

CoSCL: Cooperation of Small Continual Learners is Stronger than a Big One

Liyuan Wang^{1,2,3*}, Xingxing Zhang^{3*}, Qian Li^{1,2}, Jun Zhu^{3†}, and Yi Zhong^{1,2†}

¹ School of Life Sciences, IDG/McGovern Institute for Brain Research, Tsinghua University. ² Tsinghua-Peking Center for Life Sciences. ³ Dept. of Comp. Sci. & Tech., Institute for AI, BNRist Center, THBI Lab, Tsinghua University.
{wly19,xxzhang2020}@mails.tsinghua.edu.cn,
{liqian8,dcszj,zhongyithu}@tsinghua.edu.cn

Abstract. Continual learning requires incremental compatibility with a sequence of tasks. However, the design of model architecture remains an open question: In general, learning all tasks with a shared set of parameters suffers from severe interference between tasks; while learning each task with a dedicated parameter subspace is limited by scalability. In this work, we theoretically analyze the generalization errors for learning plasticity and memory stability in continual learning, which can be uniformly upper-bounded by (1) discrepancy between task distributions, (2) flatness of loss landscape and (3) cover of parameter space. Then, inspired by the robust biological learning system that processes sequential experiences with multiple parallel compartments, we propose Cooperation of Small Continual Learners (CoSCL) as a general strategy for continual learning. Specifically, we present an architecture with a fixed number of narrower sub-networks to learn all incremental tasks in parallel, which can naturally reduce the two errors through improving the three components of the upper bound. To strengthen this advantage, we encourage to cooperate these sub-networks by penalizing the difference of predictions made by their feature representations. With a fixed parameter budget, CoSCL can improve a variety of representative continual learning approaches by a large margin (e.g., up to 10.64% on CIFAR-100-SC, 9.33% on CIFAR-100-RS, 11.45% on CUB-200-2011 and 6.72% on Tiny-ImageNet) and achieve the new state-of-the-art performance. Our code is available at <https://github.com/lywang3081/CoSCL>.

Keywords: Continual Learning; Catastrophic Forgetting; Ensemble Model.

1 Introduction

The ability to incrementally learn a sequence of tasks is critical for artificial neural networks. Since the training data distribution is typically dynamic and unpredictable, this usually requires a careful trade-off between learning plasticity and memory stability. In general, excessive plasticity in learning new tasks

*Liyuan Wang and Xingxing Zhang contributed equally.

†Jun Zhu and Yi Zhong are corresponding authors.

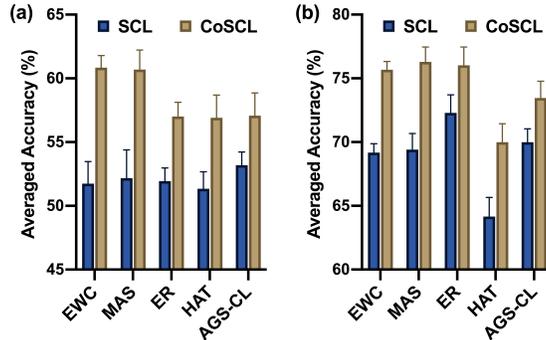


Fig. 1. Comparison of a single continual learner (SCL) and CoSCL (Ours) on (a) CIFAR-100-SC and (b) CIFAR-100-RS.

leads to the catastrophic forgetting of old tasks [25], while excessive stability in remembering old tasks limits the learning of new tasks. Most efforts in continual learning either use a single model to learn all tasks, which has to sacrifice the performance of each task to find a shared solution [33]; or allocate a dedicated parameter subspace for each task to overcome their mutual interference [14, 39], which usually lacks scalability. Recent work observed that a wider network can suffer from less catastrophic forgetting [26], while different components such as batch normalization, skip connections and pooling layers play various roles [27]. Thus, how to achieve effective continual learning in terms of model architecture remains an open question.

In contrast, the robust biological learning system applies multiple compartments (i.e, sub-networks) to process sequential experiences in parallel, and integrates their outputs in a weighted-sum fashion to guide adaptive behaviors [3, 9, 30]. This provides a promising reference for artificial neural networks.

In this work, we first theoretically analyze the generalization errors of learning plasticity and memory stability in continual learning. We identify that both aspects can be uniformly upper-bounded by (1) *discrepancy between task distributions*, (2) *flatness of loss landscape* and (3) *cover of parameter space*. Inspired by the biological strategy, we propose a novel method named Cooperation of Small Continual Learners (CoSCL). Specifically, we design an architecture with multiple narrower sub-networks¹ to learn all incremental tasks in parallel, which can naturally alleviate the both errors through improving the three components. To strengthen this advantage, we further encourage the cooperation of sub-networks by penalizing differences in the predictions of their feature representations.

With a fixed parameter budget, CoSCL can substantially boost a variety of representative continual learning strategies (e.g., up to **10.64%** on CIFAR-100-SC, **9.33%** on CIFAR-100-RS, **11.45%** on CUB-200-2011 and **6.72%** on Tiny-ImageNet, detailed in Fig. 1 and Table 1). The superior performance comes

¹In contrast to a single continual learning model with a wide network, we refer to such narrower sub-networks as “small” continual learners.

from reducing the errors in both learning plasticity and memory stability by tightening the upper bound, consistent with our theoretical analysis.

Our contributions include: (1) We present a unified form of generalization bounds for learning plasticity and memory stability in continual learning; (2) The generalization bounds suggest that the two aspects are not necessarily in conflict, but can be simultaneously enhanced in a compatible parameter space of a well-designed model architecture; (3) To achieve this goal, we draw inspirations from the biological strategy and propose to cooperate multiple (small) continual learners; (4) Extensive experiments validate the efficacy and generality of our proposal, which can be adapted to a variety of representative continual learning approaches and improve their performance by a large margin.

2 Related Work

Continual Learning requires effective learning of incremental tasks without severe catastrophic forgetting. Representative strategies include weight regularization [18, 50, 1], memory replay [34, 42, 44], parameter isolation [39, 17] and dynamic architecture [14, 48]. These strategies either learn all tasks with a single model, which have to compromise the performance of each task to obtain a shared solution [33]; or allocate parameter subspace for each task to prevent mutual interference, yet limited by scalability. Several recent work tried to improve continual learning in terms of architecture, such as by using neural architecture search [32] or learning an additional set of shared parameters [16], but to a limited extent. [33] proposed a model zoo that incrementally adds sub-networks to learn new tasks, which had to store a large amount of old training samples.

Flatness of Loss Landscape provides a conceptual explanation of generalization for deep neural networks, which is recently introduced to understand catastrophic forgetting in continual learning [29, 40, 11, 6, 23]. The core idea is that convergence to a smooth region will be more robust to (mild) parameter changes. [29, 12] analyzed that the forgetting of old tasks in continual learning can be bounded by the variation of parameters between tasks and the eigenvalues of the Hessian matrix, where the lower eigenvalues indicate a flatter curvature of the solution. [6, 40, 11] explicitly encouraged the network to find a flat minima and empirically validated its efficacy in continual learning.

Ensemble Model is a powerful architecture to improve generalization, but is still under explored in continual learning. Most current applications focus on learning each single task with a sub-network [37, 2, 47], which can be seen as a special case of dynamic architecture. The main limitation is that the total amount of parameters (resp., the storage and computational cost) might grow linearly with the number of incremental tasks. [45] proposed an efficient ensemble strategy to reduce extra parameter cost for task-specific sub-networks. Similar to ours, a *concurrent* work [13] also observed that ensemble of multiple continual learning models brings huge benefits. They further exploited recent advances of mode connectivity [28] and neural network subspace [46] to save computational cost, but had to use old training samples [13]. Besides, [20] achieved more ef-

fective weight regularization by ensemble of multiple auxiliary classifiers learned from extra out-of-distribution data (e.g., SVHN [31] for CIFAR-100 [19]).

Main Advantages of Our Work are summarized in three aspects: (1) The generalization bounds presented in our work demonstrate the *direct* link between continual learning performance and flatness of loss landscape (as well as other components). (2) We use a *fixed* number of sub-networks, which are all continual learners rather than single-task learners, and adjust their width accordingly, so no additional or growing parameters are needed. (3) We mainly focus on a restrict setting where old training samples or extra data sources are *not* needed, which is more general and realistic for continual learning.

3 Preliminary Analysis

In this section, we first introduce the problem formulation and representative continual learning strategies, and then present the generalization bounds.

3.1 Problem Formulation

Let’s consider a general setting of continual learning: A neural network with parameter θ incrementally learns T tasks, called a *continual learner*. The training set and test set of each task follow the same distribution \mathbb{D}_t ($t = 1, 2, \dots, T$), where the training set $D_t = \{(x_{t,n}, y_{t,n})\}_{n=1}^{N_t}$ includes N_t data-label pairs. For classification task, it might include one or several classes. After learning each task, the performance of all the tasks ever seen is evaluated on their test sets. Although D_t is only available when learning task t , an ideal continual learner should behave as if training them jointly. To achieve this goal, it is critical to balance learning plasticity of new tasks and memory stability of old tasks. Accordingly, the loss function for continual learning can typically be defined as

$$L_{\text{CL}}(\theta) = L_t(\theta) + \lambda \hat{L}_{1:t-1}(\theta), \quad (1)$$

where $L_t(\cdot)$ is the task-specific loss for learning task t (e.g., cross-entropy for supervised classification), and $\hat{L}_{1:t-1}(\cdot)$ provides the constraint to achieve a proper trade-off between new and old tasks. For example, $\hat{L}_{1:t-1}(\theta) = \sum_i I_{1:t-1,i}(\theta_i - \theta_{1:t-1,i}^*)^2$ for weight regularization [18, 1, 50], where $\theta_{1:t-1,i}^*$ denotes the continually-learned solution for old tasks and $I_{1:t-1}$ indicates the “importance” of each parameter. $\hat{L}_{1:t-1}(\theta) = \sum_{k=1}^{t-1} L_k(\theta; \hat{D}_k)$ for memory replay [34, 42, 44], where \hat{D}_k is an approximation of D_k through storing old training samples or learning a generative model. For parameter isolation [39, 17], $\theta = \{\bigcup_{k=1}^{t-1} \hat{\theta}_k, \hat{\theta}_{\text{free}}\}$ is dynamically isolated as multiple task-specific subspaces $\hat{\theta}_k$, while $\hat{\theta}_{\text{free}}$ denotes the “free” parameters for current and future tasks. So $\hat{L}_{1:t-1}(\theta)$ usually serves as a sparsity regularizer to save $\hat{\theta}_{\text{free}}$. For dynamic architecture [14, 48], $\theta = \{\bigcup_{k=1}^{t-1} \hat{\theta}_k, \hat{\theta}_t\}$ attempts to add a new subspace $\hat{\theta}_t$ on the basis of the previous ones, and $\hat{L}_{1:t-1}(\theta)$ should limit the amount of extra parameters.

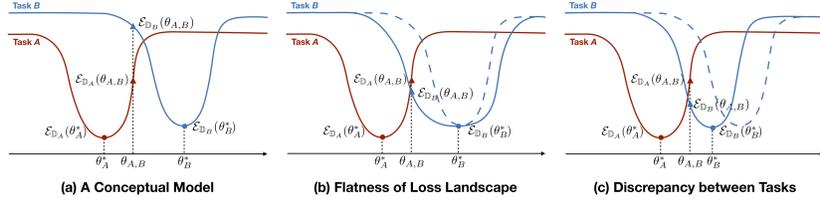


Fig. 2. A conceptual model of two tasks (the learning order of Task A and Task B does not matter). The dashed line in (b) and (c) is the original solution in (a), where finding a flatter solution or reducing the discrepancy between tasks help to mitigate the generalization errors of a shared solution $\theta_{A,B}$.

3.2 Generalization Bound for Continual Learning

Formally, the goal of continual learning is to find a solution θ in a parameter space Θ that can generalize well over a set of distribution \mathbb{D}_t and $\mathbb{D}_{1:t-1} := \{\mathbb{D}_k\}_{k=1}^{t-1}$. Let's consider a bounded loss function $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow [0, c]$ (where \mathcal{Y} denotes a label space and c is the upper bound), such that $\ell(y_1, y_2) = 0$ holds if and only if $y_1 = y_2$. Then, we can define a population loss over the distribution \mathbb{D}_t by $\mathcal{E}_{\mathbb{D}_t}(\theta) = \mathbb{E}_{(x,y) \sim \mathbb{D}_t}[\ell(f_\theta(x), y)]$, where $f_\theta(\cdot)$ is the prediction of an input parameterized by θ . Likewise, the population loss over the distribution of old tasks is defined by $\mathcal{E}_{\mathbb{D}_{1:t-1}}(\theta) = \frac{1}{t-1} \sum_{k=1}^{t-1} \mathbb{E}_{(x,y) \sim \mathbb{D}_k}[\ell(f_\theta(x), y)]$. To minimize both $\mathcal{E}_{\mathbb{D}_t}(\theta)$ and $\mathcal{E}_{\mathbb{D}_{1:t-1}}(\theta)$, a continual learning model (i.e., a continual learner) needs to minimize an empirical risk over the current training set D_t in a constrained parameter space, i.e., $\min_{\theta \in \Theta} \hat{\mathcal{E}}_{D_t}(\theta)$. Specifically, $\hat{\mathcal{E}}_{D_t}(\theta) = \frac{1}{N_t} \sum_{n=1}^{N_t} \ell(f_\theta(x_{t,n}), y_{t,n})$, and the constrained parameter space Θ depends on the previous experience carried by parameters, data, and/or task labels, so as to prevent catastrophic forgetting. Likewise, $\hat{\mathcal{E}}_{D_{1:t-1}}(\theta)$ denotes an empirical risk over the old tasks. In practice, sequential learning of each task by minimizing the empirical risk $\hat{\mathcal{E}}_{D_t}(\theta)$ in Θ can find multiple solutions, but provides significantly different generalizability on $\mathcal{E}_{\mathbb{D}_t}(\theta)$ and $\mathcal{E}_{\mathbb{D}_{1:t-1}}(\theta)$. Several recent studies suggested that a flatter solution is more robust to catastrophic forgetting [29, 40, 11, 6, 23]. To find such a flat solution, we define a robust empirical risk by the worst case of the neighborhood in parameter space as $\hat{\mathcal{E}}_{D_t}^b(\theta) := \max_{\|\Delta\| \leq b} \hat{\mathcal{E}}_{D_t}(\theta + \Delta)$ [5], where b is the radius around θ and $\|\cdot\|$ denotes the L2 norm, likewise for the old tasks as $\hat{\mathcal{E}}_{D_{1:t-1}}^b(\theta) := \max_{\|\Delta\| \leq b} \hat{\mathcal{E}}_{D_{1:t-1}}(\theta + \Delta)$. Then, solving the constrained robust empirical risk minimization, i.e., $\min_{\theta \in \Theta} \hat{\mathcal{E}}_{D_t}^b(\theta)$, will find a near solution of a flat optimum showing better generalizability. In particular, the minima found by the empirical loss $\hat{\mathcal{E}}_{D_t}(\theta)$ will also be the minima of $\hat{\mathcal{E}}_{D_t}^b(\theta)$ if the “radius” of its loss landscape is sufficiently wider than b . Intuitively, such a flat solution helps to mitigate catastrophic forgetting since it is more robust to parameter changes.

However, this connection is not sufficient. If a new task is too different from the old tasks, the parameter changes to learn it well might be much larger than the “radius” of the old minima, resulting in catastrophic forgetting. On the other hand, staying around the old minima is not a good solution for the new task,

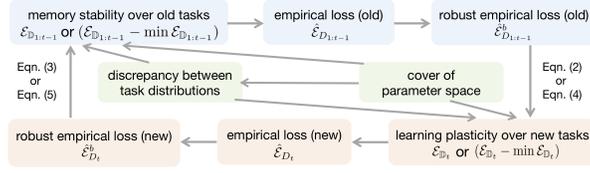


Fig. 3. Illustration of simultaneously promoting learning plasticity and memory stability in continual learning, where arrows represent the tightening process.

limiting learning plasticity. Let $\mathcal{E}_{\mathbb{D}_t}(\theta_{1:t})$ and $\mathcal{E}_{\mathbb{D}_{1:t-1}}(\theta_{1:t})$ denote the generalization errors of performing the new task and old tasks, respectively. Inspired by the PAC-Bayes theory [24] and previous work in domain generalization [5, 4], we present the upper bounds of these two errors as follows (proof in Appendix A):

Proposition 1. *Let Θ be a cover of a parameter space with VC dimension d . If $\mathbb{D}_1, \dots, \mathbb{D}_t$ are the distributions of the continually learned $1 : t$ tasks, then for any $\delta \in (0, 1)$ with probability at least $1 - \delta$, for every solution $\theta_{1:t}$ of the continually learned $1 : t$ tasks in parameter space Θ , i.e., $\theta_{1:t} \in \Theta$:*

$$\mathcal{E}_{\mathbb{D}_t}(\theta_{1:t}) < \hat{\mathcal{E}}_{D_{1:t-1}}^b(\theta_{1:t}) + \frac{1}{2(t-1)} \sum_{k=1}^{t-1} \text{Div}(\mathbb{D}_k, \mathbb{D}_t) + \sqrt{\frac{d \ln(N_{1:t-1}/d) + \ln(1/\delta)}{N_{1:t-1}}}, \quad (2)$$

$$\mathcal{E}_{\mathbb{D}_{1:t-1}}(\theta_{1:t}) < \hat{\mathcal{E}}_{D_t}^b(\theta_{1:t}) + \frac{1}{2(t-1)} \sum_{k=1}^{t-1} \text{Div}(\mathbb{D}_t, \mathbb{D}_k) + \sqrt{\frac{d \ln(N_t/d) + \ln(1/\delta)}{N_t}}, \quad (3)$$

where $\text{Div}(\mathbb{D}_i, \mathbb{D}_j) := 2 \sup_{h \in \mathcal{H}} |\mathcal{P}_{\mathbb{D}_i}(I(h)) - \mathcal{P}_{\mathbb{D}_j}(I(h))|$ is the \mathcal{H} -divergence for the distribution \mathbb{D}_i and \mathbb{D}_j ($I(h)$ is the characteristic function). $N_{1:t-1} = \sum_{k=1}^{t-1} N_k$ is the total number of training samples over all old tasks.

It can be concluded from Proposition 1 that, the generalization errors over the new task and old tasks are uniformly constrained by three components: (1) *discrepancy between task distributions*; (2) *flatness of loss landscape*; and (3) *cover of parameter space*. By the optimal solution for (robust) empirical loss, we further demonstrate that the generalization gaps of the new task and old tasks are upper bounded as follows (proof in Appendix A):

Proposition 2. *Let $\hat{\theta}_{1:t}^b$ denotes the optimal solution of the continually learned $1 : t$ tasks by robust empirical risk minimization over the current task, i.e., $\hat{\theta}_{1:t}^b = \arg \min_{\theta \in \Theta} \hat{\mathcal{E}}_{D_t}^b(\theta)$, where Θ denotes a cover of a parameter space with VC dimension d . Then for any $\delta \in (0, 1)$, with probability at least $1 - \delta$:*

$$\mathcal{E}_{\mathbb{D}_t}(\hat{\theta}_{1:t}^b) - \min_{\theta \in \Theta} \mathcal{E}_{\mathbb{D}_t}(\theta) \leq \min_{\theta \in \Theta} \hat{\mathcal{E}}_{D_{1:t-1}}^b(\theta) - \min_{\theta \in \Theta} \hat{\mathcal{E}}_{D_{1:t-1}}(\theta) + \frac{1}{t-1} \sum_{k=1}^{t-1} \text{Div}(\mathbb{D}_k, \mathbb{D}_t) + \lambda_1, \quad (4)$$

$$\mathcal{E}_{\mathbb{D}_{1:t-1}}(\hat{\theta}_{1:t}^b) - \min_{\theta \in \Theta} \mathcal{E}_{\mathbb{D}_{1:t-1}}(\theta) \leq \min_{\theta \in \Theta} \hat{\mathcal{E}}_{D_t}^b(\theta) - \min_{\theta \in \Theta} \hat{\mathcal{E}}_{D_t}(\theta) + \frac{1}{t-1} \sum_{k=1}^{t-1} \text{Div}(\mathbb{D}_t, \mathbb{D}_k) + \lambda_2, \quad (5)$$

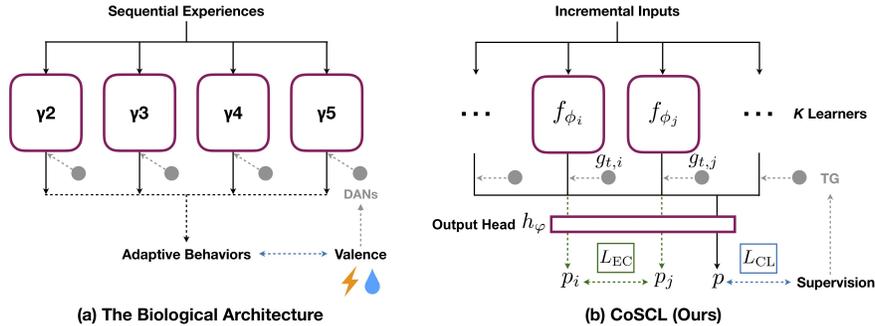


Fig. 4. (a) Fruit flies learn sequential experiences with multiple parallel compartments (γ 2-5), under the modulation of dopaminergic neurons (DANs) that convey valence. The figure is modified from [9]. (b) Inspired by the biological learning system, we propose a general strategy of cooperating multiple (small) continual learners.

where $\lambda_1 = 2\sqrt{\frac{d \ln(N_{1:t-1}/d) + \ln(2/\delta)}{N_{1:t-1}}}$, $\lambda_2 = 2\sqrt{\frac{d \ln(N_t/d) + \ln(2/\delta)}{N_t}}$, and $\text{Div}(\mathbb{D}_i, \mathbb{D}_j) := 2 \sup_{h \in \mathcal{H}} |\mathcal{P}_{\mathbb{D}_i}(I(h)) - \mathcal{P}_{\mathbb{D}_j}(I(h))|$ is the \mathcal{H} -divergence for the distribution \mathbb{D}_i and \mathbb{D}_j ($I(h)$ is the characteristic function).

Likewise, the generalization gaps over the new and old tasks are also constrained by the three components above. In particular, learning plasticity and memory stability in continual learning can be simultaneously promoted by using a more compatible parameter space, as illustrated in Fig. 3. Specifically, compatibility with the new task can facilitate a smaller robust empirical risk on the old tasks as well as improve task discrepancy, then tightening the generalization bound for learning plasticity through Eqn. (2)/Eqn. (4), and vice versa tightening the generalization bound for memory stability through Eqn. (3)/Eqn. (5).

4 Method

Unlike artificial neural networks, the robust biological learning system, such as that of fruit flies, processes sequential experiences with multiple parallel compartments (i.e., sub-networks) [3, 9]. These compartments are modulated by dopaminergic neurons (DANs) that convey valence (i.e., supervised signals), and their outputs are integrated in a weighted-sum fashion to guide adaptive behaviors [3, 9, 30] (detailed in Fig. 4, a). Inspired by this, we propose to cooperate multiple (small) continual learners as a simple yet effective method for continual learning. We present our proposal in Sec. 4.1, and validate this idea both theoretically (Sec. 4.2) and empirically (Sec. 5).

4.1 Cooperation of (Small) Continual Learners

Instead of learning all tasks with a single continual learner, we design a bio-inspired architecture to coordinate multiple continual learners. Specifically, each

continual learner is implemented with a sub-network $f_{\phi_i}(\cdot), i = 1, \dots, K$ in a parameter space for learning all incremental tasks, where the dedicated output head is removed and the output of the previous layer is weighted by a set of learnable parameters (usually a fully connected layer). Then, these outputs are fed into a shared output head $h_\varphi(\cdot)$ for prediction. For a regular classifier, this is equivalent to making predictions on a weighted-sum of feature representations, so we refer to this strategy as *feature ensemble* (FE).

When task labels are available, our architecture can more effectively incorporate task-specific information by learning an additional set of *task-adaptive gates* (TG) for each continual learner’s output. Such a gate is defined as $g_{t,i} = \sigma(s \cdot \alpha_{t,i})$ for learner i to perform task t , where $\alpha_{t,i}$ is a learnable parameter, s is a scale factor and σ denotes the sigmoid function. Therefore, the final prediction becomes $p(\cdot) = h_\varphi(\sum_{i=1}^K g_{t,i} f_{\phi_i}(\cdot))$, and all optimizable parameters include $\bar{\theta} = \{\cup_{i=1}^K \phi_i, \cup_{i=1}^K \alpha_{t,i}, \varphi\}$.

To strengthen the advantage of feature ensemble, we encourage to cooperate the continual learners by penalizing differences in the predictions of their feature representations (e.g., \mathbf{p}_i and \mathbf{p}_j). We choose the widely-used Kullback Leibler (KL) divergence and define an *ensemble cooperation* (EC) loss as

$$\begin{aligned} L_{\text{EC}}(\bar{\theta}) &= \frac{1}{K} \sum_{i=1, j \neq i}^K D_{\text{KL}}(\mathbf{p}_i || \mathbf{p}_j) = \frac{1}{K} \frac{1}{N_t} \sum_{i=1, j \neq i}^K \sum_{n=1}^{N_t} p_i(x_{t,n}) \log \frac{p_i(x_{t,n})}{p_j(x_{t,n})} \\ &= \frac{1}{K} \frac{1}{N_t} \sum_{i=1, j \neq i}^K \sum_{n=1}^{N_t} h_\varphi(g_{t,i} f_{\phi_i}(x_{t,n})) \log \frac{h_\varphi(g_{t,i} f_{\phi_i}(x_{t,n}))}{h_\varphi(g_{t,j} f_{\phi_j}(x_{t,n}))}. \end{aligned} \quad (6)$$

In practice, we reduce the sub-network width to save parameters, so we call our method “Cooperation of Small Continual Learners (CoSCL)”. Taking Eqn. (1) and Eqn. (6) together, the objective of CoSCL is defined as

$$L_{\text{CoSCL}}(\bar{\theta}) = L_{\text{CL}}(\bar{\theta}) + \gamma L_{\text{EC}}(\bar{\theta}). \quad (7)$$

4.2 Theoretical Explanation

Here we provide a theoretical explanation of how cooperating multiple continual learners can mitigate the generalization gaps in continual learning:

Proposition 3. Let $\{\Theta_i \in \mathbb{R}^r\}_{i=1}^K$ be a set of K parameter spaces ($K > 1$ in general), d_i be a VC dimension of Θ_i , and $\Theta = \cup_{i=1}^K \Theta_i$ with VC dimension d . Based on Proposition 2, for $\hat{\theta}_{1:t}^b = \arg \min_{\bar{\theta} \in \Theta} \hat{\mathcal{E}}_{D_t}^b(\bar{\theta})$, the upper bound of generalization gap is further tighter with

$$\lambda_1 = \max_{i \in [1, K]} \sqrt{\frac{d_i \ln(N_{1:t-1}/d_i) + \ln(2K/\delta)}{N_{1:t-1}}} + \sqrt{\frac{d \ln(N_{1:t-1}/d) + \ln(2/\delta)}{N_{1:t-1}}}, \quad (8)$$

$$\lambda_2 = \max_{i \in [1, K]} \sqrt{\frac{d_i \ln(N_t/d_i) + \ln(2K/\delta)}{N_t}} + \sqrt{\frac{d \ln(N_t/d) + \ln(2/\delta)}{N_t}}. \quad (9)$$

Comparing Proposition 3 and Proposition 2, we conclude that cooperating K continual learners facilitates a smaller generalization gap over the new and old tasks in continual learning than a single one. Due to the space limit, we leave more details of Proposition 3 in Appendix A, where we also analyze how a compatible parameter space of a well-designed model architecture can improve the discrepancy between task distributions, thus further tightening the generalization bounds. Next, we empirically validate our proposal as detailed below.

5 Experiment

In this section, we extensively evaluate CoSCL on visual classification tasks. All results are averaged over 5 runs with different random seeds and task orders.

Benchmark: We consider four representative continual learning benchmarks. The first two are with CIFAR-100 dataset [19], which includes 100-class colored images of the size 32×32 . All classes are split into 20 incremental tasks, based on random sequence (RS) or superclass (SC). The other two are with larger-scale datasets, randomly split into 10 incremental tasks: CUB-200-2011 [41] includes 200 classes and 11,788 bird images of the size 224×224 , and is split as 30 images per class for training while the rest for testing. Tiny-ImageNet [10] is derived from iILSVRC-2012 [36], consisting of 200-class natural images of the size 64×64 .

Implementation: We mainly focus on the task-incremental setting used in [43, 17, 6, 39] and follow their implementation for most experiments if not specified. For all the baselines, we apply a 6-layer CNN architecture for CIFAR-100-SC and CIFAR-100-RS, and an AlexNet-based architecture for CUB-200-2011 and Tiny-ImageNet.² Since our method consists of multiple continual learners, we use a similar architecture for each sub-network and accordingly reduce the width (i.e., using fewer channels) to keep the total number of parameters comparable to other baselines, so as to make the comparison as fair as possible. Then, there is an intuitive trade-off between the number and width of learners. According to our theoretical analysis in Proposition 1 and 2, the choice for the number of learners (i.e., parameter spaces) K is *independent* of the training data distribution under a limited parameter budget. Also, we empirically validate that this trade-off is only moderately sensitive (see Fig. 5). So we simply set $K = 5$ for all experiments.

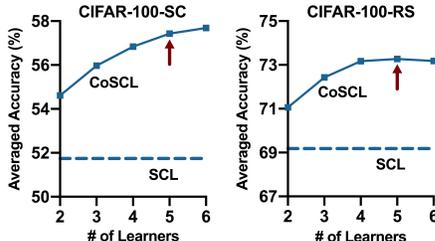


Fig. 5. Trade-off between learner number and width with a fixed parameter budget.

²A concurrent work observed that the regular CNN architecture indeed achieves better continual learning performance than more advanced architectures such as ResNet and ViT with the same amount of parameters [27].

Table 1. Averaged accuracy (%) of all the tasks learned so far in continual learning (A_t for t tasks). All results are cited from [43, 17, 6] or reproduced from their officially-released code for a fair comparison. CoSCL cooperates 5 continual learners with similar architectures as other baselines, while reducing the sub-network width accordingly to keep the total amount of parameters comparable.

Methods	CIFAR-100-SC		CIFAR-100-RS		CUB-200-2011		Tiny-ImageNet	
	A_{10}	A_{20}	A_{10}	A_{20}	A_5	A_{10}	A_5	A_{10}
SI [50]	52.20 \pm 4.37	51.97 \pm 2.07	68.72 \pm 1.11	69.21 \pm 0.77	33.08 \pm 4.05	42.03 \pm 3.06	45.61 \pm 2.05	46.00 \pm 1.13
RWALK [7]	50.51 \pm 4.53	49.62 \pm 3.28	66.02 \pm 1.89	66.90 \pm 0.29	32.56 \pm 3.76	41.94 \pm 2.35	49.69 \pm 1.47	48.12 \pm 0.96
P&C [38]	53.48 \pm 2.79	52.88 \pm 1.68	70.10 \pm 1.22	70.21 \pm 1.22	33.88 \pm 4.48	42.79 \pm 3.29	51.71 \pm 1.58	50.33 \pm 0.86
EWC [18]	52.25 \pm 2.99	51.74 \pm 1.74	68.72 \pm 0.24	69.18 \pm 0.69	32.90 \pm 2.98	42.29 \pm 2.34	50.92 \pm 1.86	48.38 \pm 0.86
w/ AFEC [43]	56.28 \pm 3.27	55.24 \pm 1.61	72.36 \pm 1.23	72.29 \pm 1.07	34.36 \pm 4.39	43.05 \pm 3.00	51.34 \pm 1.62	50.58 \pm 0.74
w/ CPR [6]	54.60 \pm 2.51	53.37 \pm 2.06	71.12 \pm 1.82	70.25 \pm 1.33	33.36 \pm 3.25	42.51 \pm 2.31	50.12 \pm 1.43	50.29 \pm 0.89
w/ CoSCL (Ours)	62.89 \pm 3.05	60.84 \pm 0.95	78.08 \pm 1.25	76.05 \pm 0.65	44.35 \pm 3.59	48.53 \pm 2.21	56.10 \pm 1.77	55.10 \pm 1.02
MAS [1]	52.76 \pm 2.85	52.18 \pm 2.22	67.60 \pm 1.85	69.41 \pm 1.27	31.68 \pm 2.37	42.56 \pm 1.84	49.69 \pm 1.50	50.20 \pm 0.82
w/ AFEC [43]	55.26 \pm 4.14	54.89 \pm 2.23	69.57 \pm 1.73	71.20 \pm 0.70	34.08 \pm 3.80	42.93 \pm 3.51	51.35 \pm 1.75	50.90 \pm 1.08
w/ CPR [6]	52.90 \pm 1.62	53.63 \pm 1.31	70.69 \pm 1.85	72.06 \pm 1.86	33.49 \pm 2.46	43.07 \pm 2.56	50.82 \pm 1.41	51.24 \pm 1.26
w/ CoSCL (Ours)	62.55 \pm 1.94	60.69 \pm 1.53	76.93 \pm 1.94	76.29 \pm 2.33	43.67 \pm 3.73	49.48 \pm 2.40	55.43 \pm 1.48	55.11 \pm 0.89

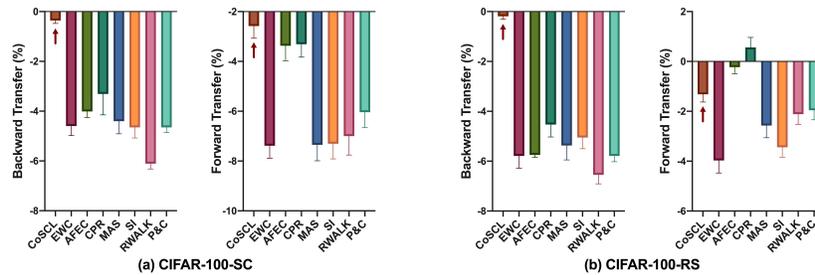


Fig. 6. Comparison of backward transfer (BWT) and forward transfer (FWT).

The learners’ training differs only in random initialization of the parameters. The implementations are further detailed in Appendix B.

Overall Performance: We first adapt CoSCL to representative continual learning strategies, including weight regularization such as EWC [18] and MAS [1], parameter isolation such as HAT [39] and AGS-CL [17], and experience replay (ER) of old training samples (20 images per class) [35]. As shown in Fig. 1, our proposal that cooperates multiple continual learners with narrower sub-networks can largely improve their performance. Then, we compare with the state-of-the-art (SOTA) methods under a realistic restriction that old training samples or additional data sources are *not* available, as detailed below.

First, we compare with the SOTA methods that can be plug-and-play with weight regularization baselines, such as AFEC [43] and CPR [6]. AFEC [43] encouraged the network parameters to resemble the optimal solution for each new task to mitigate potential negative transfer, while CPR [6] added a regularization

Table 2. Averaged accuracy (%) of architecture-based methods on CIFAR-100-RS. Here we use EWC as the default continual learning method for CoSCL.

Methods	# Param	20-split	50-split
HAT [39]	6.8M	76.96	80.46
MARK [16]	4.7M	78.31	–
BNS [32]	6.7M	–	82.39
CoSCL (Ours)	4.6M	79.43 ± 1.01	87.88 ± 1.07

term that maximized the entropy of output probability to find a flat minima. In contrast, CoSCL can more effectively improve the weight regularization baselines and achieve the new SOTA performance (detailed in Table 1).

At the same time, we consider the SOTA methods that improve continual learning in terms of architecture, such as BNS [32] and MARK [16]. BNS applied neural structure search to build a network for preventing catastrophic forgetting and promoting knowledge transfer, while MARK achieved this goal by learning an additional set of shared weights among tasks.³ With a smaller parameter budget, ours largely outperforms the two recent strong baselines (see Table 2).

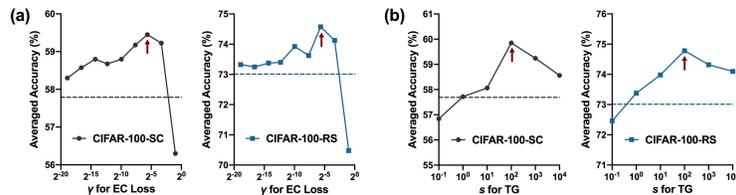
Detailed Analysis: Now, we use EWC [18] as the default continual learning method and provide a detailed analysis for the superior performance of CoSCL. First, we analyze the **knowledge transfer** among tasks by evaluating the metrics of backward transfer (BWT), which is the averaged influence of learning each new task to the old tasks, and forward transfer (FWT), which is the averaged influence of remembering the old tasks to each new task [22]. As shown in Fig. 6, CoSCL substantially improves both BWT and FWT of the default method, and in general far exceeds other representative baselines implemented in a single model. In particular, CoSCL raises BWT to almost zero, which means that catastrophic forgetting can be completely avoided. We also evaluate the **expertise** of each continual learner across tasks in Appendix C.1. The predictions made by each continual learner’s representations differ significantly and complement with each other. The functional diversity can be naturally obtained from the randomness in architecture, such as the use of dropout and a different random initialization for each learner, and is explicitly regulated by our ensemble cooperation loss (Fig. 7, a, discussed later).

Next, we present the results of an **ablation study** in Table 3. We first consider a naive baseline that averages the predictions of multiple independently-trained small continual learners, referred to as the “classifier ensemble (CE)”. However, such a naive baseline even underperforms the single continual learner (SCL). In contrast, the proposed feature ensemble (FE) of multiple small continual learners can naturally achieve a superior performance, where the ensemble cooperation loss (EC) and the task-adaptive gates (TG) bring obvious benefits by properly adjusting for functional diversity among learners and exploiting the

³They both are performed against a similar AlexNet-based architecture.

Table 3. Ablation study. A_t : averaged accuracy (%) of t tasks learned so far. TG: task-adaptive gates; EC: ensemble cooperation loss.

Methods	#Param	CIFAR-100-SC		CIFAR-100-RS	
		A_{10}	A_{20}	A_{10}	A_{20}
Single Continual Learner	837K	52.25 \pm 2.99	51.74 \pm 1.74	68.72 \pm 0.24	69.18 \pm 0.69
Classifier Ensemble	901K	50.08 \pm 1.65	43.88 \pm 0.79	66.80 \pm 1.45	55.65 \pm 0.32
Feature Ensemble	773K	58.76 \pm 3.72	57.69 \pm 1.42	73.57 \pm 0.50	73.01 \pm 1.22
Feature Ensemble + EC	773K	61.12 \pm 3.11	59.49 \pm 1.59	75.46 \pm 1.35	74.76 \pm 0.84
Feature Ensemble + TG	799K	62.01 \pm 3.36	59.85 \pm 1.77	76.11 \pm 0.98	74.78 \pm 0.41
Feature Ensemble + EC + TG	799K	62.89 \pm 3.05	60.84 \pm 0.95	78.08 \pm 1.25	76.05 \pm 0.65

**Fig. 7.** Effects of hyperparameters in CoSCL. (a) γ for ensemble cooperation (EC) loss; (b) s for task-adaptive gates (TG). The dashed lines indicate the performance w/o EC or TG in corresponding benchmarks. The arrows denote the chosen values.

additional information from task labels, respectively. Then we evaluate the effect of **hyperparameters** in Fig. 7. The hyperparameters of EC and TG are only moderately sensitive within a wide range. In this case, an appropriate (positive) strength of EC constrains the excessive diversity of predictions to improve the performance, while the continual learners will lose diversity if EC is too strong, resulting in a huge performance drop. If CoSCL cannot obtain sufficient diversity from the randomness of its architecture, the use of negative strength of EC can naturally serve this purpose, left for further work.

Moreover, we empirically validate our theoretical analysis as below. We first evaluate the **\mathcal{H} -divergence** of feature representations between tasks, which relies on the capacity of a hypothesis space to distinguish them [21]. Specifically, the \mathcal{H} -divergence can be empirically approximated by training a discriminator to distinguish if the features of input images belong to a task or not, where a larger discrimination loss indicates a smaller \mathcal{H} -divergence. As shown in Fig. 8, a, the proposed FE together with EC can largely decrease the \mathcal{H} -divergence while TG has a moderate benefit (there is a saturation effect when they are combined together). Meanwhile, we evaluate the **curvature of loss landscape** for the continually-learned solution by permuting the parameters to ten random directions [11], where the solution obtained by CoSCL enjoys a clearly flatter loss landscape than SCL (Fig. 8, b).

Taking all results together, cooperating multiple small continual learners can mitigate the discrepancy between tasks in feature space and improve flatness of

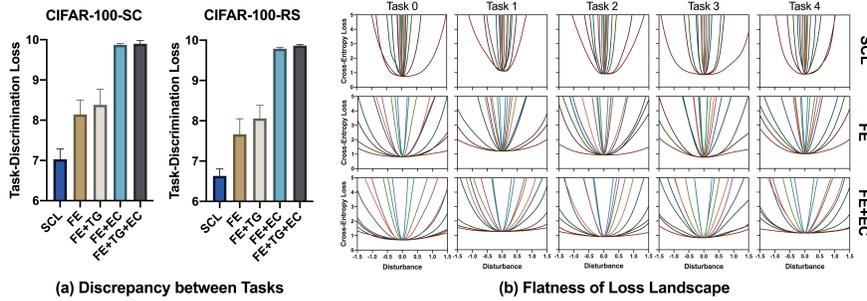


Fig. 8. Empirical validation of our theoretical analysis. (a) Task-discrimination loss in feature space. Larger loss indicates a smaller \mathcal{H} -divergence. (b) Curvature of the test loss landscape for the first five incremental tasks on CIFAR-100-SC. Each line indicates the result of a random direction.

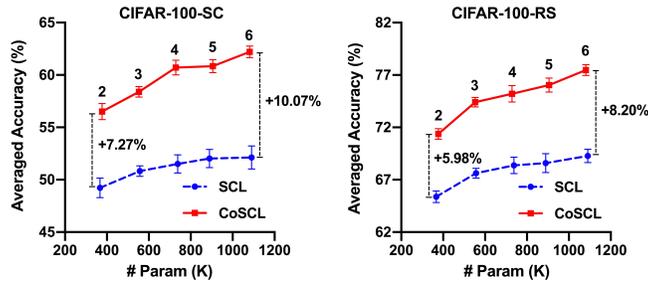


Fig. 9. Adding continual learners in CoSCL is more effective than widening the network of a single continual learner (SCL). We present the results of cooperating 2-6 continual learners with the same sub-network width, while accordingly adjust the size of SCL for a fair comparison.

the continually-learned solution (Fig. 8), thus facilitating both FWT and BWT (Fig. 6). This is consistent with our theoretical analysis, suggesting that learning plasticity and memory stability are *not* necessarily conflicting in continual learning, but can be simultaneously enhanced by a well-designed model architecture.

Adding Continual Learners is More Effective than Widening a Single Network: All of the above experiments are performed under a *fixed* parameter budget. A recent work observed that a wider network usually suffers from less catastrophic forgetting [26], providing an initial exploration of the effects of architecture in continual learning. Here we argue that adding continual learners with CoSCL is a better choice. In Fig. 9 we compare the performance of using an increasing number of continual learners (the width is the same as that used in Table 1) and accordingly widening the network of a single continual learner (SCL). It can be clearly seen that the performance gap between CoSCL and SCL is growing when more parameters are used. Therefore, CoSCL presents a promising direction for continual learning that can leverage network parameters in an efficient and scalable way.

Unsupervised Continual Learning (UCL): has the unique property of being naturally robust to catastrophic forgetting when fine-tuning on incremental unlabeled data [15, 23]. An empirical explanation is that UCL achieves a flatter loss landscape and more meaningful feature representations [23], which is consistent with our analysis. We further validate this idea by adapting CoSCL to UCL⁴, where we follow the UCL setting of [23] for CIFAR-100-RS and use a similar architecture as Table 1. As shown in Table 4, CoSCL can significantly improve the performance of UCL with two strong unsupervised learning strategies such as SimSiam [8] and BarlowTwins [49].

Table 4. Averaged accuracy (%) of unsupervised continual learning on CIFAR-100-RS. The results are reproduced from the officially-released code of [23].

Methods	SimSiam [8]	BarlowTwins [49]
Finetune	41.38 ± 0.80	63.29 ± 0.38
<i>w/</i> CoSCL	46.33 ± 0.51	74.03 ± 0.36

6 Conclusion

Numerous efforts in continual learning have been devoted to developing effective approaches based on a single model, but their efficacy might be limited by such a priori assumption. In this work, we present a unified form of generalization bounds for learning plasticity and memory stability in continual learning, consisting of three components, and demonstrate that the both aspects can be simultaneously improved by a compatible parameter space of a well-designed mode architecture. Inspired by the robust biological learning system, we propose to cooperate multiple (small) continual learners, which can naturally tighten the generalization bounds through improving the three components. Our method can substantially enhance the performance of representative continual learning strategies by improving both learning plasticity and memory stability. We hope that this work can serve as a strong baseline to stimulate new ideas for continual learning from an architecture perspective. A promising direction is to cooperate a variety of continual learning approaches with properly-designed architectures, so as to fully leverage task attributes for desired compatibility.

Acknowledgements. This work was supported by the National Key Research and Development Program of China (2017YFA0700904, 2020AAA0106000, 2020AAA0104304, 2020AAA0106302, 2021YFB2701000), NSFC Projects (Nos. 62061136001, 62106123, 62076147, U19B2034, U1811461, U19A2081, 61972224), Beijing NSF Project (No. JQ19016), BNRist (BNR2022RC01006), Tsinghua-Peking Center for Life Sciences, Tsinghua Institute for Guo Qiang, Beijing Academy of Artificial Intelligence (BAAI), Tsinghua-OPPO Joint Research Center for Future Terminal Technology, the High Performance Computing Center, Tsinghua University, and China Postdoctoral Science Foundation (Nos. 2021T140377, 2021M701892).

⁴Here we only use feature ensemble (FE) with ensemble cooperation loss (EC).

References

1. Aljundi, R., Babiloni, F., Elhoseiny, M., Rohrbach, M., Tuytelaars, T.: Memory aware synapses: Learning what (not) to forget. In: Proceedings of the European Conference on Computer Vision. pp. 139–154 (2018)
2. Aljundi, R., Chakravarty, P., Tuytelaars, T.: Expert gate: Lifelong learning with a network of experts. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3366–3375 (2017)
3. Aso, Y., Hattori, D., Yu, Y., Johnston, R.M., Iyer, N.A., Ngo, T.T., Dionne, H., Abbott, L., Axel, R., Tanimoto, H., et al.: The neuronal architecture of the mushroom body provides a logic for associative learning. *Elife* **3**, e04577 (2014)
4. Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., Vaughan, J.W.: A theory of learning from different domains. *Machine learning* **79**(1), 151–175 (2010)
5. Cha, J., Chun, S., Lee, K., Cho, H.C., Park, S., Lee, Y., Park, S.: Swad: Domain generalization by seeking flat minima. *arXiv preprint arXiv:2102.08604* (2021)
6. Cha, S., Hsu, H., Hwang, T., Calmon, F., Moon, T.: Cpr: Classifier-projection regularization for continual learning. In: Proceedings of the International Conference on Learning Representations. (2020)
7. Chaudhry, A., Dokania, P.K., Ajanthan, T., Torr, P.H.: Riemannian walk for incremental learning: Understanding forgetting and intransigence. In: Proceedings of the European Conference on Computer Vision. pp. 532–547 (2018)
8. Chen, X., He, K.: Exploring simple siamese representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15750–15758 (2021)
9. Cohn, R., Morantte, I., Ruta, V.: Coordinated and compartmentalized neuromodulation shapes sensory processing in drosophila. *Cell* **163**(7), 1742–1755 (2015)
10. Delange, M., Aljundi, R., Masana, M., Parisot, S., Jia, X., Leonardis, A., Slabaugh, G., Tuytelaars, T.: A continual learning survey: Defying forgetting in classification tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2021)
11. Deng, D., Chen, G., Hao, J., Wang, Q., Heng, P.A.: Flattening sharpness for dynamic gradient projection memory benefits continual learning. In: Proceedings of the Advances in Neural Information Processing Systems. vol. 34 (2021)
12. Dinh, L., Pascanu, R., Bengio, S., Bengio, Y.: Sharp minima can generalize for deep nets. In: Proceedings of the International Conference on Machine Learning. pp. 1019–1028. PMLR (2017)
13. Doan, T., Mirzadeh, S.I., Pineau, J., Farajtabar, M.: Efficient continual learning ensembles in neural network subspaces. *arXiv preprint arXiv:2202.09826* (2022)
14. Fernando, C., Banarse, D., Blundell, C., Zwols, Y., Ha, D., Rusu, A.A., Pritzel, A., Wierstra, D.: Pathnet: Evolution channels gradient descent in super neural networks. *arXiv preprint arXiv:1701.08734* (2017)
15. Hu, D., Lu, Q., Hong, L., Hu, H., Zhang, Y., Li, Z., Shen, A., Feng, J.: How well self-supervised pre-training performs with streaming data? *arXiv preprint arXiv:2104.12081* (2021)
16. Hurtado, J., Raymond, A., Soto, A.: Optimizing reusable knowledge for continual learning via metalearning. In: Proceedings of the Advances in Neural Information Processing Systems. vol. 34 (2021)
17. Jung, S., Ahn, H., Cha, S., Moon, T.: Continual learning with node-importance based adaptive group sparse regularization. *arXiv e-prints* pp. arXiv–2003 (2020)
18. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A.A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., et al.: Overcoming

- catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences* **114**(13), 3521–3526 (2017)
19. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images. Tech. rep., Citeseer (2009)
 20. Liu, Y., Parisot, S., Slabaugh, G., Jia, X., Leonardis, A., Tuytelaars, T.: More classifiers, less forgetting: A generic multi-classifier paradigm for incremental learning. In: *Proceedings of the European Conference on Computer Vision*. pp. 699–716. Springer (2020)
 21. Long, M., Cao, Y., Wang, J., Jordan, M.: Learning transferable features with deep adaptation networks. In: *Proceedings of the International Conference on Machine Learning*. pp. 97–105. PMLR (2015)
 22. Lopez-Paz, D., et al.: Gradient episodic memory for continual learning. In: *Proceedings of the Advances in Neural Information Processing Systems*. pp. 6467–6476 (2017)
 23. Madaan, D., Yoon, J., Li, Y., Liu, Y., Hwang, S.J.: Rethinking the representational continuity: Towards unsupervised continual learning. arXiv preprint arXiv:2110.06976 (2021)
 24. McAllester, D.A.: Pac-bayesian model averaging. In: *Proceedings of the Twelfth Annual Conference on Computational Learning Theory*. pp. 164–170 (1999)
 25. McCloskey, M., Cohen, N.J.: Catastrophic interference in connectionist networks: The sequential learning problem. In: *Psychology of Learning and Motivation*, vol. 24, pp. 109–165. Elsevier (1989)
 26. Mirzadeh, S.I., Chaudhry, A., Hu, H., Pascanu, R., Gorur, D., Farajtabar, M.: Wide neural networks forget less catastrophically. arXiv preprint arXiv:2110.11526 (2021)
 27. Mirzadeh, S.I., Chaudhry, A., Yin, D., Nguyen, T., Pascanu, R., Gorur, D., Farajtabar, M.: Architecture matters in continual learning. arXiv preprint arXiv:2202.00275 (2022)
 28. Mirzadeh, S.I., Farajtabar, M., Gorur, D., Pascanu, R., Ghasemzadeh, H.: Linear mode connectivity in multitask and continual learning. arXiv preprint arXiv:2010.04495 (2020)
 29. Mirzadeh, S.I., Farajtabar, M., Pascanu, R., Ghasemzadeh, H.: Understanding the role of training regimes in continual learning. In: *Proceedings of the Advances in Neural Information Processing Systems*. vol. 33, pp. 7308–7320 (2020)
 30. Modi, M.N., Shuai, Y., Turner, G.C.: The drosophila mushroom body: from architecture to algorithm in a learning circuit. *Annual Review of Neuroscience* **43**, 465–484 (2020)
 31. Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., Ng, A.Y.: Reading digits in natural images with unsupervised feature learning (2011)
 32. Qin, Q., Hu, W., Peng, H., Zhao, D., Liu, B.: Bns: Building network structures dynamically for continual learning. In: *Proceedings of the Advances in Neural Information Processing Systems*. vol. 34 (2021)
 33. Ramesh, R., Chaudhari, P.: Model zoo: A growing brain that learns continually. In: *NeurIPS 2021 Workshop on Distribution Shifts: Connecting Methods and Applications* (2021)
 34. Rebuffi, S.A., Kolesnikov, A., Sperl, G., Lampert, C.H.: icarl: Incremental classifier and representation learning. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2001–2010 (2017)
 35. Riemer, M., Cases, I., Ajemian, R., Liu, M., Rish, I., Tu, Y., Tesauro, G.: Learning to learn without forgetting by maximizing transfer and minimizing interference. arXiv preprint arXiv:1810.11910 (2018)

36. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *International Journal of Computer Vision* **115**(3), 211–252 (2015)
37. Rusu, A.A., Rabinowitz, N.C., Desjardins, G., Soyer, H., Kirkpatrick, J., Kavukcuoglu, K., Pascanu, R., Hadsell, R.: Progressive neural networks. arXiv preprint arXiv:1606.04671 (2016)
38. Schwarz, J., Czarnecki, W., Luketina, J., Grabska-Barwinska, A., Teh, Y.W., Pascanu, R., Hadsell, R.: Progress & compress: A scalable framework for continual learning. In: *Proceedings of the International Conference on Machine Learning*. pp. 4528–4537. PMLR (2018)
39. Serra, J., Suris, D., Miron, M., Karatzoglou, A.: Overcoming catastrophic forgetting with hard attention to the task. In: *Proceedings of the International Conference on Machine Learning*. pp. 4548–4557. PMLR (2018)
40. Shi, G., Chen, J., Zhang, W., Zhan, L.M., Wu, X.M.: Overcoming catastrophic forgetting in incremental few-shot learning by finding flat minima. In: *Proceedings of the Advances in Neural Information Processing Systems*. vol. 34 (2021)
41. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The caltech-ucsd birds-200-2011 dataset (2011)
42. Wang, L., Yang, K., Li, C., Hong, L., Li, Z., Zhu, J.: Ordisco: Effective and efficient usage of incremental unlabeled data for semi-supervised continual learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5383–5392 (2021)
43. Wang, L., Zhang, M., Jia, Z., Li, Q., Bao, C., Ma, K., Zhu, J., Zhong, Y.: Afec: Active forgetting of negative transfer in continual learning. In: *Proceedings of the Advances in Neural Information Processing Systems*. vol. 34 (2021)
44. Wang, L., Zhang, X., Yang, K., Yu, L., Li, C., Lanqing, H., Zhang, S., Li, Z., Zhong, Y., Zhu, J.: Memory replay with data compression for continual learning. In: *Proceedings of the International Conference on Learning Representations*. (2021)
45. Wen, Y., Tran, D., Ba, J.: Batchensemble: an alternative approach to efficient ensemble and lifelong learning. In: *Proceedings of the International Conference on Learning Representations*. (2020)
46. Wortsman, M., Horton, M.C., Guestrin, C., Farhadi, A., Rastegari, M.: Learning neural network subspaces. In: *Proceedings of the International Conference on Machine Learning*. pp. 11217–11227. PMLR (2021)
47. Wortsman, M., Ramanujan, V., Liu, R., Kembhavi, A., Rastegari, M., Yosinski, J., Farhadi, A.: Supermasks in superposition. In: *Proceedings of the Advances in Neural Information Processing Systems*. vol. 33, pp. 15173–15184 (2020)
48. Yan, S., Xie, J., He, X.: Der: Dynamically expandable representation for class incremental learning. arXiv preprint arXiv:2103.16788 (2021)
49. Zbontar, J., Jing, L., Misra, I., LeCun, Y., Deny, S.: Barlow twins: Self-supervised learning via redundancy reduction. In: *Proceedings of the International Conference on Machine Learning*. pp. 12310–12320. PMLR (2021)
50. Zenke, F., Poole, B., Ganguli, S.: Continual learning through synaptic intelligence. In: *Proceedings of the International Conference on Machine Learning*. pp. 3987–3995. PMLR (2017)