

Dealing with Cross-Task Class Discrimination in Online Continual Learning

Yiduo Guo¹, Bing Liu⁴, Dongyan Zhao^{1,2,3}

¹Wangxuan Institute of Computer Technology, Peking University. ²BIGAI, Beijing, China. ³National Key Laboratory of General Artificial Intelligence. ⁴Department of Computer Science, University of Illinois Chicago
 yiduo@stu.pku.edu.cn, liub@uic.edu, zhaodongyan@pku.edu.cn

Abstract

Existing continual learning (CL) research regards catastrophic forgetting (CF) as almost the only challenge. This paper argues for another challenge in class-incremental learning (CIL), which we call **cross-task class discrimination (CTCD)**, i.e., how to establish decision boundaries between the classes of the new task and old tasks with no (or limited) access to the old task data. CTCD is implicitly and partially dealt with by replay-based methods. A replay method saves a small amount of data (replay data) from previous tasks. When a batch of current task data arrives, the system jointly trains the new data and some sampled replay data. The replay data enables the system to partially learn the decision boundaries between the new classes and the old classes as the amount of the saved data is small. However, this paper argues that the replay approach also has a **dynamic training bias** issue which reduces the effectiveness of the replay data in solving the CTCD problem. A novel optimization objective with a gradient-based adaptive method is proposed to dynamically deal with the problem in the online CL process. Experimental results show that the new method achieves much better results in online CL.

1. Introduction

Continual learning (CL) learns a sequence of tasks incrementally. This work focuses on the *class incremental learning* (CIL) setting [32] in online CL. In CIL, each task consists of a set of unique classes, the sets of classes of any two different tasks are disjoint and the system has no access to the task information in testing. In online CL, the data comes gradually from a data stream. Whenever the small batch of data arrives, it is trained in one iteration. Thus, the data for each task is effectively trained in one epoch.

Existing CL papers almost regard *catastrophic forgetting* (CF) as the only issue for CL. In fact, CIL also has another major challenge. When the system learns a new task, if no data from previous tasks is available, it has no way to establish decision boundaries between new classes and old

classes in previous tasks. Even if there is no CF, the classification results will still be poor. We call this problem, **cross-task class discrimination (CTCD)**. Those approaches that do not save any previous data, e.g., *regularization-based* or *orthogonal projection-based*, do not deal with CTCD. Replay-based methods implicitly deal with CTCD to some extent because such a method uses a memory buffer \mathcal{M} to save a small amount of data (*replay data*) from old tasks. When a small batch of current task data X^{new} arrives, the system jointly trains X^{new} and some sampled replay data X^{buf} from \mathcal{M} . X^{buf} enables the system to partially learn the decision boundaries between the new classes and the old classes because the amount of the saved data is very small.

Due to the limited replay data, the **training is biased**, which reduces its ability to solve the CTCD problem. To make matters worse, the training bias also changes as more tasks are learned. This paper first shows that the problem is reflected as **gradient imbalance (GI)** on logits, i.e., higher positive gradients than negative gradients *on the logits* and vice versa. It further shows that GI is caused by two main issues. The first is **data imbalance**. Since the memory buffer size, the batch size of the new data X^{new} , and the sampled data X^{buf} from the memory buffer are all fixed, if the system has learned many tasks, the average number of samples in each previous class in X^{buf} will be much smaller than that of each class in X^{new} . This results in higher positive gradients than negative gradients on the logits of the previous classes leading to training bias and poor decision boundaries (or weak CTCD capability) between the classes of the new and old tasks. The second is **CL imbalance**, i.e., CL training focuses more on the new samples (which are harder to train as they are new) than the replayed samples (which have been seen and trained many times before). This causes further GI. This imbalance is involved (see Sec. 4.2 for details).

Some existing works [2, 42] have tried to deal with data imbalance in offline CL. For example, SSIL [2] separately calculates the cross-entropy loss of the new data and the replay data to mitigate data imbalance. But they are not from the gradient angle. The second issue of GI is more complex and has not been attempted before.

This paper proposes a novel method, called GSA (Gradient Self-Adaptation), to deal with GI (and CTCD) in online CL. GSA includes a new training objective and a *gradient-based self-adaptive loss* to compensate for the GI. The loss is dynamically controlled by two *gradient rates* which automatically measure and adapt to the dynamic GI situation. The main contributions of this paper are:

(1) It deals with the CTCD problem in online CL and proposes a new optimization framework that decomposes the problem into cross-task classification and within-task classification (see Section 5). In [22], CTCD is called *inter-task class separation*, but it uses an out-of-distribution based approach to dealing with the problem in offline CL. The paper uses a replay-based approach for online CL.

(2) It analyzes the CTCD problem from the gradient imbalance (GI) angle and finds two kinds of gradient imbalance (data imbalance and CL imbalance) (see Section 4). Based on the analysis, it proposes a *gradient-based self-adaptive loss* to compensate for the GI.

(3) Experiments in both the disjoint and long-tail online CL settings show that GSA outperforms strong baselines by a large margin (see Section 6).

2. Related Work

There are many existing CL techniques. *Regularization-based* approaches penalize changes to important parameters of old tasks [23, 50, 52]. *Replay-based* approaches save some past data and replay them in new task training [8, 12, 13, 24, 37, 44, 47, 54]. *Generative replay* builds data generators to generate pseudo old data for replaying [18, 21, 34, 38, 41, 45]. *Dynamic architectures* based approaches [1, 14, 20, 28, 36, 39, 43, 49] overcome CF by expanding or isolating parameters. Data augmentation has also been used to learn better features for CL recently [55].

Online CL and replay methods: Existing online CL methods mainly use the replay approach. ER randomly samples the buffer data [11]. MIR chooses buffer data whose loss increases the most [3]. Shapley value theory is applied in ASER for memory update/retrieval [40]. Knowledge distillation is employed in DER++ [6]. NCCL deals with CF by calibrating the network [48]. Contrastive learning is used in SCR [30]. GDumb samples and stores class-balanced data in memory [35]. However, it does not deal with CL imbalance identified in this paper. GSS diversifies the gradients of the samples in the buffer [4]. ER-AML [7] avoids CF by shielding the learned representations from drastic adaptation to accommodate new classes. OCM is based on mutual information [16]. Our work identifies two gradient imbalance issues and designed a new training strategy and a novel loss function to deal with them.

Data imbalance in CL: Several researchers have dealt with data imbalance in offline CL. SSIL [2] isolates the computation of softmax on previous and new classes to

update the model. LUCIR [19] uses the cosine normalization to calculate the predicted probability and a margin ranking loss to separate ground-truth old classes from new classes. [8] adds an additional fine-tuning stage with a small learning rate and a balanced subset of samples. BiC [46] adds a bias correction layer. [53] uses weight alignment to correct the biased weights. CCIL [33] applies a loss to balance intra-task and inter-task learning. However, many of these algorithms [8, 53] need the full data of the current task to be available upfront and/or multiple training epochs to address the data imbalance issue, which are not suitable for online CL as online CL does not have the full training data of a task available when the task arrives in the data stream.

Gradient-based CL methods: GEM [29] and A-GEM [10] rotate the current gradient when the angle between the current gradient and the gradient computed on the reference memory is obtuse. MEGA [17] uses the loss of X^{new} and X^{buf} to adjust the relative importance of learning a new task and maintaining the past knowledge. OWM [51], OGD [14], and AOP [15] avoid CF by projecting the gradients on the new task onto an orthogonal subspace of old tasks. These methods do not deal with gradient imbalance but represent a different family of CL approaches.

3. Preliminary

Problem description. Following [3], we learn a sequence of tasks, 1, 2, ..., t , ..., in the online CIL setting. In online CL, each incrementally arriving data batch is seen only once by the system. So the system effectively learns each task in one epoch. For a task t , we denote its dataset as $\{(x_k^{(t)}, y_k^{(t)})\}_{k=1}^{n^t}$, where n^t is the number of training samples in task t . $L^{(t)}$ denotes the set of classes of task t .

Model architecture and training setting. Our model F consists of a feature extractor h_θ with the parameter set θ and a classifier f_ϕ with the parameter set ϕ . It uses a replay method with a memory buffer \mathcal{M} . Whenever a small batch of new data X^{new} from the data stream is accumulated, it is trained jointly with a small batch of data X^{buf} sampled from \mathcal{M} to update the model in one training iteration. The model produces the logits $F(x; \theta, \phi) = f_\phi(h_\theta(x))$, which are used to calculate the loss or to predict in testing. Reservoir sampling is used for memory update.

Gradient imbalance (GI) on logits. We now introduce the proposed gradient-based analysis on logits and the concept of GI. In CL, each task is usually learned by minimizing the softmax cross-entropy loss, \mathcal{L}_{ce} :

$$\mathcal{L}_{ce}(o(x; \theta, \phi)) = - \sum_{i=1}^{|C_{seen}|} l_{c_i} \log(p^{c_i}), \quad p^{c_i} = \frac{e^{o_{c_i}}}{\sum_{s=1}^{|C_{seen}|} e^{o_{c_s}}} \quad (1)$$

where $|C_{seen}|$ is the number of classes that the model has seen, $l_{c_i} \in \{0, 1\}$ is the one-hot label of class c_i , and $o(x; \theta, \phi) = [o_{c_1}, o_{c_2}, \dots, o_{c_{|C_{seen}|}}]$ is the set of logit values

for input x . Given a training sample x of class c_i , the gradients on logits ($c_j \neq c_i$) are given by

$$\frac{\partial \mathcal{L}_{ce}(o(x; \theta, \phi))}{\partial o_{c_i}} = p^{c_i} - 1, \quad \frac{\partial \mathcal{L}_{ce}(o(x; \theta, \phi))}{\partial o_{c_j}} = p^{c_j}. \quad (2)$$

From Eq. 2, we see that x gives its true logit o_{c_i} a **negative gradient** and the other logits o_{c_j} **positive gradients**. As the gradient update rule for a parameter w is $w = w - lr * \nabla w$, where lr is the learning rate, the negative gradient ($p^{c_i} - 1$) results in an increase in o_{c_i} for the true class c_i and the positive gradient (p^{c_j}) results in a decrease in o_{c_j} for each wrong class c_j . Thus, *the negative gradient encourages the model to output a larger probability for the true class, and positive gradients help output lower probabilities for the wrong classes.*

However, in CL, as the model has no access to the training data of previous tasks when it learns a new task incrementally, all gradients on previous classes are positive during the new task training (there are no negative gradients) (**imbalance of positive and negative gradients**). Then the model tends to output smaller probabilities on previous classes, biasing the classification towards the new classes.

4. Gradient Imbalance (GI) in Replaying

A replay-based method actually makes the (positive and negative) gradients more balanced. There are adjustments of both ups and downs for logits for previous classes. However, this is insufficient. We discuss two reasons: *data imbalance* and *CL imbalance*. We propose some metrics first.

Metrics: When the model is learning the new task t , for a seen class c_i in any seen task (including the current task t), we define the mean of *positive gradients* ($p_k^{c_i}$) that the logits of all other classes than c_i have and the mean of *negative gradients* ($p_k^{c_i} - 1$) that class c_i 's logit has over all $n^{t,r}$ training samples (including both the data samples seen so far from the current task t and the corresponding replay data used) for this task as

$$\begin{aligned} P(t, c_i) &= \frac{\sum_{k=1}^{n^{t,r}} p_k^{c_i} \cdot \mathbb{I}(y_k \neq c_i)}{n^t} \\ N(t, c_i) &= \frac{\sum_{k=1}^{n^{t,r}} (p_k^{c_i} - 1) \cdot \mathbb{I}(y_k = c_i)}{n^t} \end{aligned} \quad (3)$$

where y_k is the label of the k th sample of task t and \mathbb{I} is the indicator function.

At the task level, for classes in any task t' ($t' \leq t$), we define the mean of *positive* and *negative gradients* that the logits of the class set $L^{(t')}$ receives in training task t respectively as

$$\begin{aligned} P(t, L^{(t')}) &= \frac{\sum_{j=1}^{|C_{seen}|} P(t, c_j) \cdot \mathbb{I}(c_j \in L^{(t')})}{\sum_{j=1}^{|C_{seen}|} 1 \cdot \mathbb{I}(c_j \in L^{(t')})} \\ N(t, L^{(t')}) &= \frac{\sum_{j=1}^{|C_{seen}|} N(t, c_j) \cdot \mathbb{I}(c_j \in L^{(t')})}{\sum_{j=1}^{|C_{seen}|} 1 \cdot \mathbb{I}(c_j \in L^{(t')})} \end{aligned} \quad (4)$$

4.1. Gradient Imbalance due to Data Imbalance

In a replay-based method, the memory buffer size is usually fixed and small and the batch size N^{buf} for the buffer batch X^{buf} and the batch size N^{new} for the new data batch X^{new} are also fixed. As the number of previous classes grows with more tasks learned, the number of sampled data for each previous class $\frac{N^{buf}}{\bigcup_{r=1}^{t-1} |L^{(r)}|}$ in X^{buf} gets smaller, but the number of samples for each new class $\frac{N^{new}}{|L^{(t)}|}$ in X^{new} of the new task remains unchanged. Then we have $\frac{N^{new}}{|L^{(t)}|} > \frac{N^{buf}}{\bigcup_{r=1}^{t-1} |L^{(r)}|}$ and the samples from previous classes and the new data classes can become highly imbalanced. The positive gradients for the previous classes can surpass their negative gradients (in absolute values). We can empirically verify this by calculating the *positive-negative* (PN) *gradient rate* with $t' \leq t$:

$$PN(t, t') = \frac{P(t, L^{(t')})}{N(t, L^{(t')})} \quad (5)$$

Empirical Verification: We introduce the setup first.

Datasets and Tasks: We conduct our experiments using two dataset settings: (1) Split CIFAR100, where we divide the CIFAR100 dataset into 10 tasks with 10 unique classes per task, and (2) Split TinyImagenet, where we divide the TinyImageNet dataset into 20 tasks with 10 different classes per task. For online CL, we run each task in one epoch.

Model, Optimizer and Batch Size: We use the full size ResNet-18 to perform our experiments. To give a more general analysis, we run with both SGD and Adam optimizers. For learning rate, we follow [3] and set it as 0.1 for the SGD optimizer. To ensure a good performance, we search and set the learning rate as 0.001 for the Adam optimizer. For batch size, we also follow [3] and set it to 10 for X^{new} and 10 for the buffer batch X^{buf} , which is randomly sampled from the memory buffer. Since ER (Experience Replay) is a basic replay method, we conduct our experiments using ER [11].

We show the PN results of the classes of a few tasks of the two datasets in Fig. 1 as more subsequent tasks are learned. We can make the following observations from Fig. 1. See plots based on SGD in Appendix 1 (see footnote 2).

(1). Sub-figures (A) and (B) in Fig. 1 show that when the model is trained with the new task t , the *PN* rates for the classes of old tasks get smaller than -1 and tend to decrease. This indicates that positive gradients surpass negative gradients (in absolute values) for the old classes. The reason is that the new data classes dominate the training, and the new data give high positive gradients to classes of the old tasks to ensure the new task data are not classified to those old classes. As more tasks are learned, the data get even more imbalanced, resulting in lower negative gradients on the old classes from the replay data of the old task t' and thus the downward trend of the curves in (A) and (B).

(2). Sub-figure (C) in Fig. 1 shows that the tendency




Figure 1. *PN* rate (*rate* in the figures) of the CIFAR100 or TinyImageNet experiments with Adam optimizers. The buffer size is 1000. We choose four different tasks in each experiment and plot their *PN* rates as subsequent tasks are learned. In (C), we report their accuracy. The plots based on the SGD optimizer are given in Appendix 1, which show the same trend.

of the test accuracy performance of the previous classes is similar to the tendency of their *PN* rates.

Dealing with Data Imbalance. We propose a sampling strategy that guarantees the same number of training samples for each class in each training iteration. Specifically, we fix $N^{new} + N^{buf} = 20$, where N^{new} and N^{buf} are the number of samples in X^{new} and X^{buf} , respectively. We use the ratio between $|L^{(t)}|$ and $\sum_{r=1}^t |L^{(r)}|$ to decide the number of samples from X^{new} (the rest are not used here), i.e., $N^{new} = \max(\text{int}(20 \cdot \frac{|L^{(t)}|}{\sum_{r=1}^t |L^{(r)}|}), 1)$ where $\text{int}(\cdot)$ returns the nearest integer of a given number and $N^{buf} = 20 - N^{new}$. We use N^{buf} to sample previous tasks’ data in the memory buffer. In this way, the number of training data for every class is approximately equal.

4.2. Gradient Imbalance due to CL Imbalance

GI still occurs even after the number of samples used in each class is balanced using the technique above. Let us call the balanced data for training X^{mix} , which includes both the replay data from previous tasks and new data from the current task. We now introduce the second cause of GI, *CL imbalance*, which is due to incremental training in CL. To explore CL imbalance, we define the *rate of accumulated positive and negative gradients* (A-*PN* rate) for a class $c_i \in L^{(t')}$ from task t' to the current task t ($t' \leq t$),

$$A\text{-}PN(t', t, c_i) = \frac{\sum_{r=t'}^t P(r, c_i)}{\sum_{r=t'}^t N(r, c_i)} \quad (6)$$

$A\text{-}PN(t', t, c_i)$ gives the accumulated gradient rate of each class over the learning process to the current task t .

Empirical Verification: We plot the A-*PN* rate of each class in Fig. 2(A) for the CIFAR100 dataset and in Fig. 2(B) for the TinyImageNet dataset after the last task is learned. We make the following observations from Fig. 2 (A), (B), (C) and some more details in the experiments:

(1). The A-*NP* rate accumulated is close to -1 or balanced when we jointly train all classes in the CIFAR100

dataset as a single task (the blue curve and its moving average yellow curve in Fig. 2 (A)). We also observe balanced gradients from detailed results (not shown here) for all classes when we learn the first task (no replay data).

(2). We ignore the accumulative part in Eq. 6 for the time being and focus on the gradient imbalance in the classes of each task alone. Based on the detailed results (not given here), we observe that as more tasks are learned, the GI for the classes of the current task being learned gets worse and worse. This explains why the classes of the last task in sub-figures (A) and (B) have the highest GI. Note, the last task has no accumulation as $t' = t$. Let us explain why.

As an example, we consider class c_i from the current task t . As more and more tasks are learned, the replay data from previous tasks in X^{mix} will contribute less and less positive gradients to c_i because the replay data may have been trained many times in the past and are well overfitted to their own classes in the past and their probability of being classified to the new class c_i is very small and hence they contribute very low positive gradients to c_i . In our experiments, we observed this and also decreased negative gradient on c_i . We plot the average cross-entropy loss of the new data batch and the replay data batch when the model learns the second task of CIFAR100 in Appendix 2. The figure shows that the average cross-entropy loss of the new data batch is much higher than that of the replay data batch in the whole training process. This is because the system has learned good features from previous tasks that makes the learning of the new class c_i easier resulting in less negative gradients. But positive gradients drop more significantly.

We now consider the accumulation part (the summations in Eq. 6). The next section will show that A-*PN* is used to dynamically adjust the loss for each class in learning to balance the gradients. By right, when we learn the current task, we only need to consider the current GI situation to adjust the loss. However, considering only one case (the current situation) is risky due to random fluctuations. That




Figure 2. A-PN rate (*rate*) for each class (x -axis) after the last task is trained. The classes are ordered by their sequence of appearances in the original data. The buffer size is 1000 and the optimizer is Adam. (C) shows the test performances (see the formula in the legend) of all learned classes. The results based on the SGD optimizer are given in Appendix 3.

is why we consider the impact on a class c_i from all subsequent tasks, which gives us a more robust gradient rate for adjusting the loss in dealing with GI in learning.

(3). Fig. 2(C) shows that higher A-PN rates (imbalanced gradients) result in lower test accuracy. The fact that the accumulated negative gradient being greater than the accumulated positive gradients in absolute values (i.e., $A-PN > -1$) makes the model biased towards the new classes.

5. The Proposed GSA Method

Our method GSA consists of two parts: (1) a new optimization framework that separately optimizes *cross-task classification* and *within-task classification*. (2) A *gradient-based self-adaptive loss* to alleviate the gradient imbalance (CL imbalance) in our framework.

5.1. Optimizing Within-class Classification and Cross-task Classification

Assume that the system has seen $n - 1$ previous classes (c_1, \dots, c_{n-1}) and there $m - n + 1$ new classes (c_n, \dots, c_m) in the current batch X^{new} , for a class, say c_n , in the batch of new data, we decompose its learning into two parts: (1) learning the decision boundaries between c_n and the other classes in the current batch X^{new} and (2) learning the decision boundaries between c_n and all previous $n - 1$ classes of old tasks in X^{buf} . For example, for a sample x_{c_n} of new class c_n from X^{new} , its original cross-entropy loss is:

$$\mathcal{L}_{ce}(x_{c_n}) = -\log\left(\frac{e^{o_{c_n}}}{\sum_{s=1}^m e^{o_{c_s}}}\right) \quad (7)$$

where o_{c_s} is the logit value of x_{c_n} for class c_s . We decompose the loss (denoted by $\mathcal{L}_{decom}(x_{c_n})$) into two parts:

$$\mathcal{L}_{decom}(x_{c_n}) = -\log\left(\frac{e^{o_{c_n}}}{\sum_{s=n}^m e^{o_{c_s}}}\right) - \log\left(\frac{e^{o_{c_n}}}{\sum_{s=1}^{n-1} e^{o_{c_s}} + e^{o_{c_n}}}\right) \quad (8)$$

The first part of \mathcal{L}_{decom} distinguishes c_n and other new classes with respect to x_{c_n} . Their gradient rates are similar.

The second part distinguishes c_n and the previous classes with respect to x_{c_n} . The negative interference from new classes to previous classes is limited in the second part. For the relationship of loss \mathcal{L}_{decom} and loss \mathcal{L}_{ce} , we have the following proposition:

Proposition 1. For a sample x_{c_n} from the new class c_n in X^{new} , the following holds

$$\mathcal{L}_{decom}(x_{c_n}) \geq \mathcal{L}_{ce}(x_{c_n}) \quad (9)$$

The proof is in Appendix 4. Similarly, for a sampled data point x_{c_j} of class c_j ($j < n$) from the buffer data X^{buf} , we define its new loss as:

$$\mathcal{L}_{decom}(x_{c_j}) = -\log\left(\frac{e^{o_{c_j}}}{\sum_{s=1}^{n-1} e^{o_{c_s}}}\right) - \log\left(\frac{e^{o_{c_j}}}{\sum_{s=n}^m e^{o_{c_s}} + e^{o_{c_j}}}\right) \quad (10)$$

In this loss, we put the previous classes (similar gradient rates) in the first term to maintain the established decision boundaries between previous classes and put the logit of the true class and the logits of the classes in the current task in the second term to establish the decision boundaries between class c_j and the new classes. The negative transfer from the new classes to previous classes is limited to the second term. We prove that $\mathcal{L}_{decom}(x_{c_j})$ is the upper bound of $\mathcal{L}_{ce}(x_{c_j})$ in Appendix 4.

Considering data balance. To estimate the first term in Eq. 8 for class c_n , we use X^{new} as the data of new classes as X^{new} is usually class-balanced.¹ For a similar reason, we use X^{buf} to estimate the first term in Eq. 10 for the previous class c_j . However, for the second term in Eq. 8, using X^{new} causes prediction bias towards new classes as X^{new} lacks samples of previous classes. Also, using X^{buf} to estimate the second term in Eq. 10 is not appropriate as the number of new classes samples in the buffer increases gradually and for a long time, the new classes are the minor classes in the buffer. To solve both problems, we use the sampling strategy proposed in the last section under “Dealing with Data

¹If this is not the case, we can sample some data of the same class from the memory buffer or directly duplicate the samples to guarantee it.

Imbalance.” Then we obtain a mixed set of new data and old data sampled from the memory buffer, called X^{mix} . X^{mix} can be regarded as samples from the uniform joint distribution for all classes, including both the new and old classes. Then our final optimization goal is:

$$\begin{aligned} \min_{F \in \mathcal{F}} & -\mathbb{E}_{x \sim X^{new}} \log\left(\frac{e^{o_{yx}}}{\sum_{s=n}^m e^{o_{c_s}}}\right) - \mathbb{E}_{x \sim X^{buf}} \log\left(\frac{e^{o_{yx}}}{\sum_{s=1}^{n-1} e^{o_{c_s}}}\right) \\ & - \mathbb{E}_{x \sim X^{mix}} (\mathbb{I}(y_x \cap S(X^{new}) \neq \emptyset) \cdot \log\left(\frac{e^{o_{yx}}}{\sum_{s=1}^{n-1} e^{o_{c_s}} + e^{o_{yx}}}\right) + \\ & \mathbb{I}(y_x \cap S(X^{new}) = \emptyset) \cdot \log\left(\frac{e^{o_{yx}}}{\sum_{s=n}^m e^{o_{c_s}} + e^{o_{yx}}}\right)) \end{aligned} \quad (11)$$

where y_x is the label of sample x and $S(X^{new})$ is the set of classes appeared in the X^{new} and \mathbb{I} is the indicator function and \mathcal{F} is the function space. The first term and second terms focus on the within-task classification goal and the last term focus on the cross-task classification goal.

5.2. Self-Adaptive Loss for CL Imbalance

Based on the analysis in Section 4, we know that even though the data balance is guaranteed, GI still exists as the gradient rate is different for different classes, which is caused by CL imbalance. We note that CL imbalance (measured by the gradient rate of each class) not only occurs between new and previous classes but also occurs among previous classes. To mitigate it, we propose *GSA-CE* (Gradient-based Self-Adaptive CE) loss \mathcal{L}_{GSA-CE} . The loss function is *dynamically adjusted* based on the gradient rates, which change as more tasks and batches are learned. Let $\{x_k, y_k\}_{k=1}^{N^{mix}}$ be the samples in X^{mix} , their \mathcal{L}_{GSA-CE} is

$$\begin{aligned} \mathcal{L}_{GSA-CE}\{x_k, y_k\}_{k=1}^{N^{mix}} &= -\frac{1}{N^{mix}} \sum_{k=1}^{N^{mix}} w_{y_k} \cdot (\mathbb{I}(y_k \cap S(X^{new}) \neq \emptyset) \cdot \\ & \log\left(\frac{e^{o_{y_k}}}{\sum_{s=1}^{n-1} e^{o_{c_s}} + e^{o_{y_k}}}\right) + \mathbb{I}(y_k \cap S(X^{new}) = \emptyset) \cdot \\ & \log\left(\frac{e^{o_{y_k}}}{\sum_{s=n}^m v_{c_s} e^{o_{c_s}} + e^{o_{y_k}}}\right)) \end{aligned} \quad (12)$$

where

$$w_{y_k} = \frac{2}{1 + e^{A \cdot PN(t', t, y_k)}}, \quad v_{c_s} = \frac{1}{-PN(t, c_s)} \quad (13)$$

where t' is the task in which class c_s appears. When the model is training task t , we incrementally update the $A \cdot PN(t', t, y_k)$ and $PN(t, y_k)$ by adding the new gradients into the current sums of all previous gradients. \mathcal{L}_{GSA-CE} has three advantages:

(1). When the new task arrives, we do not need to assume that we have the entire training set for calculating some statistics or an exemplar set like CCIL [33]. Our method is thus suitable for online CL.

(2). The model sets the values of w_{y_k} and v_{c_s} based on $A \cdot PN$ rates and PN rates automatically as they vary dynamically with the current gradient imbalance situation.

(3). $A \cdot PN$ and PN rates are class-based metrics that reflect the discriminative power of the model for all seen classes. Our loss then gives class-based weights for samples from different classes. The new loss thus automatically adjust the loss weight w_{y_k} to balance the accumulated negative gradient and the accumulated positive gradient for each class. We justify the new loss in Appendix 5.

(4) No new hyperparameter is introduced in the process.

Similarly, as the gradient rate imbalance occurs also within previous classes, for samples $\{x_k, y_k\}_{k=1}^{N^{buf}}$ from X^{buf} , their \mathcal{L}_{GSA-CE} loss is:

$$\mathcal{L}_{GSA-CE}(\{x_k, y_k\}_{k=1}^{N^{buf}}) = -\frac{\sum_{k=1}^{N^{buf}} w_{y_k} \log\left(\frac{e^{o_{y_k}}}{\sum_{s=1}^{n-1} e^{o_{c_s}}}\right)}{N^{buf}} \quad (14)$$

Considering the gradient imbalance, we replace the second term and the third term in Eq. 11 with our \mathcal{L}_{GSA-CE} losses Eq. 14 and Eq. 12 in the expectation form respectively. The final optimization objective in Eq. 11 is turned into:

$$\begin{aligned} \min_{F \in \mathcal{F}} & -\mathbb{E}_{x \sim X^{new}} \log\left(\frac{e^{o_{yx}}}{\sum_{s=n}^m e^{o_{c_s}}}\right) - \mathbb{E}_{x \sim X^{buf}} w_{y_x} \cdot \log\left(\frac{e^{o_{yx}}}{\sum_{s=1}^{n-1} e^{o_{c_s}}}\right) \\ & - \mathbb{E}_{x \sim X^{mix}} w_{y_x} \cdot (\mathbb{I}(y_x \cap S(X^{new}) \neq \emptyset) \cdot \log\left(\frac{e^{o_{yx}}}{\sum_{s=1}^{n-1} e^{o_{c_s}} + e^{o_{yx}}}\right) \\ & + \mathbb{I}(y_x \cap S(X^{new}) = \emptyset) \cdot \log\left(\frac{e^{o_{yx}}}{\sum_{s=n}^m v_{c_s} e^{o_{c_s}} + e^{o_{yx}}}\right)) \end{aligned} \quad (15)$$

6. Experiments

Evaluation data. Four image classification datasets are used. 1) **MNIST** [27] has 10 classes with 60,000/10,000 training/test examples. We created 5 disjoint tasks with 2 classes per task. 2) **CIFAR10** [25] has 10 classes with 50,000/10,000 training/test samples. We created 5 disjoint tasks with 2 classes per task. 3) **CIFAR100** [25] has 100 classes with 50,000/10,000 training/test samples. 10 disjoint tasks are created with 10 classes per task. 4) **TinyImageNet** [26] has 200 classes. We created 100 disjoint tasks with 2 classes per task. Each class has 500 training examples and 50 test examples.

Compared Baselines. GSA² is compared with 16 recent baselines: 9 online CL baselines, **AGEM**, **ER**, **MIR**, **GSS**, **ASER**, **RainBow-rt**, **ER-AML**, **GDumb**, and **SCR** and 7 offline CL baselines, **DER++**, **IL2A**, **Co²L**, **LUCIR**, **CCIL**, **BiC** and **SSIL**, as they deal with data/class imbalance and can be run in the online CL mode, i.e., training in one epoch without requiring the full data of each task to be available when the task arrives. The citations of these systems are given in Table 1 associated with their results. Note that Rainbow [5] is not a standard online CL method as it re-trains all replay-data with data augmentation for 256 epochs after each task, which is not suitable for online CL because

²Code and Appendix: <https://github.com/gydpku/GSA>

Table 1. Accuracy on the four experiment datasets with different memory buffer sizes M . All values are averages of 15 runs.

Method	MNIST (5 tasks)			CIFAR10 (5 tasks)			CIFAR100 (10 tasks)			TinyImageNet (100 tasks)		
	$M=0.1k$	$M=0.5k$	$M=1k$	$M=0.2k$	$M=0.5k$	$M=1k$	$M=1k$	$M=2k$	$M=5k$	$M=2k$	$M=4k$	$M=10k$
AGEM [10]	56.9±5.2	57.7±8.8	61.6±3.2	22.7±1.8	22.7±1.9	22.6±0.7	5.8±0.2	6.5±0.3	5.8±0.2	0.9±0.1	2.1±0.1	3.9±0.2
ER [11]	78.7±0.4	88.0±0.2	90.3±0.1	49.7±0.6	55.2±0.6	59.3±0.2	15.7±0.3	22.4±0.5	29.5±0.9	4.7±0.5	10.1±0.7	11.7±0.2
MIR [3]	79.0±0.5	88.3±0.1	91.3±1.9	37.3±0.3	40.0±0.6	41.0±0.6	15.7±0.2	19.1±0.1	24.1±0.2	6.1±0.5	11.7±0.2	13.5±0.2
GSS [4]	70.4±1.5	80.7±5.8	87.5±5.9	26.9±1.2	30.7±1.3	40.1±1.4	11.1±0.2	13.3±0.5	17.4±0.1	3.3±0.5	10.0±0.2	10.5±0.2
ASER [40]	61.6±2.1	71.0±0.6	82.1±5.9	27.8±1.0	36.2±1.2	44.7±1.2	16.4±0.3	12.2±1.9	27.1±0.3	5.3±0.3	8.2±0.2	10.3±0.4
Rainbow-rt [5]	89.1±0.3	92.1±0.1	95.0±0.3	45.2±0.4	50.6±0.3	51.5±0.8	15.4±0.5	15.9±0.3	20.4±0.3	6.6±0.3	10.1±0.3	13.1±0.5
ER-AML [7]	76.5±0.1	86.0±0.2	91.5±0.1	40.5±0.7	48.7±0.5	50.1±0.4	16.1±0.4	17.6±0.5	22.6±0.1	5.4±0.2	7.1±0.5	10.1±0.4
GDumb [35]	81.2±0.5	91.0±0.2	94.5±0.1	35.9±1.1	50.7±0.7	63.5±0.5	17.1±0.4	25.1±0.2	38.6±0.5	12.6±0.1	12.7±0.3	15.7±0.2
SCR [31]	86.2±0.5	92.8±0.3	94.6±0.1	47.2±1.7	58.2±0.5	64.1±1.2	26.5±0.2	31.6±0.5	36.5±0.2	10.6±1.1	17.2±0.1	20.4±1.1
DER++ [6]	74.4±1.1	91.5±0.2	92.1±0.2	44.2±1.1	47.9±1.5	54.7±2.2	15.3±0.2	19.7±1.5	27.0±0.7	4.5±0.3	10.1±0.3	17.6±0.5
IL2A [55]	90.2±0.1	92.7±0.1	93.9±0.1	54.7±0.5	56.0±0.4	58.2±1.2	18.2±1.2	19.7±0.5	22.4±0.2	5.5±0.7	8.1±1.2	11.6±0.4
Co ² L [9]	83.1±0.1	91.5±0.1	94.7±0.1	42.1±1.2	51.0±0.7	58.8±0.4	17.1±0.4	24.2±0.2	32.2±0.5	10.1±0.2	15.8±0.4	22.5±1.2
LUCIR [19]	73.2±0.1	87.2±0.2	90.1±0.1	27.9±1.2	33.5±0.7	37.8±0.5	8.6±1.3	19.5±0.7	16.9±0.5	7.6±0.5	9.6±0.7	12.5±0.7
CCIL [33]	86.4±0.1	92.8±0.2	94.0±0.1	50.5±0.2	55.3±0.5	59.8±0.3	18.5±0.3	19.1±0.4	20.5±0.3	5.6±0.9	7.0±0.5	15.2±0.5
BiC [46]	90.4±0.1	93.0±0.2	94.8±0.1	48.2±0.7	57.5±1.4	63.8±0.2	21.2±0.3	36.1±1.3	42.5±1.2	10.2±0.9	18.9±0.3	25.2±0.6
SSIL [2]	88.2±0.1	93.0±0.2	95.1±0.1	49.5±0.2	59.2±0.4	64.0±0.5	26.0±0.1	33.1±0.5	39.5±0.4	9.6±0.7	15.2±1.5	21.1±0.1
GSA	91.4±0.1	93.2±0.1	96.5±0.1	58.0±0.4	64.6±0.2	69.1±0.3	31.4±0.2	39.7±0.6	49.7±0.2	18.4±0.4	26.0±0.2	33.2±0.4

Table 2. Average forgetting rate. All numbers are the averages of 15 runs.

Method	MNIST (5 tasks)			CIFAR10 (5 tasks)			CIFAR100 (10 tasks)			TinyImageNet (100 tasks)		
	$M=0.1k$	$M=0.5k$	$M=1k$	$M=0.2k$	$M=0.5k$	$M=1k$	$M=1k$	$M=2k$	$M=5k$	$M=2k$	$M=4k$	$M=10k$
AGEM	32.5±5.9	30.1±4.2	32.0±2.9	36.1±3.8	43.2±4.3	48.1±3.4	78.6±2.1	77.5±1.3	78.3±1.2	73.9±0.2	78.9±0.2	74.1±0.3
ER	22.7±0.5	9.7±0.4	6.7±0.5	42.0±0.3	26.7±0.7	20.7±0.7	65.1±1.3	59.3±0.9	60.0±1.6	68.2±2.8	66.2±0.8	67.2±0.2
MIR [3]	22.3±0.5	9.0±0.5	5.7±0.9	40.0±1.6	25.9±0.7	24.5±0.5	24.5±0.3	21.4±0.3	21.0±0.1	61.1±3.2	60.9±0.3	59.5±0.3
GSS	26.1±2.2	17.8±5.22	10.5±6.7	75.5±1.5	65.9±1.6	54.9±2.0	73.4±4.2	69.3±3.1	70.9±2.9	72.8±1.2	72.6±0.4	71.5±0.2
ASER [40]	33.8±1.1	24.8±0.5	13.8±0.4	71.1±1.8	59.1±1.5	50.4±1.5	25.0±0.2	12.2±1.9	13.2±0.1	65.7±0.7	64.2±0.2	62.2±0.1
Rainbow-rt	10.1±0.1	4.7±0.4	2.4±0.5	20.4±0.3	18.1±0.4	15.3±0.8	25.5±0.5	19.3±0.4	13.3±0.5	25.5±0.6	23.2±0.3	20.0±0.4
ER-AML	23.1±0.1	9.7±0.4	6.4±0.5	50.9±0.3	40.1±0.5	34.2±0.8	51.5±0.8	49.2±0.5	38.7±0.6	47.4±0.5	43.2±0.3	41.0±0.5
GDumb	10.3±0.1	6.2±0.1	4.8±0.2	26.5±0.5	24.5±0.2	18.9±0.4	16.7±0.5	17.6±0.2	16.8±0.4	15.9±0.5	14.6±0.3	11.7±0.2
SCR	10.7±0.1	4.7±0.1	4.0±0.2	41.3±0.1	31.5±0.2	24.7±0.4	17.5±0.2	11.6±0.5	5.6±0.4	19.4±0.3	15.4±0.3	14.9±0.7
DER++ [6]	25.0±0.3	7.3±0.3	6.6±1.2	30.1±0.8	31.8±2.5	18.7±3.4	43.4±0.2	44.0±1.9	25.8±3.5	67.2±1.7	63.6±0.3	55.2±0.7
IL2A [55]	8.7±0.1	7.2±0.1	4.1±0.1	36.0±0.2	32.1±0.4	29.1±0.4	24.6±0.6	12.5±0.7	20.0±0.5	65.5±0.7	60.1±0.5	57.6±1.1
Co ² L [9]	14.7±0.2	7.1±0.1	3.1±0.1	32.0±0.1	21.0±0.3	16.9±0.2	16.9±0.4	16.6±0.6	9.9±0.7	60.5±0.5	52.5±0.9	42.5±0.8
LUCIR	20.3±0.1	8.5±0.1	7.8±0.1	63.5±0.5	55.3±0.2	46.5±1.2	60.0±0.1	47.5±0.9	44.3±0.7	46.4±0.7	42.2±0.9	37.6±0.7
CCIL	14.1±0.1	7.7±0.1	4.8±0.1	18.6±0.1	16.5±0.4	12.5±0.8	16.7±0.5	16.1±0.3	17.5±0.2	59.4±0.3	56.2±1.3	48.9±0.6
BiC	11.1±0.1	3.7±0.1	2.5±0.1	35.4±0.7	25.3±0.4	14.5±0.7	40.2±0.4	30.9±0.7	18.7±0.5	43.5±0.5	32.9±0.5	24.9±0.4
SSIL	11.3±0.1	2.7±0.1	2.8±0.1	36.0±0.7	29.6±0.4	13.5±0.4	40.1±0.5	33.9±1.2	21.7±0.8	44.4±0.7	36.6±0.7	29.0±0.7
GSA	8.1±0.1	2.5±0.1	1.4±0.1	23.5±0.2	12.6±0.4	10.0±0.3	33.2±0.6	22.8±0.4	8.7±0.3	35.5±0.3	25.8±0.4	16.9±0.6

Table 3. Ablation accuracy - average of 5 runs. Memory size M is 1k for CIFAR100 and 2k for TinyImageNet.

Dataset	no new loss	no previous loss	no cross loss	no X^{mix}	no CL imbalance
CIFAR100	30.1±0.6	12.7±0.2	29.9±0.4	30.1±0.1	29.7±0.3
TinyImageNet	17.5±0.2	4.7±0.5	17.2±0.4	17.5±0.5	16.9±0.2

it cannot be used for any-time inference. We thus removed this retraining operation and denote it as Rainbow-rt.

6.1. Architecture, Data Augmentation, Training Details and Evaluation Protocol

Architecture. For MNIST, GSA and baselines employ a fully-connected network with two hidden layers as the feature extractor h_θ , each comprising of 400 ReLU units. A linear layer of size [400, 10] is used as the classifier f_ϕ . For CIFAR10, CIFAR100, and TinyImageNet, we follow [6] and use ResNet18 (not pre-trained) as the feature extractor h_θ with around 11 million trainable parameters for GSA and

all baselines. Denoting C_{num} as the number of all classes, we employ a linear layer of size $[dim_h, C_{num}]$ as the classifier f_ϕ . For an input x , we use $F(x)$ to compute \mathcal{L}_{ce} .

Data Augmentation. To learn better features, we apply two data augmentations to each image in X^{new} and X^{mix} : *random-resized-crop* and *random-gray-scale*. For fair comparison, the same data augmentations are applied to all baselines (X^{new} and X^{buf}) to improve their performance, which results in an average of about 2% of improvement for all methods with no drop in performance for any baseline.

Training and hyperparameter settings. Like ER and many other online CL systems, GSA uses reservoir sampling for memory update. We follow [16] and train GSA with the Adam optimizer. We set the learning rate as 0.001 and fix the weight decay as 0.0001 for all settings. We follow [16] and set the batch size of X^{new} as 10 for all methods and X^{mix} as 64 for our GSA and X^{buf} also as 64 for baselines. We list other hyper-parameters in Appendix 6.

We use the official codes of baselines. Their source links

and default hyper-parameters are listed in Appendix 6. We run all methods with one epoch for each task.

Evaluation protocol. We first learn all tasks from the data stream for each dataset, and then test the final model using the test data of all tasks. We report the average accuracy of 15 random runs.

6.2. Results Analysis and Ablation Experiments

The results in Table 1 show that the best online CL baseline is SCR and the best class-imbalance baseline is BiC. Although BiC and SSIL were not originally designed for online CL, they perform well for online CL, better than all online CL baselines, which do not deal with data imbalance. This indicates that the data-balanced approaches in BiC and SSIL help. Our method GSA consistently outperforms them by very large margins as we also deal with the proposed CL imbalance. Note that OCM is not compared in the tables as it is not a competitor of GSA. GSA is actually complementary to OCM. For example, with the largest memory size for each dataset, OCM+GSA gives 96.5% on MNIST, 77.5% on CIFAR10, 53.7% on CIFAR100, and 35.7% on TinyImageNet and also outperforms OCM [16].

Forgetting rate. Table 2 shows that our GSA has substantially lower forgetting rates than baselines except GDumb, Rainbow-rt and SCR (in two datasets), but GDumb, Rainbow-rt, and SCR’s accuracy values are substantially lower than that of GSA (see Table 1). The calculation method of the average forgetting rate and **Training times** are given in Appendix 7.

Ablation. We conducted ablation experiments on GSA using two datasets, CIFAR100 and TinyImageNet. The results are given in Table 3.

(1). *Ablation study of training loss in GSA.* In the experiments “no new loss”, “no previous loss”, and “no cross loss”, we do not consider the first term (establishing the boundaries between new classes), the second term (maintaining the decision boundaries between previous classes), and the last term (establishing the cross-task boundaries) in Eq. 15 respectively. Table 3 shows that their performances are all poorer than GSA (Table 1). This is because in the first case, the model forgets the knowledge of previous classes, which causes a drastic drop in performance. The second experiment shows that our method benefits from considering all the new data in X^{new} . The third experiment shows that establishing cross-task boundaries is an important problem for CIL and can improve the overall performance further.

(2). *Ablation study of data imbalance and CL imbalance.* In the experiments “no balanced sampling,” we replace the sampling strategy for X^{mix} with random sampling from the memory buffer X^{buf} . In the experiment “no CL imbalance,” we do not use the \mathcal{L}_{GSA-CE} losses proposed in Eq. 12 and Eq. 14 and introduced into the final optimization objective in Eq. 15, but replace them with the simple cross-entropy



Figure 3. The gradient rate after the model has learned all classes of CIFAR100. A-PN is the gradient rate. The memory buffer size is 1000. Joint training means the model learns all classes together as one task in one epoch.

loss. Table 3 shows that all these incomplete GSA systems are poorer than the full GSA in Table 1. The poorer performance of “no balanced-sampling” is mainly due to data imbalance between new classes and previous classes in the buffer. Specifically, in the initial training process of the new task, the stored training data of the new task classes are usually fewer than that of the previous classes in the buffer and the cross-task class boundaries are not established well. The poorer performance of “no CL imbalance” is because the method does not consider different gradient rates across different classes, which again results in poorer performances.

GSA-CE loss mitigates the gradient imbalance. Fig. 3 shows that the GSA-CE loss clearly reduces the A-PN gradient rate imbalance than the method without it by dropping the A-NP rate that is bigger than -1 and improving the A-NP rate that is smaller than -1. And it makes the curve of A-PN rate more similar to the joint training curve.

Long-tail online CL experiments. See Appendix 8, where we will show that our method also outperforms the baselines.

7. Conclusion

This paper discussed the challenge of *cross-task class discrimination* (CTCD) and showed that the replay approach partially deals with the problem. However, the replay approach has a major bias in training, which is manifested by *gradient imbalance on the logits* and significantly limits the online CL performance and the ability to solve the CTCD problem. The paper then analyzed gradient imbalance from two perspectives: *data imbalance* introduced by replay and *CL imbalance* due to CL itself. After that, it proposed a new learning strategy and a new self-adaptive loss function to deal with the problems. Empirical evaluation demonstrated that the new approach GSA improves the baselines by large margins.

References

- [1] Davide Abati, Jakub Tomczak, Tijmen Blankevoort, Simone Calderara, Rita Cucchiara, and Babak Ehteshami Bejnordi. Conditional channel gated networks for task-aware continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3931–3940, 2020. [2](#)
- [2] 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. [1](#), [2](#), [7](#)
- [3] Rahaf Aljundi, Lucas Caccia, Eugene Belilovsky, Massimo Caccia, Min Lin, Laurent Charlin, and Tinne Tuytelaars. Online continual learning with maximally interfered retrieval. *arXiv preprint arXiv:1908.04742*, 2019. [2](#), [3](#), [7](#)
- [4] Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient based sample selection for online continual learning. *arXiv preprint arXiv:1903.08671*, 2019. [2](#), [7](#)
- [5] Jihwan Bang, Heesu Kim, YoungJoon Yoo, Jung-Woo Ha, and Jonghyun Choi. Rainbow memory: Continual learning with a memory of diverse samples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8218–8227, 2021. [6](#), [7](#)
- [6] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: a strong, simple baseline. *arXiv preprint arXiv:2004.07211*, 2020. [2](#), [7](#)
- [7] Lucas Caccia, Rahaf Aljundi, Nader Asadi, Tinne Tuytelaars, Joelle Pineau, and Eugene Belilovsky. Reducing representation drift in online continual learning. *arXiv preprint arXiv:2104.05025*, 2021. [2](#), [7](#)
- [8] Francisco M Castro, Manuel J Marín-Jiménez, Nicolás Guil, Cordelia Schmid, and Karteek Alahari. End-to-end incremental learning. In *ECCV*, 2018. [2](#)
- [9] Hyuntak Cha, Jaeho Lee, and Jinwoo Shin. Co2l: Contrastive continual learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9516–9525, October 2021. [7](#)
- [10] Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a gem. *arXiv preprint arXiv:1812.00420*, 2018. [2](#), [7](#)
- [11] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and M Ranzato. Continual learning with tiny episodic memories. In *ICML-2019*, 2019. [2](#), [3](#), [7](#)
- [12] 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. [2](#)
- [13] Cyprien de Masson d’Autume, Sebastian Ruder, Lingpeng Kong, and Dani Yogatama. Episodic memory in lifelong language learning. In *NeurIPS*, 2019. [2](#)
- [14] Mehrdad Farajtabar, Navid Azizan, Alex Mott, and Ang Li. Orthogonal gradient descent for continual learning. In *International Conference on Artificial Intelligence and Statistics*, pages 3762–3773. PMLR, 2020. [2](#)
- [15] Yiduo Guo, Wenpeng Hu, Dongyan Zhao, and Bing Liu. Adaptive orthogonal projection for batch and online continual learning. In *Proceedings of AAAI-2021*, 2022. [2](#)
- [16] Yiduo Guo, Bing Liu, and Dongyan Zhao. Online continual learning through mutual information maximization. In *International Conference on Machine Learning*, pages 8109–8126. PMLR, 2022. [2](#), [7](#), [8](#)
- [17] Yunhui Guo, Mingrui Liu, Tianbao Yang, and Tajana Rosing. Improved schemes for episodic memory-based lifelong learning. *Advances in Neural Information Processing Systems*, 33:1023–1035, 2020. [2](#)
- [18] Tyler L Hayes, Kushal Kafle, Robik Shrestha, Manoj Acharya, and Christopher Kanan. Remind your neural network to prevent catastrophic forgetting. *arXiv preprint arXiv:1910.02509*, 2019. [2](#)
- [19] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *CVPR*, pages 831–839, 2019. [2](#), [7](#)
- [20] Steven C. Y. Hung, Cheng-Hao Tu, Cheng-En Wu, Chien-Hung Chen, Yi-Ming Chan, and Chu-Song Chen. Compacting, picking and growing for unforgetting continual learning. In *NeurIPS*, 2019. [2](#)
- [21] Ronald Kemker and Christopher Kanan. FearNet: Brain-Inspired Model for Incremental Learning. In *ICLR*, 2018. [2](#)
- [22] Gyuhak Kim, Changnan Xiao, Tatsuya Konishi, Zixuan Ke, and Bing Liu. A theoretical study on solving continual learning. In *Advances in Neural Information Processing Systems*, 2022. [2](#)
- [23] 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. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017. [2](#)
- [24] Łukasz Korycki and Bartosz Krawczyk. Class-incremental experience replay for continual learning under concept drift. *arXiv preprint arXiv:2104.11861*, 2021. [2](#)
- [25] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. *Technical Report TR-2009*, University of Toronto, Toronto., 2009. [6](#)
- [26] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7:7, 2015. [6](#)
- [27] Yann LeCun, Corinna Cortes, and Christopher JC Burges. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998. [6](#)
- [28] Xilai Li, Yingbo Zhou, Tianfu Wu, Richard Socher, and Caiming Xiong. Learn to grow: A continual structure learning framework for overcoming catastrophic forgetting. In *ICML*, 2019. [2](#)
- [29] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient Episodic Memory for Continual Learning. In *NIPS*, pages 6470–6479, 2017. [2](#)
- [30] Zheda Mai, Ruiwen Li, Hyunwoo Kim, and Scott Sanner. Supervised contrastive replay: Revisiting the nearest class

- mean classifier in online class-incremental continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3589–3599, 2021. 2
- [31] Zheda Mai, Ruiwen Li, Hyunwoo Kim, and Scott Sanner. Supervised contrastive replay: Revisiting the nearest class mean classifier in online class-incremental continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 3589–3599, 2021. 7
- [32] Marc Masana, Xialei Liu, Bartłomiej Twardowski, Mikel Menta, Andrew D Bagdanov, and Joost van de Weijer. Class-incremental learning: survey and performance evaluation on image classification. *arXiv preprint arXiv:2010.15277*, 2020. 1
- [33] Sudhanshu Mittal, Silvio Galesso, and Thomas Brox. Essentials for class incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3513–3522, 2021. 2, 6, 7
- [34] Oleksiy Ostapenko, Mihai Puscas, Tassilo Klein, Patrick Jah-nichen, and Moin Nabi. Learning to remember: A synaptic plasticity driven framework for continual learning. In *CVPR*, pages 11321–11329, 2019. 2
- [35] Ameya Prabhu, Philip HS Torr, and Puneet K Dokania. Gdumb: A simple approach that questions our progress in continual learning. In *EECV*, pages 524–540, 2020. 2, 7
- [36] Jathushan Rajasegaran, Munawar Hayat, Salman Khan, Fahad Shahbaz, and Khan Ling Shao. Random path selection for incremental learning. In *NeurIPS*, 2019. 2
- [37] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *CVPR*, pages 2001–2010, 2017. 2
- [38] Ari Seff, Alex Beatson, Daniel Suo, and Han Liu. Continual learning in generative adversarial nets. *arXiv preprint arXiv:1705.08395*, 2017. 2
- [39] Joan Serrà, Dídac Surís, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. In *ICML*, 2018. 2
- [40] Dongsub Shim, Zheda Mai, Jihwan Jeong, Scott Sanner, Hyunwoo Kim, and Jongseong Jang. Online class-incremental continual learning with adversarial shapley value. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 9630–9638, 2021. 2, 7
- [41] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In *NIPS*, pages 2994–3003, 2017. 2
- [42] Albin Soutif-Cormerais, Marc Masana, Joost Van de Weijer, and Bartłomiej Twardowski. On the importance of cross-task features for class-incremental learning. *arXiv: 2106.11930*, 2021. 1
- [43] Johannes von Oswald, Christian Henning, João Sacramento, and Benjamin F Grewe. Continual learning with hypernetworks. *ICLR*, 2020. 2
- [44] Liyuan Wang, Xingxing Zhang, Kuo Yang, Longhui Yu, Chongxuan Li, Lanqing Hong, Shifeng Zhang, Zhenguo Li, Yi Zhong, and Jun Zhu. Memory replay with data compression for continual learning. *ICLR-2022*, 2022. 2
- [45] Chenshen Wu, Luis Herranz, Xialei Liu, Joost van de Weijer, Bogdan Raducanu, et al. Memory replay gans: Learning to generate new categories without forgetting. In *NIPS*, pages 5962–5972, 2018. 2
- [46] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 7
- [47] Shipeng Yan, Jiale Zhou, Jiangwei Xie, Songyang Zhang, and Xuming He. An em framework for online incremental learning of semantic segmentation. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 3052–3060, 2021. 2
- [48] Haiyan Yin, Ping Li, et al. Mitigating forgetting in online continual learning with neuron calibration. *Advances in Neural Information Processing Systems*, 34, 2021. 2
- [49] Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong Learning with Dynamically Expandable Networks. In *ICLR*, 2018. 2
- [50] Lu Yu, Bartłomiej Twardowski, Xialei Liu, Luis Herranz, Kai Wang, Yongmei Cheng, Shangling Jui, and Joost van de Weijer. Semantic drift compensation for class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6982–6991, 2020. 2
- [51] Guanxiong Zeng, Yang Chen, Bo Cui, and Shan Yu. Continual learning of context-dependent processing in neural networks. *Nature Machine Intelligence*, 1(8):364–372, 2019. 2
- [52] Juntao Zhang, Jie Zhang, Shalin Ghosh, Dawei Li, Serafettin Tasci, Larry Heck, Heming Zhang, and C.-C. Jay Kuo. Class-incremental learning via deep model consolidation. In *CVPR*, 2020. 2
- [53] Bowen Zhao, Xi Xiao, Guojun Gan, Bin Zhang, and Shu-Tao Xia. Maintaining discrimination and fairness in class incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13208–13217, 2020. 2
- [54] Hanbin Zhao, Hui Wang, Yongjian Fu, Fei Wu, and Xi Li. Memory efficient class-incremental learning for image classification. *IEEE Transactions on Neural Networks and Learning Systems*, 2021. 2
- [55] Fei Zhu, Zhen Cheng, Xu-yao Zhang, and Cheng-lin Liu. Class-incremental learning via dual augmentation. *Advances in Neural Information Processing Systems*, 34, 2021. 2, 7