Prompt-Based Exemplar Super-Compression and Regeneration for Class-Incremental Learning

Ruxiao Duan[1]
ruxiao.duan@yale.edu
Jieneng Chen[2]
jchen293@jhu.edu
Adam Kortylewski[3,4]
akortyle@mpi-inf.mpg.de
Alan Yuille[2]
ayuille1@jhu.edu
Yaoyao Liu[5]
lyy@illinois.edu

[1] Yale University
[2] Johns Hopkins University
[3] University of Freiburg
[4] Max Planck Institute for Informatics
[5] University of Illinois Urbana-Champaign

Abstract

Replay-based methods in class-incremental learning (CIL) have attained remarkable success. Despite their effectiveness, the inherent memory restriction results in saving a limited number of exemplars with poor diversity. In this paper, we introduce PESCR, a novel approach that substantially increases the quantity and enhances the diversity of exemplars based on a pre-trained general-purpose diffusion model, without fine-tuning it on target datasets or storing it in the memory buffer. Images are compressed into visual and textual prompts, which are saved instead of the original images, decreasing memory consumption by a factor of 24. In subsequent phases, diverse exemplars are regenerated by the diffusion model. We further propose partial compression and diffusion-based data augmentation to minimize the domain gap between generated exemplars and real images. PESCR significantly improves CIL performance across multiple benchmarks, e.g., 3.2% above the previous state-of-the-art on ImageNet-100. The code is available at https://github.com/KerryDRX/PESCR.

1 Introduction

Ideally, AI systems should be adaptive to changing environments, where the data are continuously observed. The AI models should be capable of learning concepts from new data while maintaining the ability to recognize the old ones. In practice, AI systems often have constrained memory budgets, because of which most of the historical data must be abandoned. However, deep AI models suffer from catastrophic forgetting when updated by abundant new data and limited historical data, as previous knowledge can be overridden by new information [33, 37]. To study how to overcome catastrophic forgetting, the class-incremental learn-

© 2025. The copyright of this document resides with its authors.
It may be distributed unchanged freely in print or electronic forms.

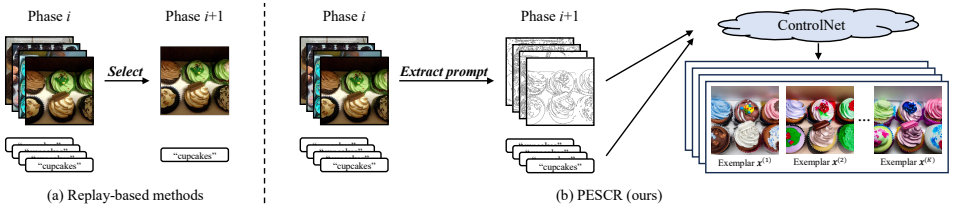

Figure 1: (a) **Traditional replay-based CIL methods** select and save only a small number of exemplars due to the memory restriction, leading to two severe issues: firstly, the relatively small size of the exemplar set compared to the new training dataset gives rise to a pronounced imbalance between old and new classes; secondly, the limited diversity of the exemplar set compared to the original training set incurs an overfitting problem. (b) **PESCR** compresses old images into visual and textual prompts (edge maps and class tags) and saves prompts to the memory. In subsequent phases, diverse high-resolution exemplars are regenerated from these prompts via an off-the-shelf pre-trained diffusion model, e.g., ControlNet. PESCR improves the quantity and diversity of exemplars without violating the memory constraint.

ing (CIL) protocol [38] is established. CIL assumes training samples from various classes are introduced to the model in phases, with previous data mostly discarded from memory.

CIL has enjoyed significant advancements [2, 7, 22, 23, 35, 44, 46, 47, 48, 53], among which replay-based methods [18, 25, 26, 27, 38, 46] stand out in terms of performance by employing a memory buffer to store a limited number of representative samples (i.e., exemplars) from former classes. During subsequent learning phases, these exemplars are revisited to help retain previously learned knowledge.

Although replay-based methods demonstrate notable effectiveness, they are still restricted by two main drawbacks. Firstly, since the exemplar set is much smaller compared to the new training data, the model is biased towards the new classes. Secondly, the poor exemplar diversity leads to overfitting on old classes. These two issues are essentially incurred by the lack of quantity and diversity of exemplars, respectively. Therefore, by tackling these two problems, all replay-based CIL methods can potentially be enhanced.

We consider these questions in CIL: is it efficient to save old-class information as RGB images? Can we compress the images into something more compact so that more information can be stored with the same memory consumption? In this paper, we propose a novel approach named ***P**rompt-based **E**xemplar **S**uper-**C**ompression and **R**egeneration* (PESCR). Instead of directly storing the previous images, we compress them into visual and textual prompts, e.g., edge maps and textual descriptions, and save these prompts in memory. In subsequent phases, diverse high-resolution exemplars are regenerated from the prompts by leveraging an off-the-shelf pre-trained diffusion model, e.g., ControlNet [49].

Compared to traditional direct replay methods, PESCR enjoys increased **quantity** and enhanced **diversity** of the exemplar set (Figure 1). Firstly, the exemplar quantity is boosted by compression: since the memory consumption of a 1-bit edge map is merely $\frac{1}{24}$ that of its 8-bit, 3-channel RGB image counterpart, 24 times more exemplars can be saved within the same memory budget. We call this process super-compression, as this compression ratio is far beyond that of existing image compression algorithms, without even compromising the image quality. Secondly, instead of relying on the images only from the dataset itself, we apply diffusion-based image generation to produce unseen samples with great diversity. This can be achieved simply by changing the random seed of the diffusion model.

However, utilizing generated images for CIL model training leads to a potentially huge domain gap between old exemplars and new data. We propose two techniques, partial compression and diffusion-based data augmentation, to mitigate this problem, enabling the CIL model to properly benefit from the synthetic exemplars without the need to fine-tune the diffusion model on the target dataset. Since the same pre-trained diffusion model can be directly downloaded from the public cloud at any time when necessary, we do not need to store the fine-tuned generator using our own memory.

Extensive experiments show that our PESCR achieves top performance on both fine-grained and coarse-grained classification datasets: Caltech-256 [14], Food-101 [5], Places-100 [52], and ImageNet-100 [12]. We fully investigate the effect of PESCR under different CIL settings and demonstrate that our approach achieves tremendous improvements compared to the state-of-the-art (SOTA) CIL method, e.g., substantially increasing the average learning from half accuracy on 11-phase ImageNet-100 by 3.2%.

Our contributions can be summarized as follows. We challenge the traditional manner of saving old class exemplars as RGB images in CIL and propose a memory-efficient data storage approach based on prompts, significantly increasing exemplar quantity with the same memory cost. A general-purpose ControlNet is employed to regenerate diverse high-resolution images from prompts during incremental stages, without fine-tuning it on our target datasets. We devise two techniques, partial compression and diffusion-based data augmentation, to alleviate the domain gap between generated exemplars and real images. Extensive experiments are conducted on four classification datasets, two CIL protocols, seven CIL methods, and three budget settings to evaluate the performance of our approach.

2 Related Work

Class-incremental learning. The goal of continual learning [10, 31] or lifelong learning [1, 8] is to develop machine learning models that can effectively adapt to and learn from data presented in a series of training stages. Recent approaches on these topics are either task-based, where each new phase introduces data from all classes but belonging to a different task [24, 40, 50], or class-based, where each phase provides data from a new set of classes of the same dataset [35, 44, 46, 47, 53]. The latter is typically called class-incremental learning (CIL). CIL methods can be grouped into three main categories: replay-based, regularization-based, and parameter-based [11, 36]. **Replay-based** methods employ a memory buffer to store information from previous classes for rehearsal in later stages. Direct replay [3, 4, 20, 21, 27, 29, 30] saves representative images from the dataset, while generative replay [15, 19, 28, 34, 40, 55] saves generators trained by previous images. These strategies are centered around selecting key samples, training generators, and effectively enhancing the classification model by utilizing a dataset that combines both exemplars and new data. **Regularization-based** methods [2, 7, 22, 23, 48] incorporate regularization terms into the loss function to reinforce previously acquired knowledge while assimilating new data. **Parameter-based** methods [35, 44, 46, 47, 53] allocate distinct model parameters to each incremental phase, aiming to avoid model forgetting that arises from overwriting parameters.

Exemplar compression. Attempts have been made to compress exemplars and reduce their memory consumption. MRDC [45] employs JPEG [43] to compress exemplars. CIM [32] identifies foreground of exemplars by CAM [51] and downsamples the background to compress images. These approaches have two weaknesses. 1) The compression is lossy, and the exemplar image quality degrades. 2) The compression ratio is data-dependent, as JPEG

compression falls short with irregular patterns and non-smooth gradients, and background downsampling is inefficient with dominating foreground regions. Our PESCR, in contrast, guarantees high-resolution exemplars and a constant compression ratio of 24, independent of the image dataset selected.

Diffusion models with spatial control. Diffusion models function by progressively deteriorating the data, introducing Gaussian noise in incremental steps, and then learning to reconstruct the data by reversing the noise introduction process [17, 41, 42]. ControlNet [49] adds spatial conditioning controls to pre-trained text-to-image diffusion models, generating high-quality images from the input text and a visual prompt, which can be edges, depth, segmentation, or human pose representations. This extraordinary ability sheds light on a new potential data format for storing old-class exemplars. Since visual prompts such as edge maps cost much less space than the original RGB images, more prompts can be saved with the same memory budget, and more high-resolution images with the same details can be regenerated by ControlNet in later stages for CIL model training.

3 Methodology

As illustrated in Figure 2, our approach generates diverse exemplars from former prompts to overcome the forgetting problem in CIL. We describe the problem setup of replay-based CIL in Section 3.1. Then, we explain how to compress images to prompts in Section 3.2. Next, we show how to regenerate the exemplars from prompts and train the CIL model with them in Section 3.3. Two techniques to reduce the domain gap are introduced in Section 3.4.

3.1 Problem Setup

Reply-based CIL has multiple phases during which the number of classes gradually increases to the maximum [13, 18, 25]. In the 1-st phase, we observe data $\mathcal{D}_1$, using them to learn an initial model Θ_1. After training, we can only store a small subset of $\mathcal{D}_1$ (i.e., exemplars denoted as $\mathcal{E}_1$) in memory used as replay samples in later phases. In the i-th phase ($i \geq 2$), we get new class data $\mathcal{D}_i$ and load exemplars $\mathcal{E}_{1:i-1} = \mathcal{E}_1 \cup \cdots \cup \mathcal{E}_{i-1}$ from the memory. Then, we initialize Θ_i with Θ_{i-1}, and train it using $\mathcal{E}_{1:i-1} \cup \mathcal{D}_i$. We evaluate the model Θ_i on a test set $\mathcal{Q}_{1:i}$ with all classes observed so far. Eventually, we select exemplars $\mathcal{E}_{1:i}$ from $\mathcal{E}_{1:i-1} \cup \mathcal{D}_i$ and save them in the memory. For the growing memory buffer setup with a budget of b units per class (where one unit corresponds to the memory cost of one image), we need to ensure that the final number of exemplars stored at this stage $|\mathcal{E}_{1:i}| \leq ib$. For the fixed memory buffer setup with a budget of B units in total, we need to have $|\mathcal{E}_{1:i}| \leq B$ in all phases.

3.2 Prompt-Based Super-Compression

The performance of replay-based methods is severely limited by the quantity and diversity of exemplars. To tackle these two issues, we compress old images into visual and textual prompts and use them to regenerate the exemplars by utilizing an off-the-shelf diffusion model. As the prompts require much less memory compared to RGB images, we are able to save a large number of prompts within the same memory budget, so that far more exemplars can be regenerated in subsequent phases.

Prompt extraction. At the end of the i-th phase, we first randomly select a training instance $(\boldsymbol{x}, y)$ from the dataset $\mathcal{D}_i$, where $\boldsymbol{x}$ denotes the image and y denotes the classification label.

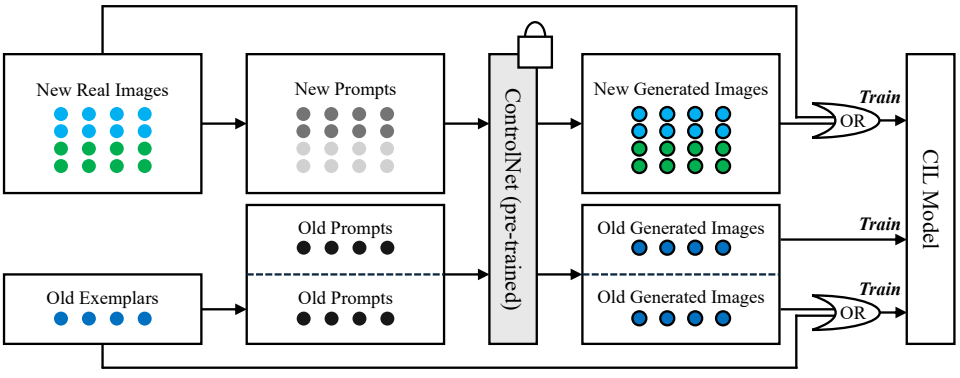

Figure 2: PESCR training process in the i-th phase ($i \geq 2$). A pre-trained ControlNet is downloaded in advance and remains frozen for image generation. Initially, we have three subsets of data: new real images of this phase, old prompts (i.e., edge maps and class tags), and old exemplars. Firstly, we transform the new real images and old exemplars into edge maps by Canny edge detection. Then, we use ControlNet to generate images from all the prompts we have. Finally, we train the CIL model with real and generated images. An image, if having both generated and real versions, will appear only once in each epoch, with a probability p of being the generated version and $1-p$ of being the real version.

Next, we extract the visual prompts. The visual prompts should preserve as many sufficient details as possible to help the CIL model retain the old-class knowledge. They should also be small enough to reduce the memory cost. We choose Canny edge maps as the visual prompts, using the classical Canny edge detector [6] to obtain the edge map $\boldsymbol{e}$:

$$\boldsymbol{e} = \text{CannyEdge}(\boldsymbol{x}). \tag{1}$$

Then, we directly save the class label t as the textual prompt. For example, if the class label is "cupcakes", then $t =$ "cupcakes". The above process is repeated to add other visual and textual prompts to the memory until the memory budget is exhausted. After that, we obtain the final prompt set $\mathcal{P}_i$ for the i-th phase, i.e., $\mathcal{P}_i = \{(\boldsymbol{e}_j, t_j)\}_{j=1}^{R_i}$, where R_i denotes the maximum number of prompts at the i-th phase to fit in the memory.

Prompt memory consumption. Storing edge maps and class tags in the memory buffer takes far less space than storing the original images. For an 8-bit RGB image, compressing it to a 1-bit edge map of equal size achieves a compression ratio of $8 \times 3 = 24$. This means that each memory unit, which can originally store 1 exemplar image, can now store 24 edge maps. The class labels, which usually contain one or two words, consume negligible memory.

3.3 Exemplar Regeneration and CIL Model Training

To regenerate high-resolution exemplars from visual and textual prompts, we leverage a pre-trained ControlNet [49], which is directly applicable without being fine-tuned on target datasets. There is no need to allocate any memory space to save ControlNet, as we are able to re-download the ControlNet model from the cloud at the beginning of each phase. This setting has its strengths and weaknesses. By avoiding saving the generator in our buffer, we

take full advantage of the memory to store old data. However, we have to sacrifice the generator's capability to fit our dataset. Consequently, the generated images might greatly differ from our images in various properties such as brightness, contrast, noise, etc., leading to a potentially huge domain gap. We propose two solutions to tackle this issue in Section 3.4.

Exemplar regeneration. In the i-th phase, we first take out a pair of visual and textual prompts $(\boldsymbol{e},t)$ from the memory $\mathcal{P}_{1:i-1} = P_1 \cup \cdots \cup \mathcal{P}_{i-1}$ of $R_{1:i-1} = \sum_{m=1}^{i-1} R_m$ prompts. Then, we resize the edge map $\boldsymbol{e}$ with nearest neighbor interpolation to meet the input size requirement of ControlNet (both height and width must be a multiple of 64). After that, we forward the prompts $(\boldsymbol{e},t)$ to ControlNet and generate K new exemplars with different random seeds:

$$\hat{\boldsymbol{x}}_k = \text{ControlNet}(\boldsymbol{e},t,s_k), \quad k \in \{1,\cdots,K\}, \tag{2}$$

where k and s_k are the index and its corresponding random seed, respectively. The generated image $\hat{\boldsymbol{x}}_k$ is then resized to the original size of $\boldsymbol{x}$ for consistency of image resolution. We repeat this operation until we finish processing the entire prompt set $\mathcal{P}_{1:i-1}$. Finally, we obtain the regenerated exemplar set $\hat{\mathcal{E}}_{1:i-1}$, which contains $R_{1:i-1}K$ exemplars.

CIL model training. Then, we combine the regenerated exemplars $\hat{\mathcal{E}}_{1:i-1}$ with the new training data $\mathcal{D}_i$ to train the CIL model:

$$\Theta_i \leftarrow \Theta_i - \gamma \nabla_{\Theta_i} \mathcal{L}(\Theta_i; \hat{\mathcal{E}}_{1:i-1} \cup \mathcal{D}_i), \tag{3}$$

where γ denotes the learning rate, and $\mathcal{L}$ is the CIL training loss. To avoid bias towards a large number of similar regenerated images, in each epoch, we randomly sample only one image from the K synthetic exemplars originating from the same base image for training.

3.4 Exemplar-Image Domain Gap Reduction

With PESCR, we can significantly increase the quantity and diversity of exemplars without breaking the memory constraint. However, the domain gap between the synthetic exemplars and real images is still a concern. Therefore, directly using the regenerated exemplars for CIL model training leads to suboptimal performance in practice. A common approach to solve this problem is to fine-tune the generative model on the target dataset. However, if the generator is updated, we have to store it at the end of each phase, taking considerable buffer memory. To circumvent this, we introduce the following two techniques to reduce the domain gap without fine-tuning ControlNet on our datasets.

Partial compression. It might be tempting to utilize all memory budget to save edge maps, obtaining 24 times the original number of exemplars. However, excessive exemplars incur a greater domain gap. Therefore, we consider only spending partial memory on saving edge maps, while leaving the rest to save RGB image exemplars as usual. At each phase i, we assume the memory budget at this phase is B_i units in total, i.e., we are allowed to save at most B_i images as exemplars normally. We set α as the compressed proportion of the dataset whose value can be adjusted depending on the CIL setting. Hence, only αB_i memory units will be allocated for edge maps, saving $24\alpha B_i$ edge maps; while the remaining $(1-\alpha)B_i$ units will be allocated for images, saving $(1-\alpha)B_i$ original images.

Herding [38] is applied to sort the training images in $\mathcal{D}_i$ by their representativeness. Then the $(1-\alpha)B_i$ most representative ones are saved directly, while the next $24\alpha B_i$ images are stored as edge maps in the buffer. This preserves the most representative information of old classes in their original form. The remaining less representative images in $\mathcal{D}_i$ are discarded.

Diffusion-based data augmentation. During CIL model training, every real image $\boldsymbol{x}$ has a certain probability of being replaced by one of its K generated copies. Before training starts at the i-th phase, for each instance $(\boldsymbol{x}, y)$ with a real image $\boldsymbol{x}$, we extract the edge map $\boldsymbol{e}$ from $\boldsymbol{x}$ and obtain K generated copies $\{\hat{\boldsymbol{x}}_k\}_{k=1}^K$ by Equation 2. In each training epoch, $\boldsymbol{x}$ has a probability p of being replaced by any one of $\{\hat{\boldsymbol{x}}_k\}_{k=1}^K$ using uniform sampling. This augmentation operation enables the model to learn from generated features and mitigates the domain gap between real and synthetic images.

Quantity–quality trade-off. By adjusting the compressed proportion α and augmentation probability p, we can control the trade-off between the quantity and quality of generated information that is learned by the CIL model. Larger α and p let the model learn from more generated exemplars more frequently, but the widened domain gap might cause performance degradation. Smaller α and p alleviate the domain gap and improve the information quality, but the exemplar quantity and learning frequency are compromised. The optimal choice of α and p depends on the CIL setting.

Overall CIL training loss. The CIL training process of the i-th phase (Equation 3) is adjusted to:

$$\Theta_i \leftarrow \Theta_i - \gamma \nabla_{\Theta_i} \mathcal{L}(\Theta_i; (\hat{\mathcal{E}}_{1:i-1} \vee \mathcal{E}_{1:i-1}) \cup (\hat{\mathcal{D}}_i \vee \mathcal{D}_i)), \tag{4}$$

where $\hat{\mathcal{E}}_{1:i-1}$ and $\mathcal{E}_{1:i-1}$ represent the regenerated subset and the real-image subset of previous exemplars, respectively. $\hat{\mathcal{D}}_i$ and $\mathcal{D}_i$ are the augmented version and the original version of the new dataset at phase i, respectively. $\vee$ denotes the logic OR operation.

4 Experiments

4.1 Experiment Settings

Datasets. We conduct CIL experiments and evaluate PESCR on four image classification datasets: Caltech-256, Food-101, Places-100, and ImageNet-100. **Caltech-256** [14] is an object recognition dataset with 30,607 images from 257 classes (256 object classes and a clutter class), each class having 80 to 827 images. We remove the clutter class and keep at most 150 images in each class by random selection to avoid extreme class imbalance. Remaining images of each class are randomly split into training (80%) and test (20%) sets. **Food-101** [5] contains 101,000 food images of 101 classes, each class with 750 training and 250 test images. **Places-100** is a subset of Places-365-Standard [52], a large-scale dataset including 365 scene categories with 3,068 to 5,000 training and 100 validation images per class. We construct the subset by randomly choosing 100 classes with seed 0. Then 3,000 training images are randomly chosen from each category. We use their validation set as the test set. **ImageNet-100** is a subset of ImageNet-1000 [12] randomly sampled with seed 1993, following [18], and each class has about 1,300 training and 50 test images.

CIL protocols. We adopt two protocols in our experiments: learning from half (LFH) and learning from scratch (LFS). LFH assumes the model is trained on half of the classes in the first phase and on the remaining classes evenly in the following N phases. LFS assumes the model learns from an equal number of classes in each of the N phases. We set N to be 5 or 10 in our experiments. The model is evaluated on all the classes observed so far at each phase, and the final average classification accuracy is reported.

Memory budget. Although MEMO [53] recently proposes a new buffer setting (with model memory taken into account), based on which MEMO attains SOTA performance by storing lighter models, we still follow the common setup, assigning a fixed number of b memory

Table 1: Average N-phase LFS and $(N+1)$-phase LFH accuracies (%) of different methods, with $b = 5$ memory units/class for Caltech-256 and $b = 20$ for other datasets.

CIL Method	Learning from Scratch (LFS)				Learning from Half (LFH)							
	Caltech-256		ImageNet-100		Caltech-256		Food-101		Places-100		ImageNet-100	
	N=5	10	5	10	5	10	5	10	5	10	5	10
iCaRL [43]	57.7	48.9	66.2	58.7	53.4	49.1	58.6	50.0	42.2	37.6	59.2	50.3
WA [50]	66.2	54.6	76.2	69.7	60.1	46.8	74.2	64.6	62.0	57.3	73.6	66.2
MEMO [53]	65.7	61.4	78.5	74.0	62.6	60.1	71.1	50.1	53.4	48.8	72.5	70.6
PODNet [13]	67.0	60.9	76.4	68.7	68.4	66.6	79.5	77.0	68.3	66.3	79.4	77.2
FOSTER [44]	41.3	36.4	81.1	78.7	62.4	60.9	81.2	78.9	69.4	68.5	81.6	78.6
CIM [32]	65.5	66.3	82.0	78.0	64.1	65.5	79.5	77.1	71.1	70.5	80.5	79.5
DER [46]	68.1	64.8	81.8	78.5	68.4	66.8	82.0	80.4	70.4	69.5	81.8	80.2
PESCR (ours)	**72.8**	**69.8**	**83.4**	**81.3**	**72.1**	**71.3**	**83.9**	**83.2**	**72.3**	**71.8**	**84.2**	**83.4**

Table 2: Average 6-phase LFH accuracies (%) of PESCR on ImageNet-100, with $b = 20$ memory units/class. Diffusion-based augmentation is applied with probability p. The numbers $R+S$ indicate R real and S synthetic exemplars/class are saved in the buffer.

p	20+0	19+24	18+48	17+72	16+96	15+120	14+144	13+168	12+192
0.0	81.8	82.4	82.4	82.5	82.8	82.6	81.6	81.4	80.8
0.1	82.4	82.9	83.4	83.4	83.5	83.4	83.8	83.4	83.5
0.2	82.8	83.0	83.3	83.6	83.9	**84.2**	83.7	83.4	83.9
0.3	82.5	83.4	83.6	84.1	83.8	84.0	**84.2**	84.0	84.0
0.4	82.9	83.1	83.5	84.0	83.7	83.9	83.8	84.0	83.6
0.5	82.4	82.2	83.1	82.8	83.5	83.5	83.4	83.7	83.4

units per class for all methods. The growing budget setting is more challenging than the fixed budget setting, where the memory in earlier phases is much more abundant. (As Caltech-256 has fewer images per class, we set $b = 5$ on Caltech-256 and $b = 20$ on other datasets by default, unless otherwise specified.)

Textual prompt extraction. Textual prompts are derived from class labels with minimal processing. For Caltech-256, we remove the prefix and suffix and replace the hyphen with a space. e.g., "063.electric-guitar-101" is changed to "electric guitar". For Food-101, we replace the underscore with a space. e.g., "apple_pie" is modified as "apple pie". For Places-100, which adopts a bi-level categorization scheme for some classes, such as "general_store/indoor", "general_store/outdoor", and "train_station/platform", we transform them to "indoor general store", "outdoor general store", and "train station platform" to ensure semantic meaningfulness. For ImageNet-100, the original class labels are used directly.

Training setup. Algorithms are evaluated on ResNet-18 [16] which is trained by 200 epochs in the first phase and 170 epochs in subsequent ones with SGD. Data augmentations include random resized cropping, horizontal flip, color jitter, and AutoAugment [9], following [44]. Hyperparameters in PyCIL [54] are adopted to implement all methods. The previous SOTA exemplar compression approach CIM [32] is also implemented by plugging into DER [46] and FOSTER [44], and we report the better result for each experiment. Unless otherwise mentioned, we incorporate PESCR into DER [46], which generally has the best performance across various settings. We choose α and p based on grid search and find that $\alpha \in [0.05, 0.3]$ and $p \in [0.2, 0.4]$ tend to work better in general. $K = 5$ synthetic copies are generated per image for diffusion-based augmentation, but we do not train the model by more epochs.

Table 3: Average 11-phase LFH accuracies (%) on ImageNet-100 with $b = 20$ memory units/class, with and without PESCR plugged in.

	iCaRL	WA	MEMO	PODNet	FOSTER	DER
Baseline	50.3	66.2	70.6	77.2	78.6	80.2
PESCR	66.4	71.7	77.6	80.6	79.4	83.4
Improvements	+16.1	+5.5	+7.0	+3.4	+0.8	+3.2

Table 4: Average 11-phase LFH accuracies (%) of DER on three datasets with b memory units/class, with and without PESCR.

	Food-101			Places-100			ImageNet-100		
	b=5	10	20	5	10	20	5	10	20
DER	77.3	79.1	80.4	67.1	68.6	69.5	77.9	79.3	80.2
DER+PESCR	80.1	81.3	83.2	70.4	70.9	71.8	81.1	81.4	83.4
Improvements	+2.8	+2.2	+2.8	+3.3	+2.3	+2.3	+3.2	+2.1	+3.2

4.2 Results and Discussion

Comparison with previous approaches. We test the CIL performance of different methods with LFS and LFH protocols, illustrating the results in Table 1. PESCR significantly enhances the baseline method DER by a large margin. In 10-phase LFS setting, PESCR improves accuracy by 5.0% and 2.8% on Caltech-256 and ImageNet-100, respectively. In 11-phase LFH setting, PESCR improves accuracy by 4.5%, 2.8%, 2.3%, and 3.2% on Caltech-256, Food-101, Places-100, and ImageNet-100, respectively.

Ablation study. We investigate the effect of partial compression (by setting $\alpha > 0$) and diffusion-based data augmentation (by setting $p > 0$) on PESCR in Table 2. These two operations respectively focus on improving quantity and diversity of exemplars. For straightforward representation, we directly show the number of real (R) and synthetic (S) exemplars per class, instead of α. The compressed ratio can be expressed as $\alpha = 1 - \frac{R}{b}$. It can be observed that when augmentation is not applied, the improvement by increasing exemplar quantity is relatively limited. When augmentation is applied, as the domain gap is reduced, the model benefits much more from additional exemplars. Jointly applying these two techniques yields the best result. However, excessive exemplars (when $\alpha > 0.3$) harm the model performance, as the model is biased towards learning from the dominating synthetic data.

PESCR with different CIL methods. We incorporate PESCR with six different CIL methods and measure the accuracy increase in Table 3. PESCR improves the performance of five methods considerably by a large margin (from 3.2% to 16.1%), and has relatively weak performance only on FOSTER [44]. As reported in [44], increasing the number of exemplars per class from 20 to 50 in FOSTER brings almost no improvement, but with more than 50 exemplars, the domain gap leads to performance degradation for PESCR, thus the accuracy increase on FOSTER is not as tremendous as in other methods.

PESCR with different memory budgets. To further understand the effectiveness of our approach under more limited memory limitations, we alter the memory budget b and quantify the improvements of PESCR in Table 4. It is remarkable that our approach can consistently enhance CIL performance even with more restricted budgets.

Number of generated copies per image. We explore the influence of increasing K (the number of generated copies per image) on PESCR. Table 5 illustrates the accuracy improvement by generating more copies, thanks to the growing sample diversity. The performance increment gradually diminishes when $K \geq 5$, due to 1) the duplicated features generated

Table 5: Average 6-phase LFH accuracies (%) of PESCR on ImageNet-100 with $b = 20$ memory units/class and K generated copies per image. Diffusion augmentation is applied with probability $p = 0.4$, and $R = 16$ real and $S = 96$ synthetic exemplars/class are saved.

K	0	1	5	10	15	20	25
Accuracy	81.8	83.1	83.7	83.4	**83.8**	83.7	83.4

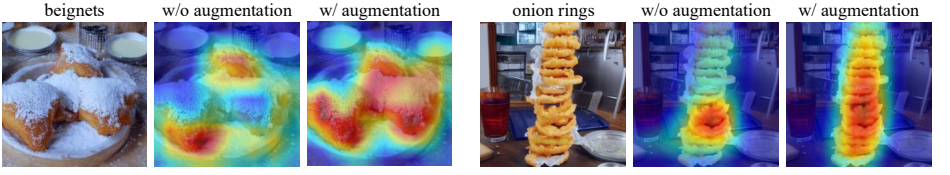

Figure 3: Two generated exemplars of Food-101 with Grad-CAM given by the classification model trained without diffusion-based augmentation and with augmentation. Activation regions given by the model trained with augmentation capture class objects more accurately.

from each edge map, 2) limited training epochs to learn from additional generations, and 3) the widened domain gap between real and superfluous synthetic images. Therefore, we use $K = 5$ copies per image for all the experiments.

Grad-CAM visualization of generated exemplars. To verify if the class objects generated by ControlNet can be successfully identified by the CIL model, we use Grad-CAM [39] to detect the image region with the most importance to the classification decision. Displayed in Figure 3 are two synthetic exemplars of Food-101 with their activation maps produced by a model trained without diffusion-based data augmentation and a model trained with augmentation. Evidently, the augmentation process is vital for the classification model to comprehensively detect the generated objects. This ensures that the model can properly benefit from the synthetic exemplars during training in subsequent stages.

5 Conclusion

We present PESCR, a novel prompt-based exemplar super-compression and regeneration approach that advances the state of replay-based class-incremental learning (CIL). By introducing a compact and expressive storage format based on visual and textual prompts, we challenge the conventional reliance on storing raw RGB images as exemplars. PESCR effectively compresses old-class information, dramatically increasing the number of stored exemplars within the same memory constraints. The prompts are later used to regenerate diverse, high-resolution synthetic images by an off-the-shelf pre-trained diffusion model, such as ControlNet, without requiring fine-tuning or additional memory storage for the generative model. To address the potential domain gap between generated and real images, we introduce two complementary strategies, partial compression and diffusion-based data augmentation, to allow PESCR to maintain model generalization and stability across phases. Comprehensive experiments demonstrate that our approach constantly improves the model performance by a large margin on numerous datasets and CIL settings.

References

[1] Rahaf Aljundi, Punarjay Chakravarty, and Tinne Tuytelaars. Expert gate: Lifelong learning with a network of experts. In *CVPR*, pages 3366–3375, 2017.

[2] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *ECCV*, pages 139–154, 2018.

[3] Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient based sample selection for online continual learning. *NeurIPS*, 32, 2019.

[4] Jihwan Bang, Heesu Kim, YoungJoon Yoo, Jung-Woo Ha, and Jonghyun Choi. Rainbow memory: Continual learning with a memory of diverse samples. In *CVPR*, pages 8218–8227, 2021.

[5] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101–mining discriminative components with random forests. In *ECCV*, pages 446–461. Springer, 2014.

[6] John F. Canny. A computational approach to edge detection. *TPAMI*, 8(6):679–698, 1986.

[7] Arslan Chaudhry, Puneet K Dokania, Thalaiyasingam Ajanthan, and Philip HS Torr. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *ECCV*, pages 532–547, 2018.

[8] Zhiyuan Chen and Bing Liu. Lifelong machine learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 12(3):1–207, 2018.

[9] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *CVPR*, pages 113–123, 2019.

[10] Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Aleš Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *arXiv*, 1909.08383, 2019.

[11] Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Ales Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. Continual learning: A comparative study on how to defy forgetting in classification tasks. *arXiv*, 1909.08383, 2019.

[12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009.

[13] Arthur Douillard, Matthieu Cord, Charles Ollion, Thomas Robert, and Eduardo Valle. Podnet: Pooled outputs distillation for small-tasks incremental learning. In *ECCV*, pages 86–102. Springer, 2020.

[14] Gregory Griffin, Alex Holub, and Pietro Perona. Caltech 256, Apr 2022.

[15] Chen He, Ruiping Wang, Shiguang Shan, and Xilin Chen. Exemplar-supported generative reproduction for class incremental learning. In *BMVC*, page 98, 2018.

[16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.

[17] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020.

[18] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *CVPR*, pages 831–839, 2019.

[19] Wenpeng Hu, Zhou Lin, Bing Liu, Chongyang Tao, Zhengwei Tao, Jinwen Ma, Dongyan Zhao, and Rui Yan. Overcoming catastrophic forgetting for continual learning via model adaptation. In *ICLR*, 2018.

[20] Ahmet Iscen, Jeffrey Zhang, Svetlana Lazebnik, and Cordelia Schmid. Memory-efficient incremental learning through feature adaptation. In *ECCV*, pages 699–715. Springer, 2020.

[21] David Isele and Akansel Cosgun. Selective experience replay for lifelong learning. In *AAAI*, volume 32, 2018.

[22] 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. *PNAS*, 114(13):3521–3526, 2017.

[23] Sang-Woo Lee, Jin-Hwa Kim, Jaehyun Jun, Jung-Woo Ha, and Byoung-Tak Zhang. Overcoming catastrophic forgetting by incremental moment matching. *NeurIPS*, 30, 2017.

[24] Zhizhong Li and Derek Hoiem. Learning without forgetting. *TPAMI*, 40(12):2935–2947, 2017.

[25] Yaoyao Liu, Yuting Su, An-An Liu, Bernt Schiele, and Qianru Sun. Mnemonics training: Multi-class incremental learning without forgetting. In *CVPR*, pages 12245–12254, 2020.

[26] Yaoyao Liu, Bernt Schiele, and Qianru Sun. Adaptive aggregation networks for class-incremental learning. In *CVPR*, pages 2544–2553, 2021.

[27] Yaoyao Liu, Bernt Schiele, and Qianru Sun. Rmm: Reinforced memory management for class-incremental learning. *NeurIPS*, 34:3478–3490, 2021.

[28] Yaoyao Liu, Qianru Sun, Xiangnan He, An-An Liu, Yuting Su, and Tat-Seng Chua. Generating face images with attributes for free. *TNNLS*, 32(6):2733–2743, 2021.

[29] Yaoyao Liu, Yingying Li, Bernt Schiele, and Qianru Sun. Online hyperparameter optimization for class-incremental learning. In *AAAI*, pages 8906–8913, 2023.

[30] Yaoyao Liu, Yingying Li, Bernt Schiele, and Qianru Sun. Wakening past concepts without past data: Class-incremental learning from online placebos. In *WACV*, pages 2215–2224, 2024.

[31] David Lopez-Paz and Marc'Aurelio Ranzato. Gradient episodic memory for continual learning. In *NIPS*, pages 6467–6476, 2017.

[32] Zilin Luo, Yaoyao Liu, Bernt Schiele, and Qianru Sun. Class-incremental exemplar compression for class-incremental learning. In *CVPR*, pages 11371–11380, 2023.

[33] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier, 1989.

[34] Grégoire Petit, Adrian Popescu, Hugo Schindler, David Picard, and Bertrand Delezoide. Fetril: Feature translation for exemplar-free class-incremental learning. In *WACV*, pages 3911–3920, 2023.

[35] Quang Pham, Chenghao Liu, and Steven Hoi. Dualnet: Continual learning, fast and slow. *NeurIPS*, 34:16131–16144, 2021.

[36] Ameya Prabhu, Philip HS Torr, and Puneet K Dokania. Gdumb: A simple approach that questions our progress in continual learning. In *ECCV*, 2020.

[37] R. Ratcliff. Connectionist models of recognition memory: Constraints imposed by learning and forgetting functions. *Psychological Review*, 97:285–308, 1990.

[38] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *CVPR*, pages 2001–2010, 2017.

[39] 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*, pages 618–626, 2017.

[40] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. *NeurIPS*, 30, 2017.

[41] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, et al. Make-a-video: Text-to-video generation without text-video data. *arXiv preprint arXiv:2209.14792*, 2022.

[42] Ruben Villegas, Mohammad Babaeizadeh, Pieter-Jan Kindermans, Hernan Moraldo, Han Zhang, Mohammad Taghi Saffar, Santiago Castro, Julius Kunze, and Dumitru Erhan. Phenaki: Variable length video generation from open domain textual description. *arXiv preprint arXiv:2210.02399*, 2022.

[43] Gregory K Wallace. The jpeg still picture compression standard. *IEEE Transactions on Consumer Electronics*, 1992.

[44] Fu-Yun Wang, Da-Wei Zhou, Han-Jia Ye, and De-Chuan Zhan. Foster: Feature boosting and compression for class-incremental learning. In *ECCV*, pages 398–414. Springer, 2022.

[45] 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. *arXiv preprint arXiv:2202.06592*, 2022.

[46] Shipeng Yan, Jiangwei Xie, and Xuming He. Der: Dynamically expandable representation for class incremental learning. In *CVPR*, pages 3014–3023, 2021.

[47] Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong learning with dynamically expandable networks. *arXiv preprint arXiv:1708.01547*, 2017.

[48] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *ICML*, pages 3987–3995. PMLR, 2017.

[49] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *ICCV*, pages 3836–3847, 2023.

[50] Bowen Zhao, Xi Xiao, Guojun Gan, Bin Zhang, and Shu-Tao Xia. Maintaining discrimination and fairness in class incremental learning. In *CVPR*, pages 13208–13217, 2020.

[51] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *CVPR*, pages 2921–2929, 2016.

[52] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Antonio Torralba, and Aude Oliva. Places: An image database for deep scene understanding. *arXiv preprint arXiv:1610.02055*, 2016.

[53] Da-Wei Zhou, Qi-Wei Wang, Han-Jia Ye, and De-Chuan Zhan. A model or 603 exemplars: Towards memory-efficient class-incremental learning. *arXiv preprint arXiv:2205.13218*, 2022.

[54] Da-Wei Zhou, Fu-Yun Wang, Han-Jia Ye, and De-Chuan Zhan. Pycil: a python toolbox for class-incremental learning. *SCIENCE CHINA Information Sciences*, 66(9):197101–, 2023. doi: https://doi.org/10.1007/s11432-022-3600-y.

[55] Fei Zhu, Xu-Yao Zhang, Chuang Wang, Fei Yin, and Cheng-Lin Liu. Prototype augmentation and self-supervision for incremental learning. In *CVPR*, pages 5871–5880, 2021.