

Create and Find Flatness: Building Flat Training Spaces in Advance for Continual Learning

Wenhang Shi^a, Yiren Chen^b, Zhe Zhao^b, Wei Lu^{a,*}, Kimmo Yan^b and Xiaoyong Du^a

^aSchool of Information and DEKE, MOE, Renmin University of China

^bTencent AI Lab

Abstract. Catastrophic forgetting remains a critical challenge in the field of continual learning, where neural networks struggle to retain prior knowledge while assimilating new information. Most existing studies emphasize mitigating this issue only when encountering new tasks, overlooking the significance of the pre-task phase. Therefore, we shift the attention to the current task learning stage, presenting a novel framework, C&F (*Create and Find Flatness*), which builds a flat training space for each task in advance. Specifically, during the learning of the current task, our framework adaptively **creates** a flat region around the minimum in the the loss landscape. Subsequently, it **finds** the parameters' importance to the current task based on their flatness degrees. When adapting the model to a new task, constraints are applied according to the flatness and a flat space is simultaneously prepared for the impending task. We theoretically demonstrate the consistency between the created and found flatness. In this manner, our framework not only accommodates ample parameter space for learning new tasks but also preserves the preceding knowledge of earlier tasks. Experimental results exhibit C&F's state-of-the-art performance as a standalone continual learning approach and its efficacy as a framework incorporating other methods. Our work is available at <https://github.com/Eric8932/Create-and-Find-Flatness>.

1 Introduction

As both the size of model parameters and the volume of training data persistently expand, recent Large Language Models (LLMs) have exhibited remarkable generalization abilities across a wide array of tasks [2]. Nevertheless, the competences and knowledge of language models remain limited to their training data, and even colossal models struggle in specific situations such as financial scenarios using private data [35] and recommendation systems using user data [15]. The high cost of retraining the model on enlarged data hinders language models to constantly acquire new knowledge. Therefore, Continual Learning (CL), which enables a single model to continually learn and adapt over time, maintains its significance in enhancing the ability of language models.

The major challenge CL faces is the phenomenon of *catastrophic forgetting* [19], where the model loses much of the previous acquired knowledge while adapting to new tasks. To tackle this issue, various approaches have been proposed, encompassing data distribution-based, architecture-based and regularization-based methods. Among these, the regularization-based method is closely related to our work, which inhibits excessive model parameter alternations to maintain

proficiency in earlier tasks [1, 12]. However, most approaches emphasize applying constraints during model optimization, neglecting the construction of an appropriate training space for continual learning in advance. That is, the current task optimization aims to identify the global optimal solution, where the loss landscape is sharp. Consequently, certain model parameters become exceedingly sensitive, implying that minor perturbations can trigger substantial performance decline in previous tasks when adapting to new ones [21]. Flat-minima based measures, which penalize sharpness and find flat local minima, are employed to mitigate the impact of weight perturbations and enhance single-task model generalization. However, they are not directly applicable to task sequences in CL [34].

Hence, we propose a novel framework, C&F (Create and Find), to build flat training spaces in advance for continual learning. Our framework leverages the geometrical properties of flat minima in both old and new task optimization processes. While training on the current task, we create a flat region around the minimum in the loss landscape. The surrounding of low-loss-value points makes the model robust to a specific range of perturbations from new tasks. Moreover, as different parameters' variations yield disparate effects on the loss, we find the unique flatness of each parameter in the current task using Fisher Information [32]. During the model's adaptation to new tasks, we restrict parameter changes within the flat region and regularize them according to distinct flatness degrees. Concurrently, we prepare the flat space for the subsequent task optimization. In the flat training spaces, the model can assimilate new concepts from new data without forgetting previously acquired knowledge.

Our main contributions encompass the following aspects: 1.Our work pioneers the utilization of flatness in crafting apt training spaces in advance for new task optimization, and we theoretically demonstrate the consistency of the flatness in the learning spaces. 2.We propose C&F, an innovative framework for continual learning, offering flexibility and seamless integration with other advanced CL techniques. 3.We conduct experiments to evaluate C&F as both a standalone method and a framework incorporating other methods. The state-of-the-art results demonstrate C&F's effectiveness.

2 Related Work

2.1 Continual Learning

Existing CL approaches could be divided into three categories: 1. Data distribution-based method maintains model's ability by aligning with the distribution of previously encountered tasks. Saving a subset of old task samples and incorporating them into new task

* Corresponding Author. Email: lu-wei@ruc.edu.cn

training is most direct, referred as *replay* [5, 16]. In addition, distillation using both old and new samples is also explored [22, 26]. Some generate pseudo-data of old tasks for replay [30, 31] or distillation [33]; 2. Architecture-based method focuses on allocating a subset of the model’s parameters for each task, optimizing the model exclusively on those parameters. Some allocate and utilize existing parameters to maintain a fixed model size [7, 18] and others freeze all model parameters and introduce additional parameters for each new task [25, 27]. However, the former imposes a limit on the model’s capacity and the latter requires expanding the storage occupied by the model as new tasks appear; 3. Regularization-based method typically incorporates an auxiliary regularization term into the loss function to impede excessive model weight alterations [1, 12, 38]. [14] constrains the current model’s output to resemble the previous model’s output when processing new data. [39] matches the magnitude of weight vectors corresponding to each task. Owing to the disparate importance of parameters, C&F implements a soft constraint on model optimization by adjusting the loss function.

2.2 Continual Pre-training of Large Language Models

Large language models have demonstrated remarkable capabilities. Enhancing their abilities towards different objectives necessitates supplementary training. Techniques such as specific domain or task data training are employed to refine and enrich model’s knowledge in specific areas [24, 11]. Instruction tuning empowers these models with the ability to follow instructions, bolstering their performance in zero/few-shot scenarios [28, 36]. Moreover, for these models to adhere to human standards, Reinforcement Learning from Human Feedback (RLHF) is deployed for alignment tuning [20, 23]. All these continual pre-training sacrifice models’ generic capabilities in favor of specific abilities, resulting in forgetting.

2.3 Flat-Minima Optimization

Flat-minima optimizers optimize both loss and sharpness to find flat local minima and leverage the surrounding low loss values to enhance generalization and scalability for a single task [3]. [8] optimizes the corrupted loss to ensure low losses around the minima and [13] introduces the concept of adaptive sharpness to achieve scale-invariant flatness. Given that the flat regions are inherently suitable for model’s continual leaning on new tasks, [21] examines the role of local minima’s geometric properties for each task in the overall degree of forgetting during continual learning. To the best of our knowledge, the only work related to our framework is F2M [29], which identifies a flat local minimum in the base task and employs it in incremental few-shot learning. Our framework extends to more general applications in continual learning with a more efficient and effective framework.

3 Methodology

In continual learning, the model must adapt to new tasks while preserving proficiency in previous ones. C&F incorporates flatness characteristics into continual learning process to address the forgetting issue (§3.1). Specifically, during current task learning, it builds flat training spaces in advance by creating flat regions and finding parameter flatness (§3.2). The new task optimization should be conducted in the last constructed space, while involving preparing a training space for the next task (§3.3). The whole training process is illustrated in Figure 1. Last, we justify theoretically the utilization of

Fisher Information as flatness indicator by demonstrating its consistency with the Create phase concerning flatness (§3.4).

3.1 Meaning of Flatness

Flat local minima means surrounding parameter points all have low loss values. Initially aimed at improving model generalization for a single task, flatness essentially enhances the model’s resilience to test set distribution shifts [8]. Continual learning is a process of encountering perpetually varying task distributions. Therefore, creating flat regions and constraining subsequent tasks to learn within them fosters the model’s robustness against continual distribution alterations. In addition, due to the low surrounding loss values around the flat minima, parameter changes do not substantially affect model’s performance. Continual learning necessitates ongoing model updates. Learning new tasks within flat regions ensures low losses, i.e., high performances, on prior tasks, addressing the forgetting issue.

3.2 Build Flat Training Spaces

3.2.1 Create Flat Regions

To ensure that the parameter points around minimum all possess low loss values, we optimize for the point with the largest loss value in the entire neighborhood. To formally state the concept, we denote the loss function $\mathcal{L} : \mathcal{W} \times \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ and \mathcal{S} represents the training set. We convert the original $L_{\mathcal{S}}(\mathbf{w})$ minimization into a MIN-MAX problem, minimizing the maximum loss value around the current point. This approach has been extensively demonstrated, both experimentally and theoretically, to be effective and efficient in creating flat regions [8, 13]. The optimization objective is as follows:

$$\min_{\mathbf{w}} \max_{\|\epsilon\|_2 \leq \rho} L_{\mathcal{S}}(\mathbf{w} + \epsilon). \quad (1)$$

where ρ denotes the range of the flat region. It is challenging to find the exact optima ϵ^* with the maximum loss value, so we employ a first-order approximation. Since ρ determines the learning space for new tasks, we aim to expand it while avoiding the problems associated with its size. To accomplish this, we further adaptively adjust ρ based on the parameter values:

$$\hat{\epsilon}(\mathbf{w}) = \rho \frac{\mathbf{w}^2 \nabla_{\mathbf{w}} L_{\mathcal{S}}(\mathbf{w})}{\|\mathbf{w} \nabla_{\mathbf{w}} L_{\mathcal{S}}(\mathbf{w})\|_2}, \quad (2)$$

where \mathbf{w}^2 is an element-wise operation and $\hat{\epsilon}$ is an approximation of ϵ^* . Finally, we use the gradient of $\mathbf{w} + \hat{\epsilon}(\mathbf{w})$ to update the current parameter point \mathbf{w}

$$\nabla_{\mathbf{w}} L_{\mathcal{S}}^{Create}(\mathbf{w}) \approx \nabla_{\mathbf{w}} L_{\mathcal{S}}(\mathbf{w}) \Big|_{\mathbf{w} + \hat{\epsilon}(\mathbf{w})}, \quad (3)$$

and the loss for Create is

$$L_C(\mathbf{w}) = L_{\mathcal{S}}(\mathbf{w} + \hat{\epsilon}(\mathbf{w})). \quad (4)$$

3.2.2 Find Parameter Flatness

Within flat regions, sharp parameters still exist. As minor fluctuations in these parameters can result in considerable changes in the loss value, they are crucial to the task and should be found and optimized sparingly. However, determining the exact flatness of each parameter is challenging, so we use Fisher Information as a rapid approximation of flatness [32]. It gauges importance by indicating the

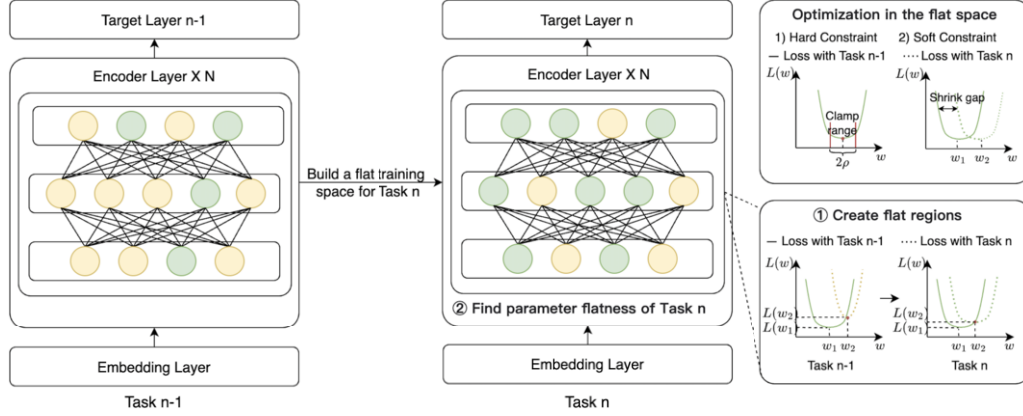


Figure 1: Task optimization in the C&F framework involves two parts: 1. Build a flat training space for the next task: 1) Create a flat region around the minimum and 2) Find parameter flatness for the current task; 2. Optimize the model in the last flat space with two constraints. Yellow and green represent sharpness and flatness, respectively. $L(w_1)$ and $L(w_2)$ denote the loss of task n-1 and n after optimization, respectively.

extent to which altering the parameter impacts the model’s output. Due to the assumption of parameter independence, we only consider the diagonal elements in the Fisher Information matrix. The computation is as follows:

$$F(\mathbf{w}_i) = \frac{1}{|S|} \sum_{j=1}^{|S|} \left(\frac{\partial \log p(y_j | x_j; \mathbf{w})}{\partial \mathbf{w}_i} \right)^2. \quad (5)$$

A smaller Fisher value suggests that the parameters are more likely situated in a flatter region, thereby permitting the model to update more substantially on them.

To enable Fisher Information to represent the flatness of each parameter across all seen tasks, it is computed for every task and accumulated using the following formula:

$$F(\mathbf{w}_i)' = \gamma * F(\mathbf{w}_i)' + F(\mathbf{w}_i). \quad (6)$$

3.3 Optimization in the Flat Spaces

The last task’s Creating and Finding flatness constructs a learning space for the current task. Optimizing parameters in it involves two corresponding constraints.

Hard Constraint. One is a hard constraint for **Create**, restricting the model to be optimized in the flat region to preserve its performance on the previous task. After each batch optimization, we clamp parameter changes to ensure they remain in the region:

$$\mathbf{w}^* - \rho * (\mathbf{w}^*)^{abs} \preceq \mathbf{w} \preceq \mathbf{w}^* + \rho * (\mathbf{w}^*)^{abs}. \quad (7)$$

since we create the flat regions adaptively, we also implement an adaptive clamping strategy based on the size of the parameters.

Soft Constraint. We apply a soft constraint for **Find** step to incorporate the found flatness into the optimization process. Specifically, we add an L2 loss term to the loss function, with the flatness serving as a unique coefficient for each parameter, as described in [12]. Due to the smaller flat range of sharp parameters, restricting their update not only sustains performance on old tasks, but also shrinks the gap between successively created regions so that most parameters stay within all regions. The loss for Find is

$$L_F(\mathbf{w}) = \sum_i F(\mathbf{w}_i)' (\mathbf{w}_i - \mathbf{w}_i^*)^2. \quad (8)$$

where \mathbf{w}^* and \mathbf{w} denote the optimal parameter vector learned in the last task and the current parameter vector, respectively.

Apart from two constraints from Create and Find, current task’s optimization involves the same creating and finding process to prepare a flat space for the next task. The overall optimization loss is

$$L_O = L_C + \lambda L_F. \quad (9)$$

where λ is the coefficient for the soft loss. The full C&F training process is shown in Algorithm 1.

Algorithm 1: The training algorithm for C&F framework

Input: Loss function $\mathcal{L} : \mathcal{W} \times \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}_+$,
Train set $\mathcal{S} \triangleq \cup_{i=1}^n \{(\mathbf{x}_i, \mathbf{y}_i)\}$, Batch size b , Learning rate α ,
Flat region range ρ , Coefficient λ , γ , Weight \mathbf{w}_0 , Task $t \geq 1$.

Output: Trained weight \mathbf{w} .

// Optimization for Task 1 and Build a Flat Space for Task 2

begin

Initialize weight $\mathbf{w}_1 := \mathbf{w}_0$.

Create a flat region around minimum

while not converged do

Sample a mini-batch \mathcal{B} of size b from \mathcal{S} ,

Compute $\hat{\epsilon}(\mathbf{w})$ with Eq.2,

Compute loss for **Create** objective with Eq.3,

Compute gradient \mathbf{g}_C ,

Update weights: $\mathbf{w}_1 = \mathbf{w}_1 - \alpha \mathbf{g}_C$.

end

Find flatness $F(\mathbf{w})$ for each parameter in \mathbf{w}_1 with Eq.5

end

for task $t = 2, \dots, M$ **do**

// Optimization for Task t and Build a Space for Task t+1

begin

Initialize weight $\mathbf{w}_t := \mathbf{w}_{t-1}$, $\mathbf{w}^* := \mathbf{w}_{t-1}$.

while not converged do

Sample a mini-batch \mathcal{B} of size b from \mathcal{S} ,

Compute gradient \mathbf{g}_C for **Create** objective,

Compute gradient \mathbf{g}_F for **Find** flatness with Eq.8,

Update weights: $\mathbf{w}_t = \mathbf{w}_t - \alpha(\mathbf{g}_C + \lambda \mathbf{g}_F)$,

Clamp the parameters \mathbf{w}_t with Eq.7.

end

Find flatness $F(\mathbf{w})$ of \mathbf{w}_t with Eq.5,

Accumulate flatness $F(\mathbf{w})'$ with Eq.6.

end

end

return \mathbf{w}_t ;

3.4 Flatness Consistency between Create and Find

The selection of Fisher Information as the flatness indicator is due to its consistency with Create concerning flatness. To prove this, we illustrate that the process of creating also minimizes the Fisher values of parameters. Firstly, min-maximization of loss optimizes both the task loss and the sharpness. According to Eq.1:

$$\begin{aligned} L_S^{Create}(\mathbf{w}) &= \max_{\|\epsilon\|_2 \leq \rho} L_S(\mathbf{w} + \epsilon) \\ &= (\max_{\|\epsilon\|_2 \leq \rho} L_S(\mathbf{w} + \epsilon) - L_S(\mathbf{w})) + L_S(\mathbf{w}), \end{aligned} \quad (10)$$

Since ϵ is around 0, the sharpness could be defined:

$$sharpness = \max_{\|\epsilon\|_2 \leq \rho} L_S(\mathbf{w} + \epsilon) - L_S(\mathbf{w}), \quad (11)$$

Substituting it back, we then have:

$$L_S^{Create}(\mathbf{w}) = sharpness + L_S(\mathbf{w}). \quad (12)$$

Therefore, in the process of create, the task loss and the sharpness are optimized simultaneously.

Then we prove that minimizing the sharpness equals to minimizing the Fisher Info. According to Eq.2, the approximately optimal $\epsilon^*(\mathbf{w})$ is:

$$\epsilon^*(\mathbf{w}) \approx \rho \frac{\nabla_{\mathbf{w}} L_S(\mathbf{w})}{\|\nabla_{\mathbf{w}} L_S(\mathbf{w})\|_2}, \quad (13)$$

Notice that we don't consider adaptively adjusting ρ here, so it is different from Eq.2. Substituting it back to Eq.11 and approximating the sharpness via a first-order Taylor expansion of $L_S(\mathbf{w} + \hat{\epsilon}(\mathbf{w}))$ w.r.t ϵ around 0, obtaining:

$$\begin{aligned} sharpness &\approx \epsilon^*(\mathbf{w})^T \nabla_{\mathbf{w}} L_S(\mathbf{w}) \\ &= \rho \|\nabla_{\mathbf{w}} L_S(\mathbf{w})\|_2, \end{aligned} \quad (14)$$

Denoting $g(\mathbf{w}) = \|\nabla_{\mathbf{w}} L_S(\mathbf{w})\|_2$ as the norm of gradient of the loss function and $I(\mathbf{w}) = \mathbb{E} [g(\mathbf{w})g(\mathbf{w})^T]$ as Fisher Information matrix. Now, let's consider the square of the Euclidean norm of the gradient:

$$\|g(\mathbf{w})\|_2^2 = g(\mathbf{w})^T g(\mathbf{w}), \quad (15)$$

Taking the expected value of both sides:

$$\mathbb{E} [\|g(\mathbf{w})\|_2^2] = \mathbb{E} [g(\mathbf{w})^T g(\mathbf{w})] = Tr(I(\mathbf{w})). \quad (16)$$

where $Tr(\bullet)$ denotes the trace of a matrix. This reveals that minimizing sharpness essentially minimizes the trace of the Fisher Information matrix. Since we compute only the diagonal elements of the matrix during the Find phase, the trace is solely interested, which represents the overall flatness for all parameters. Thus, the flatness created and found are consistent.

4 Experiment

4.1 Datasets and Experimental Setup

We validate our method on NLP tasks. For most of the experiments, we follow MBPA++ [5] and IDBR [10] to use five text classification datasets and map Amazon and Yelp tasks to the same label space, resulting in 33 classes with 4 task layers in total. Since the original datasets are too large, we randomly sample 2000 examples per class for training and validation sets to build sampled-version datasets. We conduct major experiments on the sampled-setting text classification

datasets, except when comparing to state-of-the-art methods because they are all trained and validated on the full setting. To exclude the effect of task order and sequence length, we follow IDBR [10] and construct seven task sequences of various lengths and orders. Our implementation is based on TencentPretrain [41], a pre-training toolkit extended from UER-py [40]. More specific information about datasets and task sequences is in supplemental material.

4.2 Methods for Comparison

We compare our C&F as an individual approach with the following methods: **Seq** [37], sequentially finetunes the model at each task, and **Replay** [10], stores and replays the most representative old exemplars while learning new tasks, which could be seen as the lower bound. **LwF** [14], **EWC** [12] and **MAS** [1] are three regularization-based methods. LwF limits the output of the current model to be consistent with the old model. EWC and MAS constraint parameter optimization based on their effects on loss and model output, respectively. **F2M** [29] finds a flat local minimum in the first task and constrains the following task to update around it. **ORT** [6] stores the gradients of some samples (size is 20 here) for each task to construct the orthogonal basis and projects gradient direction of the subsequent tasks to be orthogonal to the basis. **MTL** learns all the tasks simultaneously and is considered an upper-bound of continual learning.

Apart from baseline methods, we also compare with the state-of-the-art methods. **MBPA++** [5] replays the seen exemplars at training time and uses KNN to select the most representative examples for local adaptation at test time. **LAMOL** [31] converts all tasks into a Q&A format and trains a language model by learning the task and generating training samples at the same time. **IDBR** [10] disentangles text representation space into a task generic space and a task specific space. **ProgressivePrompt** [25] learns a new soft prompt for each task and sequentially concatenates it with the previously learned ones, while keeping the base model frozen.

All baseline methods except for **Seq** are equipped with replay. We follow [10] to store the most representative training samples at the end of each task's learning and replay them while learning new tasks with the same frequency. It's noted that the state-of-the-art methods except **ProgressivePrompt** also incorporate replay, even with a higher replaying frequency. The concrete implementation details could be found in supplemental material.

4.3 Model

We employ BERT-base-uncased as the pretrained model to ensure a fair comparison, as it has been utilized in all previous works. It should be noted that our C&F framework is versatile and applicable to models of any size. While it might be appealing to investigate whether CL techniques can further enhance the capabilities of recently emergent large language models, no benchmark is effective against LLMs' extensive knowledge, as they are all composed of public datasets. Additionally, the high costs associated with LLMs and the closed-source nature of the most powerful models pose challenges for comparisons on them. Consequently, we use the medium-sized model to validate the efficacy of the continual learning methods on benchmarks and select the most effective one.

5 Result and Analysis

When training on the current task, we validate with a frequency of 500 steps using validation sets consist of all seen tasks, and select the

Table 1: Summary of baseline methods’ performance and error bar on Sampled Setting datasets. Metric is averaged accuracy after training on the last task. For each order, we conduct significance test against the best reproducible model, and * means that the improvement is significant at 0.05 significance level.

Method Order	Length-3 Task Sequences				Length-5 Task Sequences			
	I	II	III	AVG	IV	V	VI	AVG
Seq	69.91	63.91	70.78	68.20(± 1.00)	68.38	67.21	69.87	68.49(± 1.08)
Replay	72.85	70.54	73.48	72.29(± 0.46)	71.68	71.84	74.37	72.63(± 0.57)
MAS [1]	73.45	72.99	73.61	73.35(± 0.38)	73.78	75.22	74.30	74.43(± 0.5)
EWC [12]	74.39	74.12	74.47	74.33(± 0.19)	75.15	74.51	75.29	74.98(± 0.25)
LwF [14]	75.17	74.21	74.57	74.65(± 0.24)	75.37	74.98	75.48	75.28(± 0.21)
F2M [29]	73.56	71.43	73.71	72.90(± 0.37)	73.08	72.49	74.61	73.39(± 0.29)
ORT [6]	73.82	71.50	73.78	72.88(± 0.44)	72.58	74.79	73.03	73.42(± 0.67)
C&F	75.89*	74.81*	75.01*	75.24 (± 0.19)	76.61*	76.46*	76.06*	76.38 (± 0.23)
MTL	75.31	75.31	75.31	75.31(± 0.08)	76.94	76.94	76.94	76.94(± 0.09)

Table 2: Summary of previous state-of-the-art methods’ performances on Full Setting datasets. Metric is averaged accuracy after training on the last task. Results of other methods are fetched from the original papers, and C&F’s results are averaged over 3 runs.

Method Order	Length-5 Task Sequences				
	IV	V	VI	VII	AVG
MBPA++ [5]	70.7	70.2	70.9	70.8	70.7
MBPA++ [31]	74.9	73.1	74.9	74.1	74.3
LAMOL [31]	76.1	76.1	77.2	76.7	76.5
IDBR [10]	75.9	76.2	76.4	76.7	76.3
ProgressivePrompt [25]	78.0	77.7	77.9	77.9	77.9
C&F	78.5	78.2	77.9	78.8	78.4

model with the best performance at the end of the task training. We report the averaged accuracy for text classification tasks on test sets after training on all tasks. Without special statement, all results are averaged over 5 runs.

5.1 Main Results

Firstly, we evaluate C&F as an individual method against baselines on sampled-setting task sequences of length-3 and length-5. As illustrated in Table 1, C&F outperforms other methods across all 6 orders and the performance enhancement becomes more pronounced as the task sequence gets longer, indicating a diminished accumulation of forgetting for C&F. Notably, when compared to flatness-based or gradient-based methods like F2M or ORT, our approach yields superior results, demonstrating that a simple employment of gradient or flatness features is insufficient to address forgetting in NLP task sequences. While the overall performance improvements might not be statistically significant, they are far from negligible. The pretrained model’s strong capabilities have raised the performance lower bound, rendering further advancements challenging, and our performance closely approaches the upper bound MTL.

In Table 2, we compare with the previous state-of-the-art methods: MBPA++ [5], LAMOL [31], IDBR [10] and ProgressivePrompt [25] on full-setting datasets. The results show that C&F achieves the best performances on 4 different orders. These findings suggest that combining continual learning and geometric properties effectively alleviates catastrophic forgetting, and the flat training spaces constructed in advance indeed aid the model’s constant adaptations.

5.2 Integration with other methods

Our proposed C&F can function not only as a stand-alone approach for continual learning, but also as a framework to integrate other approaches. We incorporate the IDBR [10] approach and conduct

experiments in the full-dataset setting as in Table 2. In addition to IDBR’s original regularization, C&F framework creates flat regions and finds parameter flatness during current task learning, and adds the corresponding constraints in optimization. Results in Table 3 show that C&F brings further performance boost to IDBR. Moreover, we extend the experiments to NER tasks, in which a new entity type is learned at a time and classification head’s dimensions are continually expanded. We integrate our framework with the state-of-the-art method ExtendNER [22], which employs knowledge distillation to replay previous entities while learning a new one using samples annotated exclusively on the new entity type. Full dataset details of NER tasks are available in the supplemental material. Results in Table 3 indicate that C&F further enhances the performance of the SOTA method on NER tasks. These findings suggest that C&F framework is versatile and CL techniques can be seamlessly and effectively introduced to improve performance further.

5.3 Analysis of CREATE

To investigate whether the created flatness helps mitigate catastrophic forgetting and provides ample learning space for new tasks, we compare C&F with F (Find), which discards Create, in terms of learning ability, forgetting and accuracy. Despite the absence of Create, F still retains the parameter clamping after each update for a fair comparison. We adopt the measure of intransigence, representing the inverse of learning ability, and forgetting from [4]. We detail the measures in supplemental material. Figure 2 displays the comparison of the two methods under sampled-setting order 4 as ρ rises and the updating regions expand. C&F consistently exhibits lower levels of forgetting, suggesting that the creation of flat regions effectively diminishes forgetting. Additionally, C&F’s lower intransigence indicates its enhanced ability to learn new tasks. Although the improvement in generalization partially stems from the created flatness, it demonstrates that the flat regions do not impede new task learning.

Table 3: Comparisons between original method and method in C&F framework. Metric for Text Classification is averaged accuracy after training on the last task. Metric for NER is the averaged F1 over seen entity types on the test set at each step and scores at each step are averaged over all class orders.

Method	Text Classification				
	IV	V	VI	VII	AVG
IDBR [10]	75.9	76.2	76.4	76.7	76.3
+C&F	77.4	77.2	77.4	77.6	77.4

Method	NER									
	CoNLL-03				OntoNotes-5.0					
	Step 1	Step 2	Step 3	Step 4	Step 1	Step 2	Step 3	Step 4	Step 5	Step 6
ExtendNER [22]	86.41	86.47	86.41	86.41	79.80	80.59	80.88	80.80	80.27	79.51
+C&F	87.10	87.20	87.23	86.98	81.24	81.58	81.68	81.41	80.98	80.21

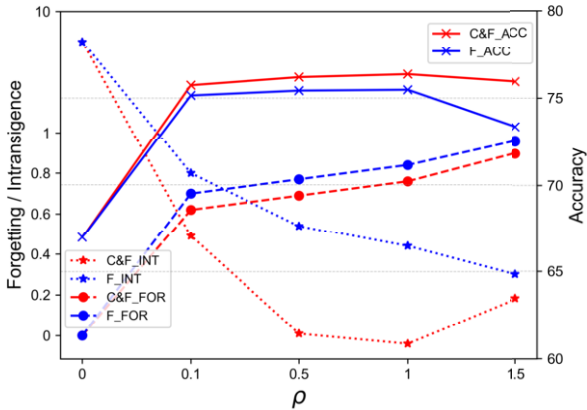


Figure 2: Comparisons of intransigence, forgetting and accuracy between C&F and F, C&F without creating, under order 4 as ρ rises. INT is short for intransigence and FOR is short for forgetting. Left y-axis is for intransigence and forgetting, and right is for accuracy.

It is worth noting that C&F’s intransigence increases when the value of ρ is high, as the flatness comes at the expense of increased loss. Overall, C&F successfully reduces forgetting and bolsters learning by creating flat regions, as evidenced by the accuracy curves as well.

5.4 Analysis of FIND

5.4.1 Overlap of Regions

Flat regions are created for each task, but the model only needs updating within the most recent region. Overlapping regions ensure that more parameters remain in areas of low task loss, crucial for maintaining performances across all tasks. Regularizing parameter optimization according to the found flatness helps shrink the gap between the regions. Experimentally, Figure 3 shows the overall model sharpness change for the first task after trained on each task by C&F and C (Create), which discards Find and the corresponding loss regularization term. We approximate the Hessian Spectrum using the Lanczos algorithm [9] and choose the largest eigenvalue to characterize sharpness, reporting the log value. As expected, the model trained with Find maintains lower sharpness for the first task as learning new tasks, implying that more parameters reside in the firstly created flat region. Thus, Find enhances the regions’ overlap degree, resulting in the construction of superior flat training spaces.

Moreover, although the creation of flat regions flattens all parameters, some parameters remain sharp with a small flat range. Over-optimizing them can cause the model to deviate from the minimum

and leave the region. Results in Figure 3 demonstrate that Find further constrains the parameters within regions by preventing excessive optimization on those sharp parameters. It experimentally verifies that the flatness found for each parameter is consistent with the created flatness.

5.4.2 Flatness Indicator

We have validated the use of Fisher Information as the flatness indicator in our framework both theoretically (§3.4) and experimentally (§5.4.1). Furthermore, we experiment with alternative indicators from MAS [1] and SI [38] and replace Fisher with them in C&F framework, respectively. We also compare with C (Create), eliminating Find and its corresponding regularization from C&F, and **Random**, randomly sampling flatness value for each parameter. Results in Table 4 underscore the superiority of Fisher Information as a flatness indicator within our C&F framework, proving that it is more conducive to constructing optimal learning spaces.

Table 4: Comparisons of different flatness indicator for Find. MAS and SI replace Fisher Information, which is originally used in C&F, with indicator utilized in MAS [1] and SI [38], respectively.

Method	IV	V	VI	AVG
C	74.70	74.06	75.32	74.70
Random	75.34	74.67	75.49	75.17
MAS	76.26	75.03	75.87	75.72
SI	76.00	75.04	75.63	75.45
C&F	76.61	76.46	76.06	76.38

5.4.3 Sparsity of Find

Given that continual learning necessitates saving every model for rollback in practical applications, we further investigate the effect of parameter sparsity in Find. The most straightforward approach is to find parameters with the least importance to the last task in each layer and solely optimize on them in the new task. Table 5 presents the results of varying sparingly optimizing ratios applied to C&F per layer. The performance gradually increases as the ratio rises, with only 50% of parameters being optimized achieving performances comparable to full optimization. Sparsity ensures more parameters located in the flat regions, especially those are sharp, which further alleviates catastrophic forgetting. Therefore, C&F with sharp parameters fixed can be a suitable choice in practice.

5.5 Time Comparison

While the C&F framework offers a distinctive approach that encompasses both the creating and finding flatness processes, it tends to

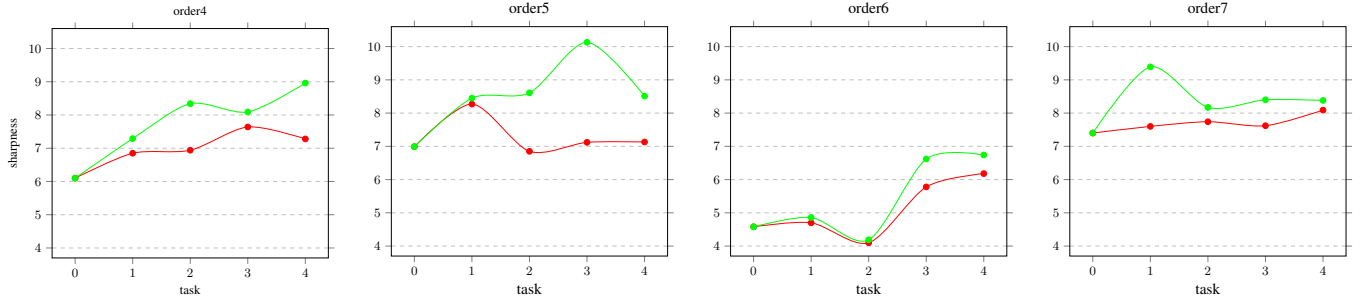


Figure 3: Comparisons of whole model’s sharpness for first task during continual learning on different orders. **Red lines** denote C&F and **green lines** denote C, C&F without Find. Sharpness is calculated by the log value of the largest eigenvalue of the Hessian Spectrum.

Table 5: Comparisons of performance with increasing sparingly updating ratio per layer. 100% is the setting of C&F.

SU_Ratio	IV	V	VI	AVG
0.1%	69.41	72.83	68.04	70.09
1%	71.26	74.77	68.60	71.54
10%	74.74	75.82	73.96	74.84
20%	75.34	75.89	74.90	75.38
50%	76.26	75.92	75.83	76.00
100%	76.61	76.46	76.06	76.38

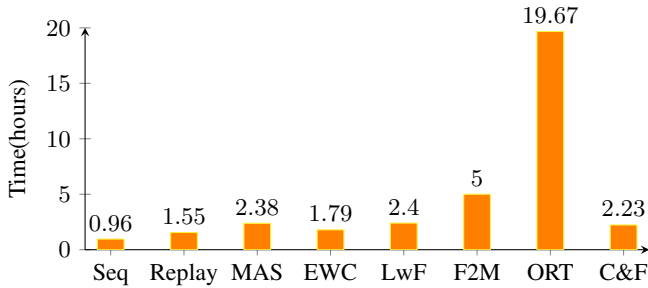


Figure 4: Time comparisons in hours for the main methods on sampled-setting order4 task sequence.

exhibit marginally lower efficiency compared to direct optimization methods. This primarily results from the additional time required for creating flat regions, which may impede the optimization speed. We provide time comparisons in Fig 4, illustrating the learning time in hours for the main methods on the sampled-setting order4 task sequence, as used in Table 1. The training settings of these methods are consistent. It’s noted that C&F’s time extension is insignificant and its training time is considerably less than that of other flatness-based or gradient-based methods like F2M or ORT. This reveals that directly incorporating gradient or flatness characteristics into continual learning may lead to substantial efficiency degradation.

5.6 Effect of C&F Components

We have devised a generic framework for continual learning, with the core of creating and finding flatness and the associated optimization constraints. We outline each component of C&F below and examine the effects of ablating them one by one. Firstly, We replay samples saved from previous tasks, which are picked out as the most representative using K-Means [17]. In addition, we create a flat region around the minimum of each task. To further mitigate catastrophic forgetting, we incorporate an L2 loss term in optimization, employ-

ing random coefficient as the flatness value for each parameter. Furthermore, we find each parameters’ flatness and use it as the coefficient of the L2 term specific to each parameter. Finally, we clamp all parameters’ optimization within the created regions. As shown in Table 6, incrementally discarding each ingredient leads to a decline in accuracy. Thus, every component in the C&F framework is essential for achieving the optimal performance.

Table 6: Effects of ablating different components in C&F.

Method	IV	V	VI	AVG
C&F	76.61	76.46	76.06	76.38
- Clamp	76.32	75.26	75.80	75.79
- Find	75.77	73.73	75.61	75.04
- L2	74.10	73.64	75.21	74.31
- Create	71.68	71.84	74.37	72.63
- Replay	68.38	67.21	69.87	68.49

6 Conclusion

In this work, we introduce C&F, an innovative framework incorporating geometrical properties of flatness to build flat training spaces in advance for continual learning. C&F **creates** a flat region around the minimum and **finds** each parameters’ flatness both with respect to the current task. During the new task’s optimization, C&F constrains parameter changes in the created flat regions and regularizes them according to the found flatness, simultaneously preparing a new flat training space for the subsequent task. As an individual continual learning method, C&F outperforms the previous state-of-the-art approaches in text classification task sequences. Moreover, when other methods are integrated into C&F framework, their performances attain further boost. In future research, we aim to explore the flatness’s applicability to other fields concerning continual learning.

Acknowledgements

The work was supported by the National Natural Science Foundation of China under Grant No. 62072458.

References

- [1] 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)*, pp. 139–154, (2018).
- [2] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuezhi Li, Scott Lundberg, et al., ‘Sparks of artificial general intelligence: Early experiments with gpt-4’, *arXiv preprint arXiv:2303.12712*, (2023).

- [3] Pratik Chaudhari, Anna Choromanska, Stefano Soatto, Yann LeCun, Carlo Baldassi, Christian Borgs, Jennifer Chayes, Levent Sagun, and Riccardo Zecchina, 'Entropy-sgd: Biasing gradient descent into wide valleys', *Journal of Statistical Mechanics: Theory and Experiment*, **2019**(12), 124018, (2019).
- [4] Arslan Chaudhry, Puneet K Dokania, Thalaiyasingam Ajanthan, and Philip HS Torr, 'Riemannian walk for incremental learning: Understanding forgetting and intransigence', in *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 532–547, (2018).
- [5] Cyprien de Masson D'Autume, Sebastian Ruder, Lingpeng Kong, and Dani Yogatama, 'Episodic memory in lifelong language learning', *Advances in Neural Information Processing Systems*, **32**, (2019).
- [6] Mehrdad Farajtabar, Navid Azizan, Alex Mott, and Ang Li, 'Orthogonal gradient descent for continual learning', in *International Conference on Artificial Intelligence and Statistics*, pp. 3762–3773. PMLR, (2020).
- [7] Chrisantha Fernando, Dylan Banarse, Charles Blundell, Yori Zwols, David Ha, Andrei A Rusu, Alexander Pritzel, and Daan Wierstra, 'Pathnet: Evolution channels gradient descent in super neural networks', *arXiv preprint arXiv:1701.08734*, (2017).
- [8] Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur, 'Sharpness-aware minimization for efficiently improving generalization', *arXiv preprint arXiv:2010.01412*, (2020).
- [9] Behrooz Ghorbani, Shankar Krishnan, and Ying Xiao, 'An investigation into neural net optimization via hessian eigenvalue density', in *International Conference on Machine Learning*, pp. 2232–2241. PMLR, (2019).
- [10] Yufan Huang, Yanzhe Zhang, Jiaao Chen, Xuezhi Wang, and Diyi Yang, 'Continual learning for text classification with information disentanglement based regularization', *arXiv preprint arXiv:2104.05489*, (2021).
- [11] Zixuan Ke, Yijia Shao, Haowei Lin, Tatsuya Konishi, Gyuhak Kim, and Bing Liu, 'Continual learning of language models', *arXiv preprint arXiv:2302.03241*, (2023).
- [12] 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).
- [13] Jungmin Kwon, Jeongseop Kim, Hyunseo Park, and In Kwon Choi, 'Asam: Adaptive sharpness-aware minimization for scale-invariant learning of deep neural networks', in *International Conference on Machine Learning*, pp. 5905–5914. PMLR, (2021).
- [14] Zhizhong Li and Derek Hoiem, 'Learning without forgetting', *IEEE transactions on pattern analysis and machine intelligence*, **40**(12), 2935–2947, (2017).
- [15] Junling Liu, Chao Liu, Renjie Lv, Kang Zhou, and Yan Zhang, 'Is chatgpt a good recommender? a preliminary study', *arXiv preprint arXiv:2304.10149*, (2023).
- [16] David Lopez-Paz and Marc Aurelio Ranzato, 'Gradient episodic memory for continual learning', *Advances in neural information processing systems*, **30**, (2017).
- [17] J MacQueen, 'Classification and analysis of multivariate observations', in *5th Berkeley Symp. Math. Statist. Probability*, pp. 281–297, (1967).
- [18] 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*, pp. 7765–7773, (2018).
- [19] Michael McCloskey and Neal J Cohen, 'Catastrophic interference in connectionist networks: The sequential learning problem', in *Psychology of learning and motivation*, volume 24, 109–165, Elsevier, (1989).
- [20] Jacob Menick, Maja Trebacz, Vladimir Mikulik, John Aslanides, Francis Song, Martin Chadwick, Mia Glaese, Susannah Young, Lucy Campbell-Gillingham, Geoffrey Irving, et al., 'Teaching language models to support answers with verified quotes', *arXiv preprint arXiv:2203.11147*, (2022).
- [21] Seyed Iman Mirzadeh, Mehrdad Farajtabar, Razvan Pascanu, and Hassan Ghasemzadeh, 'Understanding the role of training regimes in continual learning', *Advances in Neural Information Processing Systems*, **33**, 7308–7320, (2020).
- [22] Natawut Monaiikul, Giuseppe Castellucci, Simone Filice, and Oleg Rokhlenko, 'Continual learning for named entity recognition', in *Proceedings of the AAI Conference on Artificial Intelligence*, volume 35, pp. 13570–13577, (2021).
- [23] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al., 'Training language models to follow instructions with human feedback', *Advances in Neural Information Processing Systems*, **35**, 27730–27744, (2022).
- [24] Yujia Qin, Jiajie Zhang, Yankai Lin, Zhiyuan Liu, Peng Li, Maosong Sun, and Jie Zhou, 'Elle: Efficient lifelong pre-training for emerging data', *arXiv preprint arXiv:2203.06311*, (2022).
- [25] Anastasia Razdaibiedina, Yuning Mao, Rui Hou, Madian Khabza, Mike Lewis, and Amjad Almahairi, 'Progressive prompts: Continual learning for language models', *arXiv preprint arXiv:2301.12314*, (2023).
- [26] Sylvester-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*, pp. 2001–2010, (2017).
- [27] 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).
- [28] Victor Sanh, Albert Webson, Colin Raffel, Stephen H Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, et al., 'Multitask prompted training enables zero-shot task generalization', *arXiv preprint arXiv:2110.08207*, (2021).
- [29] Guangyuan Shi, Jiabin Chen, Wenlong Zhang, Li-Ming Zhan, and Xiao-Ming Wu, 'Overcoming catastrophic forgetting in incremental few-shot learning by finding flat minima', *Advances in Neural Information Processing Systems*, **34**, 6747–6761, (2021).
- [30] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim, 'Continual learning with deep generative replay', *Advances in neural information processing systems*, **30**, (2017).
- [31] Fan-Keng Sun, Cheng-Hao Ho, and Hung-Yi Lee, 'Lamol: Language modeling for lifelong language learning', *arXiv preprint arXiv:1909.03329*, (2019).
- [32] Ming Tu, Visar Berisha, Martin Woolf, Jae-sun Seo, and Yu Cao, 'Ranking the parameters of deep neural networks using the fisher information', in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 2647–2651. IEEE, (2016).
- [33] Rui Wang, Tong Yu, Handong Zhao, Sungchul Kim, Subrata Mitra, Ruiyi Zhang, and Ricardo Henao, 'Few-shot class-incremental learning for named entity recognition', in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 571–582, (2022).
- [34] Wei Wen, Yandan Wang, Feng Yan, Cong Xu, Chunpeng Wu, Yiran Chen, and Hai Li, 'Smoothout: Smoothing out sharp minima to improve generalization in deep learning', *arXiv preprint arXiv:1805.07898*, (2018).
- [35] Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhajan Kambadur, David Rosenberg, and Gideon Mann, 'Bloomberggpt: A large language model for finance', *arXiv preprint arXiv:2303.17564*, (2023).
- [36] Zhaofeng Wu, Robert L Logan IV, Pete Walsh, Akshita Bhagia, Dirk Groeneveld, Sameer Singh, and Iz Beltagy, 'Continued pre-training for better zero-and few-shot promptability', *arXiv preprint arXiv:2210.10258*, (2022).
- [37] Dani Yogatama, Cyprien de Masson d'Autume, Jerome Connor, Tomas Kocisky, Mike Chrzanowski, Lingpeng Kong, Angeliki Lazaridou, Wang Ling, Lei Yu, Chris Dyer, et al., 'Learning and evaluating general linguistic intelligence', *arXiv preprint arXiv:1901.11373*, (2019).
- [38] Friedemann Zenke, Ben Poole, and Surya Ganguli, 'Continual learning through synaptic intelligence', in *International Conference on Machine Learning*, pp. 3987–3995. PMLR, (2017).
- [39] Bowen Zhao, Xi Xiao, Guojun Gan, Bin Zhang, and Shu-Tao Xia, 'Maintaining discrimination and fairness in class incremental learning', in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13208–13217, (2020).
- [40] Zhe Zhao, Hui Chen, Jinbin Zhang, Xin Zhao, Tao Liu, Wei Lu, Xi Chen, Haotang Deng, Qi Ju, and Xiaoyong Du, 'Uer: An open-source toolkit for pre-training models', *arXiv preprint arXiv:1909.05658*, (2019).
- [41] Zhe Zhao, Yudong Li, Cheng Hou, Jing Zhao, Rong Tian, Weijie Liu, Yiren Chen, Ningyuan Sun, Haoyan Liu, Weiyan Mao, et al., 'Tencentpretrain: A scalable and flexible toolkit for pre-training models of different modalities', *arXiv preprint arXiv:2212.06385*, (2022).