

# Principal Gradient Direction and Confidence Reservoir Sampling for Continual Learning

Zhiyi Chen<sup>1</sup> and Tong Lin<sup>(✉)</sup><sup>2</sup>

<sup>1</sup> Georgia Institute of Technology, USA  
zchen798@gatech.edu

<sup>2</sup> The Key Laboratory of Machine Perception (MOE), School of EECS, Peking University; Peng Cheng Laboratory, Shenzhen, China  
lintong@pku.edu.cn

**Abstract.** Task-free online continual learning aims to alleviate catastrophic forgetting of the learner on a non-iid data stream. Experience Replay (ER) is a SOTA continual learning method, which is broadly used as the backbone algorithm for other replay-based methods. However, the training strategy of ER is too simple to take full advantage of replayed examples and its reservoir sampling strategy is also suboptimal. In this work, we propose a general proximal gradient framework so that ER can be viewed as a special case. We further propose two improvements accordingly: **Principal Gradient Direction** (PGD) and **Confidence Reservoir Sampling** (CRS). In Principal Gradient Direction, we optimize a target gradient that not only represents the major contribution of past gradients, but also retains the new knowledge of the current gradient. We then present Confidence Reservoir Sampling for maintaining a more informative memory buffer based on a margin-based metric that measures the value of stored examples. Experiments substantiate the effectiveness of both our improvements and our new algorithm consistently boosts the performance of MIR-replay, a SOTA ER-based method: our algorithm increases the average accuracy up to 7.9% and reduces forgetting up to 15.4% on four datasets.

**Keywords:** continual learning, principal gradient direction, confidence reservoir sampling

## 1 Introduction

Primates and humans can continually learn new skills and accumulate knowledge throughout their lifetime [5]. However, in machine learning, the agents hardly have a steady good performance when they learn a data stream. *Catastrophic forgetting* [10] is a common challenge when training a single neural network model on consecutive tasks: the model may perform well over the first task but suffers a serious accuracy decay along with the training process on the next tasks.

---

This work was supported by NSFC Tianyuan Fund for Mathematics (No. 12026606), and National Key RD Program of China (No. 2018AAA0100300).

2 Zhiyi Chen and Tong Lin<sup>(✉)</sup>

*Continual learning* [14], also known as *lifelong learning* [16], is a special field in machine learning that focuses on avoiding or alleviating catastrophic forgetting.

The primary setting of continual learning (CL) is the task-incremental setting [17], which assumes the stream of data can be clearly divided into sequential tasks and learnt offline. However, task-free online has received increasing attention recently, which is more practical: not only each sample can be merely observed once (single pass setting) but also the data stream is non-iid without any task information to assist the process of continual learning.

There are three major families of architecture in CL: expansion-based methods, regularization-based methods and replay-based methods. In this paper, we focus on the last one, which store the previous raw data and replay some of them when learning current data to alleviate forgetting. Experience Replay (ER) [4] is one of the most representative methods, and has been proven as a strong baseline. Because of its superior performance, ER becomes the backbone algorithm for many recent replay-based methods, such as ER-MIR [1], GSS [2], etc.

However, there is still room for improvement: on the one hand, the training strategy of ER is too simple to make full use of examples. On the other hand, reservoir sampling, which is a commonly used memory update strategy, can only ensure the equilibrium of previous samples but not good enough to maintain a more informative memory buffer. Our paper aims to tackle these defects and produces a stronger backbone algorithm for other continual learning methods based on ER.

In this paper, we firstly present a new algorithm for the training strategy called Principal Gradient Direction (PGD), which attempts to optimize a new gradient that not only represents the past data better but also retains the new knowledge of the current example. Secondly, we define a margin-based metric to measure the value of stored data and propose Confidence Reservoir Sampling (CRS), which helps to maintain a more informative memory buffer.

Under the online CL setting, our experimental results show that both of our two approaches improve ER and also boost the performance of other ER-based CL methods, such as MIR [1], which achieve the best accuracy and forgetting measure among all the replay-based methods.

## 2 Methods

In this section, we will first discuss the setup of task-free online continual learning and replay-based methods in Section 2.1, and then propose a proximal gradient framework to analyze the training strategy of ER from a new perspective in Section 2.2. Finally, we elaborate our two methods: Principal Gradient Direction and Confidence Reservoir Sampling in Section 2.3 & 2.4.

### 2.1 Setup

In task-free online continual learning setting, there is a stream of non-iid data:  $\dots, (x_t, y_t), \dots$ , which doesn't contain any task information to identify the specific

Title Suppressed Due to Excessive Length 3

task that one example belongs to. The learner  $f(\cdot; \theta^t)$  can only observe  $(x_t, y_t)$  at the  $t^{th}$  training step due to the single pass constraint.

For replay-based methods, a space-limited memory buffer  $\mathcal{M}$  can be used to store some examples to help provide information of past data. The learner should try to maximize the overall performance of all data, i.e., the average accuracy, and minimize the forgetting of past knowledge.

Many methods [1] [2] have addressed their improvements on the simple random selection used in ER, which is orthogonal to our improvements. In the following subsections, we will analyze the shortcomings of ER on training strategy and storage strategy and present our improvements in Section 2.3 & 2.4 accordingly.

## 2.2 Proximal Gradient Framework

In this subsection, we use *Proximal operator* [12], a well-studied numerical method in optimization, to build a proximal gradient framework, which is the foundation of our Principal Gradient Direction and also provides a new perspective to the training strategy of ER.

The proximal operator of a function  $f(\cdot)$  with a scalar parameter  $\lambda (> 0)$  is defined by

$$\text{prox}_{\lambda f}(v) := \arg \min_x f(x) + \frac{1}{2\lambda} \|x - v\|_2^2, \quad (1)$$

where  $x \in \mathbb{R}^n, v \in \mathbb{R}^n$  are two  $n$  dimensional vectors and  $f : \mathbb{R}^n \rightarrow \mathbb{R}$  is a closed proper convex function. Proximal operators can be interpreted as modified gradient steps:

$$\text{prox}_{\lambda f}(v) = v - \lambda \nabla M_{\lambda f}(v), \quad (2)$$

where  $M_{\lambda f}$  is a smoothed or regularized form of  $f$  termed as Moreau envelop  $M_{\lambda f}(v) := \inf_x f(x) + \frac{1}{2\lambda} \|x - v\|_2^2$ .

As shown in [9], continual learning can be formulated as a minimization problem that finds a new gradient close enough to the gradient of the new data and satisfies some constraints at the same time. In other words, the new gradient should still be beneficial to the current task and also takes the past tasks into consideration.

Based on this insight, we introduce the proximal operator into the setting of continual learning:

$$\text{prox}_{\lambda f}(g) = \arg \min_w f(w) + \frac{1}{2\lambda} \|g - w\|_2^2, \quad (3)$$

where  $g$  is the gradient vector calculated on the new data, and  $w$  is the target gradient to update the network weights.  $f(\cdot)$  is the convex function we need to design which characterizes the relation between the target gradient and gradients of past examples selected from the memory.

The training strategy of ER is simple: the learner randomly samples a small batch of past data from memory and directly uses the sampled data  $B_t$  as well as the new input data  $(x_t, y_t)$  to co-train the network. From the perspective of proximal gradient framework, the constraint function  $f(\cdot)$  of ER is the inner

4 Zhiyi Chen and Tong Lin<sup>(✉)</sup>

product of the target gradient and the average gradient of selected past data without  $\lambda$ :

$$\min_w \frac{1}{2} \|g - w\|_2^2 - \langle g_{ref}, w \rangle, \quad (4)$$

where  $g_{ref}$  is the reference gradient of  $B_t$ . The Equation (4) has an analytic solution as follows, which is the actual training strategy of ER:

$$w^* = g + g_{ref}. \quad (5)$$

However, this strategy ignores the difference of sampled examples and it also regards new data and past data equally weighted, which is suboptimal.

### 2.3 Principal Gradient Direction

A more reasonable idea of utilizing the new data and selected examples is to find a target gradient that not only represents the overall contribution of the sampled past examples, but also maintains the knowledge of new data. Such a gradient can be found in the neighbor of  $g$ , which should also follow the principal direction of all past gradients. In this way, the new gradient will not violate the past knowledge for the reason that principal direction ensures a gradient descent towards a overall decrease on losses of past examples. In addition, the gradient also promotes the memorization of new data because it is a near neighbour of  $g$ .

To find the principal direction, we attempt to minimize the sum of solid angles between the new gradient vector and the past gradients, i.e., maximize the sum of cosine value. Besides, the length of a gradient should also be taken into consideration, because the ‘‘short’’ gradient vector means that current model  $f(\cdot; \theta^t)$  can learn it well and hence is less important than a ‘‘long’’ gradient. So we apply *sigmoid* function on length of the gradient as weight. We can also set a small threshold  $\epsilon$  for the length of gradient:  $\|g_i\|_\epsilon = \max(\epsilon, \|g_i\|)$  to further decrease the impact of the short one.

Under the proximal gradient framework, we formulate a optimization problem as follows:

$$\min_w - \sum_{i=0}^K \frac{\langle w, g_i \rangle}{\|w\| \|g_i\|_\epsilon} \text{sigmoid}(\|g_i\|) + \frac{1}{2\alpha} \|w - g\|_2^2, \quad (6)$$

where  $w$  is the target gradient,  $g$  is the gradient of the new input,  $g_i$  is the gradient of the sampled past example,  $\alpha$  is a hyperparameter to balance the two parts and  $K$  is the size of sampled batch  $B_n$ .

To solve this optimization problem, we choose *Proximal Gradient Method* [12] to get an iterative solution of the proximal problem. Considering a general optimization problem:

$$\text{minimize } f(x) + h(x), \quad (7)$$

where  $f : \mathbb{R}^n \rightarrow \mathbb{R}$  and  $h : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$  are two closed proper convex functions and  $f$  is differentiable. The *Proximal Gradient Method* is formulated

Title Suppressed Due to Excessive Length 5

as follows:

$$x^{(k+1)} = \text{prox}_{\beta h}(x^{(k)} - \beta \nabla f(x^{(k)})). \quad (8)$$

As for our problem, we regard the target gradient  $w$  as the optimization variable, the principal direction term in (6) as function  $f$  and the distance constraint term as function  $h$ .

After substituting the variables and expanding the formulation of (8), we get the standard form of proximal gradient method for our optimization problem:

$$\begin{aligned} w^{(k+1)} = \arg \min_w & \frac{1}{2\beta} \|w - (w^{(k)} - \beta \nabla f(w^{(k)}))\|_2^2 \\ & + \frac{1}{2\alpha} \|w - g\|_2^2. \end{aligned} \quad (9)$$

To find the solution, we need to set the derivative of (9) to zero. Note that we can ignore the constant term, e.g.  $g^T g$ , so we can get:

$$w^{(k+1)} = \frac{\alpha(w^{(k)} - \beta \nabla f(w^{(k)})) + \beta g}{\alpha + \beta}. \quad (10)$$

For the gradient  $\nabla f(w^{(k)})$ , with the rule of derivation for fraction, the solution is:

$$\begin{aligned} \nabla f(w^{(k)}) = & - \sum_{i=0}^K \left( \frac{g_i}{\|w^{(k)}\| \|g_i\|_\epsilon} \text{sigmoid}(\|g_i\|) \right. \\ & \left. - \frac{\langle w^{(k)}, g_i \rangle w^{(k)}}{\|w^{(k)}\|_2^3 \|g_i\|_\epsilon} \text{sigmoid}(\|g_i\|) \right). \end{aligned} \quad (11)$$

Here we choose the gradient of new input data  $g$  as  $w^{(0)}$  for the reason that the new gradient should be a neighbor of  $g$ . From empirical observation, we find that just one step optimization is good enough, so an approximate solution is:

$$w^{(1)} = g - \frac{\alpha\beta}{\alpha + \beta} \nabla f(g). \quad (12)$$

We replace the fraction  $\alpha\beta/(\alpha + \beta)$  in (12) with a single hyperparameter  $\lambda$  in experiment, which makes it look like one step gradient descent from  $g$  on our principal direction function  $f$ . In practice, we can choose to group the examples averagely to decrease the number of backward propagation to obtain an appropriate computational complexity.

## 2.4 Confidence Reservoir Sampling

In this subsection, we focus on the storage strategy about how to update the memory with the new example  $(x_t, y_t)$ .

6 Zhiyi Chen and Tong Lin<sup>(✉)</sup>**Algorithm 1 Reservoir sampling**


---

```

Procedure:  $\mathcal{M}$ , mem.sz,  $t$ ,  $(x_t, y_t)$ 
if  $|\mathcal{M}| \leq \text{mem.sz}$  then
     $\mathcal{M}.\text{append}((x_t, y_t))$ 
else
     $i = \text{randint}(0, t)$ 
    if  $i \leq \text{mem.sz}$  then
         $\mathcal{M}[i] \leftarrow (x_t, y_t)$ 
    end if
end if

```

---

ER and many other replay-based methods apply reservoir sampling strategy (Algorithm 1) [18], where mem.sz is the total memory size of  $\mathcal{M}$  and  $t$  is the order number of input  $(x_t, y_t)$ .

Though this strategy can ensure the equilibrium for memory buffer, the random replacement (the blue row in Algorithm 1) still has a room for improvement considering the limited memory space. We aspire to maintain a more informative memory buffer by replacing the less useful examples, which can improve continual learning no matter which subset is selected to consolidate the past knowledge.

Just like the exploration and exploitation dilemma in reinforcement learning, the same situation also exists in online continual learning: exploration is replacing the old data with the new one to explore the new knowledge, while exploitation is keeping the old data intact. Actually, only when an example is selected, it is really exploited by the learner.

Inspired by the idea of Upper-Confidence Bound (UCB) algorithm, which balances the uncertainty and reward of a certain action to choose one from the action set, we use a similar strategy to calculate a score for each example in memory buffer and choose the appropriate one to be replaced.

The exploitation rate, denoted as  $EX$ , is the first part of the metric, which is calculated by a division from the times  $n$  that the example is selected into  $B_t$  and the age of the example  $a$ :  $EX = n/a$ . We intend to replace the highly exploited one, which is more likely to be overfitted by the learner.

Then we define *margin* [8] based on the prediction probability from the forward propagation: the output prediction  $p(x; \theta)$  on an example  $(x, y)$  is computed through a softmax activation function, and we formulate margin, denoted as  $m$ , as:

$$m := p_y(x; \theta) - \max_{y' \neq y} p_{y'}(x; \theta). \quad (13)$$

When the model makes a correct prediction, the margin of the certain input is positive, otherwise, we get a negative margin. Margin value indicates the confidence of the prediction: larger the margin is in magnitude, more confidence we have in the prediction.

At the  $t^{\text{th}}$  training step, we can first get  $m_t$  of  $(x_t, y_t)$  from model  $f(\cdot; \theta^t)$  and then  $m_{t+1}$  from the new model  $f(\cdot; \theta^{t+1})$  that executes one step gradient descent. Then we define margin increment:  $MI = m_{t+1} - m_t$ , which measures

Title Suppressed Due to Excessive Length 7

**Algorithm 2 Confidence Reservoir sampling**


---

```

Procedure:  $\mathcal{M}$ , mem.sz,  $t$ ,  $(x_t, y_t)$ 
if  $|\mathcal{M}| \leq \text{mem.sz}$  then
   $\mathcal{M}.\text{append}((x_t, y_t))$ 
else
   $i = \text{randint}(0, t)$ 
  if  $i \leq \text{mem.sz}$  then
    if Using strategy  $s_1$  then
       $j \leftarrow \max(\mathcal{S}(\mathcal{M}))$ 
    else if Using strategy  $s_2$  then
       $j \sim P(j) = \mathcal{S}_j / \sum_k \mathcal{S}_k$ 
    end if
     $\mathcal{M}[j] \leftarrow (x_t, y_t)$ 
  end if
end if

```

---

the importance of a certain example at one training step. If margin increment is large, it means that this training step has learnt the example very well, in other words, the example is simple and less informative for the model.

So we can calculate our metric, denoted as  $\mathcal{S}$ , for all the examples in memory buffer:

$$\mathcal{S} := EX + c \cdot MI, \quad (14)$$

where  $EX$  is the exploitation rate,  $MI$  is the margin increment and  $c$  is a weight hyperparameter. For a high score, the example is either over-exploited or less informative, which is more appropriate to be replaced.

We have two strategies to replace examples based on  $\mathcal{S}$ :

$s_1$  directly chooses the biggest score, and  $s_2$  replaces each example with a probability  $P(i) = \mathcal{S}_i / \sum_j \mathcal{S}_j$ , which applies to different datasets.

So far, we complete the definition of our margin-based metric and implement it on reservoir sampling as Confidence Reservoir Sampling (Algorithm 2). In this way, Confidence Reservoir Sampling not only satisfies the requirement of equal storage, but also maintains a more informative memory buffer. Note that our margin-based metric can also be extended to other storage strategy.

### 3 Experiments

In this section, we report the details of experiments and the performance of our two improvements. We apply PGD and CRS on ER and conduct ablation study. We also use the renewed backbone algorithm over MIR-replay [1] to demonstrate the effectiveness of our approaches.

#### 3.1 Datasets and Architectures

We consider four commonly used datasets:

- (1) **MNIST Split** is derived from MNIST, the famous dataset on handwritten

8 Zhiyi Chen and Tong Lin<sup>(✉)</sup>**Table 1.** Average accuracy (%) of ablation Study ( $\uparrow$ )

METHOD	ER	ER-P	ER-C	ER-PC
MNIST-S	79.8 $\pm$ 3.2	82.4 $\pm$ 2.1	81.5 $\pm$ 2.1	<b>84.0<math>\pm</math>2.3</b>
MNIST-P	79.1 $\pm$ 0.7	80.9 $\pm$ 0.3	79.9 $\pm$ 0.5	<b>81.7<math>\pm</math>0.6</b>
CIFAR10-S	30.7 $\pm$ 2.0	36.1 $\pm$ 1.8	38.5 $\pm$ 1.1	<b>40.0<math>\pm</math>2.1</b>
MINI-S	23.0 $\pm$ 1.2	25.5 $\pm$ 0.6	25.2 $\pm$ 0.8	<b>25.8<math>\pm</math>1.0</b>

**Table 2.** Forgetting measure (%) of ablation Study ( $\downarrow$ )

METHOD	ER	ER-P	ER-C	ER-PC
MNIST-S	19.2 $\pm$ 4.0	13.2 $\pm$ 3.1	17.3 $\pm$ 3.2	<b>9.6<math>\pm</math>1.9</b>
MNIST-P	4.3 $\pm$ 0.5	2.6 $\pm$ 0.5	4.0 $\pm$ 0.6	<b>2.4<math>\pm</math>0.4</b>
CIFAR10-S	63.3 $\pm$ 2.7	56.6 $\pm$ 3.7	<b>49.4<math>\pm</math>1.7</b>	49.7 $\pm$ 3.3
MINI-S	32.1 $\pm$ 2.0	25.7 $\pm$ 1.4	28.5 $\pm$ 1.1	<b>25.7<math>\pm</math>1.1</b>

digits, which directly splits 10 classes of MNIST into 5 non-overlapping different tasks.

(2) **MNIST Permutations** is also derived from MNIST, which randomly generates different pattern of pixel permutation for each task to exchange the position of the original images of MNIST. For both MNIST Split and MNIST Permutations, we use the similar benchmark setting as [9] that each task consists of 1000 examples.

(3) **CIFAR10 Split** is derived from CIFAR10, which averagely divides the whole classes in CIFAR10 into 5 tasks, where each task has 9750 samples and 250 retained for validation just as [1].

(4) **MiniImageNet Split** is derived from miniImageNet, a subset of ImageNet with 100 classes and 600 images per class, which averagely divides the whole classes into 20 tasks.

For MNIST-S and MNIST-P, all baselines use fully-connected neural networks with two hidden layers of 100 ReLU units. A smaller version of ResNet18 [6] is used for CIFAR10-S and MINI-S, which has three times less feature maps for each layer than the original ResNet18.

### 3.2 Metrics

We use *Average Accuracy* and *Forgetting Measure* [3] to evaluate the performances of the baselines over four datasets. For *Average Accuracy*, the higher the number (indicated by  $\uparrow$ ) the better is the model. For *Forgetting Measure*, the lower the number (indicated by  $\downarrow$ ) the better is the model. We run 10 times to get each result.



Title Suppressed Due to Excessive Length 9

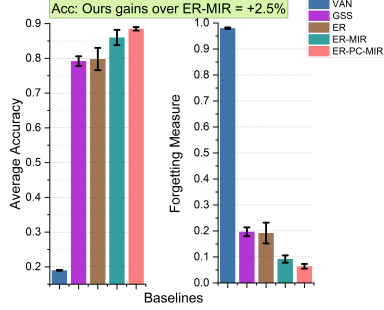


Fig. 1. Performances on MNIST-S

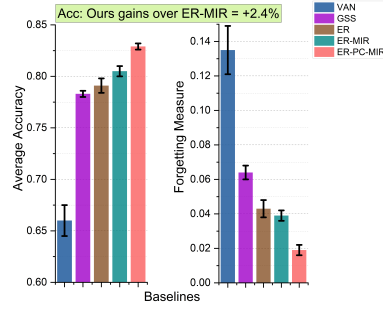


Fig. 2. Performances on MNIST-P

### 3.3 Ablation Study

We conduct ablation study on four datasets by combining our two approaches with ER, and the resulting algorithms are as follows: basic ER (noted as ER), ER pluses PGD (noted as ER-P), ER pluses CRS (noted as ER-C) and ER pluses both PGD and CRS (noted as ER-PC). We store 50 examples per class and select 10 past examples for  $B_t$  on MNIST-S, MNIST-P and CIFAR10-S while store 100 examples per class and select 20 examples on MINI-S. The results are showed in Table 1 & 2.

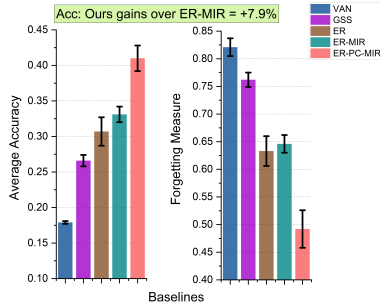
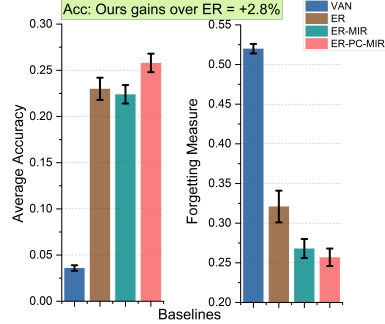
**Effectiveness of PGD and CRS** From the results, we can observe that both PGD and CRS can improve the performance of ER on all four datasets: the two methods can boost the average accuracy up to 7.7% and reduce the forgetting measure up to 13.9%. On MNIST-S and MNIST-P, whose size are relatively small and network is simpler, PGD contributes more than CRS. The situation reverses on CIFAR10-S. MINI-S has the longest task sequence (20 tasks) and the biggest input size, where our two approaches have similar contribution in average accuracy. The comparative relations are same in forgetting measure.

**Joint improvement of PGD and CRS** The results also demonstrate that PGD and CRS can always jointly render a further improvement. On all four datasets, ER-PC is the best algorithm in terms of average accuracy which outperforms ER from 2.6% to 9.3%. ER-PC also achieves least forgetting on the first three datasets, which only performs slightly worse than ER-C on CIFAR10-S.

Our aim is to produce a stronger backbone algorithm for other ER-based methods, so we use ER-PC as a renewed backbone algorithm for the following comparison.

### 3.4 Performance of ER-PC

In this subsection, we will show the performance of ER-PC, where we use it as the new backbone algorithm by overlying MIR-replay [1] on it, which is an example-selection strategy for replay and is SOTA replay-based method so far. We note the new method as **ER-PC-MIR**.

**Fig. 3.** Performances on CIFAR10-S**Fig. 4.** Performances on MINI-S

**Basic comparison** We take the following four baselines into comparison: **VAN** (a vanilla method that a single predictor for all the tasks without any continual learning strategy), **ER** [4], **ER-MIR** [1] (the basic version of MIR-replay based on ER) and **GSS** [2].

For the reason that the training time of GSS on MINI-S is unacceptable, we don't take GSS into comparison on this dataset. We also don't take GEM [9] and A-GEM [3] into comparison because they all need the task information to update the memory and train the network, which violate the task-free online CL setting. Prior works show that ER and ER-MIR outperform GEM-like algorithms. The settings of memory size are same as our ablation study. The results are reported in Figure 1 - 4.

First, ER-PC-MIR achieves the best average accuracy on all four datasets. On MNIST-S, MNIST-P and CIFAR10-S, ER-PC-MIR achieves better average accuracy than ER-MIR, the best baseline on these datasets, with improvements up to 7.9%. In MINI-S, our method is better than ER, the best baseline, with improvement 2.8%.

Second, our method also forgets least knowledge among the baselines on all four datasets: ER-PC-MIR reduces forgetting than ER-MIR with improvements from 1.1% to 15.4%. On CIFAR10-S, ER is the best baseline in terms of forgetting measure, and ER-PC-MIR is better than it with 13.6%.

The results show that our method ER-PC is a stronger backbone algorithm than vanilla ER: after combining with MIR-replay, ER-PC-MIR not only outperforms than ER-MIR, but also achieves the best performance among all other replay-based methods.

**Comparison in Different Memory Size** As MNIST-P and CIFAR10-S are two representative datasets in domain-incremental and class-incremental datasets, we run ER-MIR and ER-PC-MIR on them in different memory size. We store 100, 50, 25 and 10 examples per class, which means that the total size of memory buffer is 1000, 500, 250, 100 on two datasets. We report the average accuracy and forgetting measure in Table 3 & 4.

Title Suppressed Due to Excessive Length 11

**Table 3.** Average accuracy (%) on MNIST-P and CIFAR10-S in different memory size ( $\uparrow$ )

MNIST-P	1000	500	250	100
ER-MIR	82.7 $\pm$ 0.4	80.5 $\pm$ 0.5	77.5 $\pm$ 0.9	73.6 $\pm$ 1.0
ER-PC-MIR	<b>84.4<math>\pm</math>0.4</b>	<b>82.9<math>\pm</math>0.3</b>	<b>79.6<math>\pm</math>0.6</b>	<b>76.1<math>\pm</math>0.4</b>
CIFAR10-S	1000	500	250	100
ER-MIR	43.5 $\pm$ 1.7	33.1 $\pm$ 1.1	27.1 $\pm$ 2.3	22.0 $\pm$ 2.2
ER-PC-MIR	<b>48.9<math>\pm</math>2.5</b>	<b>41.0<math>\pm</math>1.8</b>	<b>33.7<math>\pm</math>1.9</b>	<b>26.6<math>\pm</math>3.0</b>

**Table 4.** Forgetting measure (%) on MNIST-P and CIFAR10-S in different memory size ( $\downarrow$ )

MNIST-P	1000	500	250	100
ER-MIR	2.3 $\pm$ 0.4	3.9 $\pm$ 0.3	6.0 $\pm$ 0.6	8.8 $\pm$ 0.9
ER-PC-MIR	<b>1.2<math>\pm</math>0.3</b>	<b>1.9<math>\pm</math>0.3</b>	<b>4.4<math>\pm</math>0.5</b>	<b>7.0<math>\pm</math>0.7</b>
CIFAR10-S	1000	500	250	100
ER-MIR	46.4 $\pm$ 5.1	64.6 $\pm$ 1.6	72.2 $\pm$ 4.3	77.0 $\pm$ 2.5
ER-PC-MIR	<b>36.0<math>\pm</math>4.8</b>	<b>49.2<math>\pm</math>3.4</b>	<b>54.6<math>\pm</math>3.8</b>	<b>69.1<math>\pm</math>5.3</b>

In all memory size, ER-PC-MIR consistently improves the performance of ER-MIR. ER-PC-MIR achieves more average accuracy than ER-MIR from 1.7% to 2.5% on MNIST-P. On CIFAR10-S, ER-PC-MIR gains over ER-MIR from 4.6% to 7.9% in average accuracy. The results show the reliability of our renewed backbone algorithm in different memory size.

## 4 Conclusion

In this paper, we firstly focus on the training strategy of CL and present a proximal gradient framework. Based on it, **Principal Gradient Direction** is proposed to take full advantage of replayed examples and new data. Then we pay attention to memory updating strategy: we define a new margin-based metric to measure the value of stored data and propose **Confidence Reservoir Sampling** based on it to maintain a more informative memory buffer. The experiments demonstrate that our two approaches are both beneficial and can jointly give a further improvement. After applied with PGD and CRS, the renewed backbone algorithm can boost the performance of MIR-replay and always achieves the best performance among other replay-based baselines on four datasets. On task-incremental and domain-incremental datasets, our method also consistently outperforms ER-MIR in different memory size. The experiments show that our method is a reliable and stronger backbone algorithm than vanilla ER.

12 Zhiyi Chen and Tong Lin<sup>(✉)</sup>

## References

1. Rahaf Aljundi, Lucas Caccia and Eugene Belilovsky, et al: Online Continual Learning with Maximally Interfered Retrieval. In: NeurIPS 2019a
2. Rahaf Aljundi, Min Lin, Baptiste Goujaud and Yoshua Bengio: Gradient based sample selection for online continual learning. In: NeurIPS 2019.
3. Arslan Chaudhry, Marc'Aurelio Ranzato, Marcus Rohrbach and Mohamed Elhoseiny: Efficient Lifelong Learning with A-GEM. In: ICLR 2019
4. Arslan Chaudhry, Marcus Rohrbach and Mohamed Elhoseiny, et al: Continual Learning with Tiny Episodic Memories. arXiv, abs/1902.10486, 2019b
5. Fagot, Joël and Cook, Robert G.: Evidence for large long-term memory capacities in baboons and pigeons and its implications for learning and the evolution of cognition. In: Proceedings of the National Academy of Sciences, 103, (46)17564-17567, 2006
6. Kaiming He, Xiangyu Zhang, Shaoqing Ren and Jian Sun: Deep Residual Learning for Image Recognition. In: CVPR, 2016
7. Yen-Chang Hsu, Yen-Cheng Liu and Zsolt Kira: Re-evaluating Continual Learning Scenarios: A Categorization and Case for Strong Baselines. CoRR, abs/1810.12488, 2018
8. Vladimir Koltchinskii and Dmitry Panchenko: Empirical margin distributions and bounding the generalization error of combined classifiers. In: The Annals of Statistics, 30(1): 1–50, 2002
9. David Lopez-Paz and Marc'Aurelio Ranzato: Gradient Episodic Memory for Continual Learning. In: NeurIPS, 2017
10. McCloskey, Michael and Cohen, Neal J.: Catastrophic Interference in Connectionist Networks: The Sequential Learning Problem. In: Psychology of Learning and Motivation, 1989
11. Maxime Oquab, Léon Bottou, Ivan Laptev and Josef Sivic: Learning and Transferring Mid-level Image Representations Using Convolutional Neural Networks. In: CVPR, 2014
12. N. Parikh and S. Boyd: Proximal algorithms. In: Foundations and Trends in Optimization, 1(3): 123–231, 2014
13. Matthew Riemer, Ignacio Cases and Robert Ajemian, et al: Learning to Learn without Forgetting by Maximizing Transfer and Minimizing Interference. In: ICLR, 2019
14. Mark B. Ring: Continual Learning in Reinforcement Environments. University of Texas at Austin, 1994
15. Ozan Sener and Vladlen Koltun: Multi-Task Learning as Multi-Objective Optimization. In: NeurIPS, 2018
16. Thrun and Sebastian: A Lifelong Learning Perspective for Mobile Robot Control. In: IEEE/RSJ/GI International Conference on Intelligent Robots and Systems, 1994
17. Gido M. van de Ven and Andreas S. Tolias: Three scenarios for continual learning. CoRR: abs/1904.07734, 2019
18. Jeffrey Scott Vitter: Random Sampling with a Reservoir. In: ACM Trans. Math. Softw., 11(1): 37–57, 1985