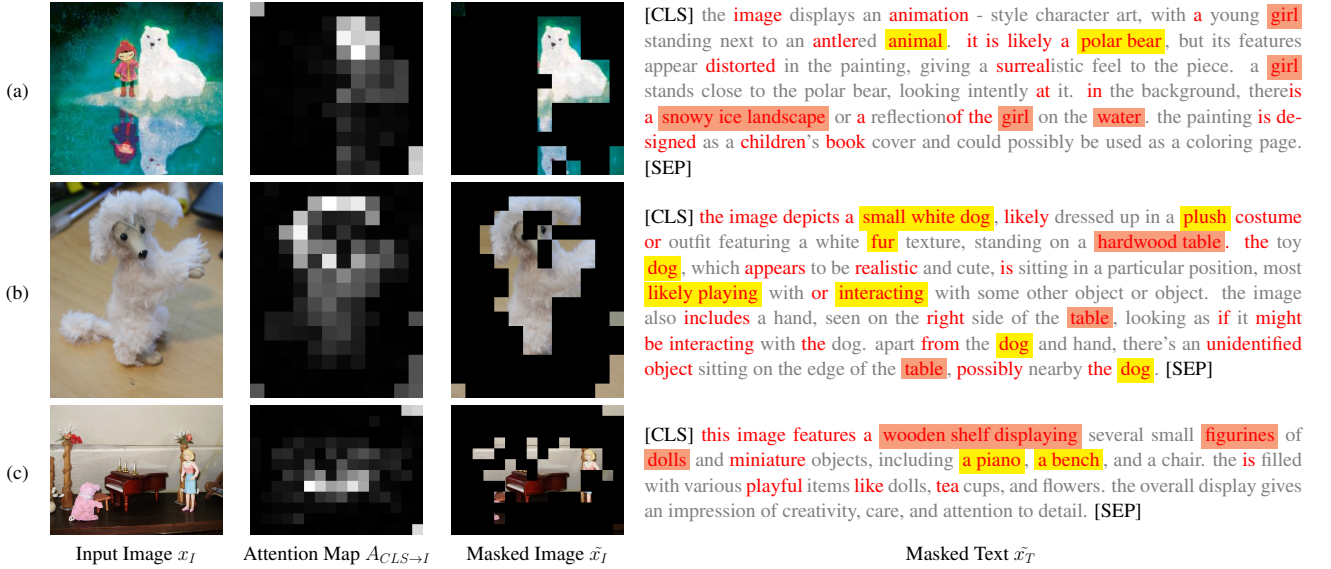


Figure 4. Examples of the masked exemplars and the corresponding attention maps from different classes, including (a) ice bear, (b) poodle, and (c) grand piano. We denote colors for the masked text in red and the discarded text in gray, while also highlighting the important words related to contextual information as well as class-related information.

Table 2. Average incremental accuracy on our MM-ImageNet-R dataset with different numbers of incremental phases  $L=5, 10, 20$ . The memory size is 5 raw multimodal samples per class.

Methods	MM-ImageNet-R		
	$L=5$	$L=10$	$L=20$
L2P [36]	68.17	69.21	66.23
DualPrompt [35]	77.66	75.57	70.47
EASE [43]	78.06	76.84	73.91
SSF	80.58	77.03	75.26
SSF + MRDC [32]	80.77	78.66	76.13
SSF + CIM [21]	80.97	78.21	77.23
SSF + Ours	<b>82.76</b>	<b>80.55</b>	<b>78.64</b>

focuses on specific few tokens; 2) CAM is based on the activation value from the last layer of the model, reflecting the response of each token to the given class; 3) GradCAM further considers the gradient of the target class flowing into the last layer of the model, highlighting the important tokens for predicting that class; 4) Random masking simply

draws some random tokens to discard them. Note that the first three design choices follow the same procedure as our proposed method (i.e., using attention maps as the masking reference) to first compute the mean of their respective maps and threshold the maps by the resultant mean to construct the final masks.

In Table 3, using the attention map in our proposed method reaches the best performance, indicating that the attention map is more suitable for reflecting the importance of input tokens. The reason is that the attention map is a more meaningful way that indicates the class token learned by accumulating the information from text and image tokens.

**Masking threshold.** The value of the masking threshold determines the amount of preserved tokens, which also affects the numbers of the stored exemplars directly. With the higher threshold, the preserved tokens are fewer and thus the number of exemplars increases, and vice versa. We validate our method on 5 different threshold values ( $\tau_I$  for

Table 3. Ablation study of adopting different design choices as the masking reference (see the first paragraph of Section 4.4 for more details). The memory size is 5 raw samples per class.

Masking Reference	Preserved Ratio	# of Exemplars	$\bar{A}$
Attention Map	0.34	14.08	<b>80.55</b>
Entropy	0.48	9.78	79.89
CAM	0.38	11.91	79.36
GradCAM	0.45	10.68	79.66
Random	0.35	14	79.58

Table 4. Analysis of the masking threshold (either  $\tau_I$  or  $\tau_T$ ), where  $\mu$  and  $\sigma$  are the mean and the standard deviation of the attention map (either  $A_{CLS \rightarrow I}$  or  $A_{I \rightarrow T}$ ).

Masking Threshold $\tau$	Preserved Ratio	# of Exemplars	$\bar{A}$
$\mu + 0.5 \times \sigma$	0.21	21.34	80.24
$\mu + 0.25 \times \sigma$	0.27	16.67	<b>81.03</b>
$\mu$	0.34	14.08	80.55
$\mu - 0.25 \times \sigma$	0.46	10.18	80.18
$\mu - 0.5 \times \sigma$	0.66	7.1	78.82
$\tau_i = 0.005, \tau_t = 0.001$	0.38	12.60	79.89
$\tau_i = 0.004, \tau_t = 0.001$	0.42	11.26	79.86

Table 5. Comparisons under the same memory size (raw samples per class) or a similar number of exemplars.

Methods	# of Exemplars	Memory Size	$\bar{A}$
SSF	5	5	77.03
	14	14	79.79
SSF + CIM	9.91	5	78.21
SSF + Ours	14.08	5	<b>80.55</b>

$A_{CLS \rightarrow I}$  and  $\tau_T$  for  $A_{I \rightarrow T}$ ) in Table 4. We find no obvious difference in performance when the threshold is equal to or higher than the attention mean, while the performance drops with a much lower threshold. This phenomenon illustrates our key motivation: replaying with fewer but more important tokens per exemplar yet saving more exemplars is the better way to alleviate the forgetting issue. We also compare our masking threshold based on the mean of the attention maps with the fixed pre-determined threshold for image  $\tau_i$  and text  $\tau_t$ , as shown in the last two rows in Table 4. We observe that even with a similar number of exemplars and preserved ratio, the model with a fixed masking threshold performs worse. This highlights that dynamically determining the masking threshold based on attention maps is more effective in preserving important regions given various input contents.

**Discussion on memory.** In Table 5, we find that under the same memory size (i.e., 5 raw multimodal samples per class), our proposed exemplar masking framework saves more exemplars in buffer for replaying the old knowledge, thus improving the performance by a large margin. Furthermore, we observe that even with a similar number of




Figure 5. Experimental results under the different constraints of the memory buffer size. Our proposed method preserves more exemplar samples under the same limited storage space and improves the baseline by a large margin.

exemplars, our method “SSF+Ours” with 14.08 masked exemplars per class still outperforms the “SSF” baseline based on conventional exemplar replay, which has 14 exemplars per class. More importantly, our memory demand for saving such amount of masked exemplars is 2.8 times smaller than the “SSF”. This observation implies that there is no need to save the entire multimodal samples for replaying old knowledge, as most regions are redundant and do not bring significant benefits. Specifically, we calculate the storage reduction for the exemplars in the first incremental phase. Given the MM-ImageNet-R dataset, the raw exemplar data occupies 3.60 MB, while the masked exemplars use only 1.29 MB, saving 64.17% (2.31 MB) of storage.

Furthermore, we compare the performance of baselines with different memory buffer sizes in Figure 5. As the CIM [21] still requires storage space for non-discriminative tokens, the number of exemplars is limited and thus cannot reach the optimal performance. In contrast, under the same memory size, our exemplar masking framework preserves more exemplars without keeping the redundant tokens and thus outperforms baselines.

## 5. Conclusions

In this paper, we propose a data-efficient exemplar masking framework with the parameter-efficient tuning method for multimodal incremental learning. Our method involves masking exemplars based on attention weights, preserving valuable discriminative tokens while discarding less important ones. This technique allows us to store a greater number of masked exemplars within the same memory buffer. In addition, we have developed a multimodal data augmentation strategy that facilitates the exchange of multimodal data within the same class, enhancing both generalization capabilities and replay performance. Moreover, as a practical application, we extend the existing image dataset to a multimodal dataset by creating captions via captioning models and refining it with large language models. Extensive experiments and ablation studies demonstrate the efficiency and effectiveness of our framework.



## References

- [1] Davide Abati, Jakub Tomczak, Tijmen Blankevoort, Simone Calderara, Rita Cucchiara, and Babak Ehteshami Bejnordi. Conditional channel gated networks for task-aware continual learning. In *CVPR*, 2020. 2
- [2] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *NeurIPS*, 2020. 1
- [3] Xusheng Cao, Haori Lu, Linlan Huang, Xialei Liu, and Ming-Ming Cheng. Generative multi-modal models are good class incremental learners. In *CVPR*, 2024. 3
- [4] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning. *arXiv preprint arXiv:2305.06500*, 2023. 2, 5
- [5] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *ICCV*, 2021. 2, 5
- [6] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *CVPR*, 2019. 2
- [7] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *ICML*, 2019. 3
- [8] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *ICLR*, 2022. 2, 3
- [9] Ahmet Iscen, Jeffrey Zhang, Svetlana Lazebnik, and Cordelia Schmid. Memory-efficient incremental learning through feature adaptation. In *ECCV*, 2020. 2
- [10] Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. Mdetrm: modulated detection for end-to-end multi-modal understanding. In *ICCV*, 2021. 1
- [11] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *ICML*, 2021. 1, 3
- [12] James Kirkpatrick et al. Overcoming catastrophic forgetting in neural networks. *PNAS*, 2017. 2
- [13] Yi-Lun Lee, Dian-Shan Chen, Chen-Yu Lee, Yi-Hsuan Tsai, and Wei-Chen Chiu. Data efficient incremental learning via attentive knowledge replay. In *IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2023. 2, 5
- [14] Yi-Lun Lee, Yi-Hsuan Tsai, Wei-Chen Chiu, and Chen-Yu Lee. Multimodal prompting with missing modalities for visual recognition. In *CVPR*, 2023. 1
- [15] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *ArXiv:2104.08691*, 2021. 2, 3
- [16] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *ArXiv:2101.00190*, 2021. 2, 3
- [17] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE TPAMI*, 2017. 2
- [18] Dongze Lian, Daquan Zhou, Jiashi Feng, and Xinchao Wang. Scaling & shifting your features: A new baseline for efficient model tuning. *NeurIPS*, 2022. 2, 3, 5
- [19] Yaoyao Liu, Yuting Su, An-An Liu, Bernt Schiele, and Qianru Sun. Mnemonics training: Multi-class incremental learning without forgetting. In *CVPR*, 2020. 2
- [20] Yaoyao Liu, Bernt Schiele, and Qianru Sun. Adaptive aggregation networks for class-incremental learning. In *CVPR*, 2021. 2
- [21] Zilin Luo, Yaoyao Liu, Bernt Schiele, and Qianru Sun. Class-incremental exemplar compression for class-incremental learning. In *CVPR*, 2023. 2, 6, 7, 8
- [22] Mengmeng Ma, Jian Ren, Long Zhao, Davide Testuggine, and Xi Peng. Are multimodal transformers robust to missing modality? In *CVPR*, 2022. 1
- [23] Hai Pham, Paul Pu Liang, Thomas Manzini, Louis-Philippe Morency, and Barnabás Póczos. Found in translation: Learning robust joint representations by cyclic translations between modalities. In *AAAI*, 2019. 1
- [24] Sunyuan Qiang, Jiayi Hou, Jun Wan, Yanyan Liang, Zhen Lei, and Du Zhang. Mixture uniform distribution modeling and asymmetric mix distillation for class incremental learning. In *AAAI*, 2023. 2
- [25] Jack W Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, et al. Scaling language models: Methods, analysis & insights from training gopher. *ArXiv:2112.11446*, 2021. 1
- [26] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. Exploring the limits of transfer learning with a unified text-to-text transformer. *JMLR*, 2020. 1
- [27] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *CVPR*, 2017. 2, 3, 5
- [28] 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 *ICCV*, 2017. 6
- [29] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In *NeurIPS*, 2017. 2
- [30] Xiaoyu Tao, Xinyuan Chang, Xiaopeng Hong, Xing Wei, and Yihong Gong. Topology-preserving class-incremental learning. In *ECCV*, 2020. 2
- [31] Fu-Yun Wang, Da-Wei Zhou, Han-Jia Ye, and De-Chuan Zhan. Foster: Feature boosting and compression for class-incremental learning. In *ECCV*, 2022. 2
- [32] Liyuan Wang, Xingxing Zhang, Kuo Yang, Longhui Yu, Chongxuan Li, Lanqing Hong, Shifeng Zhang, Zhenguo Li, Yi Zhong, and Jun Zhu. Memory replay with data compression for continual learning. In *ICLR*, 2022. 2, 6, 7

- [33] Xin Wang, Devinder Kumar, Nicolas Thome, Matthieu Cord, and Frederic Precioso. Recipe recognition with large multi-modal food dataset. In *IEEE International Conference on Multimedia & Expo (ICME) Workshops*, 2015. [5](#)
- [34] Zilong Wang, Zhaohong Wan, and Xiaojun Wan. Trans-modality: An end2end fusion method with transformer for multimodal sentiment analysis. In *ACM Web Conference (WWW)*, 2020. [1](#)
- [35] Zifeng Wang, Zizhao Zhang, Sayna Ebrahimi, Ruoxi Sun, Han Zhang, Chen-Yu Lee, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, et al. Dualprompt: Complementary prompting for rehearsal-free continual learning. In *ECCV*, 2022. [2](#), [5](#), [6](#), [7](#)
- [36] Zifeng Wang, Zizhao Zhang, Chen-Yu Lee, Han Zhang, Ruoxi Sun, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, and Tomas Pfister. Learning to prompt for continual learning. In *CVPR*, 2022. [2](#), [5](#), [6](#), [7](#)
- [37] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, Zhengyou Zhang, and Yun Fu. Incremental classifier learning with generative adversarial networks. *ArXiv:1802.00853*, 2018. [2](#)
- [38] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *CVPR*, 2019. [2](#)
- [39] Shipeng Yan, Jiangwei Xie, and Xuming He. Der: Dynamically expandable representation for class incremental learning. In *CVPR*, 2021. [2](#)
- [40] Elad Ben Zaken, Yoav Goldberg, and Shauli Ravfogel. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 2022. [2](#), [3](#)
- [41] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *ICLR*, 2018. [2](#)
- [42] Bowen Zhao, Xi Xiao, Guojun Gan, Bin Zhang, and Shu-Tao Xia. Maintaining discrimination and fairness in class incremental learning. In *CVPR*, 2020. [2](#)
- [43] Da-Wei Zhou, Hai-Long Sun, Han-Jia Ye, and De-Chuan Zhan. Expandable subspace ensemble for pre-trained model-based class-incremental learning. In *CVPR*, 2024. [6](#), [7](#)