
SparCL: Sparse Continual Learning on the Edge

Anonymous Author(s)

Affiliation

Address

email

Abstract

Existing work in continual learning (CL) focuses on mitigating catastrophic forgetting, *i.e.*, model performance deterioration on past tasks when learning a new task. However, the training efficiency of a CL system is under-investigated, which limits the real-world application of CL systems under resource-limited scenarios. In this work, we propose a novel framework called Sparse Continual Learning (SparCL), which is the first study that leverages sparsity to enable cost-effective continual learning on edge devices. SparCL achieves both training acceleration and accuracy preservation through the synergy of three aspects: *weight sparsity*, *data efficiency*, and *gradient sparsity*. Specifically, we propose task-aware dynamic masking (TDM) to learn a sparse network throughout the entire CL process, dynamic data removal (DDR) to remove less informative training data, and dynamic gradient masking (DGM) to sparsify the gradient updates. Each of them not only improves efficiency, but also further mitigates catastrophic forgetting. SparCL consistently improves the training efficiency of existing state-of-the-art (SOTA) CL methods by at most $23\times$ less training FLOPs, and, surprisingly, further improves the SOTA accuracy by at most 1.7%. SparCL also outperforms competitive baselines obtained from adapting SOTA sparse training methods to the CL setting in both efficiency and accuracy. We also evaluate the effectiveness of SparCL on a real mobile phone, further indicating the practical potential of our method. Source code will be released.

1 Introduction

The objective of Continual Learning (CL) is to enable an intelligent system to accumulate knowledge from a sequence of tasks, such that it exhibits satisfying performance on both old and new tasks [28]. Recent methods mostly focus on addressing the *catastrophic forgetting* [39] problem – learning model tends to suffer performance deterioration on previously seen tasks. However, in the real world, when the CL applications are deployed in resource-limited platforms [44] such as edge devices, the learning efficiency, w.r.t. both training speed and memory footprint, are also crucial metrics of interest, yet they are rarely explored in prior CL works.

Existing CL methods can be categorized into regularization-based [2, 28, 33, 62], rehearsal-based [8, 12, 46, 56], and architecture-based [27, 38, 48, 53, 54]. Both regularization- and rehearsal-based methods directly train a dense model, which might even be over-parametrized for the union of all tasks [18, 35]; Though several architecture-based methods [47, 53, 58] start with a sparse sub-network from the dense model, they still grow the model size progressively to learn emerging tasks. The aforementioned methods, although striving for greater performance with less forgetting, still introduce significant memory and computation overhead during the whole CL process.

Recently, another stream of work, sparse training [4, 19, 31] has emerged as a new training trend to achieve training acceleration, which embraces the promising training-on-the-edge paradigm. With sparse training, each iteration takes less time with the reduction in computation achieved by sparsity,




Figure 1: **Left:** Overview of SparCL. SparCL consists of three complementary components: task-aware dynamic masking (TDM) for weight sparsity, dynamic data removal (DDR) for data efficiency, and dynamic gradient masking (DGM) for gradient sparsity. **Right:** SparCL successfully preserves the accuracy and significantly improves efficiency over DER++ [8], one of the SOTA CL methods, with different sparsity ratios on the Split Tiny-ImageNet [16] dataset.

under the traditional i.i.d. learning setting. Inspired by these sparse training methods, we naturally think about introducing sparse training to the field of CL. A straightforward idea is to directly combine existing sparse training methods, such as SNIP [31], RigL [19], with a rehearsal buffer under the CL setting. However, these methods fail to consider key challenges in CL to mitigate catastrophic forgetting, for example, properly handling transition between tasks. As a result, these sparse training methods, though enhancing training efficiency, cause significant accuracy drop (see Section 5.2). Thus, we would like to explore a general strategy, which is orthogonal to existing CL methods, that not only leverages the idea of sparse training for efficiency, but also addresses key challenges in CL to preserve (or even improve) accuracy.

In this work, we propose *Sparse Continual Learning* (SparCL), a general framework for cost-effective continual learning, aiming at enabling practical CL on edge devices. As shown in Figure 1 (left), SparCL achieves both learning acceleration and accuracy preservation through the synergy of three aspects: *weight sparsity*, *data efficiency*, and *gradient sparsity*. Specifically, to maintain a small dynamic sparse network during the whole CL process, we develop a novel task-aware dynamic masking (TDM) strategy to keep only important weights for both the current and past tasks, with special consideration during task transitions. Moreover, we propose a dynamic data removal (DDR) scheme, which progressively removes “easy-to-learn” examples from training iterations, which further accelerates the training process and also improves accuracy of CL by balancing current and past data and keeping more informative samples in the buffer. Finally, we provide an additional dynamic gradient masking (DGM) strategy to leverage gradient sparsity for even better efficiency and knowledge preservation of learned tasks, such that only a subset of sparse weights are updated. Figure 1 (right) demonstrates that SparCL successfully preserves the accuracy and significantly improves efficiency over DER++ [8], one of the SOTA CL methods, under different sparsity ratios.

SparCL is simple in concept, compatible with various existing rehearsal-based CL methods, and efficient under practical scenarios. We conduct comprehensive experiments on multiple CL benchmarks to evaluate the effectiveness of our method. We show that SparCL works collaboratively with existing CL methods, greatly accelerates the learning process under different sparsity ratios, and even sometimes improves upon the state-of-the-art accuracy. We also establish competing baselines by combining representative sparse training methods with advanced rehearsal-based CL methods. SparCL again outperforms these baselines in terms of both efficiency and accuracy. Most importantly, we evaluate our SparCL framework on real edge devices to demonstrate the practical potential of our method. We are not aware of any prior CL works that explored this area and considered the constraints of limited resources during training.

In summary, our work makes the following contributions:

- We propose *Sparse Continual Learning* (SparCL), a general framework for cost-effective continual learning, which achieves learning acceleration through the synergy of *weight sparsity*, *data efficiency*, and *gradient sparsity*. To the best of our knowledge, our work is the first to introduce the idea of sparse training to enable efficient CL on edge devices.
- SparCL shows superior performance compared to both conventional CL methods and CL-adapted sparse training methods on all benchmark datasets, leading to at most $23\times$ less training FLOPs and, surprisingly, 1.7% improvement over SOTA accuracy.

- We evaluate SparCL on a real mobile edge device, demonstrating the practical potential of our method and also encouraging future research on CL on-the-edge. The results indicate that our framework can achieve at most $3.1\times$ training acceleration.

2 Related work

2.1 Continual Learning

The main focus in continual learning (CL) has been mitigating catastrophic forgetting. Existing methods can be classified into three major categories. *Regularization-based* methods [2, 28, 33, 62] limit updates of important parameters for the prior tasks by adding corresponding regularization terms. While these methods reduce catastrophic forgetting to some extent, their performance deteriorates under challenging settings [36], and on more complex benchmarks [46, 56]. *Rehearsal-based* methods [13, 14, 22] save examples from previous tasks into a small-sized buffer to train the model jointly with the current task. Though simple in concept, the idea of rehearsal is very effective in practice and has been adopted by many state-of-the-art methods [8, 11, 45]. *Architecture-based* methods [38, 47, 53, 54, 57] isolate existing model parameters or assign additional parameters for each task to reduce interference among tasks. As mentioned in Section 1, most of these methods use a dense model without consideration of efficiency and memory footprint, thus are not applicable to resource-limited settings. Our work, orthogonal to these methods, serves as a general framework for making these existing methods efficient and enabling a broader deployment, *e.g.*, CL on edge devices.

A limited number of works explore sparsity in CL, however, for different purposes. Several methods [37, 38, 49, 53] incorporate the idea of weight pruning [21] to allocate a sparse sub-network for each task to reduce inter-task interference. Nevertheless, these methods still reduce the full model sparsity progressively for every task and finally end up with a much denser model. On the contrary, SparCL maintains a sparse network throughout the whole CL process, introducing great efficiency and memory benefits both during training and at the output model. A recent work [15] aims at discovering lottery tickets [20] under CL, but still does not address efficiency. However, the existence of lottery tickets in CL serves as a strong justification for the outstanding performance of our SparCL.

2.2 Sparse Training

There are two main approaches for sparse training: fixed-mask sparse training and dynamic sparse training. Fixed-mask sparse training methods [31, 50, 52, 55] first apply pruning, then execute traditional training on the sparse model with the obtained fixed mask. The pre-fixed structure limits the accuracy performance, and the first stage still causes huge computation and memory consumption. To overcome these drawbacks, dynamic mask methods [4, 17, 19, 41, 42] adjust the sparsity topology during training while maintaining low memory footprint. These methods start with a sparse model structure from an untrained dense model, then combine sparse topology exploration at the given sparsity ratio with the sparse model training. Recent work [61] further considers to incorporate data efficiency into sparse training for better training accelerations. However, all prior sparse training works are focused on the traditional training setting, while CL is a more complicated and difficult scenario with inherent characteristics not explored by these works. In contrast to prior sparse training methods, our work explores a new learning paradigm that introduces sparse training into CL for efficiency and also addresses key challenges in CL, mitigating catastrophic forgetting.

3 Continual Learning Problem Setup

In supervised CL, a model f_θ learns from a sequence of tasks $\mathcal{D} = \{\mathcal{D}_1, \dots, \mathcal{D}_T\}$, where each task $\mathcal{D}_t = \{(x_{t,i}, y_{t,i})\}_{i=1}^{n_t}$ consists of input-label pairs, and each task has a disjoint set of classes. Tasks arrive sequentially, and the model must adapt to them. At the t -th step, the model gains access to data from the t -th task. However, a small fix-sized rehearsal buffer \mathcal{M} is allowed to save data from prior tasks. At test time, the easiest setting is to assume task identity is known for each coming test example, named task-incremental learning (Task-IL). If this assumption does not hold, we have the more difficult class-incremental learning (Class-IL) setting. In this work, we mainly focus on the more challenging Class-IL setting, and only report Task-IL performance for reference.

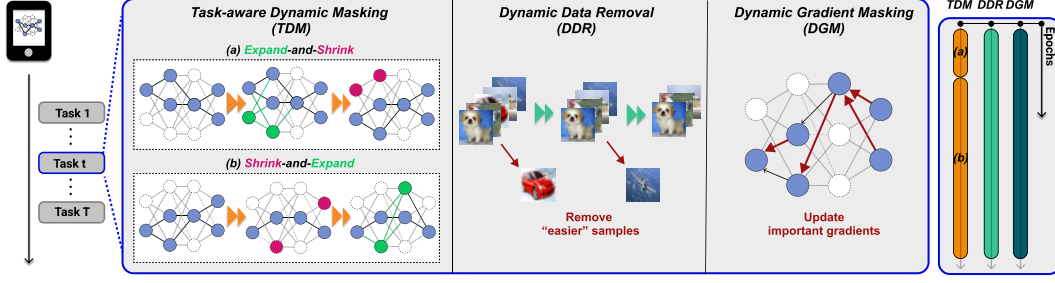


Figure 2: Illustration of the SparCL workflow. Three components work synergistically to improve training efficiency and further mitigate catastrophic forgetting for preserving accuracy.

129 The goal of conventional CL is to train a model sequentially that performs well on all tasks at test
 130 time. The main evaluation metric is average test accuracy on all tasks. In real-world resource-
 131 limited scenarios, we should further consider *training efficiency* of the model. Thus, we measure the
 132 performance of the model more comprehensively by including training FLOPs and memory footprint.

133 4 Sparse Continual Learning (SparCL)

134 Our method, Sparse Continual Learning, is a unified framework composed of three complementary
 135 components: *task-aware dynamic masking* for weight sparsity, *dynamic data removal* for data
 136 efficiency, and *dynamic gradient masking* for gradient sparsity. The entire framework is shown in
 137 Figure 2. We will illustrate each component in detail in this section.

138 4.1 Task-aware Dynamic Masking

139 To enable cost-effective CL in resource limited scenarios, SparCL is designed to maintain a dynamic
 140 structure when learning a sequence of tasks, such that it not only achieves high efficiency, but also
 141 intelligently adapts to the data stream for better performance. Specifically, we propose a strategy
 142 named *task-aware dynamic masking* (TDM), which dynamically removes less important weights
 143 and grows back unused weights for stronger representation power periodically **by maintaining a**
 144 **single binary weight mask throughout the CL process**. Different from typical sparse training work,
 145 which only leverages the weight magnitude [41] or the gradient w.r.t. data from a single training
 146 task [19, 61], TDM considers also the importance of weights w.r.t. data saved in the rehearsal buffer,
 147 as well as the switch between CL tasks.

148 Specifically, TDM strategy starts from a randomly initialized binary mask $M_\theta = M_0$, with a given
 149 sparsity constraint $\|M_\theta\|_0/\|\theta\|_0 = 1 - s$, where $s \in [0, 1]$ is the sparsity ratio. Moreover, it makes
 150 different intra- and inter-task adjustments to keep a dynamic sparse set of weights based on their
 151 continual weight importance (CWI). We summarize the process of task-aware dynamic masking in
 152 Algorithm 1 and elaborate its key components in detail below.

153 **Continual weight importance (CWI).** For a model f_θ parameterized by θ , the CWI of weight $w \in \theta$
 154 is defined as follows:

$$\text{CWI}(w) = \|w\|_1 + \alpha \left\| \frac{\partial \tilde{\mathcal{L}}(\mathcal{D}_t; \theta)}{\partial w} \right\|_1 + \beta \left\| \frac{\partial \mathcal{L}(\mathcal{M}; \theta)}{\partial w} \right\|_1, \quad (1)$$

155 where \mathcal{D}_t denotes the training data from the t -th task, \mathcal{M} is the current rehearsal buffer, and α, β are
 156 coefficients to control the influence of current and buffered data, respectively. Moreover, \mathcal{L} represents
 157 the cross-entropy loss for classification, while $\tilde{\mathcal{L}}$ is the *single-head* [1] version of the cross-entropy
 158 loss, which only considers classes from the current task **by masking out the logits of other classes**.

159 Intuitively, CWI ensures we keep (1) weights of larger magnitude for output stability, (2) weights
 160 important for the current task for learning capacity, and (3) weights important for past data to mitigate
 161 catastrophic forgetting. Moreover, inspired by the classification bias in CL [1], we use the single-head
 162 cross-entropy loss when calculating importance score w.r.t. the current task to make the importance
 163 estimation more accurate.

Algorithm 1: Task-aware Dynamic Masking (TDM)

Input: Model weight θ , number of tasks T , training epochs of the t -th task K_t , binary sparse mask M_θ , sparsity ratio s , intra-task adjustment ratio p_{intra} , inter-task adjustment ratio p_{inter} , update interval δk
Initialize: θ, M_θ , s.t. $\|M_\theta\|_0/\|\theta\|_0 = 1 - s$
for $t = 1, \dots, T$ **do**
 for $e = 1, \dots, K_T$ **do**
 if $t > 1$ **then**
 / Inter-task adjustment */*
 Expand M_θ by randomly adding unused weights,
 s.t. $\|M_\theta\|_0/\|\theta\|_0 = 1 - (s - p_{\text{inter}})$
 if $e = \delta k$ **then**
 Shrink M_θ by removing the least important weights according to equation 1,
 s.t. $\|M_\theta\|_0/\|\theta\|_0 = 1 - s$
 end
 end
 if $e \bmod \delta k = 0$ **then**
 / Intra-task adjustment */*
 Shrink M_θ by removing the least important weights according to equation 1,
 s.t. $\|M_\theta\|_0/\|\theta\|_0 = 1 - (s + p_{\text{intra}})$
 Expand M_θ by randomly adding unused weights,
 s.t. $\|M_\theta\|_0/\|\theta\|_0 = 1 - s$
 end
 Update $\theta \odot M_\theta$ via backpropagation
 end
end

164 **Intra-task adjustment.** When training the t -th task, a natural assumption is that the data distribution
165 is consistent inside the task, thus we would like to update the sparse model in a relatively stable way
166 while keeping its flexibility. Thus, in Algorithm 1, we choose to update the sparsity mask M_θ in a
167 *shrink-and-expand* way every δk epochs. We first remove p_{intra} of the weights of least CWI to retain
168 learned knowledge so far. Then we randomly select unused weights to recover the learning capacity
169 for the model and keep the sparsity ratio s unchanged.

170 **Inter-task adjustment.** When tasks switches, on the contrary, we assume data distribution shifts
171 immediately. Ideally, we would like the model to keep the knowledge learned from old tasks as much
172 as possible, and to have enough learning capacity to accommodate the new task. Thus, instead of
173 the shrink-and-expand strategy for intra-task adjustment, we follow an *expand-and-shrink* scheme.
174 Specifically, at the beginning of the $(t + 1)$ -th task, we expand the sparse model by randomly adding
175 a proportion of p_{inter} unused weights. Intuitively, the additional learning capacity *facilitates* fast
176 adoption of new knowledge and reduces interference with learned knowledge. We allow our model to
177 have smaller sparsity (*i.e.*, *larger learning capacity*) temporarily for the first δk epochs as a warm-
178 up period, and then remove the p_{inter} weights with least CWI, following the same process in the
179 intra-task case, to satisfy the sparsity constraint.

180 4.2 Dynamic Data Removal

181 In addition to weight sparsity, decreasing the amount of training data can be directly translated into
182 the saving of training time without any requirements for hardware support. Thus, we would also like
183 to explore data efficiency to reduce the training workload. Some prior CL works select informative
184 examples to construct the rehearsal buffer [3, 6, 59]. However, the main purpose of them is not
185 training acceleration, thus they either introduce excessive computational cost or consider different
186 problem settings. By considering the features of CL, we present a simple yet effective strategy,
187 *dynamic data removal* (DDR), to reduce training data for further acceleration.

188 We measure the importance of each training example by the occurrence of misclassification [51, 61]
189 during CL. In TDM, the sparse structure of our model updates periodically every δk epochs, so we
190 align our data removal process with the update of weight mask for further efficiency and training

191 stability. In Section 4.1, we have partitioned the training process for the t -th task into $N_t = K_t/\delta k$
 192 stages based on the dynamic mask update. Therefore, we gradually remove training data at the end of
 193 i -th stage, based on the following policy: 1) Calculate the total number of misclassifications $f_i(x_j)$
 194 for each training example during the i -th stage. 2) Remove a proportion of ρ_i training samples with
 195 the least number of misclassifications. Although our main purpose is to keep the “harder” examples
 196 to learn to consolidate the sparse model, we can get further benefits for better CL result. First, the
 197 removal of “easier” examples increases the probability that “harder” examples to be saved to the
 198 rehearsal buffer, given the common strategy, *e.g.* reservoir sampling [14], to buffer examples. Thus,
 199 we construct a more informative buffer in a implicit way without heavy computation. Moreover, since
 200 the buffer size is much smaller than the training set size of each task, the data from the buffer and the
 201 new task is highly imbalanced, dynamic data removal also relieves the data imbalance issue.

202 Formally, we set the data removal proportion for each task as $\rho \in [0, 1]$, and a cutoff stage, such that:

$$\sum_{i=1}^{\text{cutoff}} \rho_i = \rho, \quad \sum_{i=\text{cutoff}+1}^{N_k} \rho_i = 0 \quad (2)$$

203 The cutoff stage controls the trade-off between efficiency and accuracy: When we set the cutoff stage
 204 earlier, we reduce the training time for all the following stages; however, when the cutoff stage is set
 205 too early, the model might underfit the removed training data. Note that when we set $\rho_i = 0$ for all
 206 $i = 1, 2, \dots, N_t$ and $\text{cutoff} = N_t$, we simply recover the vanilla setting without any data efficiency
 207 considerations. In our experiments, we assume $\rho_i = \rho/\text{cutoff}$, *i.e.*, removing equal proportion of
 208 data at the end of every stage, for simplicity. We also conduct comprehensive exploration study for ρ
 209 and the selection of the cutoff stage in Section 5.3 and Appendix D.3.

210 4.3 Dynamic Gradient Masking

211 With TDM and DDR, we can already achieve bi-level efficiency during training. To further boost
 212 training efficiency, we explore sparsity in gradient and propose dynamic gradient masking (DGM)
 213 for CL. Our method focuses on reducing computational cost by only applying the most important
 214 gradients onto the corresponding unpruned model parameters via a gradient mask. The gradient mask
 215 is also dynamically updated along with the weight mask defined in Section 4.1. Intuitively, while
 216 targeting for better training efficiency, DGM also promotes the preservation of past knowledge by
 217 preventing a fraction of weights from update.

218 Formally, our goal here is to find a subset of unpruned parameters (or, equivalently, a gradient mask
 219 M_G) to update over multiple training iterations. For a model f_θ parameterized by θ , we have the
 220 corresponding gradient matrix G calculated during each iteration. To prevent the pruned weights
 221 from updating, the weight mask M_θ will be applied onto the gradient matrix G as $G \odot M_\theta$ during
 222 backpropagation. Besides the gradients of pruned weights, we in addition consider to remove less
 223 important gradients for faster training. To achieve this, we introduce the continual gradient importance
 224 (CGI) based on the CWI to measure the importance of weight gradients.

$$\text{CGI}(w) = \alpha \left\| \frac{\partial \tilde{\mathcal{L}}(\mathcal{D}_t; \theta)}{\partial w} \right\|_1 + \beta \left\| \frac{\partial \mathcal{L}(\mathcal{M}; \theta)}{\partial w} \right\|_1. \quad (3)$$

225 We remove a proportion q of non-zero gradients from G with less importance measured by CGI
 226 and we have $\|M_G\|_0/\|\theta\|_0 = 1 - (s + q)$. The gradient mask M_G is then applied onto the gradient
 227 matrix G . During the entire training process, the gradient mask M_G is updated with a fixed interval.

228 5 Experiment

229 5.1 Experiment Setting

230 **Datasets.** We evaluate our SparCL on two representative CL benchmarks, Split CIFAR-10 [29]
 231 and Split Tiny-ImageNet [16] to verify the efficacy of SparCL. In particular, we follow [8, 62] by
 232 splitting CIFAR-10 and Tiny-ImageNet into 5 and 10 tasks, each of which consists of 2 and 20 classes
 233 respectively. Dataset licensing information can be found in Appendix C.

234 **Comparing methods.** In particular, we select CL methods of all kinds including regularization-
 235 based (EWC [28], LwF [33]), architecture-based (PackNet [38], LPS [53]), and rehearsal-based

Table 1: Comparison with CL methods. SparCL consistently improves training efficiency of the corresponding CL methods while preserves (or even improves) accuracy on both class- and task-incremental settings.

Method	Sparsity	Buffer size	Split CIFAR-10			Split Tiny-ImageNet		
			Class-IL (\uparrow)	Task-IL (\uparrow)	FLOPs Train $\times 10^{15}$ (\downarrow)	Class-IL (\uparrow)	Task-IL (\uparrow)	FLOPs Train $\times 10^{16}$ (\downarrow)
EWC [28]	0.00	–	19.49 \pm 0.12	68.29 \pm 3.92	8.3	7.58 \pm 0.10	19.20 \pm 0.31	13.3
LwF [33]			19.61 \pm 0.05	63.29 \pm 2.35	8.3	8.46 \pm 0.22	15.85 \pm 0.58	13.3
PackNet [38]	0.50 [†]	–	–	93.73 \pm 0.55	5.0	–	61.88 \pm 1.01	7.3
LPS [53]			–	94.50 \pm 0.47	5.0	–	63.37 \pm 0.83	7.3
A-GEM [13]	0.00	200	20.04 \pm 0.34	83.88 \pm 1.49	11.1	8.07 \pm 0.08	22.77 \pm 0.03	17.8
iCaRL [46]			49.02 \pm 3.20	88.99 \pm 2.13	11.1	7.53 \pm 0.79	28.19 \pm 1.47	17.8
FDR [5]			30.91 \pm 2.74	91.01 \pm 0.68	13.9	8.70 \pm 0.19	40.36 \pm 0.68	22.2
ER [14]			44.79 \pm 1.86	91.19 \pm 0.94	11.1	8.49 \pm 0.16	38.17 \pm 2.00	17.8
DER++ [8]			64.88 \pm 1.17	91.92 \pm 0.60	13.9	10.96 \pm 1.17	40.87 \pm 1.16	22.2
SparCL-ER ₇₅	0.75	200	46.89 \pm 0.68	92.02 \pm 0.72	2.0	8.98 \pm 0.38	39.14 \pm 0.85	3.2
SparCL- DER++ ₇₅			66.30 \pm 0.98	94.06 \pm 0.45	2.5	12.73 \pm 0.40	42.06 \pm 0.73	4.0
SparCL-ER ₉₀	0.90		45.81 \pm 1.05	91.49 \pm 0.47	0.9	8.67 \pm 0.41	38.79 \pm 0.39	1.4
SparCL- DER++ ₉₀			65.79 \pm 1.33	93.73 \pm 0.24	1.1	12.27 \pm 1.06	41.17 \pm 1.31	1.8
SparCL-ER ₉₅	0.95		44.59 \pm 0.23	91.07 \pm 0.64	0.5	8.43 \pm 0.09	38.20 \pm 0.46	0.8
SparCL- DER++ ₉₅			65.18 \pm 1.25	92.97 \pm 0.37	0.6	10.76 \pm 0.62	40.54 \pm 0.98	1.0
A-GEM [13]	0.00	500	22.67 \pm 0.57	89.48 \pm 1.45	11.1	8.06 \pm 0.04	25.33 \pm 0.49	17.8
iCaRL [46]			47.55 \pm 3.95	88.22 \pm 2.62	11.1	9.38 \pm 1.53	31.55 \pm 3.27	17.8
FDR [5]			28.71 \pm 3.23	93.29 \pm 0.59	13.9	10.54 \pm 0.21	49.88 \pm 0.71	22.2
ER [14]			57.74 \pm 0.27	93.61 \pm 0.27	11.1	9.99 \pm 0.29	48.64 \pm 0.46	17.8
DER++ [8]			72.70 \pm 1.36	93.88 \pm 0.50	13.9	19.38 \pm 1.41	51.91 \pm 0.68	22.2
SparCL-ER ₇₅	0.75	500	60.80 \pm 0.22	93.82 \pm 0.32	2.0	10.48 \pm 0.29	50.83 \pm 0.69	3.2
SparCL- DER++ ₇₅			74.09 \pm 0.84	95.19 \pm 0.34	2.5	20.75 \pm 0.88	52.19 \pm 0.43	4.0
SparCL-ER ₉₀	0.90		59.34 \pm 0.97	93.33 \pm 0.10	0.9	10.12 \pm 0.53	49.46 \pm 1.22	1.4
SparCL- DER++ ₉₀			73.42 \pm 0.95	94.82 \pm 0.23	1.1	19.62 \pm 0.67	51.93 \pm 0.36	1.8
SparCL-ER ₉₅	0.95		57.75 \pm 0.45	92.73 \pm 0.34	0.5	9.91 \pm 0.17	48.57 \pm 0.50	0.8
SparCL- DER++ ₉₅			72.14 \pm 0.78	94.39 \pm 0.15	0.6	19.01 \pm 1.32	51.26 \pm 0.78	1.0

[†]PackNet and LPS actually have a decreased sparsity after learning every task, we use 0.50 to roughly represent the average sparsity.

(A-GEM [13], iCaRL [40], FDR [5], ER [14], **DER++** [8]) methods. Note that PackNet and LPS are only compatible with task-incremental learning. We also adapt representative sparse training methods (SNIP [31], RigL [19]) to the CL setting by combining them with **DER++** (SNIP-**DER++**, RigL-**DER++**).

Variants of our method. To show the generality of SparCL, we combine it with **DER++** (one of the SOTA CL methods), and ER (simple and widely-used) as *SparCL-**DER++*** and *SparCL-ER*, respectively. We also vary the weight sparsity ratio (0.75, 0.90, 0.95) of SparCL for a comprehensive evaluation.

Evaluation metrics. We use the average accuracy on all tasks to evaluate the performance of the final model. Moreover, we evaluate the training FLOPs [19], and memory footprint [61] (including feature map pixels and model parameters during training) to demonstrate the efficiency of each method. Please see Appendix D.1 for detailed definitions of these metrics.

Experiment details. For fair comparison, we strictly follow the settings in prior CL work [8, 26]. We sets the per task training epochs to 50 and 100 for Split CIFAR-10 and Tiny-ImageNet, respectively, with a batch size of 32. For the model architecture, We follow [8, 46] and adopt the ResNet-18 [23] without any pre-training. We also use the best hyperparameter setting reported in [8, 53] for CL methods, and in [19, 31] for CL-adapted sparse training methods. For SparCL and its competing CL-adapted sparse training methods, we adopt a uniform sparsity ratio for all convolutional layers. Please see Appendix D for other details.

5.2 Main Results

Comparison with CL methods. Table 1 summarizes the results on Split CIFAR-10 and Tiny-ImageNet, under both class-incremental (Class-IL) and task-incremental (Task-IL) settings. From Table 1, we can clearly tell that SparCL significantly improves ER and **DER++**, while also outperforms other CL baselines, in terms of training efficiency (measured in FLOPs). With higher sparsity ratio, SparCL leads to less training FLOPs. Notably, SparCL achieves 23 \times training efficiency improvement upon **DER++** with a sparsity ratio of 0.95. On the other hand, our framework also improves the average accuracy of ER and **DER++** consistently under all cases with a sparsity ratio of 0.75 and 0.90,

Table 2: Comparison with CL-adapted sparse training methods. All methods are combined with DER++ with a 500 buffer size. SparCL outperforms all methods in both accuracy and training efficiency, under all sparsity ratios. All three methods here can save 20% \sim 51% memory footprint, please see Appendix D.2 for details.

Method	Sparsity	Split CIFAR-10		Split Tiny-ImageNet	
		Class-IL (\uparrow)	FLOPs Train $\times 10^{15}$ (\downarrow)	Class-IL (\uparrow)	FLOPs Train $\times 10^{16}$ (\downarrow)
DER++ [8]	0.00	72.70 \pm 1.36	13.9	19.38 \pm 1.41	22.2
SNIP-DER++ [31]	0.90	69.82 \pm 0.72	1.6	16.13 \pm 0.61	2.5
RigL-DER++ [19]		69.86 \pm 0.59	1.6	18.36 \pm 0.49	2.5
SparCL-DER++ ₉₀		73.42 \pm 0.95	1.1	19.62 \pm 0.67	1.8
SNIP-DER++ [31]	0.95	66.07 \pm 0.91	0.9	14.76 \pm 0.52	1.5
RigL-DER++ [19]		66.53 \pm 1.13	0.9	15.88 \pm 0.63	1.5
SparCL-DER++ ₉₅		72.14 \pm 0.78	0.6	19.01 \pm 1.32	1.0

Table 3: Ablation study on Split-CIFAR10 with 0.75 sparsity ratio. All components contributes to the overall performance, in terms of both accuracy and efficiency (training FLOPs and memory footprint).

TDM	DDR	DGM	Class-IL (\uparrow)	FLOPs Train $\times 10^{15}$ (\downarrow)	Memory Footprint (\downarrow)
\times	\times	\times	72.70	13.9	169MB
\checkmark	\times	\times	73.37	3.6	101MB
\checkmark	\checkmark	\times	73.80	2.8	101MB
\checkmark	\times	\checkmark	73.97	3.3	98MB
\checkmark	\checkmark	\checkmark	74.09	2.5	98MB

and only slight performance drop when sparsity gets larger as 0.95. In particular, SparCL-DER++ with 0.75 sparsity ratio sets new SOTA accuracy, with all buffer sizes under both benchmarks. The outstanding performance of SparCL indicates that our proposed strategies successfully preserve accuracy by further mitigating catastrophic forgetting with a much sparser model. Moreover, the improvement that SparCL brings to two different existing CL methods shows the generalizability of SparCL as a unified framework, *i.e.*, it has the potential to be combined with a wide array of existing methods.

We would also like to take a closer look at PackNet and LPS, which also leverage the idea of sparsity to split the model by different tasks, a different motivation from training efficiency. Firstly, they are only compatible with the Task-IL setting, since they leverage task identity at both training and test time. Moreover, the model sparsity of these methods reduces with the increasing number of tasks, which still leads to much larger overall training FLOPs than that of SparCL. This further demonstrates the importance of keeping a sparse model without permanent expansion throughout the CL process.

Comparison with CL-adapted sparse training methods. Table 2 shows the result under the more difficult Class-IL setting. SparCL outperforms all CL-adapted sparse training methods in both accuracy and training FLOPs. The performance gap between SparCL-DER++ and other methods gets larger with a higher sparsity. SNIP- and RigL-DER++ achieve training acceleration at the cost of compromised accuracy, which suggests that keeping accuracy is a non-trivial challenge for existing sparse training methods under the CL setting. SNIP generates the static initial mask after network initialization which does not consider the structure suitability among tasks. Though RigL adopts a dynamic mask, the lack of task-aware strategy prevents it from generalizing well to the CL setting.

5.3 Effectiveness of Key Components

Ablation study. We provide a comprehensive ablation study in Table 3 using SparCL-DER++ with 0.75 sparsity on Split CIFAR10. Table 3 demonstrates that all components of our method contribute to both efficiency and accuracy improvements. Comparing row 1 and 2, we can see that the majority of FLOPs decrease results from TDM. Interestingly, TDM leads to an increase in accuracy, indicating TDM generates a sparse model that is even more suitable for learning all tasks than then full dense model. Comparing row 2 and 3, we can see that DDR indeed further accelerates training by removing less informative examples. As discussed in Section 4.2, when we remove a certain number of samples (30% here), we achieve a point where we keep as much informative samples as we need, and also balance the current and buffered data. Comparing row 2 and 4, DGM reduce both training FLOPs and memory footprint while improve the performance of the network. Finally, the last row demonstrates the collaborative performance of all components. We also show the same ablation study with 0.90 sparsity in Appendix D.4 for reference. Detail can be found in Appendix D.1.




Figure 3: Comparison between DDR and One-shot [61] data removal strategy w.r.t. different data removal proportion ρ . DDR outperforms One-shot and also achieves better accuracy when $\rho \leq 30\%$.

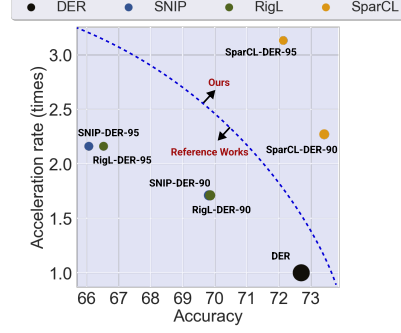


Figure 4: Comparison with CL-adapted sparse training methods in training acceleration rate and accuracy results. The radius of circles are measured by memory footprint.

Exploration on DDR. To understand the influence of the data removal proportion ρ , and the cutoff stage for each task, we show corresponding experiment results in Figure 3 and Appendix D.3, respectively. In Figure 3, we fix $\text{cutoff} = 4$, i.e., gradually removing equal number of examples every 5 epochs until epoch 20, and vary ρ from 10% to 90%. We also compare DDR with One-shot removal strategy [61], which removes all examples at once at cutoff. DDR outperforms One-shot consistently with different ρ in average accuracy. Also note that since DDR removes the examples gradually before the cutoff stage, DDR is more efficient than One-shot. When $\rho \leq 30\%$, we also observe increased accuracy of DDR compared with the baseline without removing any data. When $\rho \geq 40\%$, the accuracy gets increasingly lower for both strategies. The intuition is that when DDR removes a proper amount of data, it removes redundant information while keeps the most informative examples. Moreover, as discussed in Section 4.2, it balances the current and buffered data, while also leave informative samples in the buffer. When DDR removes too much data, it will also lose informative examples, thus the model has not learned these examples well before removal.

Exploration on DGM. We test the efficacy of DGM at different sparsity levels. Detailed exploratory experiments are shown in Appendix D.5 for reference. The results indicate that by setting the proportion q within an appropriate range, DGM can consistently improve the accuracy performance regardless of the change of weight sparsity.

5.4 Mobile Device Results

The training acceleration results are measured on the CPU of an off-the-shelf Samsung Galaxy S20 smartphone, which has the Qualcomm Snapdragon 865 mobile platform with a Qualcomm Kryo 585 Octa-core CPU. We run each test on a batch of 32 images to denote the training speed. The detail of on-mobile compiler-level optimizations for training acceleration can be found in Appendix E.1.

The acceleration results are shown in Figure 4. SparCL can achieve approximately $3.1\times$ and $2.3\times$ training acceleration with 0.95 sparsity and 0.90 sparsity, respectively. Besides, our framework can also save 51% and 48% memory footprint when the sparsity is 0.95 and 0.90. Furthermore, the obtained sparse models save the storage consumption by using compressed sparse row (CSR) storage and can be accelerated to speed up the inference on-the-edge. We provide on-mobile inference acceleration results in Appendix E.2.

6 Conclusion

This paper presents a unified framework named SparCL for efficient CL that achieves both learning acceleration and accuracy preservation. It comprises three complementary strategies: task-aware dynamic masking for weight sparsity, dynamic data removal for data efficiency, and dynamic gradient masking for gradient sparsity. Extensive experiments on standard CL benchmarks and real-world edge device evaluations demonstrate that our method significantly improves upon existing CL methods and outperforms CL-adapted sparse training methods. We discuss the limitations and potential negative social impacts of our method in Appendix A and B, respectively.

References

- [1] Hongjoon Ahn, Jihwan Kwak, Subin Lim, Hyeonsu Bang, Hyojun Kim, and Taesup Moon. Ss-il: Separated softmax for incremental learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 844–853, 2021. 4
- [2] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *ECCV*, 2018. 1, 3
- [3] Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient based sample selection for online continual learning. *Advances in neural information processing systems*, 32, 2019. 5, 20
- [4] Guillaume Bellec, David Kappel, Wolfgang Maass, and Robert Legenstein. Deep rewiring: Training very sparse deep networks. In *International Conference on Learning Representations*, 2018. 1, 3
- [5] Ari S Benjamin, David Rolnick, and Konrad Kording. Measuring and regularizing networks in function space. *arXiv preprint arXiv:1805.08289*, 2018. 7
- [6] Zalán Borsos, Mojmir Mutny, and Andreas Krause. Coresets via bilevel optimization for continual learning and streaming. *Advances in Neural Information Processing Systems*, 33:14879–14890, 2020. 5
- [7] Miles Brundage, Shahar Avin, Jack Clark, Helen Toner, Peter Eckersley, Ben Garfinkel, Allan Dafoe, Paul Scharre, Thomas Zeitzoff, Bobby Filar, et al. The malicious use of artificial intelligence: Forecasting, prevention, and mitigation. *arXiv preprint arXiv:1802.07228*, 2018. 15
- [8] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: a strong, simple baseline. In *NeurIPS*, 2020. 1, 2, 3, 6, 7, 8
- [9] Pietro Buzzega, Matteo Boschini, Angelo Porrello, and Simone Calderara. Rethinking experience replay: a bag of tricks for continual learning. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 2180–2187. IEEE, 2021. 20
- [10] Alfredo Canziani, Adam Paszke, and Eugenio Culurciello. An analysis of deep neural network models for practical applications. *arXiv preprint arXiv:1605.07678*, 2016. 16
- [11] Hyuntak Cha, Jaeho Lee, and Jinwoo Shin. Co2l: Contrastive continual learning. In *ICCV*, 2021. 3, 19
- [12] Arslan Chaudhry, Albert Gordo, Puneet Kumar Dokania, Philip Torr, and David Lopez-Paz. Using hindsight to anchor past knowledge in continual learning. *arXiv preprint arXiv:2002.08165*, 2(7), 2020. 1
- [13] Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a-gem. *arXiv preprint arXiv:1812.00420*, 2018. 3, 7
- [14] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and Marc’Aurelio Ranzato. On tiny episodic memories in continual learning. *arXiv preprint arXiv:1902.10486*, 2019. 3, 6, 7
- [15] Tianlong Chen, Zhenyu Zhang, Sijia Liu, Shiyu Chang, and Zhangyang Wang. Long live the lottery: The existence of winning tickets in lifelong learning. In *International Conference on Learning Representations*, 2020. 3
- [16] 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. Ieee, 2009. 2, 6
- [17] Tim Dettmers and Luke Zettlemoyer. Sparse networks from scratch: Faster training without losing performance. *arXiv preprint arXiv:1907.04840*, 2019. 3

- [18] Xuanyi Dong and Yi Yang. Network pruning via transformable architecture search. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 759–770, 2019. 1
- [19] Utku Evci, Trevor Gale, Jacob Menick, Pablo Samuel Castro, and Erich Elsen. Rigging the lottery: Making all tickets winners. In *International Conference on Machine Learning*, pages 2943–2952. PMLR, 2020. 1, 2, 3, 4, 7, 8, 16
- [20] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *International Conference on Learning Representations*, 2019. 3
- [21] Song Han, Jeff Pool, et al. Learning both weights and connections for efficient neural network. In *NeurIPS*, pages 1135–1143, 2015. 3
- [22] Tyler L Hayes, Nathan D Cahill, and Christopher Kanan. Memory efficient experience replay for streaming learning. In *ICRA*, 2019. 3
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 7
- [24] Yang He, Guoliang Kang, Xuanyi Dong, Yanwei Fu, and Yi Yang. Soft filter pruning for accelerating deep convolutional neural networks. *arXiv preprint arXiv:1808.06866*, 2018. 17
- [25] Yang He, Ping Liu, Ziwei Wang, Zhilan Hu, and Yi Yang. Filter pruning via geometric median for deep convolutional neural networks acceleration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4340–4349, 2019. 17
- [26] Yen-Chang Hsu, Yen-Cheng Liu, Anita Ramasamy, and Zsolt Kira. Re-evaluating continual learning scenarios: A categorization and case for strong baselines. *arXiv preprint arXiv:1810.12488*, 2018. 7
- [27] Zixuan Ke, Bing Liu, and Xingchang Huang. Continual learning of a mixed sequence of similar and dissimilar tasks. *NeurIPS*, 33, 2020. 1
- [28] 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. 1, 3, 6, 7, 19
- [29] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 6, 15
- [30] Yann LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998. 15
- [31] Namhoon Lee, Thalaiyasingam Ajanthan, and Philip HS Torr. Snip: Single-shot network pruning based on connection sensitivity. *arXiv preprint arXiv:1810.02340*, 2018. 1, 2, 3, 7, 8
- [32] Tuanhui Li, Baoyuan Wu, Yujiu Yang, Yanbo Fan, Yong Zhang, and Wei Liu. Compressing convolutional neural networks via factorized convolutional filters. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3977–3986, 2019. 17
- [33] Zhizhong Li and Derek Hoiem. Learning without forgetting. *TPAMI*, 40(12):2935–2947, 2017. 1, 3, 6, 7
- [34] Xiaolong Ma, Fu-Ming Guo, Wei Niu, Xue Lin, Jian Tang, Kaisheng Ma, Bin Ren, and Yanzhi Wang. Pconv: The missing but desirable sparsity in dnn weight pruning for real-time execution on mobile devices. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 5117–5124, 2020. 17
- [35] Xiaolong Ma, Wei Niu, Tianyun Zhang, Sijia Liu, Sheng Lin, Hongjia Li, Wujie Wen, Xiang Chen, Jian Tang, Kaisheng Ma, et al. An image enhancing pattern-based sparsity for real-time inference on mobile devices. In *European Conference on Computer Vision*, pages 629–645. Springer, 2020. 1

- [36] Zheda Mai, Ruiwen Li, Jihwan Jeong, David Quispe, Hyunwoo Kim, and Scott Sanner. Online continual learning in image classification: An empirical survey. *arXiv preprint arXiv:2101.10423*, 2021. 3
- [37] Arun Mallya, Dillon Davis, and Svetlana Lazebnik. Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In *ECCV*, pages 67–82, 2018. 3
- [38] Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *CVPR*, 2018. 1, 3, 6, 7
- [39] 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. 1
- [40] Sanket Vaibhav Mehta, Darshan Patil, Sarath Chandar, and Emma Strubell. An empirical investigation of the role of pre-training in lifelong learning. *ICML Workshop*, 2021. 7
- [41] Decebal Constantin Mocanu, Elena Mocanu, Peter Stone, Phuong H Nguyen, Madeleine Gibescu, and Antonio Liotta. Scalable training of artificial neural networks with adaptive sparse connectivity inspired by network science. *Nature communications*, 9(1):1–12, 2018. 3, 4
- [42] Hesham Mostafa and Xin Wang. Parameter efficient training of deep convolutional neural networks by dynamic sparse reparameterization. In *International Conference on Machine Learning*, pages 4646–4655. PMLR, 2019. 3
- [43] Wei Niu, Xiaolong Ma, Sheng Lin, Shihao Wang, Xuehai Qian, Xue Lin, Yanzhi Wang, and Bin Ren. Patdnn: Achieving real-time dnn execution on mobile devices with pattern-based weight pruning. *arXiv preprint arXiv:2001.00138*, 2020. 18, 20
- [44] Lorenzo Pellegrini, Vincenzo Lomonaco, Gabriele Graffieti, and Davide Maltoni. Continual learning at the edge: Real-time training on smartphone devices. *arXiv preprint arXiv:2105.13127*, 2021. 1
- [45] Quang Pham, Chenghao Liu, and Steven Hoi. Dualnet: Continual learning, fast and slow. *NeurIPS*, 34, 2021. 3
- [46] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *CVPR*, pages 2001–2010, 2017. 1, 3, 7
- [47] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016. 1, 3
- [48] Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. In *ICML*, pages 4548–4557, 2018. 1
- [49] Ghada Sokar, Decebal Constantin Mocanu, and Mykola Pechenizkiy. Spacenet: Make free space for continual learning. *Neurocomputing*, 439:1–11, 2021. 3
- [50] Hidenori Tanaka, Daniel Kunin, Daniel L Yamins, and Surya Ganguli. Pruning neural networks without any data by iteratively conserving synaptic flow. *Advances in Neural Information Processing Systems*, 33:6377–6389, 2020. 3
- [51] Mariya Toneva, Alessandro Sordoni, Remi Tachet des Combes, Adam Trischler, Yoshua Bengio, and Geoffrey J Gordon. An empirical study of example forgetting during deep neural network learning. *arXiv preprint arXiv:1812.05159*, 2018. 5
- [52] Chaoqi Wang, Guodong Zhang, and Roger Grosse. Picking winning tickets before training by preserving gradient flow. In *International Conference on Learning Representations*, 2019. 3
- [53] Zifeng Wang, Tong Jian, Kaushik Chowdhury, Yanzhi Wang, Jennifer Dy, and Stratis Ioannidis. Learn-prune-share for lifelong learning. In *ICDM*, 2020. 1, 3, 6, 7

- [54] Zifeng Wang, Zizhao Zhang, Chen-Yu Lee, Han Zhang, Ruoxi Sun, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, and Tomas Pfister. Learning to prompt for continual learning. *CVPR*, 2022. 1, 3
- [55] Paul Wimmer, Jens Mehnert, and Alexandru Condurache. Freezenet: Full performance by reduced storage costs. In *Proceedings of the Asian Conference on Computer Vision*, 2020. 3
- [56] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *CVPR*, pages 374–382, 2019. 1, 3
- [57] Shipeng Yan, Jiangwei Xie, and Xuming He. Der: Dynamically expandable representation for class incremental learning. In *CVPR*, pages 3014–3023, 2021. 3
- [58] Li Yang, Sen Lin, Junshan Zhang, and Deliang Fan. Grown: Grow only when necessary for continual learning. *arXiv preprint arXiv:2110.00908*, 2021. 1
- [59] Jaehong Yoon, Divyam Madaan, Eunho Yang, and Sung Ju Hwang. Online coreset selection for rehearsal-based continual learning. *arXiv preprint arXiv:2106.01085*, 2021. 5
- [60] Kun-Hsing Yu, Andrew L Beam, and Isaac S Kohane. Artificial intelligence in healthcare. *Nature biomedical engineering*, 2(10):719–731, 2018. 15
- [61] Geng Yuan, Xiaolong Ma, Wei Niu, Zhengang Li, Zhenglun Kong, Ning Liu, Yifan Gong, Zheng Zhan, Chaoyang He, Qing Jin, et al. Mest: Accurate and fast memory-economic sparse training framework on the edge. *Advances in Neural Information Processing Systems*, 34, 2021. 3, 4, 5, 7, 9, 16, 17
- [62] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *ICML*, 2017. 1, 3, 6

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes] The claims match the experimental results and it is expected to generalize according to the diverse experiments stated in our paper. We include all of our code, data, and models in the supplementary materials, which can reproduce our experimental results.
 - (b) Did you describe the limitations of your work? [Yes] See Section 6 and Appendix A.
 - (c) Did you discuss any potential negative societal impacts of your work? [Yes] See Section 6 and Appendix B.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes] We have read the ethics review guidelines and ensured that our paper conforms to them.
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [N/A] Our paper is based on the experimental results and we do not have any theoretical results.
 - (b) Did you include complete proofs of all theoretical results? [N/A] Our paper is based on the experimental results and we do not have any theoretical results.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See Section 5.1, Section 5.4 and we provide code to reproduce the main experimental results.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section 5.1 and Section 5.4.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Table 1, Table 2, fig 1, fig 3.

- 518 (d) Did you include the total amount of compute and the type of resources used (e.g., type
519 of GPUs, internal cluster, or cloud provider)? [Yes] See Section 5.1, Section 5.4.
- 520 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 521 (a) If your work uses existing assets, did you cite the creators? [Yes] We mentioned and
522 cited the datasets (Split CIFAR-10 and Tiny-ImageNet), and all comparing methods
523 with their paper and github in it.
- 524 (b) Did you mention the license of the assets? [Yes] The licences of used datasets/models
525 are provided in the cited references and we state them explicitly in Appendix C.
- 526 (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
527 We provide code for our proposed method in the supplement.
- 528 (d) Did you discuss whether and how consent was obtained from people whose data you're
529 using/curating? [N/A]
- 530 (e) Did you discuss whether the data you are using/curating contains personally identifiable
531 information or offensive content? [N/A]
- 532 5. If you used crowdsourcing or conducted research with human subjects...
- 533 (a) Did you include the full text of instructions given to participants and screenshots, if
534 applicable? [N/A]
- 535 (b) Did you describe any potential participant risks, with links to Institutional Review
536 Board (IRB) approvals, if applicable? [N/A]
- 537 (c) Did you include the estimated hourly wage paid to participants and the total amount
538 spent on participant compensation? [N/A]