

Class-Incremental Exemplar Compression for Class-Incremental Learning

Zilin Luo¹ Yaoyao Liu² Bernt Schiele² Qianru Sun¹

¹Singapore Management University

²Max Planck Institute for Informatics, Saarland Informatics Campus

zilin.luo.2021@phdcs.smu.edu.sg {yaoyao.liu, schiele}@mpi-inf.mpg.de qianrusun@smu.edu.sg

Abstract

Exemplar-based class-incremental learning (CIL) [36] finetunes the model with all samples of new classes but few-shot exemplars of old classes in each incremental phase, where the “few-shot” abides by the limited memory budget. In this paper, we break this “few-shot” limit based on a simple yet surprisingly effective idea: compressing exemplars by downsampling non-discriminative pixels and saving “many-shot” compressed exemplars in the memory. Without needing any manual annotation, we achieve this compression by generating 0-1 masks on discriminative pixels from class activation maps (CAM) [49]. We propose an adaptive mask generation model called class-incremental masking (CIM) to explicitly resolve two difficulties of using CAM: 1) transforming the heatmaps of CAM to 0-1 masks with an arbitrary threshold leads to a trade-off between the coverage on discriminative pixels and the quantity of exemplars, as the total memory is fixed; and 2) optimal thresholds vary for different object classes, which is particularly obvious in the dynamic environment of CIL. We optimize the CIM model alternatively with the conventional CIL model through a bilevel optimization problem [40]. We conduct extensive experiments on high-resolution CIL benchmarks including Food-101, ImageNet-100, and ImageNet-1000, and show that using the compressed exemplars by CIM can achieve a new state-of-the-art CIL accuracy, e.g., 4.8 percentage points higher than FOSTER [42] on 10-Phase ImageNet-1000. Our code is available at <https://github.com/xffzl/CIM-CIL>.

1. Introduction

Dynamic AI systems have a continual learning nature to learn new class data. They are expected to adapt to new classes while maintaining the knowledge of old classes, i.e., free from forgetting problems [31]. To evaluate this, the following protocol of class-incremental learning (CIL) was proposed by Rebuffi et al. [36]. The model training goes

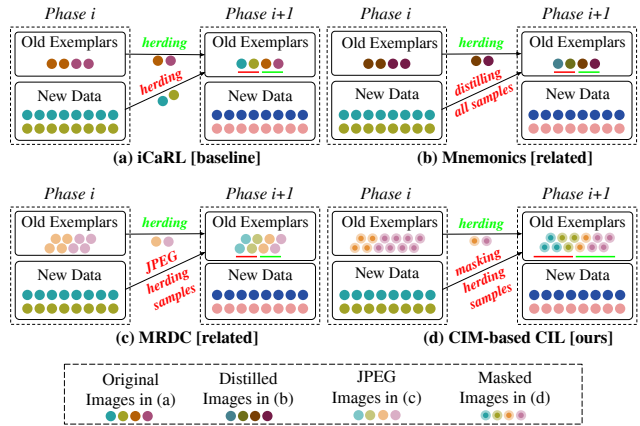


Figure 1. The phase-wise training data in different methods. (a) iCaRL [36] is the baseline method using full new class data and few-shot old class exemplars. (b) Mnemonics [27] distills all training samples into few-shot exemplars without increasing their quantity. (c) MRDC [43] compresses each exemplar uniformly into a low-resolution image using JPEG [41]. (d) Our approach based on the proposed class-incremental masking (CIM) downsamples only non-discriminative pixels in the image. The legend shows the symbols of special images generated by the methods.

through a number of phases. Each phase has new class data added and old class data discarded, and the resultant model is evaluated on the test data of all seen classes. A straightforward way to retain old class knowledge is keeping around a few old class exemplars in the memory and using them to re-train the model in subsequent phases. The number of exemplars is usually limited, e.g., 5 ~ 20 exemplars per class [12, 17, 25, 27, 36, 42, 44, 46, 48], as the total memory in CIL strictly budgeted, e.g., $2k$ exemplars.

This leads to a serious data imbalance between old and new classes, e.g., 20 per old class vs. 1.3k per new class (on ImageNet-1000 [9]), as illustrated in Figure 1a. The training is thus always dominated by new classes, and forgetting problems occur for old classes. Liu et al. [27] tried to mitigate this problem by parameterizing and distilling the exem-

plars, without increasing the number of them (Figure 1b). Wang et al. [43] traded off between the quality and quantity of exemplars by uniformly compressing exemplar images with JPEG [41] (Figure 1c). As shown in Figure 1d, our approach is also based on image compression. The idea is to downsample only non-discriminative pixels (e.g., background) and keep discriminative pixels (i.e., representative cues of foreground objects) as the original. In this way, **we do not sacrifice the discriminativeness** of exemplars when increasing their quantity. In particular, **we aim for adaptive compression** in dynamic environments of CIL, where the intuition is later phases need to be more conservative (i.e., less downsampling) as the model needs more visual cues to classify the increased number of classes.

To achieve selective and adaptive compression, we need the location labels of discriminative pixels. Without extra labeling, we automatically generate the labels by utilizing the model’s own “attention” on discriminative features, i.e., class activation maps (CAM) [49]. We take this method as a feasible baseline, and based on it, we propose an adaptive version called class-incremental masking (CIM). Specifically, for each input image (with its class label), we use its feature maps and classifier weights (corresponding to its class label) to compute a CAM by channel-wise multiplication, aggregation, and normalization. Then, we apply hard thresholding to generate a 0-1 mask.¹ We notice that when generating the masks in the dynamic environments of CIL, the optimal hyperparameters (such as the value of hard threshold and the choice of activation functions) vary for different classes as well as in different incremental phases. Our adaptive version CIM tackles this by parameterizing a mask generation model and optimizing it in an end-to-end manner across all incremental phases. In each phase, the learned CIM model adaptively generates class- and phase-specific masks. We find that the compressed exemplars based on these masks have stronger representativeness, compared to using the conventional CAM.

Technically, we have two models to optimize, i.e., the CIL model and the CIM model.² These two cannot be optimized separately as they are dependent on computation: 1) the CIM model compresses exemplars to input into the CIL model; 2) the two models share network parameters. We exploit a global bilevel optimization problem (BOP) [7, 40] to alternate their training processes at two levels. This BOP goes through all incremental training phases. In particular, for each phase, we perform a local BOP with two steps to tune the parameters of the CIM model: 1) a temporary model is trained with the compressed exemplars as input; and 2) a validation loss on the uncom-

pressed new data is computed and the gradients are back-propagated to optimize the parameters of CIM. To evaluate CIM, we conduct extensive experiments by plugging it in recent CIL methods,³ LUCIR [17], DER [46], and FOSTER [42], on three high-resolution benchmarks, Food-101 [3], ImageNet-100 [17], and ImageNet-1000 [9]. We find that using the compressed exemplars by CIM brings consistent and significant improvements, e.g., 4.2% and 4.8% higher than the SOTA method FOSTER [42], respectively, in the 5-phase and 10-phase settings of ImageNet-1000, with a total memory budget for 5k exemplars.

2. Related Work

Class-Incremental Learning (CIL). There are three main lines of work to address the catastrophic forgetting problem [30, 31] in CIL. *Regularization-based* methods apply discrepancy (between old and new models) penalization terms in their objective functions, e.g., by comparing output logits [23, 36], intermediate features [12, 17, 24, 39], and prediction heatmaps [10]. *Parameter-isolation-based* methods increase the model parameters in each new incremental phase, to prevent knowledge forgetting caused by parameter overwritten. Some of them [18, 37, 42, 45, 46] proposed to progressively expand the size of the neural network to learn new coming data. Others [1, 21, 25, 47] froze a part of network parameters (to maintain the old class knowledge) to alleviate the problem of knowledge overwriting. *Replay-based* methods assume there is a clear memory budget allowing a handful of old-class exemplars in the memory. Exemplars can be used to re-train the model in each new phase [12, 17, 26, 27, 36, 43, 44]. This re-training usually contains two steps: one step of training the model on all new class data and old class exemplars, and one step of finetuning the model with a balanced subset (i.e., using an equal number of samples per class) [12, 17, 25, 26, 46].

The replay-based methods focusing on *memory optimization* [27, 43] are closely related to our work. [27] proposed a bilevel optimization framework to distill the current new class data into exemplars before discarding them. It aims to improve the quality of exemplars without increasing the quantity. Another work [43] aimed to trade-off between the quality and quantity of exemplars by image compression using the JPEG algorithm, i.e., each exemplar is uniformly downsampled. Ours differs from these two works in three aspects. 1) Our CIM based image compression automatically segments the discriminative pixels in the exemplar and downsamples the non-discriminative pixels only. It barely weakens the representativeness of the exemplar. In contrast, the parameterization of image pixels in [27] hampers the model from capturing high-frequency (discriminative) features from the image, especially the high-resolution (e.g.,

¹Note that we do not use mask labels to do image compression because storing them is expensive. Instead, we expand the mask to a bounding box, as elaborated in Section 4.

²Note that the CIM model is actually a plug-in branch in the CIL model, which is detailed in Section 4.2.

³Using “plug-in” for evaluation is due to the fact that many baseline methods were originally evaluated in different CIL settings.

224×224) image [6]. 2) Our approach increases the diversity (i.e., quantity) of old class exemplars by reducing the memory consumption for each exemplar. In contrast, [27] keeps a fixed number of exemplars in the memory. 3) Our approach has an adaptive image compression strategy that fits well in the dynamic environments of CIL. In contrast, [43] uses uniform image compression (or uniformly increasing quality parameters from 100 to 1) without considering the properties of specific classes in each learning phase.

Class Activation Map (CAM) [49] is a simple yet effective weakly-supervised object localization method. Its model is trained with only the image-level label and can generate pixel-level masks on foreground objects. Specifically, the masks are the results of hard-thresholding the heatmaps produced by feature maps and classifier weights. Advanced CAM variants include Grad-CAM [38], ReCAM [8], AdvCAM [22], etc. Our CIM is based on the vanilla CAM because it is computationally simple and efficient.

Bilevel Optimization Problems (BOP) [7, 40] aims to solve a nested optimization problem, where the outer-level optimization is subjected to the result of the inner-level optimization. It has shown effectiveness in a wide range of machine learning areas, such as hyperparameter selection [29] and meta-learning [13]. For tackling CIL tasks, [27] leverages BOP to alternatively optimize the parameters of the CIL model and the parameterized exemplars. [25] applies BOP to learn the aggregation weights of the plastic and elastic branches in the CIL model. In our work, we use BOP to solve the optimization of the CIL model and the parameterized Class-Incremental Masking (CIM) model, where CIM is a plugin branch (in the CIL model), using few extra parameters. The process of BOP is quick and efficient.

3. Preliminary

The following is the training pipeline of standard CIL with few-shot exemplars. Assume there are N learning phases. In the 1-st phase, we load data \mathcal{D}_1 containing all training samples of c_1 classes, and use \mathcal{D}_1 to train the initial classification model (θ_1, ω_1) , where θ_1 and ω_1 denote the parameters of the feature extractor and classifier, respectively. When the training is done, we evaluate the model performance on the test samples of c_1 classes. Before the 2-nd phase, we discard most of the training samples due to the strict memory budget of CIL. In other words, we preserve only a handful of training samples \mathcal{E}_1 (i.e., exemplars) in the memory, selected from \mathcal{D}_1 . A common method for selecting exemplars is called feature herding [36] and has been used in many related works [27, 42, 43, 46]. We adopt it, too, in this work.

In the i -th phase ($i \geq 2$), we load all exemplars $\mathcal{E}_{1:i-1} = \mathcal{E}_1 \cup \dots \cup \mathcal{E}_{i-1}$ from the memory and initialize the current model (θ_i, ω_i) by the previous model $(\theta_{i-1}, \omega_{i-1})$. We use $\mathcal{E}_{1:i-1}$ and the new coming data \mathcal{D}_i (containing c_i

new classes) to train (θ_i, ω_i) . Then, we evaluate the current model using a test set of all $\sum_{j=1}^i c_j$ classes seen so far. After that, we discard most of the training samples in \mathcal{D}_i , and leave few-shot exemplars \mathcal{E}_i in the memory. It is clear that this discarding causes a strong data imbalance between old and new coming classes in the subsequent phase. In the following, we introduce our solution to this problem.

4. Methodology

As illustrated in Figure 1d, we alleviate the data imbalance problem by saving a larger number of compressed exemplars for old classes, where we leverage pixel-selective compression, i.e., downsample only non-discriminative pixels. To achieve this, **the ideal case** is that we have the pixel-level localization of foreground objects. However, **the realistic case** is that such localization labels are expensive, most CIL benchmarks do not have the labels, and it is not fair to compare with other CIL methods if using the labels. Without extra labeling, we introduce a CAM-based mask generation method, and based on it, we provide a baseline solution to pixel-selective compression in Section 4.1. **The problem** of mask generation in CIL is that the optimal generation hyperparameters such as hard thresholds are changing in the dynamic environment (with the increasing number of classes and phases). It is thus desirable to have an adaptive mask generation process. To this end, we propose class-incremental masking (CIM)—a learnable mask generation model, in Section 4.2.

4.1. CAM-based Compression Pipeline

Generating pixel-level labels for large-scale datasets, e.g., ImageNet [9], is non-trivial. Using class activation maps (CAM) is a naïve solution with little computation costs. Its key idea is to make use of the activation of the classification model itself: on the feature maps, activated pixels are more discriminative than non-activated ones for recognizing the object, where “activated” means *of high activation values and with strong correlation with the classification weights of the object*. After localizing activated pixels by CAM, we can generate a 0-1 mask on them, e.g., by hard thresholding their normalized values, and then upsample the mask to the size of the input image [49].

From CAM to 0-1 Mask. We extract CAM in the following steps. Given an image x from \mathcal{D}_i and its ground truth class label y , let $F(x; \theta_i)$ represent the feature block output by the feature extractor θ_i , and $\omega_{i,y}$ for the classification weights of class y in the classifier ω_i . The CAM of x is:

$$\mathcal{M}^{\text{CAM}} = \frac{A - \min(A)}{\max(A) - \min(A)}, \quad A = \omega_{i,y}^\top F(x; \theta_i), \quad (1)$$

where $\min(\cdot)$ and $\max(\cdot)$ operations are used for normalization. Then, we upsample \mathcal{M}^{CAM} to the size of image x and use the same notation.




Figure 2. The proposed compression pipeline. Given an image, we extract its CAM-based (baseline) or CIM-based (ours) mask, threshold it to be a 0-1 mask with a fixed threshold τ , and generate a tight bounding box (bbox) that covers all masked pixels. Then, we downsample the pixels outside the bbox and sum the downsampled image up to the masked image to generate the final compressed image.

Each value in \mathcal{M}^{CAM} denotes the activation strength of the model at a specific pixel location. Following the way of generating 0-1 masks in weakly-supervised semantic segmentation works [8, 11, 22], we apply a hard threshold τ (between 0 and 1) over all values of \mathcal{M}^{CAM} , and get the 0-1 mask \mathcal{M}^τ : $\mathcal{M}^\tau = \mathbb{I}(\mathcal{M}^{\text{CAM}} > \tau)$, where $\mathbb{I}(\cdot)$ is the indicator function. In \mathcal{M}^τ , 1s indicate the locations of discriminative pixels, e.g., foreground pixels, based on which the model makes the prediction. While 0s indicate mostly background pixels or non-discrimination pixels that can be downsampled as they make little contribution to the prediction.

After generating 0-1 masks, it is ideal to keep them in the memory as meta information of compression. However, this is not efficient or feasible in CIL. There are two reasons. 1) The space for saving image-size masks is non-negligible. Each mask pixel is a one-bit boolean value, and one mask takes around $\frac{1}{3 \times 8} = \frac{1}{24}$ of the memory of one RGB image. 2) The mask involves activated regions with irregular shapes. It is thus non-trivial to perform any standard downsampling algorithm [34] on the remaining regions.

From 0-1 Masks to Bounding Boxes (BBox). A simple workaround is to generate a tight bounding box (bbox) to cover the positions of 1s in \mathcal{M}^τ , and use the bbox for compression. Specifically, given \mathcal{M}^τ , we obtain the coordinate representation of the bounding box as:

$$\mathcal{B} = [\min h, \min w; \max h, \max w]_{(h,w): \mathcal{M}^\tau(h,w)=1}, \quad (2)$$

where h and w denote the vertical and horizontal coordinates of the 1 on \mathcal{M}^τ , respectively. We highlight that \mathcal{B} consists of four integers only and takes negligible memory overhead compared to \mathcal{M}^τ . In addition, we “reshape” the irregular shape (of the activated region) in \mathcal{M}^τ into a rectangular \mathcal{B} , so our downsampling operation on the pixels outside the rectangular becomes easy.

Compression with BBox. Given the image and its bbox on foreground, the compression is implemented by downsampling pixels outside the bbox. Specifically, as illustrated in Figure 2, we compress the image x to \tilde{x} as follows,

$$\tilde{x} = \mathcal{M}^{\mathcal{B}} \odot x + (1 - \mathcal{M}^{\mathcal{B}}) \odot x_\eta, \quad (3)$$

where $\mathcal{M}^{\mathcal{B}}$ is the binary mask according to \mathcal{B} , i.e., the val-

ues of $\mathcal{M}^{\mathcal{B}}$ are 1 inside \mathcal{B} , and 0 otherwise. x_η is the fully downsampled version of x with a downsampling ratio η ($\eta > 1$), \odot denotes the element-wise product, and $+$ denotes the element-wise addition. Both \odot and $+$ are applied independently on each RGB channel.

The memory allocated for the compressed image \tilde{x} is as follows,

$$\begin{aligned} m_{\tilde{x}} &= \frac{H_{\mathcal{B}}W_{\mathcal{B}}}{HW} + \frac{1}{\eta} \left(1 - \frac{H_{\mathcal{B}}W_{\mathcal{B}}}{HW}\right) \\ &= 1 - \left(1 - \frac{1}{\eta}\right) \cdot \left(1 - \frac{H_{\mathcal{B}}W_{\mathcal{B}}}{HW}\right), \end{aligned} \quad (4)$$

where $H_{\mathcal{B}}$ and $W_{\mathcal{B}}$ are the height and width of \mathcal{B} , respectively. H and W are the height and width of the original image x , respectively. $m_{\tilde{x}}$ is always smaller than 1 where 1 denotes the memory unit of saving one original image x . Therefore, we can save a larger number of compressed exemplars within the same memory budget. We denote the set of compressed exemplars in the i -th phase as $\tilde{\mathcal{E}}_i$.

Compression Artifacts. The above compression introduces artifacts to the compressed images, i.e., there is a resolution mutation around bounding box edges. From the perspective of spectrum analysis [5, 20], such mutation carries noisy and high-frequency components and impairs the model training in subsequent phases. We mitigate the effect of these artifacts by implementing the following data augmentation: in each training epoch, we transform a random subset of \mathcal{D}_i into compressed images with CAM-based bounding boxes using the same downsampling ratio. Using this augmentation enables the model to “simulate” the training with compressed images and learn to be invariant to compression artifacts.

4.2. Class-Incremental Masking (CIM)

Ideally, the mask generation process needs to adjust at different phases in CIL environments. The process involves two hyperparameters: the masking threshold and the choice of network activation functions. First, for the threshold, searching for its optimum (for all classes) is not trivial. Grid search is intuitive, but it is computationally expensive when the number of classes increases in CIL. Second, for the acti-

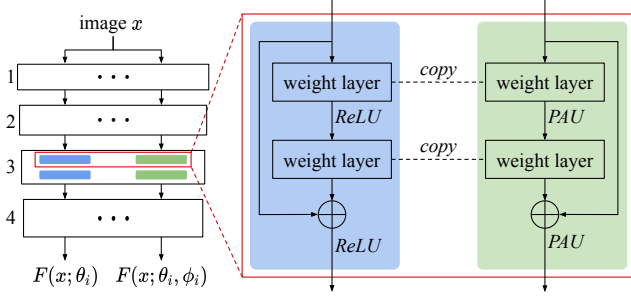


Figure 3. Our CIM installs an “extension” on the backbone (e.g., ResNet-18 [16, 36] with four blocks) by adding a learnable activation function (e.g., PAU [32]) at the position of the original activation function (i.e., ReLU [33]). Such “extension” results in a new network branch (in green), whose weight layer parameters are directly copied from the original branch (in blue).

vation function, the standard network of CIL methods uses ReLU [33, 36] and does not optimize it.

We solve this issue by applying a learnable activation function in addition to the existing ReLU function in the CIL model. Physically, we have one neural network, while logically, we have two models to learn (in each incremental phase): the conventional CIL model with ReLU activations, and the adaptive mask generation model with learnable activations. We thus call our method class-incremental masking (CIM) based CIL. In the following, we elaborate on the network design and optimization pipeline.

Network Design. Figure 3 demonstrates an example network architecture in our CIM-based CIL. The proposed CIM extends the network backbone by *logically* adding a network branch, where only the activation functions are learnable (e.g., Padé Activation Units (PAU⁴, [32])) and the parameters of weight layers are copied from the original branch. This design is motivated by the works of He et al. [15] and Bochkovskiy et al. [2], which indicate that layers with learnable activation functions can flexibly process object (localization) information at different network blocks. The difference is that we apply this flexibility to achieve adaptive mask generation for different CIL phases. We denote the CIM parameters (i.e., the parameters in learnable activation functions) in i -th phase as ϕ_i . We optimize the CIL model (θ_i, ω_i) and the CIM model ϕ_i via a global BOP, as elaborated below.

Optimization Pipeline. We demonstrate the overall optimization flow in Algorithm 1, which consists of two levels of optimization: task-level and mask-level—the former one for CIL and the latter one for CIM. Note that to maintain a unified notation, we further define $\tilde{\mathcal{E}}_0 = \emptyset$.

1) Task-level Optimization. This level aims to optimize the CIL model (θ_i, ω_i) to address the CIL task at hand. It

⁴PAU uses a rational function of given degrees m and n , i.e., $(a_0 + a_1x + \dots + a_mx^m)/(1 + b_1x + \dots + b_nx^n)$, and can be parameterized by a_0, \dots, a_m and b_1, \dots, b_n .

Algorithm 1: CIM-based CIL (i -th phase, $i \geq 1$)

Input: New data \mathcal{D}_i ; old compressed exemplars $\tilde{\mathcal{E}}_{0:i-1}$ ($\tilde{\mathcal{E}}_0 = \emptyset$); last-phase CIL model $(\theta_{i-1}, \omega_{i-1})$ and CIM model ϕ_{i-1} (θ_0, ω_0 are random parameters, ϕ_0 is set to ReLU).

Output: New compressed exemplars $\tilde{\mathcal{E}}_i$; updated CIL model (θ_i, ω_i) and CIM model ϕ_i .

- 1 Initialize (θ_i, ω_i) with $(\theta_{i-1}, \omega_{i-1})$;
 - 2 Initialize ϕ_i with ϕ_{i-1} ;
 - 3 **for** *epochs* **do**
 - 4 Train (θ_i, ω_i) using $\tilde{\mathcal{E}}_{0:i-1} \cup \mathcal{D}_i$ by Eq. 5;
 - 5 Compress \mathcal{D}_i into $\tilde{\mathcal{D}}_i^\phi$ using ϕ_i ;
 - 6 Temporarily update θ_i to θ_i^+ using $\tilde{\mathcal{E}}_{0:i-1} \cup \tilde{\mathcal{D}}_i^\phi$ by Eq. 7;
 - 7 Learn ϕ_i using \mathcal{D}_i by Eq. 9;
 - 8 **end**
 - 9 Compress \mathcal{D}_i into $\tilde{\mathcal{D}}_i$ using the learned ϕ_i by Eq. 3;
 - 10 Select exemplars $\tilde{\mathcal{E}}_i$ from $\tilde{\mathcal{D}}_i$ by, e.g., herding [36].
-

can be written as:

$$(\theta_i, \omega_i) \leftarrow (\theta_i, \omega_i) - \lambda \nabla_{(\theta, \omega)} \mathcal{L}_{\text{CIL}}(\tilde{\mathcal{E}}_{0:i-1} \cup \mathcal{D}_i; \theta_i, \omega_i), \quad (5)$$

where λ is the learning rate. We follow the implementation of CIL training loss \mathcal{L}_{CIL} in baseline methods [23, 42, 46, 50]. This means that we use different training losses when plugging CIM into different baseline methods.

2) Mask-level Optimization. This level aims to optimize the CIM model ϕ_i to produce adaptive compression masks. It is formulated as a local BOP:

$$\min_{\phi_i} [\mathcal{L}_{\text{val}}(\mathcal{D}_i; \theta_i^*, \omega_i) + \mu \mathcal{R}(\phi_i)], \quad (6a)$$

$$\text{s.t. } \theta_i^* = \arg \min_{\theta_i} \mathcal{L}_{\text{train}}(\tilde{\mathcal{E}}_{0:i-1} \cup \tilde{\mathcal{D}}_i(\phi_i); \theta_i, \omega_i). \quad (6b)$$

Eq. 6b denotes an inner-level optimization. It trains θ_i with the data $\tilde{\mathcal{D}}_i(\phi_i)$ compressed by using ϕ_i , and converges as θ_i^* . Eq. 6a denotes an outer-level optimization. It is based on the validation loss derived by θ_i^* on the original data \mathcal{D}_i . $\mathcal{R}(\phi_i)$ is a constraint representing the memory limitation and μ is its weight. In the following, we elaborate on the implementation details for the two levels.

In the inner-level optimization, we train a temporary CIL model with compressed data. Specifically, we first compress new-class data \mathcal{D}_i into $\tilde{\mathcal{D}}_i(\phi_i)$ using the generated masks by the CIM model ϕ_i . Then, we implement the inner-level optimization as a one-step gradient descent (using the CIL training loss) as:

$$\theta_i^+ \leftarrow \theta_i - \beta_1 \nabla_{\theta} \mathcal{L}_{\text{CIL}}(\tilde{\mathcal{E}}_{0:i-1} \cup \tilde{\mathcal{D}}_i(\phi_i); \theta_i, \omega_i), \quad (7)$$

where β_1 is the learning rate for θ_i .

The aim of the outer-level optimization is to optimize ϕ_i such that the temporary model (θ_i^+, ω_i) (trained with compressed data) has a low validation loss on the original data \mathcal{D}_i . To achieve this, we back-propagate the loss on the original data \mathcal{D}_i to update ϕ_i as:

$$\phi_i \leftarrow \phi_i - \beta_2 \nabla_{\phi} [\mathcal{L}_{\text{CE}}(\mathcal{D}_i; \theta_i^+, \omega_i) + \mu \mathcal{R}(\phi_i)], \quad (8)$$

where \mathcal{L}_{CE} denotes softmax cross-entropy loss and β_2 is the learning rate for ϕ_i . This trains ϕ_i to capture the most discriminative features of new-class images. The constraint $\mathcal{R}(\phi_i)$ is implemented as a ℓ_2 -regularization term on the generated mask by ϕ_i . The motivation for the regularization term is to make the mask coverage smaller, thus the compressed images take less memory.

We empirically observe that by the above optimization flow, the output activation maps by ϕ_i are easy to collapse, i.e., different images have the same map. To solve this issue, we add a cross-entropy loss term about ϕ_i in Eq. 8 to regularize it to produce image-specific activation maps:

$$\begin{aligned} \phi_i \leftarrow \phi_i - \beta_2 \nabla_{\phi} [\mathcal{L}_{\text{CE}}(\mathcal{D}_i; \theta_i^+, \omega_i) + \mu \mathcal{R}(\phi_i) \\ + \mu' \mathcal{L}_{\text{CE}}(\tilde{\mathcal{E}}_{0:i-1} \cup \mathcal{D}_i; \theta_i, \phi_i, \omega_i)], \end{aligned} \quad (9)$$

where μ' is the weight.

Limitations. Our CIM learns to generate adaptive masks for exemplar compression in CIL. It has three limitations that are left as future work. 1) It is not able to adjust any previous-phase exemplars, as the validation data (the original data of these exemplars) are not accessible anymore. 2) It introduces hundreds of activation parameters to the CIL model, although this is not a significant overhead compared to model parameters. Please check detailed overhead analyses in the supplementary materials. 3) Image compression is not that meaningful for low-resolution datasets (e.g., 32×32 CIFAR-100) It is because the memory taken by the compression parameters (e.g., the parameters of CIM) and the RGB pixels of a low-resolution image are comparable. Using the memory to save more images is more meaningful.

5. Experiments

We incorporate CIM into two baseline CIL methods (i.e., DER [46] and FOSTER [42]) and boost their model performances consistently on three datasets. Below, we introduce datasets and experiment settings (Section 5.1), followed by results and analyses (Section 5.2)

5.1. Experimental Settings

Datasets. We conduct experiments on three standard CIL benchmarks with high-resolution images. 1) **Food-101** [3] consists of 101 food categories with 750 training and 250 test samples per category. All images have a maximum side length of 512 pixels. 2) **ImageNet-1000** [9] is a large-scale

dataset with 1,000 classes and each class has around 1,300 training and 50 test samples. 3) **ImageNet-100** is a 100-class subset randomly sampled from ImageNet-1000 with a fixed NumPy [14] random seed (1993), following [17]. We provide other details of these datasets, e.g., image sizes and pre-processing methods, in the supplementary materials.

Protocols. We use two protocols: *learning from scratch* (LFS) and *learning from half* (LFH), following recent CIL works [42, 46]. In LFS, the model observes the same number of classes in all N phases, where N is optionally 5, 10, and 20. In LFH, the model is trained on half of the classes (e.g., 500 classes for ImageNet-1000) in the 1-th phase. Then, it learns the remaining classes evenly in the subsequent N phases, where N can be 5, 10, and 25. In both protocols, after the training of each phase we evaluate the resultant model on the test data of all seen classes. Our final report includes the average accuracy over all phases and the last-phase accuracy which indicates the degree of model forgetting. We run each experiment three times and report the average results.

Memory Budget. There are two memory budget⁵ settings.

1) In the “fixed” setting, we remove some old-class exemplars when new exemplars from the current phase are added in the memory to maintain the “fixed memory budget”. In this setting, we set the total memory to be 2,020 samples for Food-101 and 2,000 samples for ImageNet-100. For ImageNet-1000, we have two options—5,000 samples and 20,000 samples. 2) In the “growing” setting, a constant memory budget is allocated for each class across all phases and hence extra memory is appended when new classes come. In this setting, we set the budget to be 20 samples per class for all datasets. Following [42, 46], we apply the “fixed” setting in LFS experiments and the “growing” setting in LFH experiments.

Implementation Details. Our implementation is based on the standard deep learning library PyTorch [35] and image processing library OpenCV [4]. Following [42, 43, 46], we use an 18-layer ResNet [16] as the network backbone θ and a fully-connected layer as the classifier ω in all experiments. We use the same CIL training hyperparameters as in related works [17, 42, 46] for fair comparison: 1) there are 200 epochs in 1-st phase and 170 epochs in the subsequent phases; 2) the learning rate λ is initialized as 0.1 and decreases to zero with a cosine annealing scheduler [28]; 3) the SGD optimizer is deployed, with momentum factor set to 0.9 and weight decay set to 0.0005. For compression-related hyperparameters, we set the masking threshold τ as 0.6 and the downsampling ratio η as 4.0. To build the CIM model, we apply PAUs with degrees $m = 5$ and $n = 4$ as learnable activation layers. For the optimization of the CIM

⁵Please note that we measure the memory by the number of original images. Each compressed exemplar in CIM takes less memory than the original image (Eq. 4), resulting in more exemplars in the same memory.

Method	Learning from Scratch (LFS)						Learning from Half (LFH)					
	Food-101			ImageNet-100			Food-101			ImageNet-100		
	N=5	10	20	5	10	20	5	10	25	5	10	25
iCaRL [36]	69.66	62.18	56.70	73.90	67.06	62.36	60.13	53.42	46.87	62.53	59.88	52.97
WA [48]	70.94	63.69	58.45	74.64	68.62	63.20	63.55	57.60	52.48	65.75	63.71	58.34
PODNet [12]	68.03	61.24	47.38	72.14	63.96	53.69	75.37	70.01	65.32	75.54	74.33	68.33
AANets [25]	69.46	61.59	48.83	72.98	65.77	55.36	76.07	71.22	66.93	76.96	75.58	71.78
DER [46]	73.88	70.76	64.39	78.50	76.12	73.79	78.13	73.45	-	79.08	77.73	-
DER w/ ours	75.63	73.09	69.17	79.63	77.57	75.36	79.25	75.76	-	80.30	79.05	-
FOSTER [42]	75.03	72.72	66.73	79.93 [†]	76.55 [†]	74.49	79.08	75.07	68.08	80.07 [†]	77.54	72.40 [*]
FOSTER w/ ours	76.44	74.85	70.20	80.58	77.94	75.23	79.76	76.86	70.50	80.93	78.66	75.74

[†] The paper of FOSTER [42] did not report the numerical results for $N=5/10$ (LFS) and $N=5$ (LFH) on ImageNet-100. We run these experiments using the public code (released by authors) and report the reproduced results.

^{*} Our reproduced result (72.40) is significantly higher than the original result (69.34) reported in the paper of FOSTER.

Table 1. Average accuracies (%) of two top-performing CIL methods [42, 46] with and without our CIM-CIL plugged-in, and other four baselines [12, 25, 36, 48], on two datasets (Food-101 and ImageNet-100) and using two protocols (learning from scratch (LFS) and learning from half (LFH)). Due to the space limits, we report the 95% confidence intervals for these results in the supplementary materials.

Memory Budget	Method	N=5		N=10	
		Avg.	Last	Avg.	Last
$M = 20k$	iCaRL [36]	44.36	27.78	38.40	22.70
	WA [48]	58.37	50.62	54.10	45.66
	DER [46]	67.49	59.75	66.73	58.62
	FOSTER [42]	69.21	64.88	68.34	60.14
	FOSTER w/ ours	69.93	66.05	69.53	62.07
$M = 5k$	FOSTER	57.19	49.42	54.72	44.96
	FOSTER w/ ours	61.37	54.46	59.48	50.83

Table 2. Average and last accuracies (%) on ImageNet-1000 of FOSTER with and without our method plugged-in, and other three baselines [36, 46, 48]. We show two memory budgets, $M = 20k$ (upper block) and $M = 5k$ (lower block), in the LFS setting.

model ϕ (i.e., the mask-level optimization), we initially set β_1 as 0.1 and β_2 as 0.01 and reduce them to zero following the scheduler of λ . μ and μ' is set to 0.1 and 0.2, respectively. To smooth the training, we clip the gradient norm of ϕ to be no more than 1. We report the result of hyperparameter sensitivity analysis in the supplementary materials.

5.2. Results and Analyses

Comparing with the State-of-the-art. In Table 1, we summarize the experimental results on two datasets (Food-101 and ImageNet-100) and in two CIL protocols (LFS and LFH). From the table, we have the following observations. 1) Our CIM-based CIL consistently improves the state-of-the-art method FOSTER [42] with clear margins in all settings. E.g., our method surpasses it by an average of 1.4 percentage points on ImageNet-100, and 2.0 percentage points on Food-101. 2) Our CIM-based CIL achieves more significant improvements when N becomes larger, e.g., on ImageNet-100 (LFH), our method improves FOSTER by

Ablation Method	Food-101		ImageNet-100	
	N=10	20	10	20
1 Baseline	72.72	66.73	76.55	72.37
2 Artifact Aug.	71.38	66.03	75.63	71.45
3 Full Comp.	73.03	67.38	76.92	73.26
4 Random Acti.	73.10	67.54	76.88	73.54
5 Center Acti.	73.29	67.88	76.78	73.82
6 Class Acti.	73.76	68.65	77.21	74.67
7 Phase-wise τ	73.83	69.17	77.06	74.78
8 Joint Train	73.44	69.01	77.34	74.59
9 BOP (ours)	74.85	70.20	77.94	75.23
10 LastBlock Only	74.55	69.87	77.72	74.86
11 Fg Compressed	75.02	70.13	77.87	75.46

Table 3. Average accuracies (%) of different ablation methods. The experiments are conducted in the LFS setting.

0.9 and 3.3 percentage points when $N=5$ and $N=25$, respectively. 3) Our CIM-based CIL achieves greater improvements consistently on Food-101 (than ImageNet-100). It improves baselines by 2.1 percentage points on Food-101, while the improvement is 1.4 on ImageNet-100 ($N=10$, LFS). It shows that our method is particularly effective when the representative visual cues of a class are from some of its components, e.g., the “cream” of the class “cake”.

Table 2 shows the results on the large-scale dataset ImageNet-1000 in different memory settings ($M = 20k$ and $M = 5k$). We can see that our CIM-based CIL improves FOSTER consistently. It is impressive that it achieves more improvements in the more strict memory setting ($M = 5k$). Specifically, it boosts the average accuracy of FOSTER by 4.5 percentage points when $M = 5k$, significantly higher than that of $M = 20k$ (1.0).

Ablation Study. Table 3 shows the ablation results. First

Method	ImageNet-100			ImageNet-1000		
	N=6	11	26	6	11	26
LUCIR <i>baseline</i>	71.22	69.67	67.45	65.23	62.43	59.88
w/ Mnemonics	73.30	72.17	71.50	66.15	63.12	63.08
w/ MRDC	73.62	72.81	70.44	67.67	65.60	62.74
w/ ours	74.05	73.76	72.84	68.03	66.54	63.77

Table 4. Comparing with Mnemonics [27] and MRDC [43]. We plug each of them in baseline LUCIR [17] for fair comparison.

block: baselines. Row 1 is for the baseline FOSTER [42]. Row 2 shows the results of adding artifact augmentation (see Section 4.1). It shows directly apply this augmentation does not improve and even impair the model. *Please note the models in below blocks all use this augmentation.* *Second block: activation methods.* Rows 3-6 show the results of using different activation methods to compress exemplars. Row 3 is to downsample all pixels (i.e., no region is activated). Row 4 is to randomly select activation regions. Row 5 is to activate only the center region ($\frac{1}{4}$ of the original image), while row 6 is to use naive CAM. Comparing them to row 1, we can see that using naively compressed exemplars can improve CIL models. Row 4 outperforms rows 3-5, validating that it is more reliable to use the model’s activation to generate compressed exemplars. *Third block: optimization methods.* Rows 7-9 are on top of Row 6 and are the results of applying different optimization strategies. Row 7 is to manually select τ using a held-out set (10% of the dataset). Row 8 is to jointly train CIL and CIM models (for each input batch). Row 9 is the proposed method of using a global BOP. *Fourth block: two variants of CIM-based CIL.* Rows 10-11 are two variants of row 9. In Row 10, only the activation layers in the last block of CIM are learnable, and previous blocks use ReLU. Compared to row 9, row 10 shows slightly worse performance. Row 11 shows the version of adding a weak downsampling ($\eta' = 2.0$) on discriminative regions, based on which more compressed exemplars are saved. It results in comparable performance to row 9 but increases costs.

Comparing with Other Compression-based Methods. Table 4 shows our results comparing to two compression-based methods: Mnemonics [27] and MRDC [43]. We can see that our method consistently outperforms them in all settings. This is because our method does not sacrifice the discriminativeness of exemplars while improving the number (variance) of exemplars in a phase-adaptive manner. However, the two related methods either keep a fixed number of exemplars in the memory [27] or use uniform image compression without considering the properties of specific classes in different incremental phases [43].

Visualizations (CAM vs. CIM). Figure 4 gives two visualization examples, “Afghan hound” and “indigo bird”, each with the activation map as well as the bounding boxes.

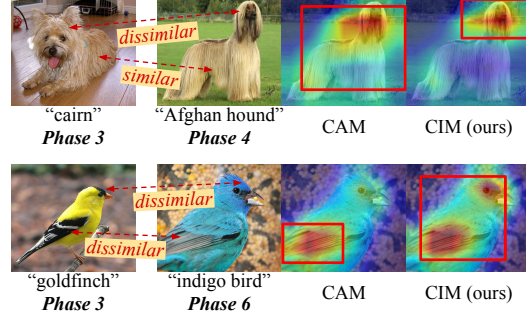


Figure 4. Visualizations (CAM vs. CIM). The experimental setting is $N=10$ (LFS) on ImageNet-100.

Metric	Small	Middle	Large
Mean of #Exemplars	39.40	38.30	34.77
Last Acc. (% , baseline)	66.13	68.40	69.93
Last Acc. (% , ours)	70.00	71.10	72.26
Improvement (%)	+3.87	+3.65	+2.33

Table 5. Results for small, middle and large objects in the setting of $N=10$ (LFS) on ImageNet-100. “Mean of #Exemplars” denotes the average number of saved exemplars by our method. The baseline (FOSTER [42]) has this number as 20 for all classes.

The first column shows their respective confusing classes appearing in earlier phases. CIM learns to focus on the discriminative (i.e., dissimilar to confusing classes) regions.

Results of Different-Size Objects. Table 5 shows the results for small, middle, and large objects. These size categorization is according to ImageNet Object Localization Challenge [19]. We calculate the bbox coverage for each class and take top 30 classes with highest coverages as “large”, rear 30 classes with lowest coverages as “small” and the rest 40 classes as “middle”. It is intriguing that our method achieves the highest improvement (over baseline) for small objects. Our explanation is that small objects benefit more from image compression (than large ones), as their images contain more background pixels to downsample.

6. Conclusions

We introduced a novel exemplar compression method for CIL, allowing us to save more representative exemplars but not increase memory budget. We achieved this compression by downsampling non-discriminative pixels of the image bounded by CAM masks. To generate adaptive masks, we proposed a novel method CIM that explicitly parameterizes a mask generation model and optimizes it in an end-to-end manner across incremental phases. Our method achieves consistent performance improvements over multiple baselines and can be taken as a flexible plug-and-play module.

Acknowledgements. The author gratefully acknowledges the support by the Lee Kong Chian (LKC) Fellowship fund awarded by Singapore Management University.

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] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. Yolov4: Optimal speed and accuracy of object detection. *arXiv*, 2020. 5
- [3] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101 – mining discriminative components with random forests. In *ECCV*, 2014. 2, 6
- [4] Gary Bradski. The opencv library. *Dr. Dobb's Journal: Software Tools for the Professional Programmer*, 2000. 6
- [5] Kenneth R Castleman. *Digital image processing*. Prentice Hall Press, 1996. 4
- [6] George Cazenavette, Tongzhou Wang, Antonio Torralba, Alexei A Efros, and Jun-Yan Zhu. Dataset distillation by matching training trajectories. In *CVPR*, 2022. 3
- [7] Can Chen, Xi Chen, Chen Ma, Zixuan Liu, and Xue Liu. Gradient-based bi-level optimization for deep learning: A survey. *arXiv*, 2022. 2, 3
- [8] Zhaozheng Chen, Tan Wang, Xiongwei Wu, Xian-Sheng Hua, Hanwang Zhang, and Qianru Sun. Class re-activation maps for weakly-supervised semantic segmentation. In *CVPR*, 2022. 3, 4
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 1, 2, 3, 6
- [10] Prithviraj Dhar, Rajat Vikram Singh, Kuan-Chuan Peng, Ziyang Wu, and Rama Chellappa. Learning without memorizing. In *CVPR*, 2019. 2
- [11] Zhang Dong, Zhang Hanwang, Tang Jinhui, Hua Xiansheng, and Sun Qianru. Causal intervention for weakly supervised semantic segmentation. In *NeurIPS*, 2020. 4
- [12] Arthur Douillard, Matthieu Cord, Charles Ollion, Thomas Robert, and Eduardo Valle. Podnet: Pooled outputs distillation for small-tasks incremental learning. In *ECCV*, 2020. 1, 2, 7
- [13] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017. 3
- [14] Charles R Harris, K Jarrod Millman, Stéfan J Van Der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J Smith, et al. Array programming with numpy. *Nature*, 2020. 6
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *ICCV*, 2015. 5
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 5, 6
- [17] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *CVPR*, 2019. 1, 2, 6, 8
- [18] Shenyang Huang, Vincent François-Lavet, and Guillaume Rabusseau. Neural architecture search for class-incremental learning. *arXiv*, 2019. 2
- [19] Imagenet object localization challenge. <https://www.kaggle.com/competitions/imagenet-object-localization-challenge/>. 8
- [20] Anil K Jain. *Fundamentals of digital image processing*. Prentice-Hall, Inc., 1989. 4
- [21] 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*, 2017. 2
- [22] Jungbeom Lee, Eunji Kim, and Sungroh Yoon. Anti-adversarially manipulated attributions for weakly and semi-supervised semantic segmentation. In *CVPR*, 2021. 3, 4
- [23] Zhizhong Li and Derek Hoiem. Learning without forgetting. *PAMI*, 2017. 2, 5
- [24] Yaoyao Liu, Yingying Li, Bernt Schiele, and Qianru Sun. Online hyperparameter optimization for class-incremental learning. In *AAAI*, 2023. 2
- [25] Yaoyao Liu, Bernt Schiele, and Qianru Sun. Adaptive aggregation networks for class-incremental learning. In *CVPR*, 2021. 1, 2, 3, 7
- [26] Yaoyao Liu, Bernt Schiele, and Qianru Sun. Rmm: Reinforced memory management for class-incremental learning. In *NeurIPS*, 2021. 2
- [27] Yaoyao Liu, Yuting Su, An-An Liu, Bernt Schiele, and Qianru Sun. Mnemonics training: Multi-class incremental learning without forgetting. In *CVPR*, 2020. 1, 2, 3, 8
- [28] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv*, 2016. 6
- [29] Dougal Maclaurin, David Duvenaud, and Ryan Adams. Gradient-based hyperparameter optimization through reversible learning. In *ICML*, 2015. 3
- [30] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of Learning and Motivation*. Elsevier, 1989. 2
- [31] Ken McRae and Phil A Hetherington. Catastrophic interference is eliminated in pretrained networks. In *Proceedings of the 15th Annual Conference of the Cognitive Science Society*, 1993. 1, 2
- [32] Alejandro Molina, Patrick Schramowski, and Kristian Kersting. Padé activation units: End-to-end learning of flexible activation functions in deep networks. In *ICLR*, 2019. 5
- [33] Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *ICML*, 2010. 5
- [34] J Anthony Parker, Robert V Kenyon, and Donald E Troxel. Comparison of interpolating methods for image resampling. *IEEE Transactions on Medical Imaging*, 1983. 4
- [35] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, 2019. 6
- [36] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. iCaRL: Incremental classifier and representation learning. In *CVPR*, 2017. 1, 2, 3, 5, 7

- [37] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv*, 2016. [2](#)
- [38] 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. [3](#)
- [39] Christian Simon, Piotr Koniusz, and Mehrtash Harandi. On learning the geodesic path for incremental learning. In *CVPR*, 2021. [2](#)
- [40] Ankur Sinha, Pekka Malo, and Kalyanmoy Deb. A review on bilevel optimization: from classical to evolutionary approaches and applications. *IEEE Transactions on Evolutionary Computation*, 2017. [1](#), [2](#), [3](#)
- [41] Gregory K Wallace. The jpeg still picture compression standard. *Communications of the ACM*, 1991. [1](#), [2](#)
- [42] Fu-Yun Wang, Da-Wei Zhou, Han-Jia Ye, and De-Chuan Zhan. Foster: Feature boosting and compression for class-incremental learning. *arXiv*, 2022. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [43] 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. [1](#), [2](#), [3](#), [6](#), [8](#)
- [44] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *CVPR*, 2019. [1](#), [2](#)
- [45] Ju Xu and Zhanxing Zhu. Reinforced continual learning. In *NeurIPS*, 2018. [2](#)
- [46] Shipeng Yan, Jiangwei Xie, and Xuming He. Der: Dynamically expandable representation for class incremental learning. In *CVPR*, 2021. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [47] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *ICML*, 2017. [2](#)
- [48] Bowen Zhao, Xi Xiao, Guojun Gan, Bin Zhang, and Shu-Tao Xia. Maintaining discrimination and fairness in class incremental learning. In *CVPR*, 2020. [1](#), [7](#)
- [49] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *CVPR*, 2016. [1](#), [2](#), [3](#)
- [50] Da-Wei Zhou, Han-Jia Ye, and De-Chuan Zhan. Co-transport for class-incremental learning. In *Proceedings of the 29th ACM International Conference on Multimedia*, 2021. [5](#)