

Joint Input and Output Coordination for Class-Incremental Learning

Shuai Wang^{1,2}, Yibing Zhan³, Yong Luo^{1,2*}, Han Hu⁴, Wei Yu^{1*},
Yonggang Wen⁵, Dacheng Tao⁵

¹Institute of Artificial Intelligence, School of Computer Science, Wuhan University, China

²Hubei LuoJia Laboratory, Wuhan, China

³JD Explore Academy, JD.com, Inc., China

⁴School of Information and Electronics, Beijing Institute of Technology, China

⁵College of Computing & Data Science, Nanyang Technological University, Singapore
wangshuai123@whu.edu.cn, zhanyibing@jd.com, luoyong@whu.edu.cn, hhu@bit.edu.cn,
yuwei@whu.edu.cn, ygwen@ntu.edu.sg, dacheng.tao@ntu.edu.sg

Abstract

Incremental learning is nontrivial due to severe catastrophic forgetting. Although storing a small amount of data on old tasks during incremental learning is a feasible solution, current strategies still do not 1) adequately address the class bias problem, and 2) alleviate the mutual interference between new and old tasks, and 3) consider the problem of class bias within tasks. This motivates us to propose a joint input and output coordination (JIOC) mechanism to address these issues. This mechanism assigns different weights to different categories of data according to the gradient of the output score, and uses knowledge distillation (KD) to reduce the mutual interference between the outputs of old and new tasks. The proposed mechanism is general and flexible, and can be incorporated into different incremental learning approaches that use memory storage. Extensive experiments show that our mechanism can significantly improve their performance.

1 Introduction

In recent years, incremental learning has attracted much attention since it can play an important role in a wide variety of fields, including unmanned driving [Santoso and Finn, 2022] and human-computer interaction [Tschandl *et al.*, 2020]. Incremental learning is nontrivial since the parameters of deep models in the old tasks are often destroyed in the process of learning new tasks. This leads to the occurrence of catastrophic forgetting [French and Chater, 2002]. How to well preserve past information and fully explore new knowledge has become a major challenge of incremental learning.

Existing incremental learning approaches mainly focus on memory storage replay [Ahn *et al.*, 2021; Li and Hoiem, 2017; Wu *et al.*, 2019; Rebuffi *et al.*, 2017; Yan *et al.*, 2021], model dynamic expansion [Serra *et al.*, 2018; Mallya and Lazebnik, 2018], and regularization constraints design [Aljundi *et al.*, 2019]. Memory store replay has been

*Corresponding authors: Yong Luo, Wei Yu.




Figure 1: An illustration of the class imbalance and mutual interference issues. The difference in the number of input data for each class between tasks and within tasks makes the weights of fully connected layers greatly biased (neuron size). The output scores of data from old tasks ($1, \dots, t-1$) on the classification heads of new task t should approximate zero, but may be much larger than zero (green solid line) after training the new task model. The output scores of data from the new task on the classification heads of old tasks may be inconsistent before (blue dotted line) and after (blue solid line) updating the old task models.

demonstrated to be very effective, and it alleviates the destruction of old task weights by storing past data or simulating human memory. However, due to the privacy restriction and limited memory, the data to be accessed from old tasks are often quite scarce. This makes incremental learning models suffer from severe inter-task class bias, or known as the class imbalance issue between old and new tasks.

There exist some recent approaches [Ahn *et al.*, 2021; Rebuffi *et al.*, 2017; Yan *et al.*, 2021] that alleviate the problem of class imbalance between old and new tasks by utilizing rescaling, balanced scoring, or softmax separating. Although these approaches can improve the performance to some extent, the problem of category imbalance still exists, since during the incremental learning progresses, the category imbalance becomes more severe as the number of sample categories continuously increase. Moreover, the mutual interference between old and new tasks has not been well addressed. That is, only the predictions in old tasks are tried to be maintained, and the output scores of old task data on the classification heads of new tasks are not well suppressed. The output consistency of new task data on old classification heads before and after updating the new task model is also not considered. Besides, none of the existing approaches deal with the

class bias within tasks. An illustration is shown in Figure 1.

In order to address these issues, we propose a joint input and output coordination (JIOC) mechanism, which enables incremental learning models to simultaneously alleviate the class imbalance and reduce the interference between the predictions of old and new tasks. Specifically, different weights are adaptively assigned to different input data according to their gradients for the output scores during the training of the new task and updating of the old task models. Then the outputs of old task data on new classification heads are explicitly suppressed and knowledge distillation (KD) [Menon *et al.*, 2021] is utilized for harmonization of the output scores based on the principle of human inductive memory [Williams, 1999; Redondo and Morris, 2011].

The main contributions are summarized as follows:

- We propose a joint input and output coordination mechanism for incremental learning. As far as we are concerned, this is the first work that simultaneously adjusts input data and output layer for incremental learning;
- We design an adaptive input weighting strategy. The samples of different classes are weighted according to their gradients of the output scores. This alleviates the class bias problem both in and between tasks.
- We develop an output coordination strategy, which maintains the outputs of new task data on the old task classification heads before and after training, and suppresses the outputs of old task data on the new task classification heads.

The proposed method is general and flexible, and can be utilized as a plug-and-play tool for existing incremental learning approaches that use memory storage. To demonstrate the effectiveness of our mechanism, we incorporate it into some recent or competitive incremental learning approaches on multiple popular datasets (CIFAR10-LT, CIFAR100-LT, CIFAR100 [Krizhevsky *et al.*, 2009], MiniImagNet [Vinyals *et al.*, 2016], TinyImageNet [Le and Yang, 2015] and Cub-200-2011 [Wah *et al.*, 2011]). The results show that we can consistently improve the existing approaches, and the relative improvement is more than 10% sometimes.

2 Related Work

2.1 Incremental Learning

Incremental learning [De Lange *et al.*, 2021] has received extensive attention in recent decades. In incremental learning, input data in new tasks are continuously used to extend the knowledge of existing models. This makes incremental learning manifest as a dynamic learning technique. An incremental learning model can be defined as one that meets the following conditions: (1) The model can learn useful knowledge from new task data; (2) The old task data that has been used to train the model does not need to be accessed or has a small amount of access; (3) It has a memory function for the knowledge that has been learned. The current study on incremental learning mainly focuses on domain incremental learning [Mirza *et al.*, 2022; Garg *et al.*, 2022; Mallya *et al.*, 2018], class-incremental learning [Ahn *et al.*, 2021; Rebuffi *et al.*, 2017; Yan *et al.*, 2021; Zhang *et al.*, 2020;

Liu *et al.*, 2021], and small sample incremental learning [Tao *et al.*, 2020; Cheraghian *et al.*, 2021].

There are many works on class-incremental learning (CIL), and most of these works overcome catastrophic forgetting by using knowledge distillation (KD) together with a small amount of old task data accessed. For example, DMC [Zhang *et al.*, 2020] utilizes separate models for the new and old classes and trains the two models by combining double distillation. SPB [Liu *et al.*, 2021] utilizes cosine classifier and reciprocal adaptive weights, and a new method of learning class-independent knowledge and multi-view knowledge is designed to balance the stability-plasticity dilemma of incremental learning.

Although the above approaches can achieve promising performance sometimes, none of them address class bias within tasks, nor adequately address class bias between old and new tasks. Therefore, we propose joint input and output coordination (JIOC) mechanism that enables incremental learning models to alleviate class imbalance and reduce interference between the predictions of old and new tasks.

2.2 Human Inductive Memory

The inductive memory method is a unique ability of human beings. It causes the memorized content to be induced according to different attributes or categories; Subsequently, these contents are memorized by different categories or attributes. As early as 1999, Williams et al. [Williams, 1999] investigated the relationship between memory for input and inductive learning of morphological rules relating to functional categories in a semi-artificial form of Italian. The ability to perform induction appears in the early age of human, while the underlying mechanisms remain unclear. Therefore, Fisher et al. [Fisher and Sloutsky, 2005] demonstrated that category- and similarity-based induction should result in different memory traces and thus different memory accuracy. Hayes et al. [Hayes *et al.*, 2013] examined the development of the relationship between inductive reasoning and visual recognition memory, and demonstrated it through two studies. Inspired by human inductive memory, Geng et al. [Geng *et al.*, 2020] proposed a Dynamic Memory Induction Network (DMIN) to further address the small-sample challenge. These examples of inductive memory inspire us to propose an output distribution coordination mechanism.

3 Method

3.1 Notations and Problem Setup

In CIL, data for new tasks are arriving constantly, which are represented as $\mathcal{D} = \{\mathcal{D}^1, \mathcal{D}^2, \dots, \mathcal{D}^t, \dots, \mathcal{D}^T\}$. The data in the t -th new task is $\mathcal{D}^t = \{(x_{i,j}^t, y_{i,j}^t)_{i=1,2,\dots,m;j=1,2,\dots,n_m}\}$, where m is the number of classes, n_m is the number of samples for the m -th category, x is the input data, and y is the corresponding data label. The number of samples may vary for different categories in the new task. When learning the t -th new task, we assume that there are a small amount of data stored for the old tasks, i.e.,

$$\mathcal{D}_{old}^t = \{(x_{i,j}^1, y_{i,j}^1), \dots, (x_{i,j}^{t-1}, y_{i,j}^{t-1})\}, \quad (1)$$

where $i = 1, 2, \dots, m$ and $j = 1, 2, \dots, n_{old}$, $n_{old} \ll n$. That is, the number of old data \mathcal{D}_{old}^t in the repository is much

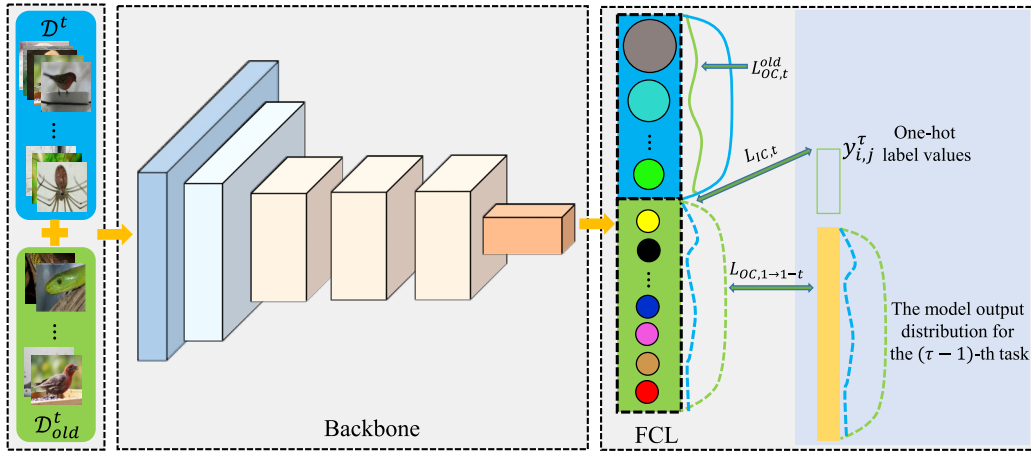


Figure 2: Overall structure of the proposed method. Firstly, the absolute gradient of the output scores is computed, based on the $\hat{p}_{i,j,k=i}^\tau$ and the $y_{i,j,k=i}^\tau$, to induce a weight for each sample, where the weights are adaptively updated during the training. Then $L_{OC,1 \rightarrow t-1}$ is employed to maintain the outputs for each old task. $L_{OC,1 \rightarrow t-1}$ is also utilized to make the outputs of new task data on old task classification heads after updating the old task models agree with those before the update. Finally, to suppress the outputs of old task data on new task classification heads, their output scores $\hat{p}_{i,j,k}^t$ are directly optimized to approach zeros (The solid blue line and solid green line represents the output distribution of the new task and old task data, respectively, on the new task classification head; The dashed blue line and dashed green line represents the output distribution of the new task and old task data, respectively, on the old task classification head).

smaller than that of \mathcal{D}^t . In CIL, a feature extractor $f(\cdot)$ (such as ResNet [He *et al.*, 2016]) and a fully connected layer (FCL) together with a *softmax* classifier is generally adopted, i.e.,

$$\mathbf{x}_{i,j}^\tau = f(x_{i,j}^\tau; \Theta), \quad (2)$$

$$\hat{\mathbf{p}}_{i,j}^\tau = \text{softmax}(\mathbf{x}_{i,j}^\tau; W), \quad (3)$$

where $\tau = \{1, 2, \dots, t\}$, Θ is the parameter of the feature extractor, W is the parameter of the classifier, and $\hat{\mathbf{p}}_{i,j}^\tau$ is a vector of output scores. When incremental learning proceeds to the t -th task, all the data in $\mathcal{D}^t \cup \mathcal{D}_{old}^t = \{(x_{i,j}^\tau, y_{i,j}^\tau), (\tau = 1, 2, \dots, t)\}$ are utilized for training, and the following cross-entropy loss is usually adopted:

$$L_{ce,t} = -\frac{1}{N_{old} + n_{new}} \sum_{i,j,k,\tau=1}^t y_{i,j,k}^\tau \log(\hat{p}_{i,j,k}^\tau), \quad (4)$$

where N_{old} is total number of stored data for old tasks, n_{new} is the total number of samples for the new task, and $\hat{p}_{i,j,k}^\tau$ is the output score at the k -th neuron.

3.2 Overview

According to the above problem setup, it can be seen that when performing incremental learning, only a limited number of samples from the old tasks will be retained. Due to the large number of samples in the new task, incremental learning suffers from the class imbalance issue between the old and new tasks. The class imbalance issue also exists within the new task, but this is ignored by existing CIL approaches [Ahn *et al.*, 2021; Rebuffi *et al.*, 2017; Yan *et al.*, 2021].

Therefore, we propose the joint input and output coordination (JIOC) mechanism, as shown in Figure 2, where we assign different weights to different input data according to their the absolute value of the gradient for output scores. In




Figure 3: A comparison of the SSIL approach (left) and the proposed output coordination (right). In SSIL, only the outputs of old task data on old classification heads are kept consistent before and after updating. We improve it by further enforcing the outputs of new task data on old classification heads to be consistent, and suppress the outputs of old task data on new classification heads.

addition, in order to prevent the mutual interference of output distributions between old and new tasks, we split the softmax layer inspired by the principle of human inductive memory. This is similar to the SSIL [Ahn *et al.*, 2021] approach, but has several significant differences, as shown in Figure 3: 1) for each of the old tasks, we utilize KD to maintain the output distribution of each task. To make the output scores of new task data on the classification heads of old tasks consistent, we also employ KD to enforce the outputs after updating the old task models agree with the scores before the update; 2) to suppress the outputs of old task data on the classification heads of new tasks, their ground-truth target values are directly set to be zero for training.

3.3 Input Coordination

As we know, the class imbalance issue may lead to significant bias in the learned weights of the fully connected layers [Li *et al.*, 2020]. Therefore, $\hat{p}_{i,j,k}^\tau$ may deviate greatly from its corresponding true value $p_{i,j,k}^\tau$, and hence it is necessary to

balance the weight of fully connected layers between tasks and within tasks.

Due to the severe bias in the weights of the fully connected layer, we propose to utilize the outputs of fully connected layer's previous layer to adjust the weights. Suppose that $\hat{\mathbf{q}}_{i,j}^\tau$ is the vector of the previous layer that outputs scores $\hat{p}_{i,j}^\tau$. The derivative of $L_{ce,t}$ w.r.t. $\hat{\mathbf{q}}_{i,j}^\tau$ (we refer to the supplementary material for the detailed calculation) can be given by:

$$\frac{\partial L_{ce,t}}{\partial \hat{\mathbf{q}}_{i,j}^\tau} = \begin{bmatrix} \hat{p}_{i,j,1}^\tau - y_{i,j,1}^\tau & \dots \\ \hat{p}_{i,j,k}^\tau - y_{i,j,k}^\tau & \dots \\ \hat{p}_{i,j,m}^\tau - y_{i,j,m}^\tau & \dots \end{bmatrix}. \quad (5)$$

Then the absolute value of the gradient of the output score for the input data when $k = i$ is:

$$\delta_{i,j,k=i}^\tau = |\hat{p}_{i,j,k=i}^\tau - y_{i,j,k=i}^\tau|. \quad (6)$$

When the number of data is large for a certain category, the model tends to bias to this category and thus the absolute value of the gradient in Eq. (6) tends to be small in the learning process. To alleviate the bias issue, we propose to regard the absolute value as the weight for the corresponding input sample and add it into the loss during the training. That is, smaller weights will be adaptively assigned to the samples of the category that has more input data, and hence the model would focus more on the category that has fewer samples.

Based on the above analysis, we utilize the absolute values $\delta_{i,j}^\tau$ of the gradient to induce a weight for each input data during the training. First of all, we incorporate the absolute value of the gradient of the input data into the traditional cross-entropy loss (Eq. (4)), i.e.,

$$L_{IC,t} = -\frac{1}{N_{old} + n_{new}} \sum_{i,j,k=i,\tau=1}^t y_{i,j,k}^\tau \delta_{i,j,k}^\tau \log(\hat{p}_{i,j,k}^\tau). \quad (7)$$

Then, we can use Eq. (7) to balance the loss of each category. In this way, the category weights of the fully connected layer can be balanced according to the absolute value $\delta_{i,j}$ of the gradient of each input data. It not only alleviates the category bias between old and new tasks in incremental learning, but also greatly reduces the within-task bias. The main procedure is summarized in Algorithm 1¹.

3.4 Output Coordination

According to the above analysis, it is necessary to keep the output distribution of the new task data \mathcal{D}^t on the old task classification heads consistently before and after updating the old task models². Also, it is necessary to suppress the output scores of the old task data \mathcal{D}_{old}^t on the classification heads of

¹In the entire algorithm pipeline, the outer loop and inner loop iterate *Total* and $\frac{m * n_m + N_{old}}{batchsize}$ times, respectively. We neglect the time complexity of Eq. (6), Eq. (7), as well as the parameter updates for Θ and W . The overall time complexity of the algorithm pipeline is $O(Total * \frac{m * n_m + N_{old}}{batchsize})$.

²During the updating of the $(t-1)$ old tasks, there are only $m \cdot (t-1)$ classification heads. This does not contain the classification heads for the t -th task.

Algorithm 1 Main procedure of input coordination.

Input: The data of the incremental learning model $\{\mathcal{D}_{old}^t, \mathcal{D}^t\}$; the feature extractor of the current model is $f(\cdot, \Theta)$; the parameter of the current fully-connected layer is W ;

Output: The updated parameters Θ and W ;

```

1: for epoch = 1; epoch < Total; epoch ++ do
2:   while batchsize loads  $\{\mathcal{D}_{old}^t, \mathcal{D}^t\}$  data do
3:     (1)  $\delta_{i,j}^\tau \leftarrow \{\hat{\mathbf{q}}_{i,j}^\tau\}$ , by using Eq. (6);
4:     (2)  $L_{IC,t} \leftarrow L_{ce,t}$  and  $\delta_{i,j}^\tau$ , by using Eq. (7);
5:     (3) According to the loss value  $L_{IC,t}$  obtained in
       the previous step, the parameters  $\Theta$  and  $W$  of the
       incremental learning model are updated.
6:   end while
7:   Return the updated  $\Theta$  and  $W$ .
8: end for
    
```

the new task (In Figure 1, this is to keep the blue solid line consistent with the dotted line, and make the green solid line approach to 0).

When the model trains the t -th task, we suppose that the output score of the data $\mathcal{D}^t \cup \mathcal{D}_{old}^t$ without going through *softmax* layer is given by $\hat{z}_{i,j,k}^\tau$. Before updating the old tasks models and training the t -th task, the output score of the data $\mathcal{D}^t \cup \mathcal{D}_{old}^t$ is $\hat{z}_{i,j,k}^\tau$. By considering the principle of human inductive memory, KD is used to enforce the output consistency of the new task data on the classification heads of each old task before and after updating the corresponding model, i.e.,

$$L_{OC,1 \rightarrow t-1} = \sum_{\tau=1}^{t-1} \left[\sum_{i,j,k} \rho_{KL}^\epsilon(\hat{z}_{i,j,k}^\tau, \hat{z}_{i,j,k}^\tau) \right], \quad (8)$$

where $\rho_{KL}^\epsilon(\cdot)$ is the distillation loss, and ϵ is a temperature scaling parameter.

The output of the old task data \mathcal{D}_{old}^t on the classification head of the new task can be adjusted according to:

$$L_{OC,t}^{old} = \frac{1}{n_{new}} \sum_{i,j,k} (\hat{p}_{i,j,k}^\tau - 0), \quad (9)$$

where $i \in \{1, \dots, m(t-1)\}$.

Although the principle of Eq. (8) is similar to the SSIL [Ahn *et al.*, 2021] approach, the output coordination mechanism proposed in this paper is different from the SSIL approach, as shown in Figure 3. Combining the output coordination loss $L_{IC,t}$ of Eq. (7), the overall loss function $L_{JIOC,t}$ of the method proposed can be obtained, i.e.,

$$L_{JIOC,t} = L_{IC,t} + \gamma_1 L_{OC,1 \rightarrow t-1} + \gamma_2 L_{OC,t}^{old}, \quad (10)$$

where $\gamma_1 \geq 0$ and $\gamma_2 \geq 0$ are trade-off hyper-parameters.

4 Experiment

4.1 Datasets and Evaluation Criteria

Datasets. In this paper, we not only validate the effectiveness of our method on unbalanced CIFAR10-LT and

Datasets	CIFAR10-LT	CIFAR100-LT
Training images	16,271	32,775
Classes	10	100
Max #{images}	5,000	500
Min #{images}	206	200
Imbalance factor	24	2.5

Table 1: The detailed information of long-tail imbalance datasets.

CIFAR100-LT datasets but also conduct corresponding validation on balanced CIFAR100 [Krizhevsky *et al.*, 2009], MiniImageNet [Vinyals *et al.*, 2016], TinyImageNet [Le and Yang, 2015], and Cub-200-2011 [Wah *et al.*, 2011] datasets. The CIFAR10 and CIFAR100 datasets both consist of 50,000 training images and 10,000 test images, with 10 and 100 categories respectively. To create unbalanced settings for the balanced datasets, we reduce the number of training samples for some classes. To ensure that our method is applicable to various settings³, we consider long-tail imbalances [Cui *et al.*, 2019], and a summarization of the dataset is reported in Table 1. The MiniImageNet dataset was excerpted from the ImageNet [Russakovsky *et al.*, 2015] dataset, and it contains 100 classes, each with 600 images of size 84×84 .

Typically, the training and test split of this dataset is 80 : 20. The TinyImageNet dataset is also a subset of the ImageNet [Russakovsky *et al.*, 2015] dataset and contains 200 classes, with each class containing 500 training images, 50 validation images, and 50 testing images. The Cub-200-2011 dataset is a bird dataset used for image classification. It covers 200 categories with a total of 11,788 images.

Evaluation Criteria. Following [Shi *et al.*, 2022], the average accuracy is used to measure the performance of the incremental learning algorithm, i.e.,

$$\bar{A} = \frac{1}{t} \sum_{\tau=1}^t A_{\tau}, \quad (11)$$

where A_{τ} is the accuracy of the τ -th task.

Baseline Protocol. The training sets of CIFAR10-LT is divided into $T = 5$ tasks, and the number of categories for each task is 2. The number of samples in the memory is fixed to be $N_{old} = \{1000, 2000\}$ during the incremental training. In the CIFAR100-LT, CIFAR100, MiniImageNet, TinyImageNet, and Cub-200-2011 datasets, the training tasks are divided into $T = 10$. The memory sizes of CIFAR100, MiniImageNet, and TinyImageNet dataset are also fixed as $N_{old} = \{1000, 2000\}$, and the memory size is chosen from $N_{old} = \{1000, 1500\}$. In addition, for the Cub-200-2011 dataset, the memory storage size fixed as $N_{old} = \{500\}$. The numbers of categories for each task of CIFAR10-LT, CIFAR100, MiniImageNet, TinyImageNet and Cub-200-2011 dataset are 10, 10, 10, 20 and 20, respectively. Regarding the memory storage samples related to our fusion algorithm, we all follow the existing algorithm [Rebuffi *et al.*, 2017; Yan *et al.*, 2021; Wang *et al.*, 2022].

³Since we use a small amount of old task data, the setup is slightly different from that of [Cui *et al.*, 2019].

Implementation details. Our method and all the compared approaches (BiC [Wu *et al.*, 2019], PODNet [Douillard *et al.*, 2020], COIL [Zhou *et al.*, 2021], SSIL [Ahn *et al.*, 2021], ICARL [Rebuffi *et al.*, 2017], DER [Yan *et al.*, 2021] and FOSTER [Wang *et al.*, 2022]) are implemented using PyCIL [Zhou *et al.*, 2023] and Pytorch [Paszke *et al.*, 2017].

On the experimental dataset, we used ResNet18 [He *et al.*, 2016] and ResNet32 as feature extractors respectively. ResNet32 is just used to further demonstrate that this mechanism can have some effectiveness in other network frameworks. In terms of parameter settings, we align with the original methods on PyCIL [Zhou *et al.*, 2023] to facilitate a fair comparison. Among these, the batch size is set to 128. Additionally, the SGD optimizer is used to gradually update the weights during incremental learning model training. The learning rate is initially set to be 0.1 and gradually decays. We run the training on two NVIDIA 3090RTX GPUs.

4.2 Results and Analysis

We incorporate the proposed JIOC strategy into existing class-incremental learning algorithms (ICARL [Rebuffi *et al.*, 2017], DER [Yan *et al.*, 2021], and FOSTER [Wang *et al.*, 2022]). The experimental results on different datasets are shown in Table 2 and Table 3.

Results on CIFAR10-LT and CIFAR100-LT. From the overall performance analysis of the Table 2, it can be seen that the ICARL, DER, and FOSTER algorithms on our created imbalanced datasets have been significantly improved. The relative improvements of ICARL_JIOC are 1.70%, 1.24%, 20.90%, 18.29%, 10.11% and 6.84% compared with the original ICARL algorithm. For the DER algorithm, our DER_JIOC improves it by 1.03%, 1.80%, 1.88%, 0.55%, 3.55% and 3.07%. In regard to the FOSTER algorithm, the relative improvements are the significant 3.51%, 3.66%, 3.02%, 10.10%, 2.20%, and 3.37% respectively. Compared with all the counterparts, the best performance is usually achieved by the proposed FOSTER_JIOC method. This not only demonstrates the effectiveness of our method on imbalanced datasets, but also further confirms its ability to alleviate catastrophic forgetting in other network frameworks.

Results on MiniImageNet, TinyImageNet, Cub-200-2011 and CIFAR100. We can observe from Table 3 that the mechanism proposed in this paper also has significant improvements on the MiniImageNet, TinyImageNet, Cub-200-2011 and CIFAR100 datasets. For example, our FOSTER_JIOC outperforms the original FOSTER algorithm by 4.28%, 6.24%, 4.07%, 2.06%, 38.60%, 6.91%, 14.90%, 1.67% and 2.82%, respectively. The best performance is also achieved by the proposed DER_JIOC and ICARL_JIOC, which are comparable. This further demonstrates that the proposed method not only alleviates catastrophic forgetting in class-imbalanced datasets but also has a forgetting-mitigation effect on normal data.

4.3 Ablation Studies


In this section, we first examine the effectiveness of the input and output coordination strategies separately, and then analyze how the output coordination strategy shows a more significant effect in mitigating forgetting compared to the SSIL.

Dataset	CIFAR10-LT		CIFAR100-LT			
Network	ResNet18				ResNet32	
T	5	5	10	10	10	10
N_{old}	1K	1.5K	1K	1.5K	1K	1.5K
Methods	Average accuracy					
BiC [Wu <i>et al.</i> , 2019]	65.46	66.48	39.91	42.88	60.03	61.64
PODNet [Douillard <i>et al.</i> , 2020]	76.00	71.23	38.38	41.33	47.15	49.21
SSIL* [Ahn <i>et al.</i> , 2021]	70.67	73.14	42.24	46.14	50.22	55.32
COIL [Zhou <i>et al.</i> , 2021]	79.60	79.64	48.98	50.72	58.31	58.31
ICARL [Rebuffi <i>et al.</i> , 2017]	70.18	73.25	38.47	42.04	50.17	54.67
ICARL_JIOC	71.37+1.19	74.16+0.91	46.51+8.04	49.73+7.69	55.24+5.07	58.41+3.74
DER [Yan <i>et al.</i> , 2021]	71.05	73.27	52.22	54.22	64.18	66.13
DER_JIOC	71.78+0.73	74.59+1.32	53.20+0.98	54.52+0.30	66.46+2.28	68.16+2.03
FOSTER [Wang <i>et al.</i> , 2022]	71.74	74.39	51.72	49.43	60.98	62.82
FOSTER_JIOC	73.95+2.52	77.20+2.72	53.28+1.56	54.42+4.99	62.32+1.34	64.94+2.12

Table 2: Results on the CIFAR10-LT and CIFAR100-LT datasets (* means our implementation).

Dataset	MiniImageNet		TinyImageNet		Cub-200-2011	CIFAR100				
Network	ResNet18								ResNet32	
T	10	10	10	10	10	10	10	10	10	10
N_{old}	1K	2K	1K	2K	0.5K	1K	2K	1K	2K	
Methods	Average accuracy									
BiC [Wu <i>et al.</i> , 2019]	52.95	52.10	52.95	54.18	29.13	34.43	45.77	58.01	61.93	
PODNet [Douillard <i>et al.</i> , 2020]	55.74	59.27	43.30	45.66	31.81	42.12	47.25	51.66	53.63	
SSIL* [Ahn <i>et al.</i> , 2021]	47.59	54.52	35.94	42.13	33.20	43.63	50.75	51.76	56.12	
COIL [Zhou <i>et al.</i> , 2021]	64.97	65.10	42.87	42.84	34.81	51.25	56.21	57.12	60.03	
ICARL [Rebuffi <i>et al.</i> , 2017]	53.43	60.90	33.63	39.24	30.10	38.29	44.32	52.20	56.35	
ICARL_JIOC	59.88+6.45	65.98+5.08	38.60+4.97	44.74+5.50	35.11+5.01	47.49+9.20	53.74+9.42	56.66+4.46	59.84+3.49	
DER [Yan <i>et al.</i> , 2021]	69.06	72.36	53.19	56.54	36.16	53.16	57.28	67.57	70.12	
DER_JIOC	70.06+1.00	73.08+0.72	56.37+3.18	57.63+1.09	38.82+2.26	56.29+3.13	59.31+2.03	70.45+2.88	71.88+1.76	
FOSTER [Wang <i>et al.</i> , 2022]	67.93	69.37	50.36	54.78	27.49	52.09	52.09	64.05	65.93	
FOSTER_JIOC	70.84+2.91	73.70+4.33	52.41+2.05	55.91+1.13	38.10+10.61	55.69+3.60	59.85+7.76	65.12+1.07	67.79+1.86	

Table 3: Results on the MiniImageNet, TinyImageNet, Cub-200, CIFAR100 datasets (* means our implementation).



(a) The results of using ResNet18 as the feature extractor on CIFAR100-LT. (b) The results of using ResNet18 as the feature extractor on CIFAR100. (c) The results of using ResNet32 as the feature extractor on CIFAR100.

Figure 4: A Compare the experimental results of SSIL with those of SSIL_OC (SSIL_OC is formed by integrating the proposed OC strategy into the SSIL algorithm. * means our implementation).

Study on the effectiveness of Input and Output Coordination. To validate the effectiveness of input and output co-

ordination strategies, we exclusively employed ResNet18 as the feature extractor on the CIFAR100-LT dataset, as shown

ICARL		Task										Avg
		1	2	3	4	5	6	7	8	9	10	
<i>AllTasks</i>	$L_{ce} + L_{KD}$	80.70	58.95	50.97	39.60	35.06	30.45	29.07	22.15	19.57	18.21	38.47
	$L_{IC} + L_{KD}$	80.70	63.50	55.80	45.18	40.84	36.88	33.46	27.21	25.24	22.79	43.16
	$L_{ce} + L_{OC}$	80.70	63.75	57.50	46.90	43.58	38.07	36.56	30.71	28.56	24.98	45.13
	L_{JIOC}	80.70	64.90	58.37	47.50	46.08	40.33	36.79	32.66	30.46	27.26	46.51
<i>NewTask</i>	$L_{ce} + L_{KD}$	80.70	61.50	76.1	70.20	81.9	70.01	74.7	73.00	77.10	71.20	73.64
	$L_{IC} + L_{KD}$	80.70	66.40	81.00	73.10	84.60	72.50	77.7	74.90	79.70	71.60	76.22
	$L_{ce} + L_{OC}$	80.70	64.40	80.30	71.70	83.20	69.70	76.50	72.40	78.80	70.60	74.83
	L_{JIOC}	80.70	66.40	80.40	70.90	84.20	71.50	80.30	73.10	78.40	72.40	75.83
<i>OldTasks</i>	$L_{ce} + L_{KD}$	—	56.40	38.40	29.40	23.35	22.52	19.15	14.89	12.38	12.32	25.42
	$L_{IC} + L_{KD}$	—	60.60	45.35	37.43	32.52	30.28	25.65	21.89	19.30	18.00	32.34
	$L_{ce} + L_{OC}$	—	63.10	46.10	38.63	33.67	31.74	29.90	24.76	22.28	19.91	34.45
	L_{JIOC}	—	63.40	47.35	39.70	36.55	34.1	29.53	26.89	24.46	22.24	36.02

Table 4: The results obtained by running with $N_{old} = 1000$, using ResNet18 as the feature extractor on the CIFAR100-LT dataset ($L_{ce} + L_{KD}$ is the loss function used by the ICARL algorithm).

in Tables 4. Furthermore, on the CIFAR100 dataset, we conducted experiments separately using ResNet18 and ResNet32 as feature extractors, as shown in Table 5 and Table 6 (we refer to the supplementary material for Table 5 and Table 6).

(1) The results are reported in Table 4, where we can see that the average accuracy of the ICARL algorithm with the proposed input coordination is 32.34 on the old tasks, 76.22 on the new task, and 43.16 overall. Compared with the original ICARL approach, the improvements are 27.22%, 3.50%, and 12.19%, respectively. The results from the old task, new task, and overall task in Table 5 and Table 6 also illustrate the competitiveness of the input coordination strategy. This demonstrates that the input coordination strategy can alleviate class imbalance in incremental learning. Besides, the ICARL algorithm only uses KD to maintain the output distribution on the old task classification heads. It does not take into account the human inductive memory mechanism for coordinating output distribution across different tasks. In Table 4, the ICARL algorithm achieved average performances of 34.45, 74.83, and 45.13 on the old tasks, new tasks, and overall, respectively. Our proposed output coordination mechanisms improved these performances by 35.52%, 1.62%, and 17.31%, respectively. It can be concluded that the input and output coordination strategy proposed in this paper yields significant improvements, whether applied to the CIFAR100 dataset, the CIFAR100-LT dataset, or different network architectures with varying depths.

(2) In the CIFAR100-LT dataset, the input coordination strategy demonstrates notable enhancements in the outcomes for old tasks (27.22%), new tasks (3.50%), and overall task (12.19%) performance when compared to the original ICARL algorithm, as shown in Table 4. Similarly, in the CIFAR100 dataset, the input coordination strategy further improves the original ICARL algorithm by 42.05%, 3.61%, and 19.51% on the old tasks, new tasks, and overall tasks, as shown in Table 5. According to the description and corresponding improvement effects of CIFAR100-LT and CIFAR100, it can be seen that the input coordination strategy has a good regulating effect on the imbalance of old

and new categories.

Experiment Comparison between Output Coordination Strategy and the SSIL. To quantitatively analyze the differences between the output coordination strategy and the SSIL, we conducted corresponding experimental results based on different feature extractors and datasets, including class-imbalanced and balanced datasets, as shown in Figure 4. From the results in Figure 4, it is evident that using the output distribution coordination strategy leads to significant improvements in each stage task on class-imbalanced datasets and in deep feature networks (ResNet32). This also indicates that the output distribution coordination strategy enables new task data to maintain consistent output distributions on the old task classification head and suppresses old task data on the new task classification head during the incremental learning process. This avoidance of interference between new and old tasks is achieved. However, SSIL does not avoid interference from old task outputs, which results in its performance being inferior to the output distribution coordination strategy.

5 Conclusion

Although the existing approaches address the class bias issue in class-incremental learning (CIL) to a certain extent by scaling and dividing the softmax layer, they all ignore the bias within the task. In addition, the mutual interference between old and new tasks has not been well resolved. Therefore, we propose a joint input and output coordination (JIOC) mechanism to enable incremental learning models to simultaneously reduce the interference between predictions for these tasks and alleviate the class imbalance issue between and within tasks. From the extensive experiments on multiple popular datasets, we observe significant improvements when incorporating the proposed mechanism into the existing CIL approaches that utilize memory storage. In the future, we intend to design more sophisticated strategies to reweight the inputs, and develop a general end-to-end framework for CIL.

Acknowledgements

This work is supported in part by the National Natural Science Foundation of China (Grant No. U23A20318), the Special Fund of Hubei LuoJia Laboratory under Grant 220100014, the Fundamental Research Funds for the Central Universities (No. 2042024kf0039), the National Research Foundation Singapore and DSO National Laboratories under the AI Singapore Programme (AISG Award No: AISG2-GC-2023-006), and the CCF-Zhipu AI Large Model Fund OF 202224.

References

- [Ahn *et al.*, 2021] 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.
- [Aljundi *et al.*, 2019] Rahaf Aljundi, Klaas Kelchtermans, and Tinne Tuytelaars. Task-free continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11254–11263, 2019.
- [Cheraghian *et al.*, 2021] Ali Cheraghian, Shafin Rahman, Sameera Ramasinghe, Pengfei Fang, Christian Simon, Lars Petersson, and Mehrtaash Harandi. Synthesized feature based few-shot class-incremental learning on a mixture of subspaces. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8661–8670, 2021.
- [Cui *et al.*, 2019] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9268–9277, 2019.
- [De Lange *et al.*, 2021] Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Aleš Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *IEEE transactions on pattern analysis and machine intelligence*, 44(7):3366–3385, 2021.
- [Douillard *et al.*, 2020] Arthur Douillard, Matthieu Cord, Charles Ollion, Thomas Robert, and Eduardo Valle. Podnet: Pooled outputs distillation for small-tasks incremental learning. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*, pages 86–102. Springer, 2020.
- [Fisher and Sloutsky, 2005] Anna V Fisher and Vladimir M Sloutsky. When induction meets memory: Evidence for gradual transition from similarity-based to category-based induction. *Child development*, 76(3):583–597, 2005.
- [French and Chater, 2002] Robert M French and Nick Chater. Using noise to compute error surfaces in connectionist networks: A novel means of reducing catastrophic forgetting. *Neural computation*, 14(7):1755–1769, 2002.
- [Garg *et al.*, 2022] Prachi Garg, Rohit Saluja, Vineeth N Balasubramanian, Chetan Arora, Anbumani Subramanian, and CV Jawahar. Multi-domain incremental learning for semantic segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 761–771, 2022.
- [Geng *et al.*, 2020] Ruiying Geng, Binhua Li, Yongbin Li, Jian Sun, and Xiaodan Zhu. Dynamic memory induction networks for few-shot text classification. *arXiv preprint arXiv:2005.05727*, 2020.
- [Hayes *et al.*, 2013] Brett K Hayes, Kristina Fritz, and Evan Heit. The relationship between memory and inductive reasoning: Does it develop? In *Developmental Psychology*, volume 44, pages 848–860, 2013.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [Krizhevsky *et al.*, 2009] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [Le and Yang, 2015] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015.
- [Li and Hoiem, 2017] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017.
- [Li *et al.*, 2020] Yu Li, Tao Wang, Bingyi Kang, Sheng Tang, Chunfeng Wang, Jintao Li, and Jiashi Feng. Overcoming classifier imbalance for long-tail object detection with balanced group softmax. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10991–11000, 2020.
- [Liu *et al.*, 2021] Yaoyao Liu, Bernt Schiele, and Qianru Sun. Adaptive aggregation networks for class-incremental learning. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 2544–2553, 2021.
- [Mallya and Lazechnik, 2018] Arun Mallya and Svetlana Lazechnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 7765–7773, 2018.
- [Mallya *et al.*, 2018] Arun Mallya, Dillon Davis, and Svetlana Lazechnik. Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 67–82, 2018.
- [Menon *et al.*, 2021] Aditya K Menon, Ankit Singh Rawat, Sashank Reddi, Seungyeon Kim, and Sanjiv Kumar. A statistical perspective on distillation. In *International Conference on Machine Learning*, pages 7632–7642. PMLR, 2021.
- [Mirza *et al.*, 2022] M Jehanzeb Mirza, Marc Masana, Horst Possegger, and Horst Bischof. An efficient domain-incremental learning approach to drive in all weather conditions. In *Proceedings of the IEEE/CVF Conference on*

- Computer Vision and Pattern Recognition*, pages 3001–3011, 2022.
- [Paszke *et al.*, 2017] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- [Rebuffi *et al.*, 2017] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2001–2010, 2017.
- [Redondo and Morris, 2011] Roger L Redondo and Richard GM Morris. Making memories last: the synaptic tagging and capture hypothesis. *Nature Reviews Neuroscience*, 12(1):17–30, 2011.
- [Russakovsky *et al.*, 2015] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115:211–252, 2015.
- [Santoso and Finn, 2022] Fendy Santoso and Anthony Finn. A data-driven cyber-physical system using deep-learning convolutional neural networks: Study on false-data injection attacks in an unmanned ground vehicle under fault-tolerant conditions. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 53(1):346–356, 2022.
- [Serra *et al.*, 2018] Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. In *International Conference on Machine Learning*, pages 4548–4557. PMLR, 2018.
- [Shi *et al.*, 2022] Yujun Shi, Kuangqi Zhou, Jian Liang, Zihang Jiang, Jiashi Feng, Philip HS Torr, Song Bai, and Vincent YF Tan. Mimicking the oracle: an initial phase decorrelation approach for class incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16722–16731, 2022.
- [Tao *et al.*, 2020] Xiaoyu Tao, Xiaopeng Hong, Xinyuan Chang, Songlin Dong, Xing Wei, and Yihong Gong. Few-shot class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12183–12192, 2020.
- [Tschandl *et al.*, 2020] Philipp Tschandl, Christoph Rinner, Zoe Apalla, Giuseppe Argenziano, Noel Codella, Allan Halpern, Monika Janda, Aimilios Lallas, Caterina Longo, Josep Malvehy, et al. Human-computer collaboration for skin cancer recognition. *Nature Medicine*, 26(8):1229–1234, 2020.
- [Vinyals *et al.*, 2016] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. *Advances in neural information processing systems*, 29, 2016.
- [Wah *et al.*, 2011] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011.
- [Wang *et al.*, 2022] Fu-Yun Wang, Da-Wei Zhou, Han-Jia Ye, and De-Chuan Zhan. Foster: Feature boosting and compression for class-incremental learning. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXV*, pages 398–414. Springer, 2022.
- [Williams, 1999] John N Williams. Memory, attention, and inductive learning. *Studies in Second Language Acquisition*, 21(1):1–48, 1999.
- [Wu *et al.*, 2019] 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*, pages 374–382, 2019.
- [Yan *et al.*, 2021] Shipeng Yan, Jiangwei Xie, and Xuming He. Der: Dynamically expandable representation for class incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3014–3023, 2021.
- [Zhang *et al.*, 2020] Junting Zhang, Jie Zhang, Shalini Ghosh, Dawei Li, Serafettin Tasci, Larry Heck, Heming Zhang, and C-C Jay Kuo. Class-incremental learning via deep model consolidation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1131–1140, 2020.
- [Zhou *et al.*, 2021] Da-Wei Zhou, Han-Jia Ye, and De-Chuan Zhan. Co-transport for class-incremental learning. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 1645–1654, 2021.
- [Zhou *et al.*, 2023] Da-Wei Zhou, Fu-Yun Wang, Han-Jia Ye, and De-Chuan Zhan. Pycil: a python toolbox for class-incremental learning. *SCIENCE CHINA Information Sciences*, 66(9):197101–, 2023.