
Continual Learning with Tiny Episodic Memories

Arslan Chaudhry¹ Marcus Rohrbach² Mohamed Elhoseiny³ Thalaiyasingam Ajanthan⁴ Puneet K. Dokania¹
Philip H. S. Torr¹ Marc’Aurelio Ranzato²

Abstract

In continual learning (CL), an agent learns from a stream of tasks leveraging prior experience to transfer knowledge to future tasks. It is an ideal framework to decrease the amount of supervision in the existing learning algorithms. But for a successful knowledge transfer, the learner needs to remember how to perform previous tasks. One way to endow the learner the ability to perform tasks seen in the past is to store a small memory, dubbed episodic memory, that stores few examples from previous tasks and then to replay these examples when training for future tasks. In this work, we empirically analyze the effectiveness of a very small episodic memory in a CL setup where each training example is only seen once. Surprisingly, across four rather different supervised learning benchmarks adapted to CL, a very simple baseline, that jointly trains on both examples from the current task as well as examples stored in the episodic memory, significantly outperforms specifically designed CL approaches with and without episodic memory. Interestingly, we find that repetitive training on even tiny memories of past tasks does not harm generalization, on the contrary, it improves it, with gains between 7% and 17% when the memory is populated with a single example per class. Code is made available to reproduce the results.

1. Introduction

The objective of continual learning (CL) is to rapidly learn new skills from a sequence of tasks leveraging the knowledge accumulated in the past. Catastrophic forgetting (McCloskey & Cohen, 1989), i.e. the inability of a model to

recall how to perform tasks seen in the past, makes such efficient adaptation extremely difficult.

This decades old problem of CL (Ring, 1997; Thrun, 1998) is now seeing a surge of interest in the research community with several methods proposed to tackle catastrophic forgetting (Rebuffi et al., 2017; Kirkpatrick et al., 2016; Zenke et al., 2017; Lee et al., 2017a; Aljundi et al., 2018; Lopez-Paz & Ranzato, 2017; Lee et al., 2017b; Chaudhry et al., 2019). In this work, we quantitatively study some of these methods (that assume a fixed network architecture) on four benchmark datasets under the following assumptions: i) each task is fully supervised, ii) each example from a task can only be seen once using the learning protocol proposed by Chaudhry et al. (2019) (see §3), and iii) the model has access to a small memory storing examples of past tasks. Restricting the size of such episodic memory is important because it makes the continual learning problem more realistic and distinct from multi-task learning where complete datasets of all the tasks are available at each step.

We empirically observe that a very simple baseline, dubbed Experience Replay (ER)¹, that jointly trains on both the examples from the current task and examples stored in the very small episodic memory not only gives superior performance over the existing state-of-the-art approaches specifically designed for CL (with and without episodic memory), but it also is computationally very efficient. We verify this finding on four rather different supervised learning benchmarks adapted for CL; Permuted MNIST, Split CIFAR, Split mini-ImageNet and Split CUB. Importantly, repetitive training on the same examples of a tiny episodic memory does not harm generalization on past tasks. In §5.5, we analyze this phenomenon and provide insights as to why directly training on the episodic memory does not have a detrimental effect in terms of generalization. Briefly, we observe that the training on the datasets of subsequent tasks acts like a data-dependent regularizer on past tasks allowing the repetitive training on tiny memory to generalize beyond the episodic memory. We further observe that methods, that do not train directly on the memory, such as GEM (Lopez-Paz & Ran-

¹University of Oxford ²Facebook AI Research ³King Abdullah University of Science and Technology ⁴Australian National University. Correspondence to: Arslan Chaudhry <arslan.chaudhry@eng.ox.ac.uk>.

¹For consistency to prior work in the literature, we will refer to this approach which trains on the episodic memory as ER, although its usage for supervised learning tasks is far less established.

zato, 2017) and A-GEM (Chaudhry et al., 2019), underfit the training data and end up not fully utilizing the beneficial effects of this implicit and data dependent regularization.

Overall, ER with tiny episodic memories offers very strong performance at a very small additional computational cost over the fine-tuning baseline. We believe that this approach will serve as a stronger baseline for the development of future CL approaches.

2. Related Work

Regularization-based CL approaches These works attempt to reduce forgetting by regularizing the objective such that it either penalizes the feature drift on already learned tasks (Li & Hoiem, 2016; Rebuffi et al., 2017) or discourages change in parameters that were important to solve past tasks (Kirkpatrick et al., 2016; Zenke et al., 2017; Chaudhry et al., 2018; Aljundi et al., 2018). The former approach relies on the storage of network activations and subsequent deployment of knowledge distillation (Hinton et al., 2014), whereas the latter approach stores a measure of parameter importance whose best case memory complexity is the same as the total number of network parameters.

Memory-Based CL approaches These approaches (Lopez-Paz & Ranzato, 2017; Riemer et al., 2019; Chaudhry et al., 2019) use episodic memory that stores a subset of data from past tasks to tackle forgetting. One approach to leverage such episodic memory is to use it to constrain the optimization such that the loss on past tasks can never increase (Lopez-Paz & Ranzato, 2017).

Experience Replay (ER) The use of ER is well established in reinforcement learning (RL) tasks (Mnih et al., 2013; 2015; Foerster et al., 2017; Rolnick et al., 2018). Isele & Cosgun (2018), for instance, explore different ways to populate a relatively large episodic memory for a continual RL setting where the learner does multiple passes over the data. In this work instead, we study supervised learning tasks with a single pass through data and a very small episodic memory. More recently, (Hayes et al., 2018; Riemer et al., 2019) used ER for supervised CL learning tasks. Hayes et al. (2018), independently, study different replay strategies in ER and show improvements over the fine-tune baseline. Our contribution is to show the improvements brought by ER, perhaps surprisingly, over the specifically designed CL approaches. We differ from (Riemer et al., 2019) in considering episodic memories of much smaller sizes. Finally, and most importantly, we extend these previous studies by analyzing *why* repetitive training on tiny memories does not lead to overfitting (§5.5).

3. Learning Framework

3.1. Protocol for Single-Pass Through the Data

We use the learning protocol proposed by Chaudhry et al. (2019). There are two streams of tasks, described by the following ordered sequences of datasets, one stream for Cross-Validation $\mathcal{D}^{CV} = \{\mathcal{D}_{-T^{CV}}, \dots, \mathcal{D}_{-1}\}$ consisting of T^{CV} tasks, and one for Evaluation $\mathcal{D}^{EV} = \{\mathcal{D}_1, \dots, \mathcal{D}_T\}$ consisting of T tasks, where $\mathcal{D}_k = \{(\mathbf{x}_i^k, t_i^k, \mathbf{y}_i^k)_{i=1}^{n_k}\}$ is the dataset of the k -th task. The sequence \mathcal{D}^{CV} contains only a handful of tasks and it is only used for cross-validation purposes. Tasks from this sequence can be replayed as many times as needed and have various degree of similarity to tasks in the training and evaluation dataset, \mathcal{D}^{EV} . The latter stream, \mathcal{D}^{EV} , instead can be played only once; the learner will observe examples in *sequence* and will be tested throughout the learning experience. The final performance will be reported on the held-out test set drawn from \mathcal{D}^{EV} .

The k -th task in any of these streams consists of $\mathcal{D}_k = \{(\mathbf{x}_i^k, t_i^k, \mathbf{y}_i^k)_{i=1}^{n_k}\}$, where each triplet constitutes an example defined by an input ($\mathbf{x}^k \in \mathcal{X}$), a task descriptor ($t^k \in \mathcal{T}$) which is an integer id in this work, and a target vector ($\mathbf{y}^k \in \mathcal{Y}^k$), where \mathcal{Y}^k is the set of labels specific to task k and $\mathcal{Y}^k \subset \mathcal{Y}$.

3.2. Metrics

We measure performance on \mathcal{D}^{EV} using two metrics, as standard practice in the literature (Lopez-Paz & Ranzato, 2017; Chaudhry et al., 2018):

Average Accuracy ($A \in [0, 1]$) Let $a_{i,j}$ be the performance of the model on the held-out test set of task ‘ j ’ after the model is trained on task ‘ i ’. The average accuracy at task T is then defined as:

$$A_T = \frac{1}{T} \sum_{j=1}^T a_{T,j} \quad (1)$$

Forgetting ($F \in [-1, 1]$) Let f_j^i be the forgetting on task ‘ j ’ after the model is trained on task ‘ i ’ which is computed as:

$$f_j^i = \max_{l \in \{1, \dots, i-1\}} a_{l,j} - a_{i,j} \quad (2)$$

The average forgetting measure at task T is then defined as:

$$F_T = \frac{1}{T-1} \sum_{j=1}^{T-1} f_j^T \quad (3)$$

4. Experience Replay

Recent works (Lopez-Paz & Ranzato, 2017; Chaudhry et al., 2019) have shown that methods relying on episodic memory have superior performance than regularization based

approaches (e.g., (Kirkpatrick et al., 2016; Zenke et al., 2017)) when using a “single-pass through the data” protocol (§3.1). While GEM (Lopez-Paz & Ranzato, 2017) and its more efficient version A-GEM (Chaudhry et al., 2019) used the episodic memory as a mean to project gradients, here we drastically simplify the optimization problem and, similar to Riemer et al. (2019) and Hayes et al. (2018), directly train on the the examples stored in a very small memory, resulting in better performance and more efficient learning.

The overall training procedure is given in Alg. 1. Compared to the simplest baseline model that merely fine-tunes the parameters on the new task starting from the previous task parameter vector, ER makes two modifications. First, it has an episodic memory which is updated at every time step, line 8. Second, it doubles the size of the minibatch used to compute the gradient descent parameter update by stacking the actual minibatch of examples from the current task with a minibatch of examples taken at random from the memory, line 7. As we shall see in our empirical validation, these two simple modifications yield much better generalization and substantially limit forgetting, while incurring in a negligible additional computational cost on modern GPU devices. Next, we explain the difference between the direct (ER) and indirect (A-GEM) training on episodic memory from the optimization perspective.

A-GEM vs ER: Let us assume that B_n is a mini-batch of size K from the current task t and $B_{\mathcal{M}}$ is the same size mini-batch from a very small episodic memory \mathcal{M} . Furthermore, following the notation from (Chaudhry et al., 2019), let g be the gradient computed with mini-batch B_n and g_{ref} be the gradient computed with $B_{\mathcal{M}}$. In A-GEM, if $g^T g_{ref} \geq 0$, then the current task gradient g is directly used for optimization whereas if $g^T g_{ref} < 0$, g is projected such that $g^T g_{ref} = 0$. Refer to Eq. 11 in (Chaudhry et al., 2019) for the exact form of projection. In ER instead, since both mini-batches are used in the optimization step, the *average* of g and g_{ref} is used. It may seem a bit counter-intuitive that, even though ER repetitively trains on \mathcal{M} , it is still able to generalize to previous tasks beyond the episodic memory. We investigate this question in §5.5.

Since we study the usage of tiny episodic memories, the sample that the learner selects to populate the memory becomes crucial, see line 8 of the algorithm. For this, we describe various strategies to write into the memory. All these strategies assume access to a *continuous* stream of data and a small episodic memory, which rules out approaches relying on the temporary storage of all the examples seen so far. This restriction is consistent with our definition of CL: a learning experience through a stream of data under the constraint of a fixed and small sized memory and limited compute budget.

Algorithm 1 Experience Replay for Continual Learning.

```

1: procedure ER( $\mathcal{D}$ , mem_sz, batch_sz, lr)
2:    $\mathcal{M} \leftarrow \{\} * \text{mem\_sz}$   $\triangleright$  Allocate memory buffer of size mem_sz
3:    $n \leftarrow 0$   $\triangleright$  Number of training examples seen in the continuum
4:   for  $t \in \{1, \dots, T\}$  do
5:     for  $B_n \overset{K}{\sim} \mathcal{D}_t$  do  $\triangleright$  Sample without replacement a mini-batch of
        size  $K$  from task  $t$ 
6:        $B_{\mathcal{M}} \overset{K}{\sim} \mathcal{M}$   $\triangleright$  Sample a mini-batch from  $\mathcal{M}$ 
7:        $\theta \leftarrow \text{SGD}(B_n \cup B_{\mathcal{M}}, \theta, \text{lr})$   $\triangleright$  Single gradient step
        to update the parameters by stacking current minibatch with minibatch from memory
8:        $\mathcal{M} \leftarrow \text{UpdateMemory}(\text{mem\_sz}, t, n, B_n)$ 
         $\triangleright$  Memory update, see §4
9:        $n \leftarrow n + \text{batch\_sz}$   $\triangleright$  Counter update
10:  return  $\theta, \mathcal{M}$ 

```

Reservoir Sampling: Similarly to Riemer et al. (2019), Reservoir sampling (Vitter, 1985) takes as input a stream of data of unknown length and returns a random subset of items from that stream. If ‘ n ’ is the number of points observed so far and ‘mem_sz’ is the size of the reservoir (sampling buffer), this selection strategy samples each data point with a probability $\frac{\text{mem_sz}}{n}$. The routine to update the memory is given in Appendix Alg. 2.

Ring Buffer: Similarly to Lopez-Paz & Ranzato (2017), for each task, the ring buffer strategy allocates as many equally sized FIFO buffers as there are classes. If C is the total number of classes across all tasks, and mem_sz is the total size of episodic memory, each stack has a buffer of size $\frac{\text{mem_sz}}{C}$. As shown in Appendix Alg. 3, the memory stores the last few observations from each class. Unlike reservoir sampling, samples from older tasks do not change throughout training, leading to potentially stronger overfitting. Also, at early stages of training the memory is not fully utilized since each stack has a constant size throughout training. However, this simple sampling strategy guarantees equal representation of all classes in the memory, which is particularly important when the memory is tiny.

k-Means: For each class, we use online k-Means to estimate the k centroids in feature space, using the representation before the last classification layer. We then store in the memory the input examples whose feature representation is the closest to such centroids, see Appendix Alg. 4. This memory writing strategy has similar benefits and drawbacks of ring buffer, except that it has potentially better coverage of the feature space in L2 sense.

Mean of Features (MoF): Similarly to Rebuffi et al. (2017), for each class we compute a running estimate of the average feature vector just before the classification layer and

store examples whose feature representation is closest to the average feature vector (see details in Appendix Alg. 5). This writing strategy has the same balancing guarantees of ring buffer and k-means, but it populates the memory differently. Instead of populating the memory at random or using k-Means, it puts examples that are closest to the mode in feature space.

5. Experiments

In this section, we review the benchmark datasets used in our evaluation, as well as the architectures and the baselines we compared against. We then report the results we obtained using episodic memory and experience replay (ER). Finally, we conclude with a brief analysis investigating generalization when using ER on tiny memories. Code to reproduce the results is available at <https://github.com/facebookresearch/agem>.

5.1. Datasets

We consider four commonly used benchmarks in CL literature. **Permuted MNIST** (Kirkpatrick et al., 2016) is a variant of MNIST (LeCun, 1998) dataset of handwritten digits where each task has a certain random permutation of the input pixels which is applied to all the images of that task. Our Permuted MNIST benchmark consists of a total of 23 tasks.

Split CIFAR (Zenke et al., 2017) consists of splitting the original CIFAR-100 dataset (Krizhevsky & Hinton, 2009) into 20 disjoint subsets, each of which is considered as a separate task. Each task has 5 classes that are randomly sampled *without* replacement from the total of 100 classes.

Similarly to Split CIFAR, **Split miniImageNet** is constructed by splitting miniImageNet (Vinyals et al., 2016), a subset of ImageNet with a total of 100 classes and 600 images per class, to 20 disjoint subsets.

Finally, **Split CUB** (Chaudhry et al., 2019) is an incremental version of the fine-grained image classification dataset CUB (Wah et al., 2011) of 200 bird categories split into 20 disjoint subsets of classes.

In all cases, \mathcal{D}^{CV} consists of 3 tasks while \mathcal{D}^{EV} contains the remaining tasks. As described in § 3.2, we report metrics on \mathcal{D}^{EV} after doing a single training pass over each task in the sequence. The hyper-parameters selected via cross-validation on \mathcal{D}^{CV} are reported in Appendix Tab. 8.

5.2. Architectures

For MNIST, we use a fully-connected network with two hidden layers of 256 ReLU units each. For CIFAR and miniImageNet, a reduced ResNet18, similar to Lopez-Paz & Ranzato (2017), is used and a standard ResNet18 with

ImageNet pretraining is used for CUB. The input integer task id is used to select a task specific classifier head, and the network is trained via cross-entropy loss.

For a given dataset stream, all baselines use the same architecture, and all baselines are optimized via stochastic gradient descent with a mini-batch size equal to 10. The size of the mini-batch sampled from the episodic memory is also set to 10 irrespective of the size of the episodic buffer.

5.3. Baselines

We compare against the following baselines:

- **FINETUNE**, a model trained continually without any regularization and episodic memory, with parameters of a new task initialized from the parameters of the previous task.
- **EWC** (Kirkpatrick et al., 2016), a regularization-based approach that avoids catastrophic forgetting by limiting the learning of parameters critical to the performance of past tasks, as measured by the Fisher information matrix (FIM). In particular, we compute the FIM as a moving average similar to EWC++ in (Chaudhry et al., 2018) and online EWC in (Schwarz et al., 2018).
- **A-GEM** (Chaudhry et al., 2019), a model that uses episodic memory as an optimization constraint to avoid catastrophic forgetting. Since GEM (Lopez-Paz & Ranzato, 2017) and A-GEM have similar performance, we only consider the latter in our experiments due to its computational efficiency.
- **MER** (Riemer et al., 2019), a model that also leverages an episodic memory and uses a loss that approximates the dot products of the gradients of current and previous tasks to avoid forgetting. To make the experimental setting more comparable (in terms of SGD updates) to the other methods, we set the number of inner gradient steps to 1 for each outer Reptile (Nichol & Schulman, 2018) meta-update with the mini-batch size of 10.

5.4. Results

In the first experiment, we measured average accuracy at the end of the learning experience on \mathcal{D}^{EV} as a function of the size of the memory (detailed numerical results are provided in Appendix Tabs 3,4,5,6). From the results in Fig. 1, we can make several observations.

First, methods using *ER greatly outperform* not only the baseline approaches that do not have episodic memory (FINETUNE and EWC) but also state-of-the-art approaches relying on episodic memory of the same size (A-GEM and MER). Moreover, the ER variants outperform even when the episodic memory is very small. For instance, on CIFAR the

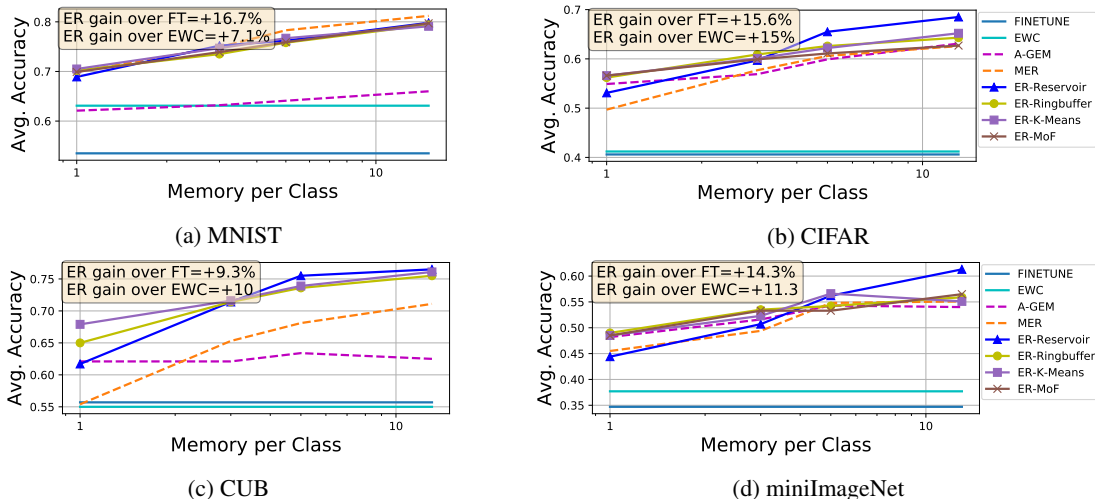


Figure 1: Average accuracy as a function of episodic memory size. The box shows the gain in average accuracy of ER-RINGBUFFER over FINETUNE and EWC baselines when only 1 sample per class is used. The performance is averaged over 5 runs. Uncertainty estimates are provided in Appendix Tabs 3,4,5,6.




Figure 2: Evolution of average accuracy (A_k) as new tasks are learned in Split CIFAR. The memory has only 85 slots (in average 1 slot per class). The vertical bar marks where the hybrid approach switches from reservoir to ring buffer strategy. The hybrid approach works better than both reservoir (once more tasks arrive) and ring buffer (initially, when the memory is otherwise not well utilized). The orange curve is a variant of ring buffer that utilizes the full memory at all times, by reducing the ring buffer size of observed classes as new classes arrive. Overall, the proposed hybrid approach works at least as good as the other approaches throughout the whole learning experience. (Averaged over 3 runs).

gain over A-GEM brought by ER is 1.7% when the memory only stores 1 example per class, and more than 5% when the memory stores 13 examples per class. This finding might seem quite surprising as repetitive training on a very small episodic memory may potentially lead to overfitting on the examples stored in the memory. We will investigate this finding in more depth in §5.5. In the same setting, the gain compared to methods that do not use memory (FINETUNE and EWC) is 15% and about 28% when using a single example per class and 13 examples per class, respectively.

Second and not surprisingly, average accuracy increases with the memory size, and does not saturate at 13 examples per class which is our self-imposed limit.

Third, experience replay based on reservoir sampling works the best across the board except when the memory size is

very small (less than 3 examples per class). Empirically we observed that as more and more tasks arrive and the size of the memory per class shrinks, reservoir sampling often ends up evicting some of the earlier classes from the memory, thereby inducing higher forgetting.

Fourth, when the memory is tiny, sampling methods that by construction guarantee a balanced number of samples per class, work the best (even better than reservoir sampling). All methods that have this property, ring buffer, k-Means and Mean of Features, have a rather similar performance which is substantially better than the reservoir sampling. For instance, on CIFAR, with one example per class in the memory, ER with reservoir sampling is 3.5% worse than ER K-Means, but ER K-Means, ER Ring Buffer and ER MoF are all within 0.5% from each other (see Appendix Tab. 4 for numerical values). These findings are further confirmed by looking at the evolution of the average accuracy (Appendix Fig. 5 left) as new tasks arrive when the memory can store at most one example per class.

The better performance of strategies like ring buffer for tiny episodic memories, and reservoir sampling for bigger episodic memories, suggests a *hybrid* approach, whereby the writing strategy relies on reservoir sampling till some classes have too few samples stored in the memory. At that point, the writing strategy switches to the ring buffer scheme which guarantees a minimum number of examples for each class. For instance, in the experiment of Fig. 2 the memory budget consists of only 85 memory slots, an average of 1 sample per class by the end of the learning experience (as there are 17 tasks and 5 classes per task). The learner switches from reservoir sampling to ring buffer once it observes that any of the classes seen in the past has only one sample left in the

Methods	Forgetting			
	MNIST	CIFAR	CUB	miniImageNet
FINETUNE	0.29	0.27	0.13	0.26
EWC	0.18	0.27	0.14	0.21
A-GEM	0.21	0.14	0.09	0.13
MER	0.14	0.19	0.10	0.15
ER-RINGBUFFER (ours)	0.12	0.13	0.03	0.12

Table 1: *Forgetting when using a tiny episodic memory of single example per class.*

memory. When the switch happens (marked by a red vertical line in the figure), the learner only keeps randomly picked $\min(n, \frac{|\mathcal{M}|}{K})$ examples per class in the memory, where n is the number of examples of class c in the memory and K are the total number of classes observed so far. The overwriting happens opportunistically, removing examples from over-represented classes as new classes are observed. Fig. 2 shows that when the number of tasks is small, the hybrid version enjoys the high accuracy of reservoir sampling. As more tasks arrive and the memory per task shrinks, the hybrid scheme achieves superior performance than reservoir (and at least similar to ring buffer).

Finally, experience replay methods are not only outperforming all other approaches in terms of accuracy (and lower forgetting as reported in Tab. 1), but also in terms of compute time. Tab. 2 reports training time on both Split CIFAR and Split CUB, using ring buffer as a use case since all other ER methods have the same computational complexity. We observe that ER adds only a slight overhead compared to the finetuning baseline, but it is much cheaper than stronger baseline methods like A-GEM and MER.

5.5. Analysis

The strong performance of experience replay methods which directly learn using the examples stored in the small episodic memory may be surprising. In fact, Lopez-Paz & Ranzato (2017) discounted this repetitive training on the memory option by saying: “Obviously, minimizing the loss at the current example together with [the loss on the episodic memory] results in overfitting to the examples stored in [the memory]”. How can the repeated training over the same very small handful of examples possibly generalize?

To investigate this matter we conducted an additional experiment. For simplicity, we consider only two tasks, T_1 and T_2 , and study the generalization performance on T_1 as we train on T_2 . We denote by \mathcal{D}_2 the training set of T_2 and by \mathcal{M}_1 the memory storing examples from T_1 ’s training set. Our hypothesis is that although direct training on the examples in \mathcal{M}_1 (in addition to those coming from \mathcal{D}_2) does indeed lead to strong memorization of \mathcal{M}_1 (as measured by nearly zero cross-entropy loss on \mathcal{M}_1), such training is still overall

Methods	Training Time [s]	
	CIFAR	CUB
FINETUNE	87	194
EWC	159	235
A-GEM	230	510
MER	755	277
ER-RINGBUFFER (ours)	116	255

Table 2: *Learning Time on \mathcal{D}^{EV} [s]*

beneficial in terms of generalization on the original task T_1 because the joint learning with the examples of the current task T_2 acts as a strong, albeit implicit and data-dependent, regularizer for T_1 .

To validate this hypothesis, we consider the MNIST Rotations dataset (Lopez-Paz & Ranzato, 2017), where each task has digits rotated by a certain degree, a setting that enables fine control over the relatedness between the tasks. The architecture is the same as for Permuted MNIST, with only 10 memory slots, one for each class of T_1 . First, we verified that the loss on \mathcal{M}_1 quickly drops to nearly 0 as the model is trained using both \mathcal{M}_1 and \mathcal{D}_2 . As expected, the model achieves a perfect performance on the examples in the memory, which is not true for methods like A-GEM which make less direct use of the memory (see Appendix Tab. 7). We then verified that only training on \mathcal{M}_1 without \mathcal{D}_2 , yields strong overfitting to the examples in the memory and poor generalization performance, with a mere average accuracy of 40% on T_1 from the initial 85% which was obtained just after training on T_1 . If we only train on \mathcal{D}_2 without using \mathcal{M}_1 (same as FINETUNE baseline), we also observed overfitting to \mathcal{D}_2 as long as T_2 and T_1 are sufficiently unrelated, Fig. 3(b) and 3(c).

When the two tasks are closely related instead (difference of rotation angles less than 20 degrees), we observe that even without the memory, generalization on T_1 improves as we train on T_2 because of positive transfer from the related task, see red curve in Fig. 3(a). However, when we train on both \mathcal{D}_2 and \mathcal{M}_1 , generalization on T_1 is better than FINETUNE baseline, *i.e.*, training with \mathcal{D}_2 only, regardless of the degree of relatedness between the two tasks, as shown by the green curves in Fig. 3.

These findings suggest that while the model essentially memorizes the examples in the memory, this does not necessarily have a detrimental effect in terms of generalization as long as such learning is performed in conjunction with the examples of T_2 . Moreover, there are two major axes controlling this regularizer: the number of examples in T_2 and the relatedness between the tasks. The former sets the strength of the regularizer. The latter, as measured by the accuracy on T_1 when training only on \mathcal{D}_2 , controls its effectiveness.




Figure 3: *Analysis on MNIST Rotation: Test accuracy on Task 1 as a function of the training iterations over Task 2. The blue curves are the accuracy when the model is trained using only \mathcal{M}_1 . The red curves are the accuracy when the model is trained using only \mathcal{D}_2 , the training set of Task 2. The green curves are the accuracy when in addition to \mathcal{D}_2 , the model uses the memory from Task 1, \mathcal{M}_1 (experience replay). (Averaged over 3 runs).*

When T_1 and T_2 are closely related, Fig. 3(a), training on \mathcal{D}_2 prevents overfitting to \mathcal{M}_1 by providing a data-dependent regularization that, even by itself, produces positive transfer. When T_1 and T_2 are somewhat related, Fig. 3(b), training on \mathcal{D}_2 still improves generalization on T_1 albeit to a much lesser extent. However, when the tasks are almost adversarial to each other as an upside down 2 may look like a 5, the resulting regularization becomes even harmful, Fig. 3(c). In this case, accuracy drops from 40% (training only on \mathcal{M}_1) to 30% (training on both \mathcal{M}_1 and \mathcal{D}_2).

One remaining question related to generalization is how ER relates to A-GEM (Chaudhry et al., 2019) and whether A-GEM overfits even less? The answer is positive. As shown in Appendix Tab. 7, A-GEM’s accuracy on the memory examples does not reach 100% even after having processed 1000 samples. Interestingly, accuracy on the training set is lower than ER suggesting that the more constrained weight updates of A-GEM make it actually underfit. This underfitting prevents A-GEM from reaping the full regularization benefits brought by training on the data of subsequent tasks.

Conclusions

In this work we studied ER methods for supervised CL tasks. Our empirical analysis on several benchmark streams of data shows that ER methods even with a tiny episodic memory offer a very large performance boost at a very marginal increase of computational cost compared to the finetuning baseline. We also studied various ways to populate the memory and proposed a hybrid approach that strikes a good trade-off between randomizing the examples in the memory while keeping enough representatives for each class.

Our study also sheds light into a very interesting phenomenon: memorization (zero cross-entropy loss) of tiny memories is useful for generalization because training on subsequent tasks acts like a data dependent regularizer. Overall, we hope the CL community will adopt experience replay methods as a baseline, given their strong empirical performance, efficiency and simplicity of implementation.

There are several avenues of future work. For instance, we would like to investigate what are the optimal inputs that

best mitigate expected forgetting and optimal strategies to remove samples from the memory when it is entirely filled up.

Acknowledgments

This work was supported by the Rhodes Trust, ERC grant ERC-2012-AdG 321162-HELIOS, EPSRC grant Seebibyte EP/M013774/1 and EPSRC/MURI grant EP/N019474/1. We would also like to acknowledge the Royal Academy of Engineering and FiveAI.

References

- Aljundi, R., Babiloni, F., Elhoseiny, M., Rohrbach, M., and Tuytelaars, T. Memory aware synapses: Learning what (not) to forget. In *ECCV*, 2018.
- Chaudhry, A., Dokania, P. K., Ajanthan, T., and Torr, P. H. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *ECCV*, 2018.
- Chaudhry, A., Ranzato, M., Rohrbach, M., and Elhoseiny, M. Efficient lifelong learning with a-gem. In *ICLR*, 2019.
- Foerster, J., Nardelli, N., Farquhar, G., Afouras, T., Torr, P. H., Kohli, P., and Whiteson, S. Stabilising experience replay for deep multi-agent reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1146–1155. JMLR. org, 2017.
- Hayes, T. L., Cahill, N. D., and Kanan, C. Memory efficient experience replay for streaming learning. *arXiv preprint arXiv:1809.05922*, 2018.
- Hinton, G., Vinyals, O., and Dean, J. Distilling the knowledge in a neural network. In *NIPS*, 2014.
- Isele, D. and Cosgun, A. Selective experience replay for lifelong learning. *arXiv preprint arXiv:1802.10269*, 2018.
- Kirkpatrick, J., Pascanu, R., Rabinowitz, N. C., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., Hassabis, D., Clopath, C., Kumaran, D., and Hadsell, R. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, 2016.
- Krizhevsky, A. and Hinton, G. Learning multiple layers of features from tiny images. <https://www.cs.toronto.edu/~kriz/cifar.html>, 2009.
- LeCun, Y. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- Lee, J., Yun, J., Hwang, S., and Yang, E. Lifelong learning with dynamically expandable networks. *arXiv preprint arXiv:1708.01547*, 2017a.
- Lee, S.-W., Kim, J.-H., Ha, J.-W., and Zhang, B.-T. Overcoming catastrophic forgetting by incremental moment matching. In *NIPS*, 2017b.
- Li, Z. and Hoiem, D. Learning without forgetting. In *European Conference on Computer Vision*, pp. 614–629, 2016.
- Lopez-Paz, D. and Ranzato, M. Gradient episodic memory for continuum learning. In *NIPS*, 2017.
- McCloskey, M. and Cohen, N. J. Catastrophic interference in connectionist networks: The sequential learning problem. *Psychology of learning and motivation*, 24:109–165, 1989.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529, 2015.
- Nichol, A. and Schulman, J. Reptile: a scalable metalearning algorithm. *arXiv preprint arXiv:1803.02999*, 2018, 2018.
- Rebuffi, S.-V., Kolesnikov, A., and Lampert, C. H. iCaRL: Incremental classifier and representation learning. In *CVPR*, 2017.
- Riemer, M., Cases, I., Ajemian, R., Liu, M., Rish, I., Tu, Y., and Tesauro, G. Learning to learn without forgetting by maximizing transfer and minimizing interference. In *ICLR*, 2019.
- Ring, M. B. Child: A first step towards continual learning. *Machine Learning*, 28(1):77–104, 1997.
- Rolnick, D., Ahuja, A., Schwarz, J., Lillicrap, T. P., and Wayne, G. Experience replay for continual learning. *CoRR*, abs/1811.11682, 2018. URL <http://arxiv.org/abs/1811.11682>.
- Schwarz, J., Luketina, J., Czarnecki, W. M., Grabska-Barwinska, A., Teh, Y. W., Pascanu, R., and Hadsell, R. Progress and compress: A scalable framework for continual learning. In *International Conference in Machine Learning*, 2018.
- Thrun, S. Lifelong learning algorithms. In *Learning to learn*, pp. 181–209. Springer, 1998.
- Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al. Matching networks for one shot learning. In *Advances in neural information processing systems*, pp. 3630–3638, 2016.
- Vitter, J. S. Random sampling with a reservoir. *ACM Transactions on Mathematical Software (TOMS)*, 11(1): 37–57, 1985.
- Wah, C., Branson, S., Welinder, P., Perona, P., and Belongie, S. The caltech-ucsd birds-200-2011 dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011.

Zenke, F., Poole, B., and Ganguli, S. Continual learning through synaptic intelligence. In *ICML*, 2017.

Appendix

In §A, we provide algorithms for different memory update strategies described in §4 of the main paper. The detailed results of the experiments which were used to generate Fig. 1 and Tab. 1 in the main paper are provided in §B. The analysis conducted in §5.5 of the main paper is further described in §C. Finally, in §D, we list the hyper-parameters used for each of the baselines across all the datasets.

A. Memory Update Algorithms

Here we provide the algorithms to write into memory as discussed in §4 of the main paper.

Algorithm 2 Reservoir sampling update. `mem.sz` is the number of examples the memory can store, t is the task id, n is the number of examples observed so far in the data stream, and B is the input mini-batch.

```

1: procedure UPDATEMEMORY(mem.sz,  $t$ ,  $n$ ,  $B$ )
2:    $j \leftarrow 0$ 
3:   for  $(\mathbf{x}, y)$  in  $B$  do
4:      $M \leftarrow |\mathcal{M}|$   $\triangleright$  Number of samples currently stored in the memory
5:     if  $M < \text{mem.sz}$  then
6:        $\mathcal{M}.\text{append}(\mathbf{x}, y, t)$ 
7:     else
8:        $i = \text{randint}(0, n + j)$ 
9:       if  $i < \text{mem.sz}$  then
10:         $\mathcal{M}[i] \leftarrow (\mathbf{x}, y, t)$   $\triangleright$  Overwrite memory slot.
11:       $j \leftarrow j + 1$ 
12:   return  $\mathcal{M}$ 
    
```

Algorithm 3 Ring buffer.

```

1: procedure UPDATEMEMORY(mem.sz,  $t$ ,  $n$ ,  $B$ )
2:   for  $(\mathbf{x}, y)$  in  $B$  do
3:     # Assume FIFO stacks  $\mathcal{M}[t][y]$  of fixed size are already initialized
4:      $\mathcal{M}[t][y].\text{append}(\mathbf{x})$ 
5:   return  $\mathcal{M}$ 
    
```

Algorithm 4 K-Means. Memory is populated using samples closest (in feature space) to sequential K-Means centroids.

```

1: procedure UPDATEMEMORY(mem.sz,  $t$ ,  $n$ ,  $B$ )
2:   # Assume array  $\mathcal{M}[t][y]$  of fixed size is already initialized
3:   # Assume K centroids  $c_j$  are already initialized
4:   # Assume cluster counters  $n_j$  are already initialized to 0
5:   for  $(\mathbf{x}, y)$  in  $B$  do
6:      $j \leftarrow \text{argmin}_{j \in \{1, \dots, K\}} \|\phi_\theta(\mathbf{x}) - c_j\|$ 
7:      $n_j \leftarrow n_j + 1$ 
8:      $c_j \leftarrow c_j + \frac{1}{n_j} * (\phi_\theta(\mathbf{x}) - c_j)$ 
9:      $d = \|\phi_\theta(\mathbf{x}) - c_j\|$ 
10:    if  $d < \mathcal{M}[t][y][j].\text{get\_dst}()$  then  $\triangleright$  Store the current
        example if it is closer to the centroid
11:       $\mathcal{M}[t][y][j] \leftarrow (\mathbf{x}, d)$ 
12:   return  $\mathcal{M}$ 
    
```

Algorithm 5 Mean of Features. Store examples that are closest to the running average feature vector.

```

1: procedure UPDATEMEMORY(mem.sz,  $t$ ,  $n$ ,  $B$ )
2:   # Assume heaps  $\mathcal{M}[t][y]$  of fixed size are already initialized
3:   # Assume average features  $f[t][y]$  are already initialized
4:   # Assume moving average decay hyper-parameter ( $\alpha$ ) is given
5:   for  $(\mathbf{x}, y)$  in  $B$  do
6:      $f[t][y] \leftarrow \alpha * f[t][y] + (1 - \alpha) * \phi_\theta(\mathbf{x})$ 
7:      $d = \|\phi_\theta(\mathbf{x}) - f[t][y]\|$ 
8:     if  $\mathcal{M}[t][y].\text{find\_max}() > d$  then  $\triangleright$  Store the current
        example if it is closer to the center
9:        $\mathcal{M}[t][y].\text{delete\_max}()$ 
10:       $\mathcal{M}[t][y].\text{insert}(\mathbf{x}; d)$ 
11:   return  $\mathcal{M}$ 
    
```

B. Detailed Results

Here we describe the detailed results used to generate the Fig. 1 in the main paper. In addition we also report the forgetting metric (3). Note that the MULTI-TASK baseline does not follow the definition of continual learning as it keeps the dataset of all the tasks around at every step.

C. Further Analysis

In Tab. 7, we provide train, memory and test set performance on both the ER-RINGBUFFER and A-GEM with two different configurations of tasks; similar tasks (10° rotation), dissimilar tasks (90° rotation). It can be seen from the table, and as argued in the §5.5 of the main paper that ER-RINGBUFFER always achieves the perfect performance on the memory. To achieve the same effect with A-GEM one has to train for more iterations.

D. Hyper-parameter Selection

Table 3: **Permuted MNIST**: Performance (average accuracy (left column) and forgetting (right column)) for different number of samples per class. The average accuracy numbers from the this table are used to generate Fig. 1 in §5.4 of the main paper.

Methods	Episodic Memory (Samples Per Class)							
	Average Accuracy [$A_T(\%)$]				Forgetting [F_T]			
	1	3	5	15	1	3	5	15
A-GEM	62.1 (± 1.39)	63.2 (± 1.47)	64.1 (± 0.74)	66.0 (± 1.78)	0.21 (± 0.01)	0.20 (± 0.01)	0.19 (± 0.01)	0.17 (± 0.02)
MER	69.9 (± 0.40)	74.9 (± 0.49)	78.3 (± 0.19)	81.2 (± 0.28)	0.14 (± 0.01)	0.09 (± 0.01)	0.06 (± 0.01)	0.03 (± 0.01)
ER-RINGBUFFER	70.2 (± 0.56)	73.5 (± 0.43)	75.8 (± 0.24)	79.4 (± 0.43)	0.12 (± 0.01)	0.09 (± 0.01)	0.07 (± 0.01)	0.04 (± 0.01)
ER-MOF	69.9 (± 0.68)	73.9 (± 0.64)	75.9 (± 0.21)	79.7 (± 0.19)	0.13 (± 0.01)	0.09 (± 0.01)	0.07 (± 0.01)	0.04 (± 0.01)
ER-K-MEANS	70.5 (± 0.42)	74.7 (± 0.62)	76.7 (± 0.51)	79.1 (± 0.32)	0.12 (± 0.01)	0.08 (± 0.01)	0.06 (± 0.01)	0.04 (± 0.01)
ER-RESERVOIR	68.9 (± 0.89)	75.2 (± 0.33)	76.2 (± 0.38)	79.8 (± 0.26)	0.15 (± 0.01)	0.08 (± 0.01)	0.07 (± 0.01)	0.04 (± 0.01)
FINETUNE	53.5 (± 1.46)	-	-	-	0.29 (± 0.01)	-	-	-
EWC	63.1 (± 1.40)	-	-	-	0.18 (± 0.01)	-	-	-
MULTI-TASK	83				-			

Table 4: **Split CIFAR**: Performance (average accuracy (left column) and forgetting (right column)) for different number of samples per class. The average accuracy numbers from the this table are used to generate Fig. 1 in §5.4 of the main paper.

Methods	Episodic Memory (Samples Per Class)							
	Average Accuracy [$A_T(\%)$]				Forgetting [F_T]			
	1	3	5	13	1	3	5	13
A-GEM	54.9 (± 2.92)	56.9 (± 3.45)	59.9 (± 2.64)	63.1 (± 1.24)	0.14 (± 0.03)	0.13 (± 0.03)	0.10 (± 0.02)	0.07 (± 0.01)
MER	49.7 (± 2.97)	57.7 (± 2.59)	60.6 (± 2.09)	62.6 (± 1.48)	0.19 (± 0.03)	0.11 (± 0.01)	0.09 (± 0.02)	0.07 (± 0.01)
ER-RINGBUFFER	56.2 (± 1.93)	60.9 (± 1.44)	62.6 (± 1.77)	64.3 (± 1.84)	0.13 (± 0.01)	0.09 (± 0.01)	0.08 (± 0.02)	0.06 (± 0.01)
ER-MOF	56.6 (± 2.09)	59.9 (± 1.25)	61.1 (± 1.62)	62.7 (± 0.63)	0.12 (± 0.01)	0.10 (± 0.01)	0.08 (± 0.01)	0.07 (± 0.01)
ER-K-MEANS	56.6 (± 1.40)	60.1 (± 1.41)	62.2 (± 1.20)	65.2 (± 1.81)	0.13 (± 0.01)	0.09 (± 0.01)	0.07 (± 0.01)	0.04 (± 0.01)
ER-RESERVOIR	53.1 (± 2.66)	59.7 (± 3.87)	65.5 (± 1.99)	68.5 (± 0.65)	0.19 (± 0.02)	0.12 (± 0.03)	0.09 (± 0.02)	0.05 (± 0.01)
FINETUNE	40.6 (± 3.83)	-	-	-	0.27 (± 0.04)	-	-	-
EWC	41.2 (± 2.67)	-	-	-	0.27 (± 0.02)	-	-	-
MULTI-TASK	68.3				-			

Table 5: **miniImageNet**: Performance (average accuracy (left column) and forgetting (right column)) for different number of samples per class. The average accuracy numbers from the this table are used to generate Fig. 1 in §5.4 of the main paper.

Methods	Episodic Memory (Samples Per Class)							
	Average Accuracy [$A_T(\%)$]				Forgetting [F_T]			
	1	3	5	13	1	3	5	13
A-GEM	48.2 (± 2.49)	51.6 (± 2.69)	54.3 (± 1.56)	54 (± 3.63)	0.13 (± 0.02)	0.10 (± 0.02)	0.08 (± 0.01)	0.09 (± 0.03)
MER	45.5 (± 1.49)	49.4 (± 3.43)	54.8 (± 1.79)	55.1 (± 2.91)	0.15 (± 0.01)	0.12 (± 0.02)	0.07 (± 0.01)	0.07 (± 0.01)
ER-RINGBUFFER	49.0 (± 2.61)	53.5 (± 1.42)	54.2 (± 3.23)	55.9 (± 4.05)	0.12 (± 0.02)	0.07 (± 0.02)	0.08 (± 0.02)	0.06 (± 0.03)
ER-MOF	48.5 (± 1.72)	53.3 (± 2.80)	53.3 (± 3.11)	56.5 (± 1.92)	0.12 (± 0.01)	0.08 (± 0.01)	0.08 (± 0.02)	0.05 (± 0.02)
ER-K-MEANS	48.5 (± 0.35)	52.3 (± 3.12)	56.6 (± 2.48)	55.1 (± 1.86)	0.12 (± 0.02)	0.09 (± 0.02)	0.06 (± 0.01)	0.06 (± 0.01)
ER-RESERVOIR	44.4 (± 3.22)	50.7 (± 3.36)	56.2 (± 4.12)	61.3 (± 6.72)	0.17 (± 0.02)	0.12 (± 0.03)	0.07 (± 0.04)	0.04 (± 0.06)
FINETUNE	34.7 (± 2.69)	-	-	-	0.26 (± 0.03)	-	-	-
EWC	37.7 (± 3.29)	-	-	-	0.21 (± 0.03)	-	-	-
MULTI-TASK	62.4				-			

Table 6: **CUB**: Performance (average accuracy (left column) and forgetting (right column)) for different number of samples per class. The average accuracy numbers from the this table are used to generate Fig. 1 in §5.4 of the main paper.

Methods	Episodic Memory (Samples Per Class)							
	Average Accuracy [$A_T(\%)$]				Forgetting [F_T]			
	1	3	5	10	1	3	5	10
A-GEM	62.1 (± 1.28)	62.1 (± 1.87)	63.4 (± 2.33)	62.5 (± 2.34)	0.09 (± 0.01)	0.08 (± 0.02)	0.07 (± 0.01)	0.08 (± 0.02)
MER	55.4 (± 1.03)	65.3 (± 1.68)	68.1 (± 1.61)	71.1 (± 0.93)	0.10 (± 0.01)	0.04 (± 0.01)	0.03 (± 0.01)	0.03 (± 0.01)
ER-RINGBUFFER	65.0 (± 0.96)	71.4 (± 1.53)	73.6 (± 1.57)	75.5 (± 1.84)	0.03 (± 0.01)	0.01 (± 0.01)	0.01 (± 0.01)	0.02 (± 0.01)
ER-K-MEANS	67.9 (± 0.87)	71.6 (± 1.56)	73.9 (± 2.01)	76.1 (± 1.74)	0.02 (± 0.01)	0.02 (± 0.01)	0.02 (± 0.01)	0.01 (± 0.01)
ER-RESERVOIR	61.7 (± 0.62)	71.4 (± 2.57)	75.5 (± 1.92)	76.5 (± 1.56)	0.09 (± 0.01)	0.04 (± 0.01)	0.02 (± 0.01)	0.03 (± 0.02)
FINETUNE	55.7 (± 2.22)	-	-	-	0.13 (± 0.03)	-	-	-
EWC	55.0 (± 2.34)	-	-	-	0.14 (± 0.02)	-	-	-
MULTI-TASK	65.6				-			




Figure 4: **MNIST**: Evolution of average accuracy (A_k) as new tasks are learned when ‘1’ and ‘15’ samples per class are used.




Figure 5: **CIFAR**: Evolution of average accuracy (A_k) as new tasks are learned when using ‘1’ and ‘13’ samples per class. The performance is averaged over 5 runs. Uncertainty estimates are provided in Tabs 3, 4, 5, 6.




Figure 6: **miniImageNet**: Evolution of average accuracy (A_k) as new tasks are learned when ‘1’ and ‘13’ samples per class are used.



 Figure 7: *CUB*: Evolution of average accuracy (A_k) as new tasks are learned when ‘1’ and ‘10’ samples per class are used.

 Table 7: *MNIST Rotation* Performance of task 1 after training on task 2.

Task 2 Samples	Rotation Angle											
	10°						90°					
	ER-RINGBUFFER			A-GEM			ER-RINGBUFFER			A-GEM		
	Train	Mem	Test	Train	Mem	Test	Train	Mem	Test	Train	Mem	Test
1000	85.6	1	86.2	81.5	86.6	82.5	68.7	1	69.4	51.7	73.3	52.1
20000	91.4	1	91.6	91.4	1	91.5	32.7	1	33.4	31.6	1	33.0

 Table 8: Hyper-parameters selection on the four benchmark datasets. ‘lr’ is the learning rate, ‘ λ ’ is the synaptic strength for EWC, ‘ γ ’ is the with in batch meta-learning rate for MER, ‘s’ is current example learning rate multiplier for MER.

Methods	MNIST	CIFAR	CUB	miniImageNet
FINETUNE	lr (0.1)	lr (0.03)	lr (0.03)	lr (0.03)
EWC	lr (0.1), λ (10)	lr (0.03), λ (10)	lr (0.03), λ (10)	lr (0.03), λ (10)
A-GEM	lr (0.1)	lr (0.03)	lr (0.03)	lr (0.03)
MER	lr (0.03), γ (0.1), s (10)	lr (0.03), γ (0.1), s (5)	lr (0.1), γ (0.1), s (5)	lr (0.03), γ (0.1), s (5)
ER-RESERVOIR	lr (0.1)	lr (0.1)	lr (0.03)	lr (0.1)
ER-[OTHERS]	lr (0.1)	lr (0.03)	lr (0.03)	lr (0.03)