Flattening Sharpness for Dynamic Gradient Projection Memory Benefits Continual Learning

Anonymous Author(s)
Affiliation
Address
email

Abstract

The backpropagation networks are notably susceptible to catastrophic forgetting, where networks tend to forget previously learned skills upon learning new ones. To address such the 'sensitivity-stability' dilemma, most previous efforts have been contributed to minimizing the empirical risk with different parameter regularization terms and episodic memory, but rarely exploring the usages of the *weight loss landscape*. In this paper, we investigate the relationship between the weight loss landscape and sensitivity-stability in the continual learning scenario, based on which, we propose a novel method, *Flattening Sharpness for Dynamic Gradient Projection Memory (FS-DGPM)*. In particular, we introduce a soft weight to represent the importance of each basis representing past tasks in GPM, which can be adaptively learned during the learning process, so that less important bases can be dynamically released to improve the sensitivity of new skill learning. We further introduce *Flattening Sharpness (FS)* to reduce the generalization gap by explicitly regulating the flatness of the weight loss landscape of all seen tasks. As demonstrated empirically, our proposed method consistently outperforms baselines with the superior ability to learn new skills while alleviating forgetting effectively.

1 Introduction

Humans and animals have the ability to continually learn new knowledge without forgetting their previously learned ones through mediating a rich set of neurocognitive mechanisms [39, 13, 37]. This ability, often known as continual learning or lifelong learning [27], is crucial for computational systems, such as deep neural networks (DNNs), which are required to sequentially learn and deal with multiple tasks when implemented in the dynamically changing environment. Continual learning remains a long-standing challenge for DNNs since these networks are typically trained with stationary training batches by stochastic gradient descent methods [17], which generally leads to an abrupt performance decrease on previously learned tasks as new tasks are learned. To address such *catastrophic forgetting*, we can brutally retrain an oracle network on the entire dataset containing all tasks to capture dynamic changes in the data distribution, but this methodology is obviously too inefficient to hinder the learning of novel data in real time.

During the last few years, lots of research efforts have been devoted to improving the *stability* of DNNs on old tasks while keeping *sensitive* to new information. The first intuitive idea is to introduce an independent branch for each new task while freezing the old task parameters to preserve the old knowledge [31, 43, 41, 23, 33, 20]. However, in this way, the network will inevitably become redundant as the task number continually increases. As presented in the neurocognitive works [39, 13], the reactivation of neuronal activity patterns, representing old memories, plays an important role in the continual learning of humans [37]. Thus, forgetness can be effectively mitigated by training a single network for new tasks by considering diverse information stored in the memory, including the

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

original training samples of old tasks [29, 4, 11], the gradients induced from old tasks [7] and the feature subspace representing old tasks [32]. However, their continual learning performance is still limited because DNNs can easily overfit the limited information stored in the small-size memory.

The overfitting problem of DNNs is often attributed to the complex loss landscape containing multiple local optima, and the sharpness of the loss landscape has been widely used to characterize the generalization gap in standard training scenarios from both theoretical and empirical perspectives [25, 19, 40, 8]. While this characterization has inspired new approaches for model training with better generalization, practical algorithms that especially seek out flatter minima to effectively address the 'sensitivity-stability' dilemma for continual learning have thus far been elusive. In this paper, our first contribution is to characterize the weight loss landscape for the continual learning scenario and identify that a flatter loss landscape with lower loss value often leads to better continual learning performance, as shown in Figure 1 and Figure 3.

Further, based on our characterization of the weight loss landscape, we find that the recently proposed Gradient Projection Memory (GPM) method [32] maintains the lowest loss value on old tasks among the previously proposed methods by taking gradient steps orthogonal to the subspace representing old tasks. However, its loss landscape on newly learned tasks is the sharpest due to the lack of sufficient subspace left for new task learning. To improve the network's sensitivity, our second contribution is to predict the importance of bases spanning the subspace for old tasks, so that less important bases can be dynamically released. In particular, we introduce a soft weight to indicate the bases importance, which can be dynamically adjusted by combining the *Flattening Sharpness (FS)* to minimize the loss value and loss sharpness simultaneously. Intuitively, a basis will be regarded as important for preserving old knowledge if the gradients induced by new tasks and old ones are aligned in the opposite direction on that basis. As demonstrated through extensive experiments, our proposed method can consistently outperform the state-of-the-art methods [15, 28, 22, 29, 4, 11, 32] by a notable margin across a range of widely used benchmark datasets.

2 Related Work

In this section, we briefly survey the representative works of continual learning by highlighting their contributions. To simplify our presentation, this section is organized by dividing the representative works into three categories, parameter isolation, regularization-based, memory-based methods.

Parameter isolation methods address forgetting by assigning a different subset of network parameters to each task. Without restrictions on network architecture, new neurons or layers or modules can be added for new tasks, while the previous task parameters can be frozen or copied to preserve old knowledge. For instance, Progressive Neural Network (PGN) [31] freezes the parameters trained with previous knowledge while expands the architecture by allocating new sub-networks with fixed capacity for new tasks. Dynamically Expandable Networks (DEN) [43] selectively retrains or expands network capacity by splitting/duplicating important units on new tasks. Reinforced Continual Learning (RCL) [41] uses reinforcement learning strategy to adaptively expand the network of each layer, while [20] use neural architecture search to find optimal network structures for each sequential task. Alternatively, with the architecture remaining static, a fixed part is allocated to each task. During the training of a new task, previous task parts are masked out to prevent interference. The mask sets are imposed at parameter level [23], or unit level [33]. PackNet [23] uses iterative pruning to fully restrict gradient updates on important weights via a binary mask, whereas HAT [33] limits the update of important units recognized by the hard attention mask through gradient backpropagation.

Regularization-based methods introduce an additional regularization term in the loss function to consolidate previous knowledge without using replay. This involves using knowledge distillation [21, 12] or penalizing changes in weights deemed important for previous tasks [15, 44, 26, 2] to reduce forgetting. There are many ways to measure the importance. Elastic Weight Consolidation (EWC) [15] identifies important weights based on the diagonal values of Fisher information matrix after training, while Synaptic Intelligence (SI) [44] calculates them online and over the entire learning trajectory in parameter space. Memory Aware Synapses (MAS) [2] estimates importance weights in an unsupervised manner, while Variational Continual Learning (VCL) [26] introduces a variational framework that spawned some Bayesian-based works [30, 6, 1, 5]. For example, [30] recursively uses a Gaussian Laplace approximation of the Hessian to approximate the posterior after every task, [6] adjusts the learning rate according to the uncertainty defined by the probability distribution of the

network weights. [1] introduces an interpretation of node-wise uncertainty on the Kullback-Leibler (KL) divergence term of the variational lower bound for Gaussian mean-field approximation.

Our method mainly follows **memory-based** methods, which mitigate forgetting based on information extracted from old tasks or based on a generative model to generate pseudo samples. For example, iCaRL [28] selects and stores samples closest to the feature mean of each class. ER [4, 29] suggests reservoir sampling under the limited and fixed budget for replay buffer. Deep Generative Replay (DGR) [34] trains a deep generative model in the Generative Adversarial Network (GAN) framework [9] to simulate past data. These previous task samples are mainly reused as model inputs for replay in the above methods. However, replay might be prone to overfitting the subset of stored samples. Alternatively, samples stored in memory can also be used to constrain the optimization of the new task loss to prevent previous task interference, thereby leaving more leeway for backward and forward transfer. Gradient Episodic Memory (GEM) [22] projects the estimated gradients in the feasible region, which is outlined by previous task gradients calculated from the episodic memory samples. Averaged-GEM (A-GEM) [3] relaxes the projection to a direction that is estimated from samples randomly selected from memory. [10] proposes a unified view of episodic memory-based continual learning methods, including GEM and A-GEM, and improves performance over these methods by using a loss-balancing update scheme. A few other works have utilized gradient information to protect previous knowledge. [29, 11] combines experience replay with optimization-based meta-learning to enforce gradient alignment between examples to learn parameters that interfere based on future gradients as low as possible. GPM [32] minimizes the interference between sequential tasks by ensuring that gradient updates only occur in directions orthogonal to the input of previous tasks.

3 The Weight Loss Landscape of Continual Learning

In this section, we first introduce our formulation of continual learning, and then characterize the weight loss landscape for the continual learning scenario from stability and sensitivity. Finally, some insights combining the weight loss landscape and continual learning are provided.

3.1 Problem Formulation

Throughout the paper, we denote scalars as a, vectors as $\boldsymbol{a}$, matrices as $\boldsymbol{A}$, and sets as $\mathcal{A}$. We consider a supervised learning setup where T tasks are sequentially learned from their training data. Each task has an identical task descriptor, $\tau \in \{1, 2, \ldots, T\}$, with its dataset $\mathcal{D}_\tau = \{\boldsymbol{x}_{i,\tau}, y_{i,\tau}\}_{i=1}^{n_\tau}$ containing n_τ samples randomly generated from a latent distribution $\mathscr{D}_\tau$. At any time-step during the learning process, we minimize the empirical risk of the model on all t tasks seen so far, with just limited size of memory $\mathcal{M}$ to summarize the training data of previous tasks $\{\mathcal{D}_\tau\}_{\tau=1}^{t-1}$. To simplify the notation, we denote $L_{\mathcal{A}}(\boldsymbol{w}) = \frac{1}{|\mathcal{A}|}\sum_{(\boldsymbol{x},y)\in\mathcal{A}}\left[\ell\left(f_{\boldsymbol{w}}\left(\boldsymbol{x}\right), y\right)\right]$ as the average empirical loss for the set $\mathcal{A}$, where $\ell(\cdot,\cdot)$ is an arbitrary loss function (*e.g.* the cross-entropy (CE) loss), $|\mathcal{A}|$ is the sample size of the set $\mathcal{A}$, and $f_{\boldsymbol{w}}$ is the DNN with weight $\boldsymbol{w}$. Our final goal is to find an optimal parameter $\boldsymbol{w}$, which minimizes the overall risk $\sum_{\tau=1}^{T} L_{\mathscr{D}_\tau}(\boldsymbol{w})$ for all tasks.

3.2 Connection of Weight Loss Landscape and Continual Learning

After learning a new task, we visualize the weight loss landscape of each task seen so far by plotting changes in its training loss when moving the weights $\boldsymbol{w}$ in a random direction $\boldsymbol{d}$ with magnitude α following [19]:

$$g_t(\alpha) = L_{\mathcal{D}_t}(\boldsymbol{w} + \alpha\boldsymbol{d}) = \frac{1}{|\mathcal{D}_t|}\sum_{(\boldsymbol{x},y)\in\mathcal{D}_t}\ell\left(f_{\boldsymbol{w}+\alpha\boldsymbol{d}}\left(\boldsymbol{x}\right), y\right), \tag{1}$$

where $\mathcal{D}_t$ is the training set for the t-th task previously learned. To eliminate the scaling invariance of DNNs, $\boldsymbol{d}$ is sampled from a Gaussian distribution and filter normalized by $\boldsymbol{d}_{l,j} \leftarrow \frac{\boldsymbol{d}_{l,j}}{\|\boldsymbol{d}_{l,j}\|_F}\|\boldsymbol{w}_{l,j}\|_F$, where $\boldsymbol{d}_{l,j}$ represents the j-th filter at l-th layer of $\boldsymbol{d}$, and $\|\cdot\|_F$ denotes the Frobenius norm. The key differnce to previous work [24] lies on the random directions used for visualization of each task. [24] only consider one task and plot the loss landscape along the directions of eigenvectors of its Fisher information matix, which only reflects some of the relationship between forgetting and landscape. Considering $\boldsymbol{d}$ is randomly selected, we repeat the visualization 10 times with different $\boldsymbol{d}$.

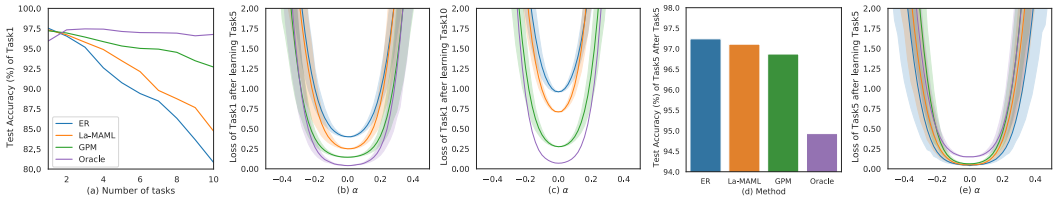

Figure 1: The connection between the weight loss landscape and continual learning is investigated on four methods. (a)-(c) shows the stability of the first task. (a) is the test accuracy change curve of the first task; (b) and (c) are the weight loss landscape of the first task after learning the fifth task and all ten tasks. (d)-(e) shows the sensitivity of the fifth task, which are the test accuracy and the weight loss landscape of the fifth task after just learning the fifth task. The shape of the weight loss landscape obtained using ten different random filter-normalized directions. ("Task$i$" is abbr. of the i-th task)

We first study the *stability* of the networks by plotting the changes of the weight loss landscape for the first task after new task learning. In particular, We use the previously proposed ER[4], La-MAML[11], and GPM[32] to train a MLP network with two hidden layers on the Permuted MNIST (PMNIST) [18] dataset that contains 10 tasks. We also retrain the network on the entire dataset contain all passed tasks as an Oracle network. Early stopping is used to halt the training with up to 10 epochs for each task based on the validation loss as proposed in [33]. As shown in Figure 1(a), all three continual learning methods lose their stability as learning new tasks. It can be observed from Figure 1(b)-(c) that the weight loss landscape becomes sharper and loss value increases simultaneously, when the testing accuracy of the first task continually decreases. We further evaluate the *sensitivity* of the network by observing the performance of the fifth task just after it has just been learned. As shown in Figure 1(d)-(e), ER shows the best learning capability compared with other methods with the lowest loss value and the flattest loss landscape. Thus, based on these empirical findings, we assume that lower loss value with a flatter neighbor often leads to better continual learning performance.

3.3 A Case Study of Flattening Sharpness for Vanilla ER

In this part, we further validate our above assumption by Flattening Sharpness for vanilla ER (FS-ER). Compared with ER that looks for a solution $\boldsymbol{w}$ that jointly minimizes the training loss of current task data and memory data, FS-ER seeks out a solution with both low loss and flat neighbor by minimizing the maximal loss in the neighbor around the parameter value. The schematic of the Flattening Sharpness (FS) is shown in Figure 2. We introduce the adversarial weight perturbation (orange dashed line) to explicitly flatten the weight loss landscape via injecting the worst-case weight perturbation, which is calculated from the current task data and past task data sampled from the replay buffer (Refer to Appendix C.1 and the next section for more details). Figure 3 respectively shows the weight loss landscapes of the fifth task after just learning the fifth task and all tasks. We find that FS-ER successfully gets a solution with a lower loss value and flatter landscape, either after the fifth task has just been learned or all ten tasks have been learned. The average testing accuracy of all tasks using FS-ER is **90.44%**, significantly higher than ER (**86.16%**), which means that flattening sharpness does benefit continual learning.

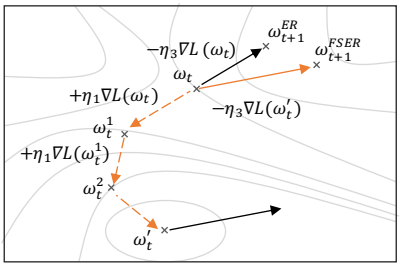

Figure 2: Schematic of FS-ER update. The dashed line and the solid line indicate the gradient ascent and descent, respectively. Orange denotes the actual update of the parameter $\boldsymbol{w}$.

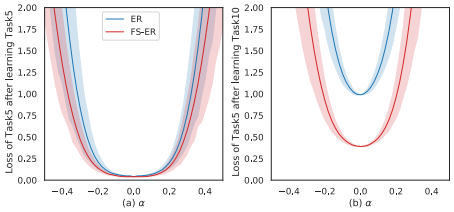

Figure 3: Landscape of the fifth task after just learning (a) the fifth task and (b) all ten tasks.

4 Flattening Sharpness for Dynamic Gradient Projection Memory

Our work is motivated by GPM [32], which achieves excellent stability on old tasks from theoretical and empirical perspectives. As shown in Figure 1, GPM achieves the lowest testing accuracy on the new task among all three practical continual learning methods, but shows less sensitivity to new task learning. To address this issue, we propose *Flattening Sharpness for Dynamic GPM (FS-DGPM)*, which dynamically adjusts the gradient subspace representing the past tasks to improve the sensitivity to the new task, while ensuring stability of the previous tasks. In particular, we let $\boldsymbol{M} = [\boldsymbol{u}_1, \boldsymbol{u}_2, \cdots, \boldsymbol{u}_k]$ denote the bases matrix that spans the gradient subspace of the previous task, $\boldsymbol{\Lambda} = \text{diag}[\lambda_1, \lambda_2, ..., \lambda_k]$ be the diagonal matrix with its i-th diagonal element $\lambda_i \in [0, 1]$ indicating the importance of each basis, and k is the number of bases.

4.1 Sharpness Evaluation

Comparing with the classical strategy that perturbs weight in the entire space [38, 14, 40], we focus on characterizing the weight loss landscape on the new task with respect to the important subspace representing old task. The important subspace can be effectively calculated based on the examples sampled from replay buffer $\mathcal{M}$ after each task training. Then, we can find the worst case by maximizing the training loss of the network on the new task in this subspace. Formally, the sharpness of the loss landscape around the solution $\boldsymbol{w}$ in the old parameter space can be predicted as,

$$\max_{\boldsymbol{v} \in \mathcal{V}} \quad L_{\hat{\mathcal{D}}_t}\left(\boldsymbol{w} + \boldsymbol{v}\right), \tag{2}$$

where $\mathcal{V}$ denotes the subspace spanned by $\boldsymbol{M}$ and $\boldsymbol{\Lambda}$, and $\hat{\mathcal{D}}_t$ denotes the batch samples of the current t-th task. As shown in Eq. (2), the high value can be obtained when the network fails to learn the new task (*sensitivity*) and the new task learning seriously interferes with the past tasks learning (*stability*). Based on the gradient method, the adversarial weight perturbation $\boldsymbol{v}$ can be solved as,

$$\boldsymbol{v} \leftarrow \boldsymbol{v} + \eta_1 \boldsymbol{M}\boldsymbol{\Lambda}\boldsymbol{M}^T \left(\nabla_{(\boldsymbol{w}+\boldsymbol{v})} L_{\hat{\mathcal{D}}_t}(\boldsymbol{w} + \boldsymbol{v})\right), \tag{3}$$

where η_1 is the update step size. Note that $\boldsymbol{v}$ is initialized as $\mathbf{0}$ and layer-wise updated. As shown in Appendix E, two-step for $\boldsymbol{v}$ (default settings) are enough to get good improvements.

4.2 Dynamic Gradient Projection Memory

After obtaining the adversarial weight perturbation $\mathbf{v}$, we can further update the bases importance matrix $\boldsymbol{\Lambda} = \text{diag}[\lambda_1, \lambda_2, ..., \lambda_k]$ by jointly considering the current task batch $\hat{\mathcal{D}}_t$ and the batch $\hat{\mathcal{M}}$ sampled from the replay buffer $\mathcal{M}$ as following,

$$\lambda_i \leftarrow \lambda_i - \eta_2 \left(\nabla_{\lambda_i} L_{\hat{\mathcal{D}}_t \cup \hat{\mathcal{M}}}(\boldsymbol{w} + \boldsymbol{v})\right), \tag{4}$$

where the sigmoid function is used here to constrain the importance value λ_i between 0 and 1. In addition, the second term on the right side can be approximated by the first-order Taylor expansion as,

$$\nabla_{\lambda_i} L_{\hat{\mathcal{D}}_t \cup \hat{\mathcal{M}}}(\boldsymbol{w} + \boldsymbol{v}) \approx \eta_1 \left(\nabla_{\boldsymbol{w}} L_{\hat{\mathcal{D}}_t}\left(\boldsymbol{w}\right)\right)^T \boldsymbol{u}_i \boldsymbol{u}_i^T \left(\nabla_{\boldsymbol{w}} L_{\hat{\mathcal{D}}_t \cup \hat{\mathcal{M}}}\left(\boldsymbol{w}\right)\right). \tag{5}$$

The above equation characterizes the relationship between the gradients induced by the current task and the old tasks based on the basis $\boldsymbol{u}_i$. As illustrated in the Figure 4, this equation implies that when the projections of two gradients on the basis $\boldsymbol{u}_i$ are aligned in the same direction, the gradient of λ_i will be positive, and when there is interference, the gradient will be negative. The positive (negative) gradient will decrease (increase) the importance λ_i, thereby releasing (tightening) the update limit of the new task on the corresponding basis $\boldsymbol{u}_i$. We provide the full derivation in the Appendix A.2. Note that the GPM is dynamically adjusted from the second task.

4.3 Weight Updating

Finally, we update the model parameters by minimizing the worst performance of $f_{\boldsymbol{w}+\boldsymbol{v}}$ with respect to $\boldsymbol{w}$, while adjusting the update magnitude of $\boldsymbol{w}$ on each basis based on its importance to alleviate forgetting. More concretely, the parameter $\boldsymbol{w}$ will be updated to:

$$\boldsymbol{w} \leftarrow \boldsymbol{w} - \eta_3 \left(\boldsymbol{I} - \boldsymbol{M}\boldsymbol{\Lambda}\boldsymbol{M}^T\right) \nabla_{\boldsymbol{w}} L_{\hat{\mathcal{D}}_t \cup \hat{\mathcal{M}}}(\boldsymbol{w} + \boldsymbol{v}). \tag{6}$$

Algorithm 1 FS-DGPM (Flattening Sharpness for Dynamic Gradient Projection Memory)

Input: Network weight $\boldsymbol{w}$, loss function ℓ, learning rate η_3, FS step size η_1, FS steps K, Soft weight step size η_2, .
Initializing $\mathcal{M} \leftarrow \{\}$, $\boldsymbol{M} \leftarrow \boldsymbol{I}, \boldsymbol{\Lambda} \leftarrow \boldsymbol{I}$
for $t = 1, 2, \cdots, T$ **do**
  for $ep = 1, 2, \cdots, num_{epochs}$ **do**
    for batch $\hat{\mathcal{D}}_t \subseteq \mathcal{D}_t$ **do**
      $\hat{\mathcal{M}} \leftarrow$ Sample $(\mathcal{M})$
      for $k = 1, \cdots, K$ **do**
        $\boldsymbol{v} \leftarrow \boldsymbol{v} + \eta_1 \boldsymbol{M}\boldsymbol{\Lambda}\boldsymbol{M}^T \left(\nabla_{(\boldsymbol{w}+\boldsymbol{v})} L_{\hat{\mathcal{D}}}(\boldsymbol{w}+\boldsymbol{v})\right)$
      end for
      if $t \geq 2$ **then**
        $\boldsymbol{\Lambda} \leftarrow \boldsymbol{\Lambda} - \eta_2 \left(\nabla_{\boldsymbol{\Lambda}} L_{\hat{\mathcal{D}}_t \cup \hat{\mathcal{M}}}(\boldsymbol{w}+\boldsymbol{v})\right)$
      end if
      $\boldsymbol{w} \leftarrow \boldsymbol{w} - \eta_3 \left(\boldsymbol{I} - \boldsymbol{M}\boldsymbol{\Lambda}\boldsymbol{M}^T\right) \nabla_{\boldsymbol{w}} L_{\hat{\mathcal{D}}_t \cup \hat{\mathcal{M}}}(\boldsymbol{w}+\boldsymbol{v})$
      Push $\hat{\mathcal{D}}_t$ to $\mathcal{M}$ with reservior sampling
    end for
  end for
  $\boldsymbol{M} \leftarrow$ UpdateGPM $(\mathcal{M})$ ▷ see Appendix Alg. 2
end for

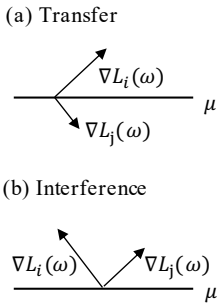

Figure 4: A depiction of transfer (a) and interference (b) in the basis $\boldsymbol{u}$ of gradient space.

Note that the optimization is performed over the model parameters $\boldsymbol{w}$, whereas the objective is computed using perturbed model $f_{\boldsymbol{w}+\boldsymbol{v}}$. In addition, we update the replay buffer $\mathcal{M}$ with reservoir sampling as in [29], and then use Singular Value Decomposition (SVD) to recalculate $\boldsymbol{M}$ based on the sampling data in the replay buffer after learning one task following [32]. Comparing with [32], we calculate the important bases in the entire gradient space and use them to replace the bases calculated last time. The complete pseudo-code of FS-DGPM is outlined in the Algorithm 1.

4.4 Theoretical Understanding

We further provide a theoretical view on why landscape can characterize the continual learning performance and why our proposed FS-DGPM works. To simplify our explanation, we only consider two tasks, which contains the training sets $\mathcal{D}_1$ and $\mathcal{D}_2$ sampled from the distributions $\mathscr{D}_1$ and $\mathscr{D}_2$, respectively. Based on the previous works on PAC-Bayes bound [25, 19, 40, 8], given a "prior" distribution P (a common assumption is zero mean, σ^2 variance Gaussian distribution) over the weights, the expected error of the classifier for the continual learning scenario can be bounded with probability at least $1-\delta$ over the draw of n training data:

$$\begin{aligned}\min_{\Delta \boldsymbol{w}} \mathbb{E}_{\boldsymbol{u}}\left[L_{\mathscr{D}_1 \cup \mathscr{D}_2}(\boldsymbol{w}+\Delta\boldsymbol{w}+\boldsymbol{u})\right] \leq & \min_{\Delta \boldsymbol{w} \in \mathcal{V}^C} \mathbb{E}_{\boldsymbol{u}}\left[L_{\mathscr{D}_1}(\boldsymbol{w}+\Delta\boldsymbol{w}+\boldsymbol{u})\right] + \min_{\Delta \boldsymbol{w} \in \mathcal{V}^C} L_{\mathcal{D}_2}(\boldsymbol{w}+\Delta\boldsymbol{w}) \\ & + \min_{\Delta \boldsymbol{w} \in \mathcal{V}^C} [\max_{\boldsymbol{u} \in \mathcal{V}} L_{\mathcal{D}_2}(\boldsymbol{w}+\Delta\boldsymbol{w}+\boldsymbol{u}) - L_{\mathcal{D}_2}(\boldsymbol{w}+\Delta\boldsymbol{w})] \\ & + \min_{\Delta \boldsymbol{w} \in \mathcal{V}^C} 4\sqrt{\frac{1}{n} KL(\boldsymbol{w}+\Delta\boldsymbol{w}+\boldsymbol{u}||P) + \ln\frac{2n}{\delta}},\end{aligned} \tag{7}$$

where $\Delta\boldsymbol{w}$ is the update based on the previously optimal solution $\boldsymbol{w}$ learned on the old task $\mathcal{D}_1$ when learning the new one $\mathcal{D}_2$, and $\boldsymbol{u}$ is often chosen as a zero mean spherical Gaussian perturbation with variance σ^2 in every direction. Let $\Delta\boldsymbol{w} \in \mathcal{V}^C$, then $\Delta\boldsymbol{w}$ lies in the complementary space of the important space representing the old task $\mathscr{D}_1$, so that $\mathbb{E}_{\boldsymbol{u}}\left[L_{\mathscr{D}_1}(\boldsymbol{w}+\Delta\boldsymbol{w}+\boldsymbol{u})\right]$ does not increase too much compared with the previously minimized $\mathbb{E}_{\boldsymbol{u}}\left[L_{\mathscr{D}_1}(\boldsymbol{w}+\boldsymbol{u})\right]$. The second term denotes the empirical loss on the second task and the third term represents the sharpness of the weight loss landscape around the $\boldsymbol{w}+\Delta\boldsymbol{w}$. Since we have constrained $\Delta\boldsymbol{w} \in \mathcal{V}^C$, then it is natural to assume $\boldsymbol{u} \in \mathcal{V}$, so that $\Delta\boldsymbol{w}+\boldsymbol{u}$ will cover the full space. Thus, our FS-DGPM exactly optimizes the worst-case of the flatness of weight loss landscape to control the PAC-Bayes bound, which theoretically justifies both lower loss value and flatter landscape lead to better continual learning performance, and why our proposed FS-DGPM works.

5 Experiments

In this section, we conduct extensive experiments to compare the performance of our proposed FS-DGPM model with the state-of-the-art methods on widely used continual learning benchmark datasets. Additional results and more details about the datasets, experiment setup, baselines, and model architectures are presented in the Appendix D.

5.1 Experimental Setup

Datasets: We evaluate our algorithm on four image classification datasets: **Permuted MNIST** (PMNIST) [18], **CIFAR-100 Split** [16], **CIFAR-100 Superclass** [42] and **TinyImageNet** [35]. The PMNIST dataset is a variant of the MNIST dataset, in which each task applies a fixed random pixel permutation to the original dataset. The PMNIST benchmark dataset consists of 20 tasks, and each contains only 1000 samples from 10 different classes [11]. The CIFAR-100 Split is constructed by randomly dividing 100 classes of CIFAR-100 into 10 tasks with 10 classes per task. The CIFAR-100 Superclass is divided into 20 tasks according to the 20 superclasses of the CIFAR-100 dataset, and each superclass contains 5 different but semantically related classes. Whereas, TinyImageNet is constructed by splitting 200 classes into 40 5-way classification tasks. In our experiments, we do not use any data augmentation. The dataset statistics are given in Appendix D.1.

Network Architecture: For PMNIST, we use a fully connected network with two hidden layers of 100 units each following [22]. For experiments of CIFAR-100 Split and CIFAR-100 Superclass, we use a 5-layer AlexNet and LeNet architecture similar to [32] respectively. For TinyImageNet, we use the same network architecture as [11], which consists of 4 conv layers and 3 fully connected layers. In PMNIST, all tasks share the final classifier layer, while other experiments use a multi-head incremental setting, that is, each task has a separate classifier. Refer to Appendix D.2 for more details.

Baselines: We compare our method against multiple methods described below. (1) **EWC** [15], a regularization-based method that uses the diagonal of Fisher information to identify important weights; (2) **ICARL** [28], a memory-based method that uses knowledge-distillation and episodic memory to reduce forgetting; (3) **GEM** [22], another memory-based method that uses the gradient of episodic memory to constrain optimization to prevent forgetting; (4) **ER** [4], a simple and competitive method based on reservoir sampling; (5) **La-MAML** [11] and (6) **GPM** [32] are memory-based methods inspired by optimization-based meta-learning and based on gradient orthogonal constraints, respectively; (7) **Multitask** is an oracle baseline that all tasks are learned jointly using the entire dataset at once in a single network. Multitask is not a continual learning strategy but serves as an upper bound on average test accuracy on all tasks.

Training Details: All baselines and our method use stochastic gradient descent (SGD) for training. For each task in PMNIST and TinyImageNet, we train the network in 1 and 10 epochs, respectively, with the batch size as 10. These experimental settings are the same as La-MAML [11], so that we directly compare with their reported results. In the CIFAR-100 Split and CIFAR-100 Superclass experiments, we use the early termination strategy to train up to 50 epochs for each task, which is based on the validation loss as proposed in [33]. For both datasets, the batch size is set to 64. The replay buffer size of PMNIST, CIFAR-100 Split, CIFAR-100 Superclass, and TinyImageNet are 200, 1000, 1000, and 400, respectively. Details about the experimental setting and the hyperparameters considered for each baseline are provided in Appendix D.4.

Metrics: We evaluate the continual learning performance by the *average accuracy* (ACC) and *backward transfer* (BWT) [22, 3, 4], formulated as following,

$$ACC = \frac{1}{T}\sum_{i=1}^{T} R_{T,i}, \quad BWT = \frac{1}{T-1}\sum_{i=1}^{T-1} R_{T,i} - R_{i,i},$$

where T is the total number of learned sequential tasks, $R_{i,j}$ is the test classification accuracy of the model on j-th task after learning the last sample from i-th task. ACC is the average test classification accuracy of all tasks, bigger is better. BWT is the interference of new learning on the past knowledge. More specifically, negative BWT implies (catastrophic) forgetting whereas positive BWT indicates learning new task increases the performance on some preceding tasks.

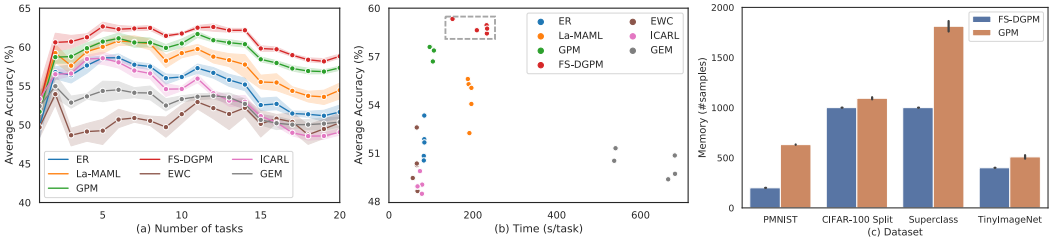

Figure 5: (a) Average accuracy as a function of the number of tasks trained on 20-Split CIFAR-100 Superclass. (b) Training time per task on 20-Split CIFAR-100 Superclass. (c) Memory usage on four datasets. ("Superclass" is abbr. of CIFAR-100 Superclass).

Table 1: Experimental results on 10-Split CIFAR-100, 20-Split CIFAR-100 Superclass and 40-Split TinyImageNet in 50 epochs. Each experiment is run with 5 seeds. † and * denotes results reported by [11] and [42] respectively.

	CIFAR-100 Split		**CIFAR-100 Superclass**		**TinyImageNet**	
Method	ACC(%)	BWT(%)	ACC(%)	BWT(%)	ACC(%)	BWT(%)
EWC	72.77 ± 0.45	-3.59 ± 0.55	50.26 ± 1.48	-7.87 ± 1.63	-	-
GEM	70.15 ± 0.34	-8.61 ± 0.42	50.35 ± 0.80	-9.50 ± 0.85	50.57 ± 0.61†	-20.50 ± 0.10†
ICARL	53.50 ± 0.81	-20.44 ± 0.82	49.05 ± 0.51	-11.24 ± 0.27	54.77 ± 0.32†	-3.93 ± 0.55†
ER	70.07 ± 0.35	-7.70 ± 0.59	51.64 ± 1.09	-7.86 ± 0.89	48.32 ± 1.51†	-19.86 ± 0.70†
La-MAML	71.37 ± 0.67	-5.39 ± 0.53	54.44 ± 1.36	-6.65 ± 0.85	66.90 ± 1.65†	-9.13 ± 0.90†
GPM	73.18 ± 0.52	**-1.17 ± 0.27**	57.33 ± 0.37	**-0.37 ± 0.12**	67.39 ± 0.47	**1.45 ± 0.22**
FS-DGPM	**74.33 ± 0.31**	-2.71 ± 0.17	**58.81 ± 0.34**	-2.97 ± 0.35	**70.41 ± 1.30**	-2.11 ± 0.84
Multitask	79.75 ± 0.38	-	61.00 ± 0.20*	-	77.10 ± 1.06†	-

5.2 Results and Discussion

PMNIST: We first evaluate our algorithm for 20 sequential PMNIST tasks with only 1000 samples per task in a single-head setting. From the results, as shown in Table 2, we see that our method (FS-DGPM) achieves the best average accuracy (**76.96%** ± **0.77**). Moreover, we achieve the least amount of forgetting except GPM, which is essentially a trade-off in accuracy to minimize forgetting. As shown in Figure 5(c), we only use about **31%** of the final memory size of GPM and achieve ∼ **2.5%** better accuracy.

Table 2: Experimental results (mean ± std in 5 runs) on PMNIST in single-epoch.

	PMNIST	
Method	ACC(%)	BWT(%)
EWC	62.25 ± 1.44	-15.22 ± 1.25
GEM	61.82 ± 0.85	-15.58 ± 1.17
ER	68.31 ± 0.56	-13.91 ± 0.67
La-MAML	75.98 ± 0.60	-10.21 ± 0.90
GPM	74.54 ± 0.36	**-7.17 ± 0.51**
FS-DGPM	**76.96 ± 0.51**	-7.45 ± 0.30
Multitask	86.54 ± 1.74	-

CIFAR-100 and TinyImageNet: Next, we use a multi-head setting to evaluate our algorithm under the more challenging visual classification benchmarks. Table 1 reports all results of these experiments. We outperform all baselines on three datasets, with achieving the best average accuracy **74.33%**, **58.81%** and **70.41%**. In these experiments, we observe that GPM is a strong baseline with the least forgetting. At the same time, we highlight that our method achieves the highest accuracy on all datasets and the second-lowest forgetting after GPM. Figure 5(a) shows the process of performance changing with the number of tasks on the CIFAR-100 Superclass. We consistently see the superior performance of our method at any stage of model evolution. It is also worth emphasizing that although our method requires more time for training and evaluation than GPM, it has lower memory usage and better test accuracy (See Figure 5(b)-(c)). As noted by [3] and [36], EWC performs poorly without multiple passes over the datasets, and GEM is not very effective under the single-head variants. These situations have also been observed in our experiments.

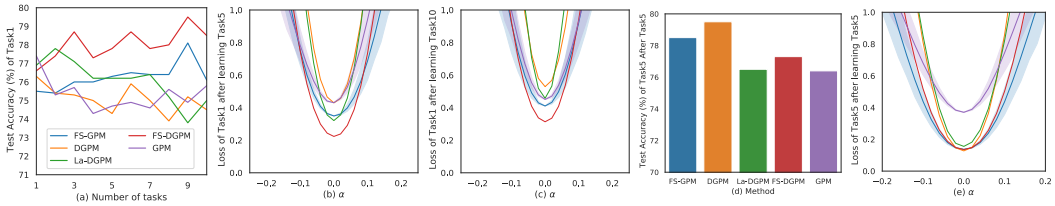

Figure 6: The ablation study implemented on CIFAR-100 Split with 10 tasks. (a)-(c) shows the stability of the first task. (a) is the test accuracy change curve of the first task; (b) and (c) are the weight loss landscape of the first task after learning the fifth task and all ten tasks. (d)-(e) shows the sensitivity of the fifth task, which are the test accuracy and the weight loss landscape of the fifth task after just learning the fifth task. The shape of the weight loss landscape obtained using ten different random filter-normalized directions. ("Task$i$" is abbr. of the i-th task)

5.3 Ablation studies on FS-DGPM

We further investigate our model performance with an ablation study and summarize it in Table 3. We respectively ablate the effects of flattening sharpness and dynamically adjusting the soft weight for bases. We refer to them as DGPM and FS-GPM. We also construct an ablation referred to as La-DGPM (Look-ahead DGPM), where the adversarial weight perturbation $\boldsymbol{v}$ in Eq. (3) is changed to the direction of gradient descent. At the same time, we also change the sign in Eq. (4) to ensure that the soft weight of the basis is reduced when the gradients are aligned in the same direction. From the results, shown in Table 3, we observe that flattening sharpness does benefits GPM, with $\sim$ **1.0%** improvement over GPM on both datasets. We can further observe through Figure 6 that all landscapes of FS-DGPM have lower loss values and flatter neighbors than DGPM and La-DGPM on the CIFAR-100 Split experiments. In addition, we see that DGPM performs well in learning new tasks, but it also leads to forgetting previous tasks. This situation can be efficiently alleviated by flattening sharpness. Hence, FS is indeed a powerful method worthy of being widely adopted for continual learning scenarios.

Table 3: The ablation study results on CIFAR-100 Split and Superclass. Each experiment is run with 5 seeds.

	CIFAR-100 Split		CIFAR-100 Superclass	
Method	ACC(%)	BWT(%)	ACC(%)	BWT(%)
FS-DGPM	**74.33** $\pm$ 0.31	**-2.71** $\pm$ 0.17	**58.81** $\pm$ 0.34	-2.97 $\pm$ 0.35
La-DGPM	73.74 $\pm$ 0.61	-3.05 $\pm$ 0.73	58.18 $\pm$ 0.41	**-2.41** $\pm$ 0.39
FS-GPM	73.96 $\pm$ 0.44	-3.12 $\pm$ 0.43	58.61 $\pm$ 0.53	-2.79 $\pm$ 0.30
DGPM	73.78 $\pm$ 0.32	-3.67 $\pm$ 0.42	56.78 $\pm$ 0.49	-2.44 $\pm$ 0.40
GPM	73.18 $\pm$ 0.52	**-1.17** $\pm$ 0.27	57.33 $\pm$ 0.37	**-0.37** $\pm$ 0.12

6 Conclusion

In this paper, we explore the weight loss landscape to characterize the well-known 'sensitivity-stability' dilemma faced by continual learning algorithms, and find that lower loss value with flatter neighbor often leads to better continual learning performance. Based on this finding, we propose our FS-DGPM algorithm, which introduces a soft weight to represent the importance of each basis representing past tasks in GPM, so that less important bases can be dynamically released to improve the sensitivity of new skill learning. Flattening Sharpness (FS) is also introduced here to reduce the generalization gap by explicitly regulating the flatness of the weight loss landscape of all tasks seen so far. The evaluation of various image classification tasks with different network architectures and the comparison with some state-of-the-art algorithms show the effectiveness of our method in achieving high classification performance while alleviating forgetting. Although our method theoretically and empirically demonstrates the advantages of introducing FS into continual learning, whether there exists a closer upper bound for the continual learning performance remains an unresolved problem and left for our future exploration.

References

[1] Hongjoon Ahn, Sungmin Cha, Donggyu Lee, and Taesup Moon. Uncertainty-based continual learning with adaptive regularization. *arXiv preprint arXiv:1905.11614*, 2019.

[2] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 139–154, 2018.

[3] Arslan Chaudhry, Marc'Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a-gem. *arXiv preprint arXiv:1812.00420*, 2018.

[4] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and M Ranzato. Continual learning with tiny episodic memories. 2019.

[5] Yu Chen, Tom Diethe, and Neil Lawrence. Facilitating bayesian continual learning by natural gradients and stein gradients. *arXiv preprint arXiv:1904.10644*, 2019.

[6] Sayna Ebrahimi, Mohamed Elhoseiny, Trevor Darrell, and Marcus Rohrbach. Uncertainty-guided continual learning with bayesian neural networks. *arXiv preprint arXiv:1906.02425*, 2019.

[7] Mehrdad Farajtabar, Navid Azizan, Alex Mott, and Ang Li. Orthogonal gradient descent for continual learning. In *International Conference on Artificial Intelligence and Statistics*, pages 3762–3773. PMLR, 2020.

[8] Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. *arXiv preprint arXiv:2010.01412*, 2020.

[9] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *arXiv preprint arXiv:1406.2661*, 2014.

[10] Yunhui Guo, Mingrui Liu, Tianbao Yang, and Tajana Rosing. Improved schemes for episodic memory-based lifelong learning. *Advances in Neural Information Processing Systems*, 33, 2020.

[11] Gunshi Gupta, Karmesh Yadav, and Liam Paull. La-maml: Look-ahead meta learning for continual learning. *arXiv preprint arXiv:2007.13904*, 2020.

[12] Wenpeng Hu, Zhou Lin, Bing Liu, Chongyang Tao, Zhengwei Tao, Jinwen Ma, Dongyan Zhao, and Rui Yan. Overcoming catastrophic forgetting for continual learning via model adaptation. In *International Conference on Learning Representations*, 2018.

[13] Daoyun Ji and Matthew A Wilson. Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nature neuroscience*, 10(1):100–107, 2007.

[14] Mohammad Khan, Didrik Nielsen, Voot Tangkaratt, Wu Lin, Yarin Gal, and Akash Srivastava. Fast and scalable bayesian deep learning by weight-perturbation in adam. In *International Conference on Machine Learning*, pages 2611–2620. PMLR, 2018.

[15] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.

[16] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.

[17] Quoc V Le, Jiquan Ngiam, Adam Coates, Ahbik Lahiri, Bobby Prochnow, and Andrew Y Ng. On optimization methods for deep learning. In *ICML*, 2011.

[18] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

[19] Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape of neural nets. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 6391–6401, 2018.

[20] Xilai Li, Yingbo Zhou, Tianfu Wu, Richard Socher, and Caiming Xiong. Learn to grow: A continual structure learning framework for overcoming catastrophic forgetting. In *International Conference on Machine Learning*, pages 3925–3934. PMLR, 2019.

[21] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017.

[22] David Lopez-Paz and Marc'Aurelio Ranzato. Gradient episodic memory for continual learning. *arXiv preprint arXiv:1706.08840*, 2017.

[23] Arun Mallya and Svetlana Lazebnik. 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.

[24] Seyed Iman Mirzadeh, Mehrdad Farajtabar, Razvan Pascanu, and Hassan Ghasemzadeh. Understanding the role of training regimes in continual learning. *arXiv preprint arXiv:2006.06958*, 2020.

[25] Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nathan Srebro. Exploring generalization in deep learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 5949–5958, 2017.

[26] Cuong V Nguyen, Yingzhen Li, Thang D Bui, and Richard E Turner. Variational continual learning. *arXiv preprint arXiv:1710.10628*, 2017.

[27] German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71, 2019.

[28] 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.

[29] Matthew Riemer, Ignacio Cases, Robert Ajemian, Miao Liu, Irina Rish, Yuhai Tu, and Gerald Tesauro. Learning to learn without forgetting by maximizing transfer and minimizing interference. *arXiv preprint arXiv:1810.11910*, 2018.

[30] Hippolyt Ritter, Aleksandar Botev, and David Barber. Online structured laplace approximations for overcoming catastrophic forgetting. *arXiv preprint arXiv:1805.07810*, 2018.

[31] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016.

[32] Gobinda Saha, Isha Garg, and Kaushik Roy. Gradient projection memory for continual learning. *arXiv preprint arXiv:2103.09762*, 2021.

[33] 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.

[34] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. *arXiv preprint arXiv:1705.08690*, 2017.

[35] Stanford. Tiny imagenet. *CS 231N*, Available online at http://cs231n.stanford.edu/tiny-imagenet-200.zip.

[36] Binh Tang and David S Matteson. Graph-based continual learning. *arXiv preprint arXiv:2007.04813*, 2020.

[37] Gido M van de Ven, Hava T Siegelmann, and Andreas S Tolias. Brain-inspired replay for continual learning with artificial neural networks. *Nature communications*, 11(1):1–14, 2020.

[38] Yeming Wen, Paul Vicol, Jimmy Ba, Dustin Tran, and Roger Grosse. Flipout: Efficient pseudo-independent weight perturbations on mini-batches. *arXiv preprint arXiv:1803.04386*, 2018.

[39] Matthew A Wilson and Bruce L McNaughton. Reactivation of hippocampal ensemble memories during sleep. *Science*, 265(5172):676–679, 1994.

[40] Dongxian Wu, Shu-Tao Xia, and Yisen Wang. Adversarial weight perturbation helps robust generalization. *Advances in Neural Information Processing Systems*, 33, 2020.

[41] Ju Xu and Zhanxing Zhu. Reinforced continual learning. *arXiv preprint arXiv:1805.12369*, 2018.

[42] Jaehong Yoon, Saehoon Kim, Eunho Yang, and Sung Ju Hwang. Scalable and order-robust continual learning with additive parameter decomposition. *arXiv preprint arXiv:1902.09432*, 2019.

[43] Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong learning with dynamically expandable networks. *arXiv preprint arXiv:1708.01547*, 2017.

[44] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, pages 3987–3995. PMLR, 2017.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] Refer to Section 6 for details.
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See Appendix D for more details.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See above.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See above.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See above.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [No] We will add the license as our code is released.
   (c) Did you include any new assets either in the supplemental material or as a URL? [No] We plan to release our code after the blind review process finishes.
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]