

Full Length Article

PVBF: A framework for mitigating parameter variation imbalance in online continual learning

Zelin Tao ^{a,d}, Hao Deng ^{a,d,*}, Mingqing Liu ^b, Lijun Zhang ^c, Shengjie Zhao ^{a,d}

^a School of Computer Science and Technology, Tongji University, Shanghai, 201804, China

^b College of Electronic and Information Engineering, Tongji University, Shanghai, 201804, China

^c National Key Laboratory for Novel Software Technology, Nanjing University, Nanjing, 210023, China

^d Engineering Research Center of Key Software Technologies for Smart City Perception and Planning, Ministry of Education, Shanghai, 201804, China

ARTICLE INFO

Keywords:

Online continual learning

Catastrophic forgetting

ABSTRACT

Online continual learning (OCL), which enables AI systems to adaptively learn from non-stationary data streams, is commonly achieved using experience replay (ER)-based methods that retain knowledge by replaying stored past during training. However, these methods face challenges of prediction bias, stemming from deviations in parameter update directions during task transitions. This paper identifies parameter variation imbalance as a critical factor contributing to prediction bias in ER-based OCL. Specifically, using the proposed parameter variation evaluation method, we highlight two types of imbalance: correlation-induced imbalance, where certain parameters are disproportionately updated across tasks, and layer-wise imbalance, where output layer parameters update faster than those in preceding layers. To mitigate the above imbalances, we propose the Parameter Variation Balancing Framework (PVBF), which incorporates: (1) a novel method to compute parameter correlations with previous tasks based on parameter variations, (2) an encourage-and-consolidate (E&C) method utilizing parameter correlations to perform gradient adjustments across all parameters during training, (3) a dual-layer copy weights with reinit (D-CWR) strategy to slowly update output layer parameters for frequently occurring sample categories. Experiments on short and long task sequences demonstrate that PVBF significantly reduces prediction bias and improves OCL performance, achieving up to 47% higher accuracy compared to existing ER-based methods.

1. Introduction

Learning constitutes the cornerstone of intelligent systems, enabling their adaptation to dynamic environments. Humans exemplify this adaptability through their ability to continuously acquire and integrate new knowledge while retaining prior experiences (Wang, Zhang, Su, & Zhu, 2024). In contrast, artificial intelligence systems, particularly those based on deep neural networks (DNNs), face a significant limitation known as catastrophic forgetting, where the acquisition of new information leads to the erosion of previously learned knowledge (Van de Ven & Tolias, 2019), thereby hindering their capacity for sequential learning. To mitigate catastrophic forgetting and enable human-like learning capabilities, *continual learning* (CL) methods have been extensively studied (Chaudhry, Dokania, Ajanthan, & Torr, 2018; Kirkpatrick et al., 2017; Li & Hoiem, 2017). Due to privacy concerns or resource constraints (Mai et al., 2022), training data in more realistic online environments is presented as dynamic, one-pass streaming data (Soutif-Cormerais et al., 2023; Wang et al., 2024).

Thus, CL models need to not only address catastrophic forgetting but also tackle the challenge of insufficient training, which can lead to prediction bias. Existing effective methods primarily employ strategies such as replaying past samples or representations, commonly referred to as experience replay (ER)-based methods (Chaudhry et al., 2019; Isele & Cosgun, 2018; Pellegrini, Graffieti, Lomonaco, & Maltoni, 2020; Rebuffi, Kolesnikov, Sperl, & Lampert, 2017). In recent years, research on ER-based methods has been considered crucial for enabling DNNs to continuously acquire new knowledge in dynamic and non-stationary sequential tasks. However, existing ER-based methods still suffer from significant prediction bias when applied to online continual learning (OCL) in sequential tasks. To address this issue, we propose an ER-based framework with bias-correction strategies to improve the OCL performance of the models.

Prediction bias is strongly linked to parameter variations during training (Sun, 2019). In OCL scenarios, the non-stationary nature of sequential tasks poses additional challenges, leading to more intricate

* Corresponding author at: School of Computer Science and Technology, Tongji University, Shanghai, 201804, China.

E-mail address: denghao1984@tongji.edu.cn (H. Deng).

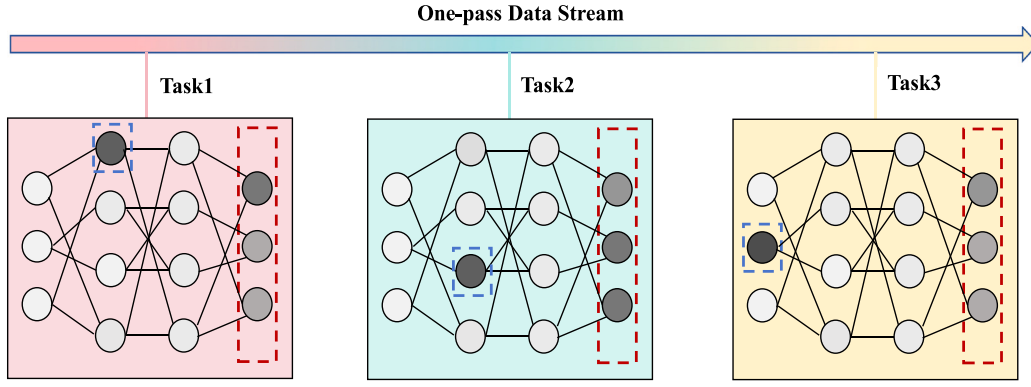


Fig. 1. Schematic diagram of parameter variation imbalance in OCL. An exemplary four-layer network is trained through a data stream with three tasks. The shading of the neurons indicates the magnitude of their variations during the training of the task, with darker colors representing greater variations. The blue dashed boxes present the correlation-induced imbalance, where a small subset of parameters exhibits significantly larger variations during certain tasks. The red dashed boxes highlight the issue of layer-wise imbalance, where the average variation of the final linear layer in the network is significantly higher than that of adjacent layers.

and uneven parameter variations compared to conventional training settings (Castro, Marin-Jimenez, Guil, Schmid, & Alahari, 2018). Existing ER-based methods primarily aim to mitigate task interference or preserve learned knowledge (Buzzega, Boschini, Porrello, Abati, & Calderara, 2020; Caccia et al., 2022), but they lack explicit evaluation of parameter variations. As a result, these methods fail to effectively link parameter updates with task-specific requirements, limiting their ability to guide precise gradient adjustments and address prediction bias. In this work, as illustrated in Fig. 1, we investigate the issue of parameter variation imbalance in OCL and identify two key phenomena. First, during training on OCL tasks, subsets of parameters undergo disproportionately large variations due to their strong correlations with task-specific memorization. We term this the correlation-induced imbalance, where these parameters are improperly optimized for individual tasks, leading to uneven learning and interference across tasks. Second, the parameters of the output classifier exhibit significantly larger variations compared to adjacent layers, an issue we denote as the layer-wise imbalance. Our findings highlight that these imbalances, arising from uneven gradient updates across parameters, are critical contributors to prediction bias in OCL scenarios and underscore the need for systematic investigations into these dynamics.

For the correlation-induced imbalance issue, several CL approaches have investigated it, particularly from the perspective of parameter isolation. These methods suggest that different parameters of neural networks exhibit varying levels of correlation for memorizing previous tasks (Fernando et al., 2017; Mallya & Lazebnik, 2018). With CL parameter isolation methods, masks are typically used to freeze or adjust parameters closely related to previous learning tasks (Hurtado, Raymond, & Soto, 2021; Mallya, Davis, & Lazebnik, 2018), demonstrating significant effectiveness in improving overall model accuracy. However, these methods require allocating distinct sets of parameters for each task, leading to increased storage consumption as the number of tasks grows, which makes them less suitable for the dynamic, non-stationary sequence tasks targeted by OCL. In contrast, ER-based methods integrate new and old knowledge by mixing input samples to correct the overall parameter variation direction, effectively mitigating prediction bias while maintaining relatively low resource consumption. However, parameter updates in ER-based OCL methods may still deviate from the optimal direction during task switches due to the non-independent and identically distributed (IID) nature of the learning samples. One of the underlying causes is the insufficient understanding of the correlation between parameter variations and previous task memory. Existing work lacks in-depth investigation of identifying crucial parameters closely related to specific tasks and developing effective strategies to prevent these parameters from being improperly updated.

For the layer-wise imbalance issue, it stems from the observation that during the short task sequence OCL training process, parameters

in the output classifier, i.e., the output layer parameters, tend to update more rapidly compared to other parameters (Chrysakis & Moens, 2023). This phenomenon can lead to a significant imbalance in parameter variations which further leads to prediction bias in the model (Wu et al., 2019). Existing research primarily focuses on improving the output layer individually based on experience replay (Pellegrini et al., 2020), e.g., adding a bias correction layer after the output layer (Wu et al., 2019), using a surrogate classifier instead of the output layer during training (Chrysakis & Moens, 2023), and developing unique update strategies for the output layer (Lomonaco & Maltoni, 2017). For instance, CWR* is a commonly adopted strategy for prediction bias correction by performing frequency-based magnitude adjustment in the output layer (Lomonaco, Maltoni, Pellegrini, et al., 2020). However, the first two methods inevitably increase the network complexity and the associated training costs, while the latter may introduce additional time expenditure and lead to suboptimal results. Designing an effective update strategy for the output layer that better adapts to dynamically non-stationary tasks remains a challenge.

In this paper, we propose a Parameter Variation Balancing Framework (PVBf), which updates overall model parameters and output layer parameters at different stages to address the above two issues in OCL. To mitigate the correlation-induced imbalance, PVBf introduces Parameter Correlation Calculation (ParamCC) to quantify the correlation between each parameter and previous tasks. Building on this correlation measure, we propose Encourage and Consolidate (E&C), a strategy that assigns adaptive gradient descent coefficients to parameters. This approach encourages the network to rapidly update parameters with low correlation to previous tasks while consolidating the memory of parameters with high correlation, thereby mitigating the prediction bias caused by the correlation-induced imbalance. To tackle layer-wise imbalance, we propose Dual-layer Copy Weights with Reinit (D-CWR), an improvement over the CWR* strategy. Inspired by human memory mechanisms (sensory memory, short-term memory, and long-term memory), D-CWR employs a two-stage consolidation process (Nairne & Neath, 2003), which can effectively mitigate the prediction bias caused by the output classifier in short task sequence OCL scenarios. The main contributions of this paper can be summarized as follows.

- It provides an evaluation method of parameter variations by capturing parameters' relative changes during training. On this basis, we identify parameter variation imbalance in OCL through two aspects: correlation-induced imbalance and layer-wise imbalance, which provides prerequisites for optimized learning strategy design.
- It proposes a two-phase approach to mitigate correlation-induced imbalance at the overall parameter level of neural networks.

First, Parameter Correlation Calculate (ParamCC) quantifies the correlation between parameter variations and the memory of previously learned tasks. Then, Encourage and Consolidate (E&C) strategy adjusts gradients by considering correlations as divisors to fine-tune the parameter updates. Together, parameter variation direction can be effectively corrected without huge network complexity increase.

- It proposes a bio-inspired Dual-layer Copy Weights with Re-init (D-CWR) method to mitigate layer-wise imbalance. Drawing inspiration from the human memory mechanism, D-CWR employs memory consolidations at two layers, i.e., from sensory to short-term and from short-term to long-term memory, which further avoids forgetting due to rapid parameter updates in the output classifier over existing methods.
- It presents a Parameter Variation Balancing Framework (PVBf) by integrating ParamCC, E&C and D-CWR. Experiments conducted in both short and long task sequence OCL scenarios show that PVBf achieves an average accuracy improvement of 31%–47% over the ER method. Especially, it reaches 97.5% of the IID method's accuracy using only 500 replay samples on the MiniImageNet dataset. PVBf also exhibits general outperformance in offline CL scenarios among classical methods.

2. Related work

In this section, we briefly review three key categories of research relevant to the current work: traditional CL methods, ER-based methods, and output correction methods. The strengths and weaknesses of each method are briefly discussed.

2.1. CL methods

Contemporary CL methodologies predominantly fall into three categories (Van de Ven & Tolias, 2019). Regularization-based approaches penalize parameter updates to enforce convergence within a shared representation space across diverse tasks (Chaudhry, Dokania, et al., 2018; Kirkpatrick et al., 2017; Zenke, Poole, & Ganguli, 2017). These approaches tackle the problem of catastrophic forgetting by constraining the parameter update methods. However, they face challenges in balancing between stability and plasticity, often resulting in suboptimal performance and high computational cost. Memory-based strategies integrate prior task knowledge through sample or representation-based memory adjustments during training (Bonnicelli, Boschini, Porrello, Spampinato, & Calderara, 2022; Chaudhry, Ranzato, Rohrbach, & Elhoseiny, 2018; Lopez-Paz & Ranzato, 2017; Riemer et al., 2018; Shin, Lee, Kim, & Kim, 2017; Wang, Zhou, Ye, & Zhan, 2022). While these methods effectively retain valuable knowledge across tasks, a capability that has been validated in question generation studies (Yuan, Yin, et al., 2022), they require storing a subset of training samples and may introduce prediction bias to some extent due to issues with the training data distribution. Dynamic structures-based approaches adapt network architectures to ensure task-specific parameter isolation, accommodating the integration of novel tasks (Li & Hoiem, 2017; Lomonaco et al., 2020; Von Oswald, Henning, Grewe, & Sacramento, 2019; Yoon, Yang, Lee, & Hwang, 2017). These approaches address the issue of interference by isolating task-specific knowledge, but they struggle with computational efficiency and spatial scalability as the number of tasks increases. Recently, prompt-based methods have also been included in CL, demonstrating promising performance in domains such as automatic program repair (Yuan, Zhang, et al., 2022). These approaches often rely heavily on large-scale pretrained models. Thus, methods in OCL often combine with memory-based strategies. However, challenges remain in striking a balance between memory usage, computational efficiency, and maintaining performance on non-stationary data streams.

2.2. ER-based methods

Existing effective approaches in OCL typically adopt strategies that involve replaying samples or representations. To mitigate catastrophic forgetting caused by task changes, iCaRL (Rebuffi et al., 2017) combines distillation loss with binary cross-entropy, classifying samples based on nearest-class prototypes computed from buffered data representations, which is suitable for class incremental learning scenarios where each task is sufficiently trained. However, in OCL scenarios, iCaRL often underperforms due to insufficient training on newer tasks. ER (Chaudhry et al., 2019) employs a fixed-size replay buffer, randomly replaying a subset of samples. Despite its simplicity, ER faces challenges in maintaining performance when learning from both replayed samples and data stream samples simultaneously. To address this issue, GDumb (Prabhu, Torr, & Dokania, 2020) maintains a class-balanced memory pool and trains the model exclusively on these samples, although the size of the memory pool often constrains its effectiveness. MIR (Aljundi, Belilovsky, et al., 2019) introduces an alternative improvement to ER by selecting samples that maximize the increase in loss during replay. This method further reduces prediction bias, albeit at the cost of increased computational burden. For loss calculation, DER++ (Buzzega et al., 2020) employs distillation loss on logits to enforce consistency over time, while ACE (Caccia et al., 2022) mitigates sudden representation changes using an asymmetric update rule. Although DER++ and ACE are promising in stabilizing learned knowledge, they fail to resolve the issue of spatially or structurally imbalanced parameter updates across the network, which can lead to significant prediction bias.

2.3. Output correction methods

In OCL scenarios, prediction bias is closely related to rapid updates of the output classifier during backpropagation (Wu et al., 2019). OBC (Chrysakis & Moens, 2023) independently optimizes output classifier to correct significant prediction bias during training. AR1* (Lomonaco et al., 2020) combines latent replay methods with optimized output classifier updates from CWR* to enhance performance. The CWR* method aligns the parameters in the output classifier with individual categories. It adjusts the update magnitude of these parameters based on the ratio of the frequency of data from a particular class in past occurrences to its frequency in a single training iteration, thereby correcting prediction bias. These strategies address the bias in the output classifier but do not simultaneously correct the bias in other parameters, leaving room for further improvement.

Although the three types of methods discussed above do not directly address the issue of parameter variation imbalance, they provide foundational ideas that have informed our approach to solving this problem. ER effectively mitigates catastrophic forgetting caused by non-IID data distributions in streaming tasks, while asymmetric cross entropy (ACE) helps achieve more accurate and balanced gradient updates. These strategies contribute to balancing parameter variations during OCL training to some extent, and thus, we integrate these concepts into our proposed framework. Furthermore, previous work on special handling of the output classifier has inspired our approach to addressing the issue of layer-wise imbalance.

3. Methodology

In this section, a detailed explanation of parameter variation imbalance is first presented, including its two specific forms and their manifestations in the OCL scenario. The proposed PVBf is then introduced, along with its applicable learning settings. Finally, detailed design and implementations of innovative methods incorporated into PVBf, namely ParamCC, E&C, and D-CWR, are described, respectively.

3.1. Parameter variation evaluate

In OCL, model parameters continuously evolve with the training of various tasks, and the ultimate result of this variation determines the model's adaptability across the entire online data stream. Therefore, we aim to calculate the variations of its parameters upon the completion of training for each task in the OCL context, revealing the patterns of parameter variation.

To record the parameters optimized through task k (where $k \in \{1, \dots, K\}$ and K denotes the total number of tasks), we capture the model state at the moment of the first occurrence of $(k + 1)$ -th task denoted as $\theta_{1,k}, \dots, \theta_{M,k}$ (where M represents the total number of parameters in the neural network). To model the variation of parameters, we first define the variation in parameter ranked m ($m \in \{1, \dots, M\}$) between task k and $k - 1$ as $\delta_{m,k}$, which is calculated using the Manhattan distance as

$$\delta_{m,k} = |\theta_{m,k} - \theta_{m,k-1}|, \quad (1)$$

where $|\cdot|$ denotes the Manhattan distance operator.

Due to significant differences in gradient updates between different parts of the network, $\delta_{m,k}$ can vary considerably for different m . Consequently, relying solely on its numerical variation may not accurately capture the correlation between parameters and tasks. To address this issue, we first standardize $\delta_{m,k}$ for each parameter while preserving relative changes, thereby mitigating the impact of numerical imbalance:

$$\delta'_{m,k} = S(\delta_{m,k}), \quad (2)$$

where S is a customized standardization function, and $\delta'_{m,k}$ is defined as the relative change of parameter m on task k .

Here we introduce three exemplary standardization functions utilized in identifying parameter variance imbalance. The first is the relative ratio (RR) function denoted by $RR(\cdot)$, which is expressed as

$$\bar{\delta}_k = \frac{\sum_{m=1}^M \delta_{m,k}}{M}, \quad RR(\delta_{m,k}) = \frac{\delta_{m,k}}{\bar{\delta}_k}, \quad (3)$$

where $\bar{\delta}_k$ represents the mean parameter variation for task k . This standardization enables $\delta'_{m,k}$ to capture relative parameter variation patterns without being affected by the absolute values of gradient updates. Additionally, by using a proportional form, this method mitigates the impact of different network architectures while intuitively highlighting imbalances in parameter variations, making it easier to identify parameters with disproportionately large or irregular updates.

As an alternative to the RR approach, Z-score (ZS) standardization method focuses on scaling the parameter variations with respect to their mean and standard deviation. The ZS standardization denoted by $ZS(\cdot)$ is defined as

$$\sigma_k = \sqrt{\frac{1}{M} \sum_{m=1}^M (\delta_{m,k} - \bar{\delta}_k)^2}, \quad ZS(\delta_{m,k}) = \frac{\delta_{m,k} - \bar{\delta}_k}{\sigma_k}, \quad (4)$$

where σ_k represents the standard deviation of the parameter variations for task k . ZS method accounts for both the central tendency and the spread of variations. Unlike the RR approach, the ZS method does not rely on task-specific variation proportions but instead standardizes the variations based on statistical properties, making it particularly useful for highly heterogeneous variations or those with significant outliers.

The last standardization method is robust scaler (RS), which is also particularly advantageous for handling significant outliers or skewed distributions of parameter variations. The RS method standardizes the variations based on the median and interquartile range (IQR), which is defined as

$$RS(\delta_{m,k}) = \frac{\delta_{m,k} - \text{Median}(\delta_k)}{IQR(\delta_k)}, \quad (5)$$

where $\text{Median}(\delta_k)$ represents the median and $IQR(\delta_k)$ is the interquartile range (the difference between the 75th and 25th percentiles) of

parameter variations for task k . The RS method is highly resilient to outliers, as the median and IQR are less sensitive to extreme values than the mean, ensuring robust parameter scaling in noisy or heavy-tailed distributions.

Using the above standardization methods for obtaining $\delta'_{m,k}$, we recorded the relative changes at the end of 1 ~ 4 task in a sequence of short tasks (a total of 5 tasks) during continual training on Cifar10 (Krizhevsky, Hinton, et al., 2009) with ER method (Chaudhry et al., 2019) using a backbone of reduced-Resnet18 (Aljundi, Belilovsky, et al., 2019). Through numerical analysis, RR intuitively demonstrates the relative differences in parameter updates, highlighting potential imbalances in both the overall and hierarchical patterns of parameter variations, while maintaining minimal computational complexity. Therefore, for all the analysis and experiments presented below, we specifically used the RR standardization method. The results obtained using the other two standardization methods are provided in Appendix A. More importantly, our observations reveal two distinct imbalances within the OCL scenario, as shown in Figs. 2 and 3.

Correlation-induced Imbalance. Fig. 2 illustrates the uneven distribution of parameter updates during OCL training, where a small fraction of parameters undergo significant changes, with some updates exceeding 64 times the mean, while over 65% of parameters exhibit updates smaller than the mean ($\delta'_{m,k} < 1$). This indicates that most parameters remain inactive, contributing minimally to task learning. In contrast, a smaller subset of parameters which experience substantial updates correlates closely with task memory. Such imbalances suggest inefficient parameter utilization, where inactive parameters are underutilized, and heavily updated parameters are at risk of causing catastrophic forgetting. Therefore, subsequent training should focus on better engaging inactive parameters while safeguarding those critical for task retention.

Layer-wise Imbalance. Fig. 3 demonstrates that in CL, the output classifier, i.e. the last linear layer, exhibits consistently higher parameter update rates compared to earlier feature extraction layers. This occurs because backpropagation begins at the classifier, which directly addresses prediction bias, particularly when input samples are imbalanced. Consequently, the classifier undergoes more frequent and larger updates, amplifying prediction biases and potentially destabilizing learning. To mitigate this, balancing updates across layers or restricting output classifier updates can reduce this disparity and improve the overall stability of the model.

3.2. Overview of parameter variation balancing framework

To address the two parameter variation imbalance issues outlined above, we propose PVBFB designed for OCL scenarios. As in Fig. 4, PVBFB primarily consists of two balancing strategies, E&C and D-CWR to mitigate imbalanced parameter variations in both the overall and the output layer level during training, thereby enhancing the network's adaptability to non-stationary data streams. PVBFB builds upon an ER framework, which retains a memory buffer storing part samples from previous tasks to aid in knowledge retention. For each neuron in each task, ParamCC is proposed to calculate correlation between parameter and previous tasks by capturing the variations in parameter weights during the training process. First, it employs a Manhattan distance-based metric to monitor relative parameter changes after each OCL task. Then, these changes are normalized using Min-Max normalization before and after task transitions to ensure a consistent correlation measurement (Jiawei & Micheline, 2006). To reduce time and space overhead, only the maximum correlation value of each parameter is retained throughout training. Based on these correlation values, E&C is then applied to update the gradients, encouraging low-correlation parameters to acquire new knowledge while consolidating high-correlation parameters to preserve previously learned knowledge. At the end of each task's training, the D-CWR strategy is applied to the output classifier. Parameters obtained through gradient descent are

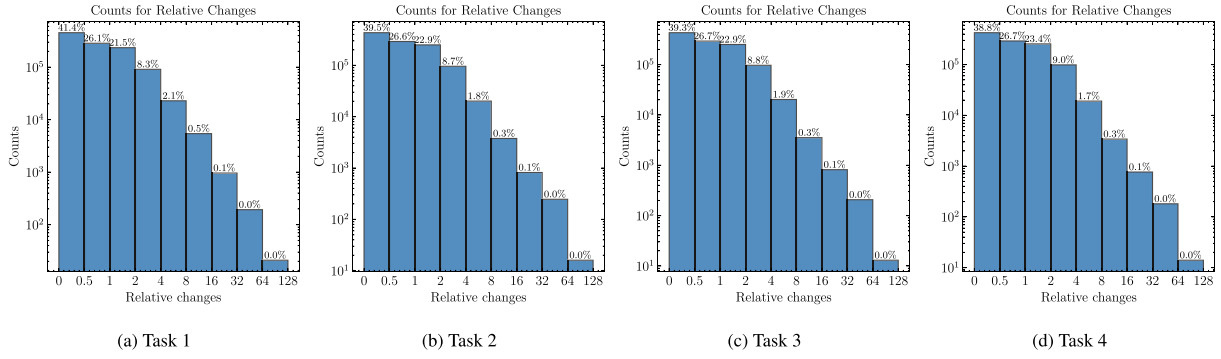


Fig. 2. Neuron counts for different relative changes $\delta'_{m,k}$.




Fig. 3. Average relative changes in parameters across different layers.

treated as sensory memory. Class-specific knowledge is then selectively transferred to a short-term memory pool through the first consolidation, followed by further consolidation into a long-term memory pool with a defined probability. Finally, the knowledge stored in long-term memory is used for prediction. This strategy slows the update rate of output layer parameters to reduce the prediction bias caused by rapid updates in this layer.

Besides, our study focuses on CL in dynamic and non-stationary sequential tasks. Firstly, we adopt the CL setting, where a model with parameters Θ must generalize well to test data without full access to previous training samples (Caccia et al., 2022). Each task k consists of training data $D_k = \{X_k, Y_k\}$, where X_k presents the input data and Y_k the corresponding labels. To ensure task diversity, we follow the classic disjoint label setting, where label sets of different tasks do not overlap ($Y_i \cap Y_j = \emptyset$ for $i \neq j$). Building upon this, we further explore OCL setting, where training process is divided into several time steps, with each time step corresponding to a single batch of training data. At each time step t , the model receives a new batch of training data (X_t^{in}, Y_t^{in}). Over time, the data distribution shifts as new tasks emerge. Specifically, at certain time steps $\{t_\alpha, t_\beta, \dots\} \in T$, the data transition from one task to another, which means that the continuous two sets of samples that the learner receives may belong to different tasks, namely $\{X_{t_{\alpha-1}}^{in}, Y_{t_{\alpha-1}}^{in}\} \in D_k$ and $\{X_{t_\alpha}^{in}, Y_{t_\alpha}^{in}\} \in D_{k+1}$. The training tasks occur sequentially, and task identifiers are not provided during evaluation. In this non-stationary setting, maintaining the neural network's classification performance across tasks is particularly challenging. We posit that preserving more task-specific network parameters' memory during learning while mitigating prediction bias from parameter updates during new task learning is critical to overcoming this challenge.

We adopt ACE to compute the loss, which separately handles data from the data stream and the memory buffer (Caccia et al., 2022). The total loss is given by

$$\mathcal{L}(X_t^{bf} \cup X_t^{in}) = \mathcal{L}^{re}(X_t^{bf}, Y_t^{old} \cup Y_t^{curr}) + \mathcal{L}^{in}(X_t^{in}, Y_t^{curr}), \quad (6)$$

Algorithm 1 Parameter Variation Balancing Framework (PVBf)

Input the hyperparameter α, β , Learning rate l .
Initialize network parameters Θ , gradients g , and memory buffer size M
Initialize C, M^s, M^c, M^l and P to zeroes
for each task $k \in \{1, \dots, K\}$ **do**
 store network parameters Θ in θ_{k-1}
 for each training batch at timestep t (B_t) **do**
 receive $X_t^{in} \sim B_t$ from input stream, $X_t^{re} \sim M$ from memory buffer
 Calculate $\mathcal{L}(X_t^{bf} \cup X_t^{in})$
 $g \leftarrow SGD(\nabla \mathcal{L}, \Theta)$ //calculate gradients
 if $k > 1$ $g' \leftarrow g, C$ ▷ section 3.4
 $\Theta \leftarrow g', l$
 end for
 $\delta' \leftarrow \Theta, \theta_{k-1}$ ▷ section 3.2
 $C \leftarrow \delta', \alpha, \beta$ ▷ section 3.3
 update M^s, M^c, M^l and P using D-CWR ▷ section 3.5
end for

where \mathcal{L} denotes the total loss function, \mathcal{L}^{in} and \mathcal{L}^{re} represent the cross-entropy loss of the data stream and data obtained from the memory buffer, X_t^{bf} and X_t^{in} represent the subsets of training samples at time step t , Y_t^{old} denotes the classes encountered up to that point, and Y_t^{curr} refers to the current training classes. Finally, a complete workflow of PVBf is provided in Algorithm 1. In the following, we will detail key designs of ParamCC, E&C strategy, and D-CWR strategy, respectively.

3.3. Parameter correlation calculate

To reduce the impact of gradient updates on absolute parameter variations and clarify the relationship between parameter variations and their correlation with previous tasks, Min-Max normalization can

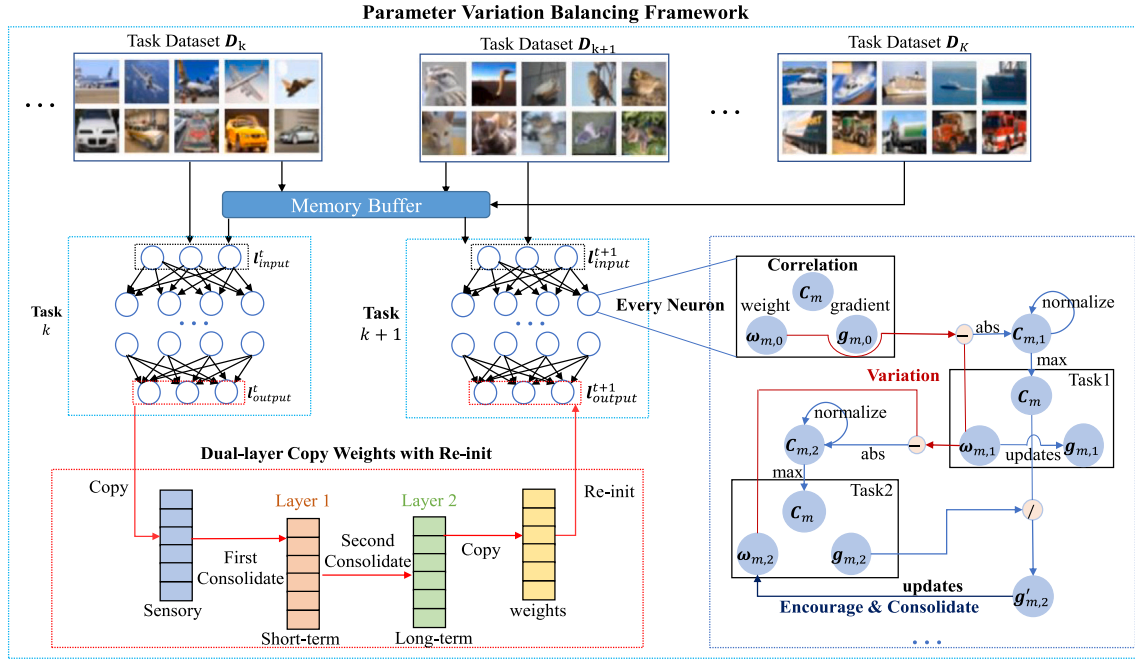


Fig. 4. Overall illustration of PVBf. PVBf is built upon a conventional experience replay framework and incorporates two main strategies for balancing parameter variations. First, the Encourage and Consolidate (E&C) strategy dynamically adjusts all neurons in the network during training. This strategy leverages the parameter correlations obtained through the Parameter Correlation Calculate (ParamCC) method during training to adjust the gradients computed by the Stochastic Gradient Decent (SGD) process. Second, the Dual-layer Copy Weights with Re-init (D-CWR) strategy is specifically applied to the output classifier, which progressively reinforces knowledge of individual class categories.

be used to standardize parameter correlations within a fixed range. Based on this idea, we introduce ParamCC, a method for evaluating the association between specific neural network parameters and previously learned tasks. The correlation of each parameter θ_m to the memory of task k can be expressed as follows:

$$C_{m,k} = \frac{\delta'_{m,k} - \min(V_k)}{\max(V_k) - \min(V_k)} \cdot (\beta - \alpha) + \alpha, \quad (7)$$

where $C_{m,k}$ denotes the correlation of parameter m with the memory of task k , and $V_k = \{\delta'_{1,k}, \dots, \delta'_{M,k}\}$ represents the set of relative changes. The hyperparameters α and β define the range of correlation for the parameters with the tasks, with $C_{m,k}$ being normalized to the interval $[\alpha, \beta]$.

By applying the aforementioned methods, we can derive a correlation measure $C_{m,k}$ for each parameter in the network with each previously trained task. To efficiently utilize this correlation measure and reduce computational and spatial overhead, we maintain only the maximum correlation of each parameter with the previous tasks throughout the training process as

$$C_m = \max_{r \in \{1, \dots, k-1\}} C_{m,r}, \quad (8)$$

where C_m represents the correlation of parameter m with all previously experienced tasks, and its value can be computed in real time during the training process. The variable k denotes the task currently being trained.

3.4. Encourage and consolidate

To mitigate catastrophic forgetting, we further integrate parameter correlation into the training process and propose a gradient correction method called E&C. The core idea is to encourage rapid updates for parameters with low correlation to previous tasks while consolidating those with high correlation by decelerating their update rates. For time step t , the gradient $g_{m,t}$ obtained by neuron m through gradient descent undergoes the following correction process:

$$g'_{m,t} = \frac{g_{m,t}}{C_m}, \quad (9)$$

where $g'_{m,t}$ represents the adjusted gradient for neuron m at time step t , taking into account the correlation of the parameters to previously learned tasks. This correction enables rapid updates for parameters with low correlation to previous tasks, thereby encouraging them to acquire new knowledge. Conversely, it ensures that parameters with high correlation to previous tasks are updated more cautiously, thereby consolidating their memory of those tasks. During the training process of online continual learning, we utilize the corrected parameters $\{g'_{1,t}, \dots, g'_{M,t}\}$ for parameter updates. This E&C approach facilitates a balance between learning new tasks and retaining knowledge of old tasks, ultimately enhancing the model's adaptability and mitigating catastrophic forgetting.

3.5. Dual-layer copy weights with Re-init

To minimize the phenomenon of parameter variation imbalance, in addition to the E&C method designed for overall model-level corrections, PVBf incorporates a specialized correction method specifically targeting the output classifier. Inspired by the interactions among sensory memory, short-term memory, and long-term memory (denoted as s , c , l) in the complementary learning systems (CLS) framework (McClelland, McNaughton, & O'Reilly, 1995), we propose an improved approach that enhances memory consolidation and reduces forgetting. Specifically, we introduce a D-CWR method that simulates the memory process after copying weights, which involves two consolidation layers: the first layer simulates the hippocampal process of converting sensory memory into short-term memory, while the second layer models the neocortical process of transferring short-term memory into long-term memory. In the following, we first define the concept of sensory memory in the context of the output classifier, followed by an explanation of the short-term and long-term consolidation processes.

We consider that the sensory memory of the output classifier originates from the knowledge acquired during the training process, particularly the model parameter weights obtained through gradient descent. In experiments conducted in Maltoni and Lomonaco (2019), a technique called mean-shift effectively normalizes the output classification layer model parameters, maintaining them within a certain range while

emphasizing the features of individual classes, which inspires our design of sensory memory. We define the sensory memory of the output classifier as the parameter weights assigned to each class after each training iteration, subtracted by the mean of the weights across all classes that appear in that iteration, specifically as follows:

$$M_{j,t}^s = \omega_{j,t} - \frac{\sum_{k \in S_t} \omega_{k,t}}{|S_t|}, \quad (10)$$

where $M_{j,t}^s$ represents the sensory memory of the model for class j , $\omega_{j,t}$ denotes the j th parameter of the output classifier, and S_t represents the set of sample labels input to the model, at time step t , operator $|\cdot|$ indicates calculating the cardinality of the set. Note that the mapping of the output classifier parameters to each class is determined by the properties of the softmax activation function. The softmax function connecting the final linear layer transforms the output layer parameters into classification probabilities for each class.

In CLS, the hippocampus is responsible for converting sensory memory into short-term memory by selectively processing certain information. The first layer of the D-CWR consolidation strategy simulates this process of the hippocampus. When sensory memory is received, D-CWR randomly selects one of the following two actions to execute: (i) Integrating the current sensory memory with the existing short-term memory; (ii) Retaining the previous short-term memory while withholding processing the current sensory memory. The model handles the short-term memory $M_{j,t}^c$ for class j at time step t , which can be calculated as

$$M_{j,t}^c = \begin{cases} M_{j,t-1}^c, & \text{if } \epsilon \geq p \\ \frac{M_{j,t-1}^c \cdot \eta_{j,t}^c + M_{j,t}^s}{\eta_{j,t}^c + 1}, & \text{if } \epsilon < p, \end{cases} \quad (11)$$

where ϵ is a random variable uniformly distributed in the interval $[0, 1]$, p is a hyperparameter representing the probability that the model converts sensory memory into short-term memory, and $\eta_{j,t}^c$ denotes the short-term memory retention coefficient for class j at time step t .

The neocortex plays a key role in converting short-term memory into long-term memory, with more frequent review of a particular class leading to stronger memory consolidation. The second layer of the D-CWR consolidation strategy simulates this process of the neocortex by integrating short-term memory with existing long-term memory according to a specified retention ratio. The long-term memory $M_{j,t}^l$ for class j at time step t is expressed as

$$M_{j,t}^l = \frac{M_{j,t-1}^l \cdot \eta_{j,t}^l + M_{j,t}^c}{\eta_{j,t}^l + 1}, \quad (12)$$

where $\eta_{j,t}^l$ denotes the long-term memory retention coefficient for class j at time step t , $\eta_{j,t}^l$ can be calculated alongside the short-term memory retention coefficient $\eta_{j,t}^c$ for class j at time step t as

$$\eta_{j,t}^c = \frac{P_{j,t}}{U_{j,t}}; \eta_{j,t}^l = \sqrt{\frac{P_{j,t}}{U_{j,t}}}, \quad (13)$$

with $P_{j,t}$ denotes the number of times class j has appeared during the model training process at time step t , and $U_{j,t}$ represents the number of occurrences of class j in the label set S_t . Notably, the retention coefficients $\eta_{j,t}^c$ and $\eta_{j,t}^l$ are designed to naturally adapt to the continual growth of the datastream. Since $P_{j,t}$ and $U_{j,t}$ are incrementally updated during training, the memory consolidation process automatically reflects the relative importance and frequency of each class over time. This mechanism allows D-CWR to effectively manage the evolving distribution of sensory, short-term, and long-term memory without requiring large storage or revisiting of past data.

To implement the D-CWR strategy, we present the algorithm that integrates the two consolidation stages, which begins by initializing memory vectors for sensory, short-term, and long-term memory. For each training step, it computes the sensory memory of the output classifier and updates short-term memory based on a probabilistic

decision determined by a random variable ϵ and a hyperparameter p . The short-term memory is subsequently consolidated into long-term memory using class-specific retention coefficients. This process ensures that the classifier's memory is continuously updated, preserving important knowledge while mitigating forgetting. The detailed steps of this memory consolidation procedure are summarized in Algorithm 2.

Algorithm 2 Dual-Layer Copy Weights with Re-init (D-CWR)

Initialize Initialize vector M^s , M^c , M^l and P to zeroes

for each training time step t batch samples $S_t(X_t^{in} \cup X_t^{re})$ **do**

 Denote ω_j as the j -th parameter of the output classifier

$\bar{\omega} \leftarrow \sum_{k \in S_t} \omega_k / |S_t|$

for each category j appearing in the batch S_t **do**

$M_j^s \leftarrow \omega_j - \bar{\omega}$

 Obtain the occurrences of class j in S_t as U_j .

$\eta_j^c \leftarrow \frac{P_j}{U_j}$

$\eta_j^l \leftarrow \sqrt{\frac{P_j}{U_j}}$

 Perform one of the following operations based on a random variable ϵ and hyperparameter p :

$\epsilon \geq p$: Retain M_j^c

$\epsilon < p$: $M_j^c \leftarrow \frac{M_j^c \cdot \eta_j^c + M_j^s}{\eta_j^c + 1}$

$M_j^l \leftarrow \frac{M_j^l \cdot \eta_j^l + M_j^c}{\eta_j^l + 1}$.

$P_j \leftarrow P_j + U_j$.

end for

$\omega \leftarrow M^l$

end for

With the above methods, we address the issue of correlation-induced imbalance and layer-wise imbalance in OCL from two perspectives: the overall model parameters and the output classifier parameters. This results in a parameter variation balancing framework that enables network models to better adapt to non-stationary data streams. Theoretically, the ParamCC and E&C strategies effectively capture the correlations between parameter variations and the knowledge of each task. By utilizing these correlations, they dynamically adjust the magnitude of parameter variations during subsequent training, thereby balancing the retention of new and old knowledge. These two strategies further enhance the performance of ER-based methods in continual learning. Not only are they highly effective in online scenarios, but they also show some benefits in offline settings. On the other hand, the D-CWR method primarily mitigates the layer-wise imbalance, proving effective in counteracting forgetting when the output classifier undergoes rapid updates.

4. Experiments

In this section, we first compare PVBF with competitive methods based on ER in two OCL scenarios: short task sequences and long task sequences. Next, we compare PVBF with other canonical ER-based methods in an offline CL scenario. The design of these experiments aims to validate the following performance metrics: (i) PVBF performs well across various OCL scenarios; (ii) PVBF components effectively mitigate the imbalance in parameter changes, thereby alleviating forgetting in OCL scenarios; (iii) Despite being specifically designed for OCL scenarios, PVBF still demonstrates strong performance when applied to offline CL settings.

4.1. OCL experiments

We first evaluate the performance of PVBF in the OCL setting, where samples from each task are sequentially presented to the model in a one-shot manner. As such, this setup places a high demand on




Fig. 5. ACC results for short task sequence OCL experiments (PVBF-D: PVBF without D-CWR).

the model's ability to adapt to non-stationary data streams. Only CL algorithms specifically designed for such challenges can effectively mitigate catastrophic forgetting in this setting. An effective OCL algorithm should efficiently learn the essential knowledge required to complete tasks with high accuracy while minimizing computational overhead.

4.1.1. Datasets and setup

We conducted experiments to evaluate OCL for image classification tasks. As with other OCL experiments (Caccia et al., 2022; Chrysakis & Moens, 2023), all benchmarks were evaluated in a single-head setting. In each of the datasets mentioned below, the model performs a classification task across N classes, where N depends on the specific dataset.

We selected three representative image datasets to conduct these OCL experiments: Split Cifar10, Split Cifar100 (Krizhevsky et al., 2009), and Split MiniImagenet (Vinyals, Blundell, Lillicrap, Wierstra, et al., 2016). These datasets are commonly used for evaluating OCL, with a focus on image classification tasks. The details on task division of the three adopted datasets are illustrated below:

Split Cifar10: This dataset consists of 60,000 images across 10 classes. It is one of the most widely adopted datasets in OCL research, often used as the standard benchmark for tasks with short task sequences. For our experiments, we partitioned the dataset into 5 tasks, each containing 2 classes.

Split Cifar100: The Cifar100 dataset includes 60,000 images spread across 100 classes. In our experiments, it was partitioned into 20 tasks, with each task containing 5 classes. This dataset presents a more complex and diverse challenge, making it suitable for evaluating OCL with longer task sequences.

Split MiniImagenet: This dataset consists of 100 classes, with each class containing 600 images. We partitioned it into 20 tasks, each with 5 classes. Like Cifar100, MiniImagenet is highly challenging for OCL due to its more diverse categories and longer task sequences.

4.1.2. Baselines

In our evaluation, we focus on replay-based methods due to their demonstrated effectiveness in the OCL setting, evidenced by previous studies (Aljundi, Belilovsky, et al., 2019; Caccia et al., 2022; Chaudhry et al., 2019; Ji, Henriques, Tuytelaars, & Vedaldi, 2020). Within this context, high-efficiency learning is necessary for all methods under consideration. To ensure a fair comparison, we have implemented uniform buffer management across all methods, employing Reservoir Sampling (Vitter, 1985) to retain or discard samples. We have selected several classic and state-of-the-art OCL baselines for our analysis, including ER (Chaudhry et al., 2019), Incremental Classifier and Representation Learning (iCaRL) (Rebuffi et al., 2017), Maximally Interfered Retrieval (MIR) (Aljundi, Belilovsky, et al., 2019), Dark Experience Replay (DER++) (Buzzega et al., 2020), Experience Replay with Asymmetric Cross-Entropy (ER-ACE) (Caccia et al., 2022), and Architect-and-Replay (AR1*) (Lomonaco et al., 2020). Additionally, we include an

independent and IID scenario (Caccia et al., 2022) as a benchmark, where the learner is trained on the dataset with a single pass, treating all classes as if they were part of a single task. This baseline variant is designed to have a similar computational budget to that of the replay methods.

4.1.3. Architectures and hyperparameters

Like previous researches (Buzzega et al., 2020; Caccia et al., 2022; Chrysakis & Moens, 2023), we use a reduced ResNet-18 network for all the datasets above, and all methods are applied to the same backbone. This choice enables us to fairly compare various methods, as these parameters directly affect the model's parameter size and the ability to classification tasks. Hyperparameters used in the experiments are presented in the following.

Learning Rate. The distance parameter shifts in the training process is directly affected by the learning rate. So in our experiments, for each dataset we set a fixed learning rate. Specifically, for Cifar10, we set the learning rate fixed at 0.1, for Cifar100 and MiniImagenet, we set the learning rate fixed at 0.01.

Memory Buffer Size. In the context of OCL scenarios, the performance of ER-based methods is closely tied to the hyperparameter of memory buffer size (MS). To ensure a fair comparison among various methods, it is necessary to evaluate their performance across different memory buffer sizes. Therefore, on the Cifar10 dataset, we varied this hyperparameter across three values: 20, 100, and 500. We conducted comparisons of selected methods under these three different settings. Conversely, on the Cifar100/MiniImageNet datasets, we fixed this hyperparameter at 500 to emphasize fairness in comparison.

Other Hyperparameters. Following previous studies (Aljundi, Belilovsky, et al., 2019; Caccia et al., 2022; Chaudhry, Ranzato, et al., 2018), for DER++, we set the hyperparameter $\alpha = 0.1$ and $\beta = 0.5$. The way to realizing masking loss for ER-ACE is the same as the author proposed in Buzzega et al. (2020). For our PVBF, we set normalization hyperparameter $\alpha = 0.5$, $\beta = 2.0$, and the probability for D-CWR p fixed at 0.9. A detailed analysis of the sensitivity of the hyperparameter p is provided in Appendix D.

4.1.4. Metrics

We utilized average per-task accuracy (ACC \uparrow , higher is better) alongside the forgetting ratio (FR \downarrow , lower is better) metric (Chaudhry, Dokania, et al., 2018), to assess the continual learning capabilities of various algorithm models in OCL scenarios. The ACC metric comprehensively assesses the learning capability of algorithmic models in OCL scenarios, while the FR metric evaluates their ability to consolidate memory during the training process. Let $a_{i,j}$ be the model's accuracy on task i on the test dataset after being trained on task j , ACC and FR are defined as

$$ACC = \frac{1}{K} \sum_{i=1}^K a_{i,i} \quad (14)$$

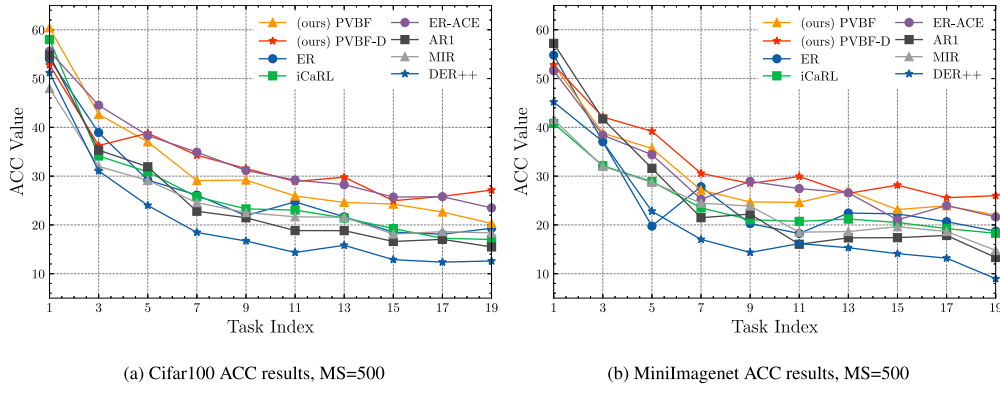


Fig. 6. ACC results for long task sequence OCL experiments (PVBf-D: PVBf without D-CWR).

Table 1

Short task sequence experiment results on split Cifar10. We conducted experiments in three scenarios: $MS = 20$, $MS = 100$, $MS = 500$. All entries are 95%-confidence intervals over 15 runs.

Method	MS=20		MS=100		MS=500	
	ACC	FR	ACC	FR	ACC	FR
IID	68.0 ± 1.1	–	68.0 ± 1.1	–	68.0 ± 1.1	–
ER	26.5 ± 1.3	50.4 ± 4.2	37.9 ± 1.5	28.1 ± 1.8	46.2 ± 1.7	18.9 ± 3.6
iCaRL	40.0 ± 1.1	38.5 ± 1.6	42.4 ± 1.1	32.3 ± 1.5	44.1 ± 1.0	30.6 ± 1.4
MIR	28.8 ± 1.4	46.7 ± 2.5	46.7 ± 0.9	18.8 ± 2.5	47.0 ± 2.0	9.1 ± 2.8
DER++	30.8 ± 1.8	33.7 ± 4.3	39.8 ± 1.3	23.3 ± 2.1	47.8 ± 1.6	13.9 ± 2.8
ER-ACE	37.7 ± 0.8	28.8 ± 2.7	47.4 ± 1.2	20.0 ± 2.6	53.3 ± 1.4	15.1 ± 1.7
AR1*	24.5 ± 1.4	75.9 ± 1.8	38.9 ± 1.8	53.0 ± 2.5	51.6 ± 2.5	36.4 ± 3.7
(ours) PVBf w/o D-CWR	37.9 ± 1.0	26.6 ± 2.0	47.1 ± 1.3	19.0 ± 1.7	54.8 ± 1.7	12.7 ± 2.0
(ours) PVBf	41.2 ± 0.8	24.0 ± 2.0	50.8 ± 1.4	18.1 ± 1.8	59.2 ± 1.8	10.6 ± 1.0

$$FR = \frac{1}{K-1} \sum_{i=1}^{K-1} \max_{k \in \{1 \dots K\}} a_{i,k} - a_{i,K} \quad (15)$$

4.1.5. Short task sequence OCL experiments

First as a standard short task sequence setting (Caccia et al., 2022), we apply our methods on Cifar10 datasets, along with the baselines. The results are shown in Table 1 and Fig. 5.

In experiments on Cifar10 under standard OCL scenarios, varying memory buffer sizes reveal distinct impacts on performance metrics. The IID method reveals the upper bound of learning accuracy for network models on this dataset (where no forgetting issue exists). As in Table 1, In scenarios with $MS=20$, $MS=100$, and $MS=500$, our PVBf demonstrates a significant accuracy advantage over other baselines, achieving an average 39.2% improvement in ACC compared to the ER method. Furthermore, our approach is able to further reduce the FR metric based on ER-ACE and AR1*, and has the lowest FR metrics in $MS = 20$ and $MS = 100$ settings, demonstrating more effective consolidating memory of past knowledge.

The results of the short task sequence experiments not only demonstrate that our PVBf exhibits strong adaptability in this setting, but also validate our hypothesis regarding the parameter variation imbalance. By comparing the PVBf without D-CWR to the ER-ACE method, we observe that the PVBf without D-CWR consistently achieves better performance across $MS = 20$, and $MS = 500$ settings, and lower FR in all settings, indicating that the E&C strategy alleviates the correlation-induced imbalance. Furthermore, when the D-CWR strategy is applied, the PVBf framework shows significantly lower FR and higher ACC, which can be attributed to the severe layer-wise imbalance in the short-task sequence OCL setting, and our D-CWR strategy effectively mitigates this imbalance.

4.1.6. Long task sequence OCL experiments

Similar experiments are conducted with the split Cifar100 and split MiniImagenet dataset with 20 tasks. The results are shown in Table 2 and Fig. 6.

Table 2

Long task sequence experiment results on split Cifar100 and split MiniImagenet. All entries are 95%-confidence intervals over 15 runs.

Method	Cifar100		MiniImagenet	
	ACC	FR	ACC	FR
IID	28.3 ± 0.9	–	24.3 ± 1.6	–
ER	17.5 ± 0.5	50.4 ± 0.7	18.1 ± 0.6	40.8 ± 1.1
iCaRL	17.1 ± 0.3	22.1 ± 0.3	16.8 ± 0.3	15.8 ± 0.6
MIR	18.1 ± 0.5	47.9 ± 0.7	18.9 ± 0.6	38.3 ± 0.9
DER++	10.9 ± 0.6	61.7 ± 0.7	10.3 ± 0.7	51.6 ± 1.0
ER-ACE	24.0 ± 0.7	10.1 ± 0.7	23.2 ± 0.7	8.9 ± 0.8
AR1*	14.6 ± 0.4	58.1 ± 0.7	15.7 ± 0.5	47.9 ± 0.6
(ours) PVBf w/o D-CWR	25.8 ± 0.7	10.3 ± 0.8	23.7 ± 0.8	9.2 ± 0.9
(ours) PVBf	21.7 ± 0.7	11.4 ± 0.6	20.9 ± 0.4	9.3 ± 0.5

The experimental results demonstrate that our PVBf without D-CWR outperforms all baseline methods on both the Cifar100 and MiniImagenet datasets. As in Table 2, particularly on the split Cifar100 dataset, the PVBf without D-CWR shows a 47% improvement in ACC over the ER method, approaching the model accuracy upper bound indicated by the IID method. On the split MiniImagenet dataset, the PVBf without D-CWR also performs well, showing improvements over the previously best-performing ER-ACE method. Additionally, the FR metric for the PVBf without D-CWR is relatively low in both datasets, very close to the best ER-ACE results, but for the ACC metric, PVBf without D-CWR has a clear advantage. This confirms that the evaluation of parameter importance and overall gradient correction in the PVBf without D-CWR are effective for long task sequences in the OCL scenario.

The OCL settings with long task sequences reveal significant differences in parameter variation imbalance compared to those with short task sequences. D-CWR may reduce the performance of PVBf, and the AR1* method, which only uses output classifier optimization, also performs poorly. We hypothesize that this is due to the




Fig. 7. Average relative changes in parameters across different layers in a long task sequence OCL experiment on Minilmagenet applying PVBF.

Table 3

Ablation experiment results on split Cifar10. All entries are 95%-confidence intervals over 15 runs.

PVBF w/o	E&C	D-CWR	ACE	–
ACC	55.9 ± 1.1	54.8 ± 1.7	51.3 ± 2.7	59.2 ± 1.8
FR	13.3 ± 1.6	12.7 ± 2.0	41.8 ± 3.1	10.6 ± 1.0

increased complexity and more frequent changes in tasks within the long task sequence scenario, which requires the output classifier to update more rapidly. In this scenario, as illustrated in Fig. 7, the relative changes in the parameters of the output classifier (linear layer) gradually decrease as the number of tasks increases, ultimately leading to a diminished effect of layer-wise imbalance. In other words, in these settings, correlation-induced imbalance is the primary contributor to parameter variation imbalance. Our PVBF without D-CWR primarily focuses on optimizing parameter dependency imbalance, which allows it to achieve performance close to that of the IID method.

4.1.7. Ablation study

Based on the results of the ablation study presented in Table 3, several key observations can be made regarding the performance of PVBF on the split Cifar10 dataset. The results of the ablation experiments demonstrate significant impacts of removing different modules on the PVBF. Removing the E&C and D-CWR modules decreased the ACC from 59.2% to 55.9% and 54.8%, indicating their substantial contributions to accuracy. Simultaneously, the FR values for these modules increased from 10.6% to 13.3% and 12.7%, illustrating their effectiveness in reducing forgetting. This implies that it is essential to correct prediction bias both holistically and specifically in the output classifier. Removing the ACE module caused the ACC to drop from 59.2% to 51.3%, and FR to increase from 10.6 to 41.8, highlighting the importance of combining ACE into PVBF to mitigate catastrophic forgetting.

4.1.8. Random task order experiments

In dynamic learning scenario, task execution orders may be adjusted in real-time due to external environmental factors, introducing uncertainty into the OCL process. Thus, we further conducted experiments to evaluate the stability and effectiveness of the PVBF framework under conditions of randomized task order.

To simulate this scenario, we performed experiments on the split Cifar10 dataset by varying MS and randomly shuffling the task sequence at each run. We evaluated the model using ACC and FR as performance metrics. Each configuration was executed 15 times to ensure statistical robustness.

Table 4 and Fig. 8 report the results under three settings: $MS = 20$, $MS = 100$, and $MS = 500$. The results indicate that the proposed PVBF method consistently achieves the highest average accuracy across all memory settings and maintains a lower forgetting rate compared to other strong baselines, including ER, iCaRL, DER++, and ER-ACE.

Furthermore, even without the D-CWR module, the ablated version (PVBF w/o D-CWR) remains competitive, indicating that our parameter variation balancing strategy is inherently robust to varying task orders.

Furthermore, as shown in Tables 1 and 4, one of the key challenges posed by random task orders in OCL lies in their significant impact under low replay buffer settings (e.g., $MS = 20$). In such cases, the uncertainty in task sequence tends to result in a higher FR, indicating that the model forgets a larger portion of previously learned knowledge, which in turn leads to a notable drop in classification accuracy. Under this challenging condition, our PVBF framework consistently achieves the lowest FR and the highest ACC among all compared methods, demonstrating strong stability and adaptability in learning from uncertain task sequences.

4.2. Offline CL experiments

To compare with more ER-based CL methods designed for offline scenarios and to further validate the stable performance of PVBF in more CL settings, we applied PVBF and other baselines to the same offline CL setting. The offline setting typically includes two specific learning modes: class-incremental (class-IL) and task-incremental (task-IL). The primary difference between these two learning modes lies in the evaluation process. In the class-IL learning mode, task identifiers are not provided to the model, and the classifier must accurately predict the class of validation samples among all seen classes. In contrast, for the task-IL learning mode, task identifiers are provided to the model, and the classifier only needs to classify samples within the corresponding task. In the offline CL setting, we use the same evaluation metrics as in the OCL scenario.

4.2.1. Datasets and setup

In reference to the standard research on offline continual learning (CL) in Buzzega et al. (2020), we conduct experiments on the widely used Split-Cifar10 dataset, addressing both class-IL and task-IL scenarios. These two scenarios respectively assess the model's memory capabilities for the entire dataset and for individual tasks, making them highly representative in the context of offline CL. Similar to the OCL scenario, we divide the Split-Cifar10 dataset into five tasks, with each task requiring the model to classify samples from two distinct classes.

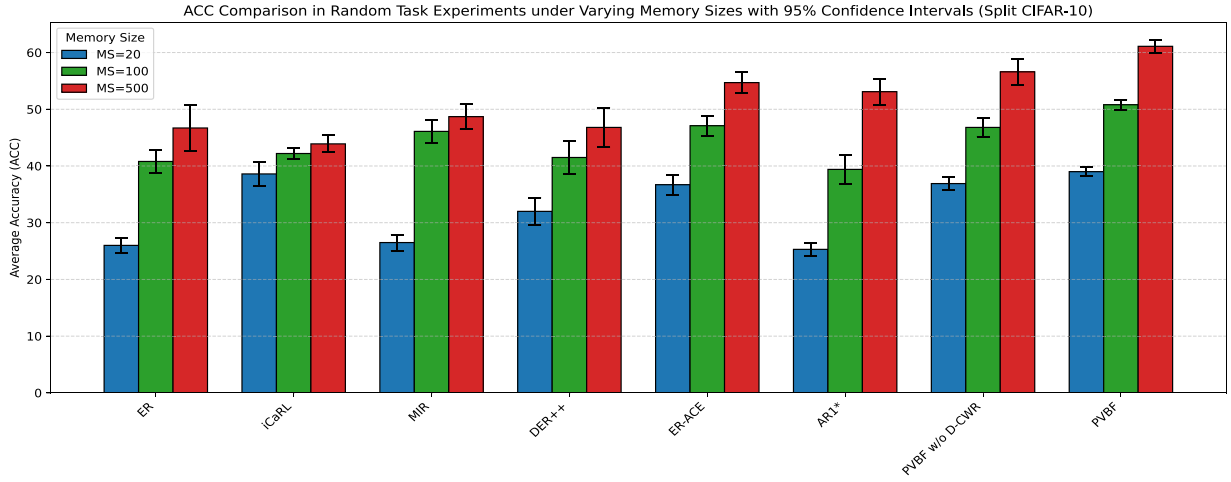
4.2.2. Baselines

Regarding the selection of baselines, we introduce seven existing ER-based methods from offline CL experiments (ER (Riemer et al., 2018), GEM (Lopez-Paz & Ranzato, 2017), iCaRL (Rebuffi et al., 2017), FDR (Benjamin, Rolnick, & Kording, 2019), GSS (Aljundi, Lin, Goujaud, & Bengio, 2019), HAL (Chaudhry, Gordo, Dokania, Torr, & Lopez-Paz, 2021), and DER++), as well as ER-ACE, which has shown strong performance in OCL scenarios. Similar to the OCL scenario, for the nine methods, including PVBF, we employ reservoir sampling for data sampling. To ensure a fair comparison across methods, we set the

Table 4

Experimental results of short task sequence with random task orders on split Cifar10. We conducted experiments in three scenarios: $MS = 20$, $MS = 100$, $MS = 500$. All entries are 95%-confidence intervals over 15 runs.

Method	MS=20		MS=100		MS=500	
	ACC	FR	ACC	FR	ACC	FR
ER	26.0 \pm 1.3	57.9 \pm 3.4	40.8 \pm 2.1	28.7 \pm 4.6	46.7 \pm 4.1	18.7 \pm 5.6
iCaRL	38.6 \pm 2.1	40.7 \pm 2.9	42.2 \pm 1.0	30.7 \pm 1.5	43.9 \pm 1.5	30.3 \pm 2.2
MIR	26.5 \pm 1.4	53.4 \pm 3.7	46.1 \pm 2.0	22.5 \pm 2.3	48.7 \pm 2.2	10.5 \pm 1.9
DER++	32.0 \pm 2.4	45.3 \pm 7.0	41.5 \pm 2.9	19.3 \pm 5.1	46.8 \pm 3.5	15.8 \pm 4.1
ER-ACE	36.7 \pm 1.8	39.8 \pm 3.0	47.1 \pm 1.8	23.9 \pm 1.8	54.7 \pm 1.9	16.2 \pm 2.5
ARI*	25.3 \pm 1.2	75.5 \pm 2.0	39.4 \pm 2.5	53.2 \pm 3.7	53.1 \pm 2.3	37.1 \pm 3.0
(ours) PVBF w/o D-CWR	36.9 \pm 1.2	37.1 \pm 2.7	46.8 \pm 1.6	24.4 \pm 2.5	56.6 \pm 2.3	14.7 \pm 2.8
(ours) PVBF	39.0 \pm 0.8	35.7 \pm 2.0	50.8 \pm 0.9	24.3 \pm 2.8	61.1 \pm 1.1	13.5 \pm 1.5

**Fig. 8.** ACC results of random task order experiments.**Table 5**

Offline experiment results on split Cifar10. MS represents the memory buffer size. We conducted experiments in three scenarios: $MS = 200$, $MS = 500$, $MS = 5120$. Results by SGD, ER, GEM, iCaRL, FDR, GSS, HAL, and DER++ are cited from Buzzega et al. (2020), all other entries are 95%-confidence intervals over 15 runs.

Method	MS=200		MS=500		MS=5120	
	ACC	FR	ACC	FR	ACC	FR
Class-IL						
SGD	19.62 \pm 0.05	96.39 \pm 0.12	19.62 \pm 0.05	96.39 \pm 0.12	19.62 \pm 0.05	96.39 \pm 0.12
ER	44.79 \pm 1.86	61.24 \pm 2.62	57.74 \pm 0.27	45.35 \pm 0.07	82.57 \pm 0.52	13.99 \pm 1.12
GEM	25.54 \pm 0.76	82.61 \pm 1.60	26.20 \pm 2.26	74.31 \pm 4.62	25.26 \pm 3.46	75.27 \pm 4.41
iCaRL	49.02 \pm 3.20	28.72 \pm 0.49	47.55 \pm 3.95	25.71 \pm 1.10	55.07 \pm 1.55	24.94 \pm 0.14
FDR	30.91 \pm 2.74	86.40 \pm 0.67	28.71 \pm 3.23	85.62 \pm 0.36	19.70 \pm 0.07	96.64 \pm 0.19
GSS	39.07 \pm 5.59	75.25 \pm 4.07	49.73 \pm 4.78	62.88 \pm 2.67	67.27 \pm 4.27	58.11 \pm 9.12
HAL	32.36 \pm 2.70	69.11 \pm 4.21	41.79 \pm 4.46	62.21 \pm 4.34	59.12 \pm 4.41	27.19 \pm 7.53
DER++	64.88 \pm 1.17	32.59 \pm 2.32	72.70 \pm 1.36	22.38 \pm 4.41	85.24 \pm 0.49	7.27 \pm 0.84
ER-ACE	63.48 \pm 1.24	16.43 \pm 1.65	69.95 \pm 1.18	12.58 \pm 0.79	83.06 \pm 0.62	6.14 \pm 0.88
(ours) PBVF	65.27 \pm 1.18	22.11 \pm 1.95	71.30 \pm 1.01	14.49 \pm 1.25	83.45 \pm 0.67	6.12 \pm 0.61
Task-IL						
SGD	61.02 \pm 3.33	46.24 \pm 2.12	61.02 \pm 3.33	46.24 \pm 2.12	61.02 \pm 3.33	46.24 \pm 2.12
ER	91.19 \pm 0.94	7.08 \pm 0.64	93.61 \pm 0.27	3.54 \pm 0.35	96.98 \pm 0.17	0.27 \pm 0.06
GEM	90.44 \pm 0.94	9.27 \pm 2.07	92.16 \pm 0.69	9.12 \pm 0.21	25.26 \pm 3.46	6.91 \pm 2.33
iCaRL	88.99 \pm 2.13	2.63 \pm 3.48	88.22 \pm 2.62	2.66 \pm 2.47	92.23 \pm 0.84	1.59 \pm 0.57
FDR	91.01 \pm 0.68	7.36 \pm 0.03	93.29 \pm 0.59	4.80 \pm 0.00	94.32 \pm 0.97	1.93 \pm 0.48
GSS	88.80 \pm 2.89	8.56 \pm 1.78	91.02 \pm 1.57	7.73 \pm 3.99	94.19 \pm 1.15	7.71 \pm 2.31
HAL	82.51 \pm 3.20	12.26 \pm 0.02	84.54 \pm 2.36	5.41 \pm 1.10	88.51 \pm 3.32	5.21 \pm 0.50
DER++	91.92 \pm 0.60	5.16 \pm 0.21	93.88 \pm 0.50	4.66 \pm 1.15	96.12 \pm 0.21	1.18 \pm 0.19
ER-ACE	92.74 \pm 1.18	5.59 \pm 0.90	94.35 \pm 0.92	3.58 \pm 0.72	97.25 \pm 0.48	0.78 \pm 0.16
(ours) PBVF	93.96 \pm 1.06	4.80 \pm 1.03	95.62 \pm 0.71	2.68 \pm 0.49	97.12 \pm 0.50	0.76 \pm 0.12

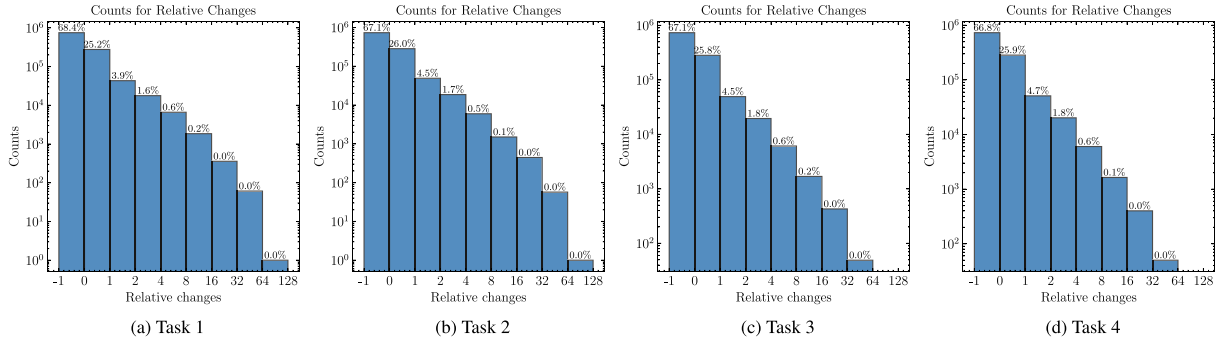

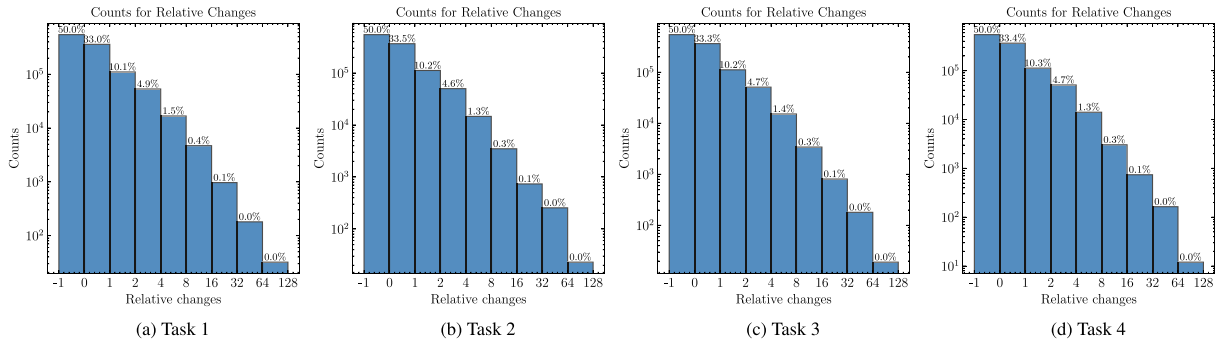
Fig. A.9. Neuron counts for different relative changes $\delta'_{m,k}$ (ZS method).

Fig. A.10. Average relative changes in parameters across different layers (ZS method).

Fig. A.11. Neuron counts for different relative changes $\delta'_{m,k}$ (RS method).

same memory buffer size for all methods within the same experimental setting. Simultaneously, to highlight the memory advantages offered by continual learning algorithms, we also include experimental results for the SGD method. This approach does not utilize any experience replay, and instead performs standard stochastic gradient descent training on each task sequentially, with the same number of training epochs as the other algorithms.

4.2.3. Training

During training, as we adopted part of the offline CL experimental data from prior work, we ensured consistency across all conditions that influence model performance. Specifically, we adhered to the following characteristics of the Mammoth experiment framework (Boschini, Bonicelli, Buzzega, Porrello, & Calderara, 2022), which is designed for offline CL experiments: (i) For all training processes, we used the same SGD optimizer with a fixed batch size of 32 and fixed epochs 50; (ii) All methods were applied to a common, unpretrained ResNet-18 based backbone; (iii) Experiments were conducted across three different memory buffer sizes: 200, 500, and 5120. For the seven classic baselines (ER, GEM, iCaRL, FDR, GSS, HAL, and DER++), we referenced their optimal performance results under previously validated

hyperparameter settings for this configuration. For ER-ACE and PVBF, we conducted experiments with a fixed learning rate of 0.03. The hyperparameter settings for PVBF were consistent with those used in the OCL setting, and all experimental results are fully reproducible.

4.2.4. Results

As shown in Table 5, experiments conducted in the offline CL settings demonstrate that PVBF exhibits stable performance across both class-IL and task-IL scenarios. In the class-IL setting, PVBF achieves the highest ACC and the second-best FR when the memory buffer size $MS = 200$. For $MS = 500$ and $MS = 5120$, PVBF consistently shows the second-highest ACC and either the lowest or second-lowest FR. In the task-IL scenario, PVBF achieves optimal performance for $MS = 200$ and $MS = 500$, surpassing all baselines. For $MS = 5120$, PVBF's performance is on par with ER-ACE. Overall, across both scenarios, PVBF demonstrates a competitive advantage over existing state-of-the-art methods in the offline CL setting. Notably, particularly in settings with a small memory buffer size and in task-IL scenarios, PVBF retains a clear advantage over DER++, which also demonstrates excellent performance in the offline CL setting.




Fig. A.12. Average relative changes in parameters across different layers in average (RS method).

4.2.5. Analysis

In the offline CL setting, compared to the OCL setting, the model is able to undergo more extensive training, which alleviates the issue of parameter variation imbalance. Particularly, in the case of $MS = 5120$, the model can store a sufficient number of samples during training, allowing it to effectively revisit knowledge from previous tasks. In this scenario, the issue of parameter variation imbalance is mitigated under the conditions of ample sample and training data. As a result, the experimental outcomes of ER-based methods eventually converge with those of the ER method. When the number of samples that the model can store is smaller, the issue of parameter variation imbalance becomes more pronounced. In such cases, PVBF can alleviate this issue to some extent, leading to better performance compared to other methods. In the task-IL scenario, the correlation-induced imbalance correction strategy employed by PVBF enables the model to retain knowledge of specific tasks more vividly, which contributes to PVBF's superior performance.

5. Conclusion

Mitigating catastrophic forgetting in OCL requires addressing the prediction bias caused by parameter variation imbalance. To solve this issue, this paper proposes a Parameter Variation Balancing Framework (PVBF), which mitigates the correlation-induced imbalance and layer-wise imbalance respectively. The paper validates PVBF through classification experiments on both short and long task sequence OCL settings, as well as offline settings. The results show that PVBF achieves an average accuracy improvement of 31%–47% over the ER method. On the MiniImageNet dataset, it attains 97.5% of the IID method's accuracy using only 500 replay samples. In offline CL settings, PVBF consistently outperforms classical CL methods, demonstrating its significant advantages.

CRedit authorship contribution statement

Zelin Tao: Visualization, Methodology, Formal analysis, Software, Writing – original draft, Validation, Investigation. **Hao Deng:** Writing – original draft, Project administration, Conceptualization, Supervision, Funding acquisition, Writing – review & editing. **Mingqing Liu:** Writing – review & editing, Visualization, Writing – original draft. **Lijun Zhang:** Writing – review & editing, Visualization. **Shengjie Zhao:** Supervision, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Table B.6

Short task sequence experiment results on split Cifar10 using PVBF with different standardization methods. All entries are 95%-confidence intervals over 15 runs.

Method	MS=20		MS=100		MS=500	
	ACC	FR	ACC	FR	ACC	FR
PBVF-RR	41.2 ± 0.8	24.0 ± 2.0	50.8 ± 1.4	18.1 ± 1.8	59.2 ± 1.8	10.6 ± 1.0
PBVF-ZS	40.4 ± 1.6	24.4 ± 2.0	51.6 ± 1.0	14.8 ± 1.2	59.0 ± 2.1	10.1 ± 2.5
PBVF-RS	42.0 ± 1.1	24.2 ± 2.2	49.5 ± 1.6	18.2 ± 1.8	56.3 ± 1.8	13.3 ± 1.2

Acknowledgments

This work was supported in part by the National Key Research and Development Project under Grant 2023YFC3806002, in part by the National Natural Science Foundation of China under Grant U23A20382 and 62371342, in part by the Natural Science Foundation of Shanghai under Grant 23ZR1467300.

Appendix A. Relative changes analysis of different standardization methods

We also use the ZS and RS method for obtaining $\delta'_{m,k}$, we recorded the relative changes at the end of 1 ~ 4 task in a sequence of short tasks (a total of 5 tasks) during continual training on Cifar10 (Krizhevsky et al., 2009) with ER method (Chaudhry et al., 2019) using a backbone of reduced-Resnet18 (Aljundi, Belilovsky, et al., 2019).

From the results in Fig. A.9 using ZS normalization, we observe that in all tasks, 66%–68% of the parameters exhibit relative changes below the mean, and over 90% of the parameters have relative changes within one standard deviation. This further supports our hypothesis regarding correlation-induced imbalance, where only a small subset of parameters show a strong correlation within a task. Similarly, the results in Fig. A.11 using RS normalization corroborate this observation, with the number of parameters exhibiting relative changes greater than 2 being less than 7% after normalization.

Similarly, the results shown in Fig. A.10, which present the average relative changes of parameters at different layers using ZS normalization, and those in Fig. A.12, which show the corresponding results with RS normalization, further support our hypothesis regarding layer-wise imbalance. In both cases, the final linear layer (i.e., the output classifier) exhibits significantly higher average relative changes in parameters compared to the adjacent layers.

Appendix B. OCL experiments results of different standardization methods

In this section, we present the experimental results of the PVBF framework under different standardization methods. The results are evaluated across multiple task sequences with varying sequence lengths and datasets. We report the ACC and FR for each method.




Fig. C.13. Time consuming Results of different methods applying on split cifar10.




Fig. D.14. The hyperparameter sensitivity heatmap for the probability parameter p in the D-CWR method, where the baseline value of p is set to 0.55. Each cell denotes the relative change in ACC corresponding to $p = 0.55 + 0.05 \times \text{column_index} + 0.01 \times \text{row_index}$, computed as the deviation from the ACC obtained under the baseline setting of $p = 0.90$.

Table B.7

Long task sequence experiment results on split Cifar100 and split MiniImagenet using PVBF with different standardization methods. All entries are 95%-confidence intervals over 15 runs.

Method	Cifar100		MiniImagenet	
	ACC	FR	ACC	FR
PVBF-RR	21.7 ± 0.7	11.4 ± 0.6	20.9 ± 0.4	9.3 ± 0.5
PVBF-ZS	22.4 ± 0.5	11.0 ± 0.5	20.8 ± 0.5	10.4 ± 0.7
PVBF-RS	22.8 ± 0.7	11.2 ± 0.6	21.8 ± 0.5	9.8 ± 0.5

Short Task Sequence Results. Table B.6 summarizes the results for the short task sequence experiment conducted on the split Cifar10 dataset. The results highlight the performance of PVBF with three standardization methods: RR, ZS, and RS. Notably, the PVBF-RR method exhibits the highest ACC when $MS = 500$ and lowest FR when $MS = 20$. The PVBF-ZS method achieves the best ACC at 51.6 ± 1.0 when $MS = 100$, and lowest FR when $MS = 100$ and $MS = 500$. The PVBF-RS method achieves best ACC when $MS = 20$. From the results of the short task sequence, both PVBF-RR and PVBF-ZS demonstrate comparable performance, outperforming PVBF-RS.

Long Task Sequence Results. Table B.7 presents the results from the long task sequence experiment conducted on the split Cifar100 and split MiniImagenet datasets. Again, we evaluate the performance of PVBF with three standardization methods: RR, ZS, and RS. In the case

of Cifar100, PVBF-RS achieves the highest ACC of 22.8, outperforming both PVBF-RR and PVBF-ZS, and PVBF-ZS achieves the lowest FR. On the MiniImagenet dataset, PVBF-RS achieves the highest ACC of 21.8, while PVBF-RR shows a lower forgetting rate of 9.8. In contrast to the short task sequence results, PVBF-RS outperforms the other two methods in the long task sequence.

Considering both OCL scenarios, all three standardization methods exhibit similar performance on the OCL task. From a practical perspective, RR has the lowest computational complexity among the three methods. Therefore, we primarily use the RR method in the main discussions.

Appendix C. Time consuming

In the context of online continual learning (OCL), evaluating the time consumption of different methods is crucial, as maintaining a low computational cost allows models to adapt more rapidly to dynamically evolving data streams. To rigorously assess the runtime efficiency of each method, we conducted time consumption experiments under a unified hardware configuration (CPU: 16 vCPUs, Intel(R) Xeon(R) Gold 6430; GPU: NVIDIA RTX 4090 with 24 GB VRAM) and on the same dataset (split Cifar10). The results are shown in Fig. C.13.

Time consuming of ParamCC. For ParamCC, since computations are only required at task transitions, it does not affect the training efficiency of the model. We independently measured the time cost of this computation process, and found that the time consumed per execution is less than 1 ms, which is negligible compared to the model's overall training time.

Time consuming of PVBF. As demonstrated by the experimental results, PVBF exhibits comparable computational time with comparative approaches, introducing less than 2 s of additional overhead compared to ER. Through systematic ablation experiments on E&C and D-CWR components, we observe that the time complexity of PVBF primarily stems from the D-CWR module, as removing this component reduces its computational time to a level comparable with ER-ACE. This indicates that the computational complexity introduced by the E&C method remains strictly bounded. Furthermore, the PVBF-E&C variant (containing only D-CWR and paramCC components) demonstrates merely 0.4 s of additional computation time compared to ER-ACE, confirming that its temporal overhead remains within practical acceptance thresholds for real-world deployment.

Appendix D. Hyperparameter sensitivity

For the probabilistic hyperparameter p in D-CWR, its value significantly impacts the PVBF model's overall ability to mitigate forgetting.

To determine the optimal value, we conducted a hyperparameter sensitivity analysis by testing p over the range of 0.61 to 1 with 0.01 increments. This evaluation was performed on split Cifar10 under a short-task sequential OCL scenario with a fixed memory size (MS) of 500. As illustrated in Fig. D.14, the model achieves stable ACCs when p ranges between 0.86 and 0.90. Based on these findings, we empirically set $p = 0.9$ for subsequent experiments in this study.

References

- Aljundi, R., Belilovsky, E., Tuytelaars, T., Charlin, L., Caccia, M., Lin, M., et al. (2019). Online continual learning with maximal interfered retrieval. *Advances in Neural Information Processing Systems*, 32.
- Aljundi, R., Lin, M., Goujaud, B., & Bengio, Y. (2019). Gradient based sample selection for online continual learning. *Advances in Neural Information Processing Systems*, 32.
- Benjamin, A., Rolnick, D., & Kording, K. (2019). Measuring and regularizing networks in function space. In *International conference on learning representations*.
- Bonicelli, L., Boschini, M., Porrello, A., Spampinato, C., & Calderara, S. (2022). On the effectiveness of Lipschitz-driven rehearsal in continual learning. *Advances in Neural Information Processing Systems*, 35, 31886–31901.
- Boschini, M., Bonicelli, L., Buzzega, P., Porrello, A., & Calderara, S. (2022). Class-incremental continual learning into the extended der-verse. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(5), 5497–5512.
- Buzzega, P., Boschini, M., Porrello, A., Abati, D., & Calderara, S. (2020). Dark experience for general continual learning: a strong, simple baseline. *Advances in Neural Information Processing Systems*, 33, 15920–15930.
- Caccia, L., Aljundi, R., Asadi, N., Tuytelaars, T., Pineau, J., & Belilovsky, E. (2022). New insights on reducing abrupt representation change in online continual learning. In *International conference on learning representations*.
- Castro, F. M., Marin-Jimenez, M. J., Guil, N., Schmid, C., & Alahari, K. (2018). End-to-end incremental learning. In *Proceedings of the European conference on computer vision*.
- Chaudhry, A., Dokania, P. K., Ajanthan, T., & Torr, P. H. (2018). Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *Proceedings of the European conference on computer vision* (pp. 532–547).
- Chaudhry, A., Gordo, A., Dokania, P., Torr, P., & Lopez-Paz, D. (2021). Using hindsight to anchor past knowledge in continual learning. In *Proceedings of the AAAI conference on artificial intelligence*: vol. 35, no. 8, (pp. 6993–7001).
- Chaudhry, A., Ranzato, M., Rohrbach, M., & Elhoseiny, M. (2018). Efficient lifelong learning with a-gem. *arXiv preprint arXiv:1812.00420*.
- Chaudhry, A., Rohrbach, M., Elhoseiny, M., Ajanthan, T., Dokania, P., Torr, P., et al. (2019). Continual learning with tiny episodic memories. In *Workshop on multi-task and lifelong reinforcement learning*.
- Chrysakis, A., & Moens, M.-F. (2023). Online bias correction for task-free continual learning. *ICLR 2023 At OpenReview*.
- Fernando, C., Banarse, D., Blundell, C., Zwols, Y., Ha, D., Rusu, A. A., et al. (2017). Pathnet: Evolution channels gradient descent in super neural networks. *arXiv preprint arXiv:1701.08734*.
- Hurtado, J., Raymond, A., & Soto, A. (2021). Optimizing reusable knowledge for continual learning via metalearning. *Advances in Neural Information Processing Systems*, 34, 14150–14162.
- Isele, D., & Cosgun, A. (2018). Selective experience replay for lifelong learning. In *Proceedings of the AAAI conference on artificial intelligence*: vol. 32, no. 1.
- Ji, X., Henriques, J., Tuytelaars, T., & Vedaldi, A. (2020). Automatic recall machines: Internal replay, continual learning and the brain. *arXiv preprint arXiv:2006.12323*.
- Jiawei, H., & Micheline, K. (2006). *Data mining: concepts and techniques*. Morgan kaufmann.
- Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., et al. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13), 3521–3526.
- Krizhevsky, A., Hinton, G., et al. (2009). Learning multiple layers of features from tiny images.
- Li, Z., & Hoiem, D. (2017). Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12), 2935–2947.
- Lomonaco, V., & Maltoni, D. (2017). Core50: a new dataset and benchmark for continuous object recognition. In *Conference on robot learning* (pp. 17–26). PMLR.
- Lomonaco, V., Maltoni, D., Pellegrini, L., et al. (2020). Rehearsal-free continual learning over small non-IID batches. In *CVPR workshops*: vol. 1, no. 2, (p. 3).
- Lopez-Paz, D., & Ranzato, M. (2017). Gradient episodic memory for continual learning. *Advances in Neural Information Processing Systems*, 30.
- Mai, Z., Li, R., Jeong, J., Quispe, D., Kim, H., & Sanner, S. (2022). Online continual learning in image classification: An empirical survey. *Neurocomputing*, 469, 28–51.
- Mallya, A., Davis, D., & Lazebnik, S. (2018). Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In *Proceedings of the European conference on computer vision* (pp. 67–82).
- Mallya, A., & Lazebnik, S. (2018). Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 7765–7773).
- Maltoni, D., & Lomonaco, V. (2019). Continuous learning in single-incremental-task scenarios. *Neural Networks*, 116, 56–73.
- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102(3), 419.
- Nairne, J. S., & Neath, I. (2003). Sensory and working memory. In *Comprehensive handbook of psychology*: vol. 4, (pp. 423–444). Wiley New York.
- Pellegrini, L., Graffieti, G., Lomonaco, V., & Maltoni, D. (2020). Latent replay for real-time continual learning. In *2020 IEEE/RSJ international conference on intelligent robots and systems* (pp. 10203–10209). IEEE.
- Prabhu, A., Torr, P. H., & Dokania, P. K. (2020). Gdumb: a simple approach that questions our progress in continual learning. In *Computer vision—ECCV 2020: 16th European conference, glasgow, UK, August 23–28, 2020, proceedings, part II 16* (pp. 524–540). Springer.
- Rebuffi, S.-A., Kolesnikov, A., Sperl, G., & Lampert, C. H. (2017). Icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 2001–2010).
- Riemer, M., Cases, I., Ajemian, R., Liu, M., Rish, I., Tu, Y., et al. (2018). Learning to learn without forgetting by maximizing transfer and minimizing interference. *arXiv preprint arXiv:1810.11910*.
- Shin, H., Lee, J. K., Kim, J., & Kim, J. (2017). Continual learning with deep generative replay. *Advances in Neural Information Processing Systems*, 30.
- Soutif-Cormerais, A., Carta, A., Cossu, A., Hurtado, J., Lomonaco, V., Van de Weijer, J., et al. (2023). A comprehensive empirical evaluation on online continual learning. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 3518–3528).
- Sun, R. (2019). Optimization for deep learning: theory and algorithms. *arXiv preprint arXiv:1912.08957*.
- Van de Ven, G. M., & Tolias, A. S. (2019). Three scenarios for continual learning. *arXiv preprint arXiv:1904.07734*.
- Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al. (2016). Matching networks for one shot learning. *Advances in Neural Information Processing Systems*, 29.
- Vitter, J. S. (1985). Random sampling with a reservoir. *ACM Transactions on Mathematical Software (TOMS)*, 11(1), 37–57.
- Von Oswald, J., Henning, C., Grewe, B. F., & Sacramento, J. (2019). Continual learning with hypernetworks. *arXiv preprint arXiv:1906.00695*.
- Wang, L., Zhang, X., Su, H., & Zhu, J. (2024). A comprehensive survey of continual learning: theory, method and application. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Wang, F.-Y., Zhou, D.-W., Ye, H.-J., & Zhan, D.-C. (2022). Foster: Feature boosting and compression for class-incremental learning. In *European conference on computer vision* (pp. 398–414). Springer.
- Wu, Y., Chen, Y., Wang, L., Ye, Y., Liu, Z., Guo, Y., et al. (2019). Large scale incremental learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 374–382).
- Yoon, J., Yang, E., Lee, J., & Hwang, S. J. (2017). Lifelong learning with dynamically expandable networks. *arXiv preprint arXiv:1708.01547*.
- Yuan, W., Yin, H., He, T., Chen, T., Wang, Q., & Cui, L. (2022). Unified question generation with continual lifelong learning. In *Proceedings of the ACM web conference 2022* (pp. 871–881).
- Yuan, W., Zhang, Q., He, T., Fang, C., Hung, N. Q. V., Hao, X., et al. (2022). Circle: continual repair across programming languages. In *Proceedings of the 31st ACM SIGSOFT international symposium on software testing and analysis* (pp. 678–690).
- Zenke, F., Poole, B., & Ganguli, S. (2017). Continual learning through synaptic intelligence. In *International conference on machine learning* (pp. 3987–3995). PMLR.