

# DATA: Decomposed Attention-based Task Adaptation for Rehearsal-Free Continual Learning

Anonymous ACL submission

## Abstract

Continual learning (CL) is essential for Large Language Models (LLMs) to adapt to evolving real-world demands, yet they are susceptible to catastrophic forgetting (CF). While traditional CF solutions rely on expensive data rehearsal, recent rehearsal-free methods employ model-based and regularization-based strategies to address this issue. However, these approaches often neglect the model’s plasticity, which is crucial to achieving optimal performance on newly learned tasks. Consequently, a key challenge in CL is striking a balance between preserving plasticity and mitigating CF. To tackle this challenge, we propose the **Decomposed Attention-based Task Adaptation (DATA)**, which explicitly decouples and learns both task-specific and task-shared knowledge using high-rank and low-rank task adapters (e.g., LoRAs). For new tasks, DATA dynamically adjusts the weights of adapters of different ranks based on their relevance and distinction from previous tasks, allowing the model to acquire new task-specific skills while effectively retaining previously learned knowledge. Specifically, we implement a decomposed component weighting strategy comprising learnable components that collectively generate attention-based weights, allowing the model to integrate and utilize diverse knowledge from each DATA. Extensive experiments on three widely used benchmarks demonstrate that our proposed method achieves state-of-the-art performance. Notably, our approach significantly enhances model plasticity and mitigates CF by extending learnable components and employing stochastic restoration during training iterations. Our code will be available at <https://anonymous.4open.science/r/DATA1>.

## 1 Introduction

Continual learning (CL) (Wang et al., 2023a; Zhou et al., 2024) is essential for Large Language Models (LLMs) (Dubey et al., 2024; Yang et al., 2024)

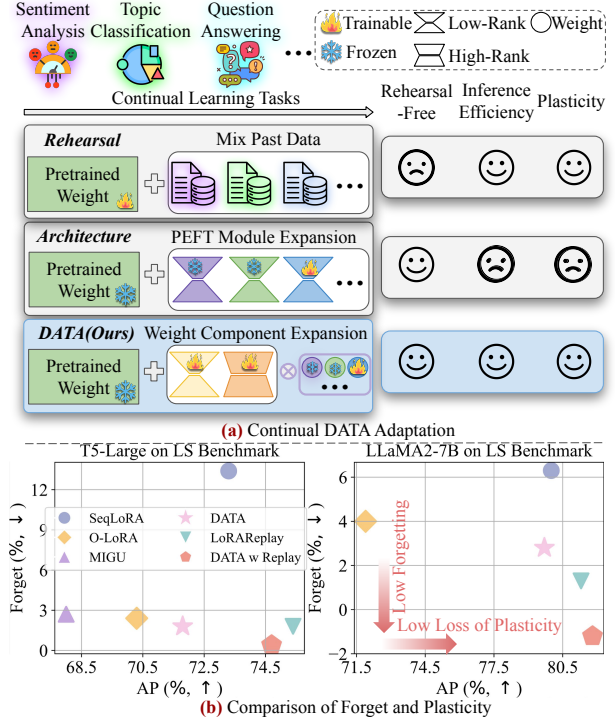


Figure 1: (a) Comparisons of previous CL methods with **DATA**. *Rehearsal-free* indicates that the methods do not require storing sample features from previous tasks. *Inference Efficiency* denotes the computational efficiency during the inference phase. *Plasticity* is the ability to adapt effectively to new tasks. (b) Comparison of **AP** and **Forget** (Sec. 4.1) across different CL methods.

to continuously evolve, adapting to real-world demands and progressively improve regarding a series of new tasks. Unlike traditional supervised learning, which relies on independent and identically distributed (i.i.d.) data, CL focuses on the dynamic and evolving demands of real-world applications (Zhai et al., 2023; Wu et al., 2024), such as domain-specific adaptations (Roziere et al., 2023) and alignment with human preferences (Ouyang et al., 2022). Departing from the i.i.d. assumption presents two major challenges: 1) **Catastrophic Forgetting (CF)** — the tendency of a model to over-

write previously learned knowledge when acquiring new tasks (McCloskey and Cohen, 1989). 2) **Loss of Plasticity**—the model’s diminishing ability to adapt effectively to distribution shifts over time (Dohare et al., 2021).

As illustrated in Figure 1 (a), an ideal CL method leveraging LLMs should possess three essential properties. A commonly used strategy in CL is the preservation (Wang et al., 2024b) or synthesis (Sun et al., 2019) of past training data, referred to as **rehearsal** (or *replay*). However, this approach presents notable challenges, particularly in scenarios involving *sensitive user data* where long-term storage is impractical (Touvron et al., 2023). Furthermore, rehearsal methods require *increased storage capacity and computational resources* during the learning process (Scialom et al., 2022). Architecture-based methods (Wang et al., 2024a) aim to mitigate interference between new and previously learned tasks by dynamically expanding the model’s capacity or isolating existing model weights. Nevertheless, these approaches often necessitate training separate expert models for each task, which in turn requires selecting (Zhao et al., 2024) or combining (Wang et al., 2023b) modules during testing. This can limit their ability to generalize effectively to long sequence tasks. Moreover, as shown in Figure 1 (b), these methods mitigate CF by *sacrificing new task accuracy* (i.e., insufficient *plasticity*), prompting us to explore innovative perspectives for enhancing CL capacity.

To address these limitations while maintaining plasticity and mitigating CF, we propose a **Decomposed Attention-based Task Adaptation** (DATA), as depicted in Figure 1 (a), which explicitly manages high-rank and low-rank adapters for task-specific and task-shared knowledge during the CL. For new tasks, DATA dynamically adjusts the weights of different ranks adapters based on their relevance and distinction from previous tasks, enabling the model to acquire new task-specific skills while preserving previously learned capabilities. To extract knowledge across tasks more effectively, we introduce a *decomposed component weighting* strategy comprising learnable components that jointly generate attention-based weights, allowing dynamic fusion of knowledge from each rank DATA. During inference, low-rank and high-rank DATA can be reparameterized and projected into LLMs without increasing parameters, preserving the model plasticity (Liu et al., 2023). We use *stochastic restoration* (Wang et al., 2022; Dohare

et al., 2024) technique to further protect source knowledge and mitigate CF.

We conduct extensive experiments to evaluate DATA on Standard CL Benchmark (Zhang et al., 2015), Long Sequence Benchmark (Razdaibiedina et al., 2023) and TRACE (Wang et al., 2023c) across T5 model (Raffel et al., 2019), Llama2 (Touvron et al., 2023), Llama3.1 (Dubey et al., 2024) and Qwen2.5 (Yang et al., 2024). DATA achieves state-of-the-art performance on both public benchmarks and in generalizing to unseen tasks. Notably, DATA can be integrated with existing mainstream methods to further enhance the capabilities of CL. In summary, our contributions are as follows:

- We propose a method DATA for rehearsal-free CL of LLMs that balances plasticity and mitigates CF by explicitly managing both *task-specific* and *task-shared* knowledge.
- We explore different task representations of high- and low-rank adapters, decomposing knowledge into shared and specific components. This approach allows for better modeling of the relevance and uniqueness of new tasks concerning previous ones.
- To address various distribution shifts for each target sample, we introduce a *decomposed component weighting* strategy for DATA, which dynamically fuses knowledge from low-rank and high-rank adapters, enhancing their task representations.
- Experiments on three benchmarks show that our approach outperforms most state-of-the-art techniques, significantly alleviating CF and improving model plasticity.

## 2 Related Work

**Continual Learning (CL).** CL aims to effectively acquire new task knowledge during ongoing training while preserving knowledge from previously learned tasks. Traditional CL methods can be broadly classified into three categories: *rehearsal-based*, *regularization-based*, and *architecture-based* approaches. Rehearsal-based methods mitigate catastrophic forgetting (CF) by selectively retaining samples (Tiware et al., 2021; Wang et al., 2024b; He et al., 2024) or pseudo-generative examples (Shin et al., 2017; Sun et al., 2019) from prior tasks. In contrast, regularization-based methods employ quadratic regularization

terms to constrain the update of weights critical to previous tasks (Zhu et al., 2024; Du et al., 2024), thereby stabilizing learned knowledge. Architecture-based methods introduce extra task-specific parameters for each new task (Wang et al., 2023b; Zhao et al., 2024), ensuring that previously acquired knowledge remains intact. These methods typically rely on access to old task data or involve discrete component learning, which presents significant challenges for the CL. In contrast, DATA leverages representations from diverse rank spaces and harnesses the inherent capabilities of LLMs to enable more effective CL (Wang et al., 2022).

**Parameter Efficient Fine-Tuning (PEFT).** LLMs have shown remarkable adaptability across a wide range of downstream tasks. However, traditional full fine-tuning for each task can incur *significant computational and storage overheads*, often resulting in overfitting (Wu et al., 2025). To mitigate these challenges, researchers have increasingly turned to PEFT methods, which can achieve comparable or even superior performance and generalization by fine-tuning only a small subset of parameters. For instance, *adapters* (Houlsby et al., 2019) incorporate additional modules into various layers of the Transformer architecture, while *prompt-tuning* (Qin and Eisner, 2021) and *prefix-tuning* (Li and Liang, 2021) introduce learnable soft tokens at different layers of the Transformer input. *Low-rank adaptation* (LoRA) (Hu et al., 2022) inserts low-rank branches into pre-trained weights and fine-tunes only those branches. Although these common PEFT techniques offer certain advantages, they are generally confined to static single-task or multi-task learning scenarios. In contrast, DATA leveraging the effectiveness and efficiency of LoRA enables continual adaptation by dynamically combining LoRAs of two different ranks (Liu et al., 2023) utilizing attention-based weights.

### 3 Methodology

**Preliminary.** Continual learning (CL) seeks to tackle challenges within ongoing sequences, specifically addressing how to adapt to new tasks without forgetting previously learned knowledge. Formally, suppose there are  $\mathcal{N}$  sequential tasks  $\{\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_{\mathcal{N}}\}$ , where each task  $\mathcal{T}_i = \{\mathbf{x}_n^i, \mathbf{y}_n^i\}_{n=1}^{\mathcal{N}_i}$  consists of  $\mathcal{N}_i$  training examples, with  $\mathbf{x}_n^i$  representing the input and  $\mathbf{y}_n^i$  is corresponding label. Let  $\mathcal{L}_i(\cdot)$  denote the empirical risk on the  $i$ -th task  $\mathcal{T}_i$ , and  $f_\theta$  represent the model with

parameters  $\theta$ . The goal is for the model to perform well on both the current task  $\mathcal{T}_j$  and all previously learned tasks  $\mathcal{T}_i$ , where  $i \in \{1, 2, \dots, j-1\}$ . The objective function of CL can be expressed as:

$$\mathcal{L}_{\text{CL}}(\theta) = \frac{1}{j} \sum_{i=1}^j \mathcal{L}_i(\theta) \quad (1)$$

where  $j$  is the index of the current training task.

#### 3.1 Motivation

Recent advances have highlighted the effectiveness of using adapters with varying ranks to model different tasks, offering a promising solution to the challenges of CL. This motivates us to further explore and validate the core principles behind utilizing both low-rank and high-rank adapters in CL.

To further explore this phenomenon, we conduct a t-SNE distribution analysis (van der Maaten and Hinton, 2008) to examine the feature distributions of Order1 (4 tasks, Sec. 4.1) within the 11th transformer block of LLaMA2-7B. The results as presented in Figure 2 (a) reveal that low-rank adapters maintain relatively consistent feature distributions across different target tasks. This consistency suggests that the low-rank embedding space effectively mitigates the effects of dynamic distribution shifts by concentrating on extracting task-shared knowledge. Conversely, high-rank adapters display noticeable variability across tasks. This variability underscores the enhanced ability of high-rank adapters to aggregate features within a single task, highlighting their suitability for capturing the unique data distribution of each target task.

Furthermore, we utilize the *H-divergence* metric to assess the representation consistency of adapters on Order 4 (15 tasks). A relatively small inter-task divergence indicates that feature representations remain stable and are less affected by cross-task shifts (Ganin et al., 2015), while a small intra-task divergence within a given task indicates that the model better understands the current distribution. We compare the divergence values obtained from the source model alone, injecting low-rank adapters and injecting high-rank adapters, as shown in Figure 2). Compared to the source model and high-rank adapters, low-rank adapters produce feature representations with lower inter-task divergence, particularly when addressing mid-range target tasks or significant task shifts between adjacent tasks (i.e., tasks 6-10). These findings underscore the efficacy of low-rank adapters in acquiring long-term,






Figure 2: (a) We perform a t-SNE distribution analysis of different adapter representations on Order 1(4 tasks). The low-rank branch shows a consistent distribution across the target tasks and the high-rank branch exhibits substantial distribution differences across the target tasks. (b) We calculate the divergence of different branches in Order 4 (15 tasks). In comparison to the source model, low-rank adapters effectively alleviates inter-task divergence across all 14 task transitions, while the high-rank adapters significantly enhances intra-task feature aggregation.

task-shared knowledge during the CL. Furthermore, high-rank adapters effectively reduce intra-task divergence across nearly all tasks, indicating their superior capability in adapting to the current task distribution and extracting task-specific knowledge in consecutive target tasks.

In summary, the structure of low-rank DATA reduces feature redundancy, placing the model in an underfitting state and enhancing its plasticity. This design enables the model to acquire general information across consecutive target tasks and effectively extract task-shared knowledge. In contrast, high-rank DATA aligns more closely with the target data distribution, allowing it to focus on learning task-specific knowledge and mitigating CF.

### 3.2 High- and Low-Rank Adapter

The above observations prompted us to introduce high- and low-rank Adapters into the model, aiming to simultaneously adapt current task distribution while maintaining task-shared knowledge.

Similar to LoRA, DATA’s design principle is simple and effective, as illustrated in Figure 3 (b). The architecture consists of three sub-branches: the central branch, which originates from the original model’s linear layer, and the right and left branches, which are bottleneck structures representing high-rank and low-rank DATA, respectively. Specifically, the right branch (high-rank) consists of an up-projection layer with parameters  $\mathbf{A}_2 \in \mathbb{R}^{d \times d_h}$ , and a down-projection layer with parameters  $\mathbf{B}_2 \in \mathbb{R}^{d_h \times d}$ , where  $d_h$  (e.g.,  $d_h = 32$ ) represents the middle dimension of the high-rank feature. The same principle applies to the low-rank branch, except it satisfies  $d_l < d_h$ . For an input  $\mathbf{x}$ , the resulting high-rank features  $\mathbf{f}_h$  and low-rank features  $\mathbf{f}_l$  from the DATA are formulated as:

$$\mathbf{f}_h = \mathbf{A}_2 \cdot (\mathbf{B}_2 \cdot \mathbf{x}); \mathbf{f}_l = \mathbf{A}_1 \cdot (\mathbf{B}_1 \cdot \mathbf{x}) \quad (2)$$

The two-branch bottleneck is connected to the output feature of the original linear  $\mathbf{f}_o$  through a residual connection, with scale weights  $\lambda_h$  and  $\lambda_l$ . The fusion of knowledge  $\mathbf{f}_x$  is then expressed as:

$$\mathbf{f}_x = \mathbf{f}_o + \lambda_h \times \mathbf{f}_h + \lambda_l \times \mathbf{f}_l \quad (3)$$

The task knowledge scale weights  $\lambda_h$  and  $\lambda_l$  can be adaptively computed through *decomposed component weighting* strategy as detailed in Sec. 3.3. During the inference process, the distinct task representations of the DATA are reparameterized and projected into the original model, enabling dynamic knowledge adjustments.

### 3.3 Decomposed Component Weighting

To effectively mitigate CF across various tasks and samples, it is crucial to extract and manage different types of knowledge. While the distinct structure of low-rank and high-rank DATA facilitates learning diverse task representations, the continual adaptation process also necessitates the normalization of knowledge fusion weights. This normalization ensures the effective capture of relevant task-specific knowledge while preserving long-term task-specific knowledge.

As illustrated in Figure 3 (a), we draw inspiration from (Smith et al., 2022) and propose a *decomposed component weighting* strategy. Specifically, we introduce a set of weight components that form a decomposed weight through weighted summation, which is then passed to the corresponding DATA layer. Furthermore, in new tasks, the weights will essentially reuse previously acquired knowledge about past tasks instead of initializing new task weights from scratch:

$$\lambda = \sum_m \alpha_m \mathbf{W}_m \quad (4)$$

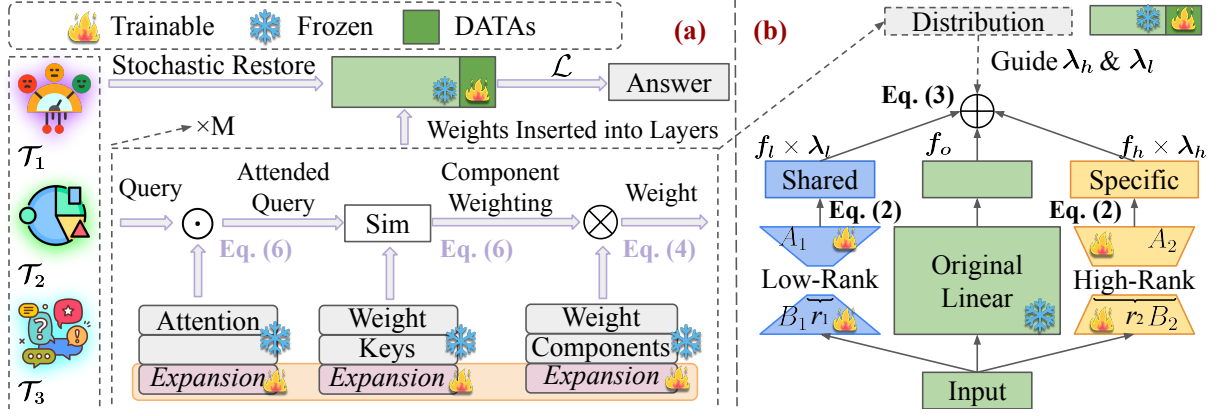


Figure 3: **Overview of the DATA framework.** (a) We introduce a novel *decomposed component weighting* strategy for generating **attention-based weights**, parameterized by a set of *extended* weight components, each associated with a corresponding key and attention vector. Only the weighting and adapter parameters are trainable which is *parameter efficient* and no training data is stored for the replay which is *memory-efficient* and *privacy-preserving*. (b) We integrate low-rank and high-rank DATA into the linear layers of the pre-trained model guided by the generated weights, allowing for the dynamic fusion of knowledge from each DATA with different task representations.

where  $W \in \mathbb{R}^{L_w \times d \times M}$  represents our set of weight components, with  $M$  denoting the length of the set (i.e., the extra axis of additional capacity),  $L_w$  is the weight length (chosen as a hyperparameter) and  $\alpha$  is a *weighting vector* that determines the contribution of each component, which is the cosine similarity  $\gamma$  between a query  $q(x)$  and keys:

$$\alpha = \gamma(q(x), \mathbf{K}) \quad (5)$$

where  $\mathbf{K} \in \mathbb{R}^{d \times M} = \{\mathbf{K}_1, \mathbf{K}_2, \dots, \mathbf{K}_M\}$  contains keys corresponding to weight components. The intuition behind this is that the contribution of each weight component  $\mathbf{W}_m$  to the final weight  $\lambda$  is weighted according to the similarity.

To address the challenge of query matching, we introduce an additional component: *attention*. Alongside  $\mathbf{K}_m$ , each weight component  $\mathbf{W}_m$  is paired with a corresponding attention vector  $\mathbf{A}_m$ , enabling the query to focus on specific features within the high-dimensional query. We utilize a simple feature selection attention mechanism, where the query vector is element-wise multiplied by the attention vector to generate an *attended query* which is then used to compute the similarity with the key. Specifically, our updated approach to producing the weighting vector is as follows:

$$\alpha = \gamma(q(x) \odot \mathbf{A}, \mathbf{K}) \quad (6)$$

where  $\mathbf{A} \in \mathbb{R}^{d \times M} = \{\mathbf{A}_1, \mathbf{A}_2, \dots, \mathbf{A}_M\}$  represents the learnable attention vectors corresponding to each weight component, and  $\odot$  denotes the element-wise (Hadamard) product. It is important to highlight that our attention vectors act

as learnable feature weightings rather than input-conditioned modules. We observe that this fixed representation with its simple design is less susceptible to forgetting, similar to the behavior of our weight component keys.

### 3.4 Expansion & Orthogonality

The key to alleviating CF is to prevent overwriting the knowledge acquired in previous tasks. When encountering a new task, we freeze the existing components and expand the collection by updating only the new components. This is illustrated at the bottom of Figure 3, where the existing parameters are *frozen*, and only the newly added parameters are *trainable*. Specifically, for task  $\mathcal{T}_t$ , we learn  $\frac{M}{N}$  components, where  $N$  denotes the number of tasks and  $M$  is the hyperparameter, while the previously learned  $\frac{(t-1) \cdot M}{N}$  components are kept frozen. This extension is achieved through our attention-based component weighting scheme, ensuring that the expansion of parameters does not affect the computation of the weights  $\alpha$  corresponding to the previously learned components.

To further mitigate CF, we introduce orthogonality constraints for  $\mathbf{W}$ ,  $\mathbf{K}$ , and  $\mathbf{A}$ . The underlying intuition is that interference is minimized when vectors are orthogonal. For instance, we aim to prevent the keys and weights learned in task  $\mathcal{T}_2$  from influencing data from task  $\mathcal{T}_1$ . To achieve this, we initialize the vectors orthogonally and incorporate a simple orthogonal penalty loss as:

$$\mathcal{L}_{\text{ortho}}(B) = \|BB^\top - I\|_2^2 \quad (7)$$

Methods		Standard CL Benchmark (SC)			Long Sequence Benchmark (LS)			TRACE		
		FP $\uparrow$	AP $\uparrow$	Forget $\downarrow$	FP $\uparrow$	AP $\uparrow$	Forget $\downarrow$	FP $\uparrow$	AP $\uparrow$	Forget $\downarrow$
T5-Large	SeqLoRA	70.7 $\pm$ .39	<b>76.7</b> $\pm$ .43	6.0	59.9 $\pm$ .56	73.3 $\pm$ .28	13.4	12.1 $\pm$ .82	44.5 $\pm$ .94	32.4
	LoRAReplay	73.3 $\pm$ .42	76.6 $\pm$ .51	3.3	73.6 $\pm$ .36	<b>75.4</b> $\pm$ .59	1.8	34.0 $\pm$ .62	<b>46.8</b> $\pm$ .63	12.8
	O-LoRA (Wang et al., 2023b)	72.0 $\pm$ .63	74.4 $\pm$ .47	2.4	67.9 $\pm$ .82	70.3 $\pm$ .65	2.4	-	-	-
	+ MIGU (Du et al., 2024)	71.6 $\pm$ .45	73.9 $\pm$ .67	2.3	65.3 $\pm$ .35	68.0 $\pm$ .47	2.7	-	-	-
	<b>DATA (ours)</b>	72.7 $\pm$ .58	74.8 $\pm$ .68	2.1	70.0 $\pm$ .44	71.8 $\pm$ .36	1.8	16.7 $\pm$ .21	41.3 $\pm$ .58	24.6
	+ Replay	<b>73.9</b> $\pm$ .24	74.5 $\pm$ .75	<b>0.6</b>	<b>74.3</b> $\pm$ .35	74.7 $\pm$ .91	<b>0.4</b>	<b>36.5</b> $\pm$ .32	45.2 $\pm$ .61	<b>8.7</b>
LLaMA2-7B	SeqLoRA	74.9 $\pm$ .42	<b>80.7</b> $\pm$ .38	5.8	73.7 $\pm$ .66	80.0 $\pm$ .65	6.3	64.1 $\pm$ .49	77.9 $\pm$ .54	13.8
	LoRAReplay	79.2 $\pm$ .53	<b>80.7</b> $\pm$ .61	1.5	80.0 $\pm$ .45	<b>81.3</b> $\pm$ .57	1.3	71.9 $\pm$ .62	<b>78.6</b> $\pm$ .75	6.7
	O-LoRA (Wang et al., 2023b)	76.4 $\pm$ .74	78.7 $\pm$ .59	2.3	67.9 $\pm$ .74	71.9 $\pm$ .49	4.0	35.0 $\pm$ .34	47.0 $\pm$ .42	12.0
	<b>DATA (ours)</b>	79.8 $\pm$ .54	80.3 $\pm$ .49	0.5	76.9 $\pm$ .35	79.7 $\pm$ .36	2.8	66.0 $\pm$ .33	74.6 $\pm$ .38	8.6
	+ Replay	<b>80.0</b> $\pm$ .24	80.4 $\pm$ .45	<b>0.4</b>	<b>81.8</b> $\pm$ .61	80.6 $\pm$ .84	<b>-1.2</b>	<b>73.1</b> $\pm$ .46	77.5 $\pm$ .98	<b>4.4</b>
LLaMA3.1-8B	SeqLoRA	79.6 $\pm$ .62	80.8 $\pm$ .49	5.8	74.8 $\pm$ .58	83.8 $\pm$ .52	9.0	65.1 $\pm$ .69	82.4 $\pm$ .54	17.3
	LoRAReplay	80.3 $\pm$ .71	<b>80.9</b> $\pm$ .56	0.6	82.0 $\pm$ .69	<b>85.0</b> $\pm$ .75	3.0	<b>78.7</b> $\pm$ .83	<b>85.7</b> $\pm$ .68	7.0
	O-LoRA (Wang et al., 2023b)	72.3 $\pm$ .86	73.9 $\pm$ .68	1.6	71.4 $\pm$ .64	74.8 $\pm$ .64	3.7	36.7 $\pm$ .57	50.1 $\pm$ .37	13.4
	<b>DATA (ours)</b>	<b>80.9</b> $\pm$ .40	80.6 $\pm$ .35	-0.3	80.0 $\pm$ .53	82.3 $\pm$ .48	2.3	72.7 $\pm$ .94	80.4 $\pm$ .88	7.7
	+ Replay	80.8 $\pm$ .33	80.4 $\pm$ .47	<b>-0.4</b>	<b>82.2</b> $\pm$ .66	82.6 $\pm$ .77	<b>0.4</b>	77.6 $\pm$ .37	81.0 $\pm$ .72	<b>3.4</b>

Table 1: Performance of baselines and ours **DATA** on standard CL benchmark (Order 1,2,3) and long sequence benchmark (Order 4,5,6) and TRACE (Order 7). **Bold** indicates the best in each setting. We report the mean and standard deviation of results with 3 different runs.

where  $B$  is an arbitrary matrix,  $I$  is identity matrix, and  $\|\cdot\|_2$  denotes Frobenius norm, which measures the squared sum of all entries of the matrix.

Additionally, we propose a stochastic restoration method to restore knowledge from the source pre-trained model, thereby enhancing plasticity. The update of the weight  $W$  at step  $t$  is as follows:

$$Mask \sim \text{Bernoulli}(p) \quad (8)$$

$$W_{t+1} = Mask \odot W_0 + (1 - Mask) \odot W_{t+1} \quad (9)$$

where  $p$  is a small restoration probability, and  $Mask$  is a mask tensor of the same shape as  $W_{t+1}$ . The mask determines which elements within  $W_{t+1}$  will be restored to the source weights  $W_0$ .

### 3.5 Full Optimization

Combining all objectives, our full optimization is:

$$\min_{\mathbf{W}_n, \mathbf{K}_n, \mathbf{A}_n} \mathcal{L}_{\text{CL}}(f_{\theta}, \mathbf{W}, \mathbf{K}, \mathbf{A}, \lambda_h, \lambda_l(\mathbf{x}), \mathbf{y}) + \beta(\mathcal{L}_{\text{ortho}}(\mathbf{W}) + \mathcal{L}_{\text{ortho}}(\mathbf{K}) + \mathcal{L}_{\text{ortho}}(\mathbf{A})) \quad (10)$$

where  $\mathbf{W}_n, \mathbf{K}_n, \mathbf{A}_n$  refer to the weight components and corresponding keys/attention vectors that are unfrozen and trained during task  $\mathcal{T}_n$  and  $\beta$  is a hyperparameter balancing the orthogonality loss.

## 4 Experiment

In Section 4.4, we compare our method with other SOTA methods on three public benchmarks. In Section 4.6, we further evaluate the task generalization ability of the proposed method on unseen target tasks. Comprehensive ablation studies are conducted in Section 4.7.

### 4.1 Datasets

**Standard CL Benchmark (SC)** is a CL benchmark for language models, which consists of five text classification datasets introduced by (Zhang et al., 2015). We follow Wang et al. (2023b) to pick four datasets (AG News, Amazon reviews, DBpedia and Yahoo Answers) and shuffle the tasks into three different orders to form orders 1, 2, and 3.

**Long Sequence Benchmark (LS)** is an extended version of the standard CL benchmark with 15 datasets (five classification tasks, nine GLUE and SuperGLUE tasks, and the IMDB dataset) (Razdaibiedina et al., 2023). Following Razdaibiedina et al. (2023), we select 1,000 random samples for training each task and hold out 500 samples per class for validation and testing. Similarly, we shuffle them to form orders 4, 5, and 6.

**TRACE** is a CL benchmark for LLMs that includes 8 datasets that cover multichoice QA, multilingual capabilities, code generation, and mathematical reasoning (Wang et al., 2023c).

**Metrics.** We adopt the following three metrics to quantify various performances: 1) **FP** =  $\frac{1}{N} \sum_{j=1}^N a_N^{T_j}$  is the average zero-shot performance across all  $N$  tasks after tuning on the final  $N$ -th task. Here,  $a_m^q$  denotes the zero-shot performance on task  $q$  after sequentially tuning the  $m$ -th task, and  $T_j$  refers to the  $j$ -th task in the sequence. 2) **AP** =  $\frac{1}{N} \sum_{j=1}^N a_j^{T_j}$  is the average zero-shot performance when learning each  $j$ -th task, which measures the plasticity of the model. 3) **Forget** = **AP** - **FP** is calculated as the differ-




Figure 4: The shifts in CL methods with **FP** and **Forget** on LS Order 4. DATA prevents the shift (blue bar) and thus mitigates forgetting (orange line).

ence between **AP** and **FP**, as commonly used in previous studies (Wu et al., 2022; Jiang et al., 2025) to quantify forgetting. More detailed information on the datasets about orders and evaluation metrics is presented in the Appendix B.1.

## 4.2 Baselines

We compare our method with the following baselines: 1) **SeqLoRA**: Trains fixed-size LoRA on a sequence of tasks. 2) **LoRAReplay**: Trains new tasks on LoRA with mixing a 2% past task. 3) **O-LoRA** (Wang et al., 2023b): Learns tasks in different LoRA subspaces that are kept orthogonal to each other and sums all LoRA weights up at testing time. 4) **MIGU** (Du et al., 2024): Only updates the model parameters with large magnitudes.

## 4.3 Implementations

We implement DATA with the LLaMA3.1-8B (Dubey et al., 2024), LLaMA2-7B (Touvron et al., 2023), Qwen2.5-7B (Yang et al., 2024) and T5-Large (Raffel et al., 2019). All experiments are conducted on 2 A100 GPUs with 80GB using LLaMA-Factory (Zheng et al., 2024). All experimental results are reported as the average of 3 runs. Please refer to the Appendix B for more detailed settings.

## 4.4 Main Results

To demonstrate the effectiveness of the proposed DATA method, we perform experiments on three CL benchmarks, as summarized in Table 1. Following O-LoRA (Wang et al., 2023b), we present the results of three independent runs with different task orders on the two previous CL benchmarks. Detailed results for each order and each task within a specific order are provided in Appendix C.

**Our DATA significantly reduces the CF of abilities and knowledge in CL.** Traditional CL approaches often result in relatively high levels of

forgetting. While O-LoRA demonstrates slightly better performance yet still experiences over 12% forgetting on the TRACE benchmark in **FP**. In contrast, DATA consistently and significantly mitigates this issue by reducing forgetting across various models and benchmarks. For instance, compared to O-LoRA, DATA increases **FP** by an average of 14.4 and decreases **Forget** by an average of 2.1 in LLaMA2. Furthermore, compared to replay methods, DATA further increases **FP** by an average of 1.3 and reduces **Forget** by an average of 1.1, highlighting its compatibility with other methods to effectively improve capabilities. DATA provides a viable solution for addressing the problem of commonsense forgetting in existing models.

**Our DATA method does not excessively compromise the plasticity needed to learn new tasks.** Compared to other CL methods, this technique demonstrates a negligible reduction in the **FP** and **AP** metric particularly in longer sequence tasks (LS, 15 tasks) and more challenging task instruction generalizations (TRACE, 8tasks), signifying a well-maintained balance between plasticity and stability. DATA outperforms O-LoRA in LLaMA2 by an average of 12.3 in **AP**, while the difference with SeqLoRA, known for not restricting parameter updates and theoretically offering the most optimal plasticity, is less than 1. This indicates that DATA is more suitable for practical applications in LLMs.

**Longer task sequences and more challenging task instructions contribute to increased forgetting.** A comparative analysis of the **Forget** metrics for LS and TRACE with SC reveals a significant rise in forgetting. For example, in the LLaMA3.1-8B model, the decline in performance escalates from 5.8 in SC to 9.0 in LC, and further to 17.3 in TRACE. Furthermore, across all tasks, the forgetting scores **Forget** for both LLaMA2 and LLaMA3 remain below 20, demonstrating that, compared to



Methods	MMLU	BBH	GSM8K	AGIEval	FP
Zero-Shot	65.65	62.12	56.33	17.72	-
SeqLoRA	63.58	<b>11.90</b>	0.00	<b>20.60</b>	79.92
LoRAReplay	60.24	5.99	1.82	10.69	80.13
O-LoRA	62.79	6.31	1.56	13.87	71.83
<b>DATA</b>	<b>64.47</b>	10.42	<b>3.63</b>	16.94	<b>81.39</b>

Table 2: Task generalization comparisons on unseen tasks based the LLaMA3.1-8B after training in Order 1.

the T5 model, LLMs exhibit a stronger capacity to mitigate CF (Jiang et al., 2025).

**Forgetting in CL is model-dependent.** A comparison of **Forget** across different models reveals that stronger foundational models tend to mitigate CF more effectively. For instance, on the LS benchmark, SeqLoRA with LLaMA2-7B exhibits 2.7 less **Forget** than SeqLoRA with LLaMA3.1-8B and 0.8 less **Forget** than LoRAReplay. This suggests that forgetting is influenced not only by the nature of the tasks but also significantly by model-related factors such as size, architecture, and the diversity of pre-training data.

#### 4.5 Performance Shifts

In Figure 4, we visually depict the performance changes observed during the training process of DATA compared to other CL methods. Our proposed method clearly demonstrates its effectiveness in mitigating forgetting and enhancing plasticity. By maintaining a stable structure during the fine-tuning phase, DATA ensures consistent performance on general tasks. In sequential task transitions, it is evident that DATA demonstrates superior **FP** and reduce **Forget** compared to other methods. The line shifts further illustrate that LLaMA3.1-8B exhibits less forgetting than LLaMA2-7B, emphasizing that **forgetting is model-dependent**.

#### 4.6 Task Generalization

We select four benchmarks to evaluate the cross-task generalization capability of DATA, an essential dimension for assessing CL algorithms. As demonstrated in Table 2, DATA effectively balances generalization and CL ability by efficiently extracting task-shared knowledge using low-rank adapters. This finding suggests that actively promoting task-shared knowledge and source capabilities between different tasks is beneficial. However, all methods experience a notable decline in performance on reasoning tasks (GSM8K and BBH), which may be attributed to the simple classification tasks potentially impairing the ability to perform generative

	DATA <sub>h</sub>	DATA <sub>l</sub>	Weight	Attention	Ortho.	Rest.	FP ↑
E <sub>1</sub>	-	-	-	-	-	-	73.0
E <sub>2</sub>	✓	-	-	-	-	-	73.7
E <sub>3</sub>	-	✓	-	-	-	-	75.4
E <sub>4</sub>	✓	✓	-	-	-	-	77.6
E <sub>5</sub>	✓	✓	-	-	-	✓	78.1
E <sub>6</sub>	✓	✓	✓	-	-	✓	78.6
E <sub>7</sub>	✓	✓	✓	✓	-	✓	79.2
E <sub>8</sub>	✓	✓	✓	✓	✓	✓	<b>79.5</b>

Table 3: Ablation studies on different components.

reasoning required for instruction following.

#### 4.7 Ablation Study

We conduct an ablation study in Order 1 using LLaMA2-7B to evaluate the contributions of various components in our method DATA, including high-rank DATA (DATA<sub>h</sub>), low-rank DATA (DATA<sub>l</sub>), decomposed component weighting strategy, attention keys, orthogonality regularization (Ortho.), and stochastic restoration (Rest.). As shown in Table 3 (E<sub>2</sub>), incorporating high-rank DATA increases the **FP** by 0.7, indicating that high-rank features can effectively extract more task-specific knowledge for adaptation to the target task. In E<sub>3</sub>, low-rank DATA improves performance by 2.4 compared to E<sub>1</sub>. The overall improvement in E<sub>4</sub> reaches 4.6, suggesting that both types of DATA complement each other during continual adaptation. The stochastic restoration in E<sub>5</sub> allows the model to maintain plasticity in CL, enhancing its adaptability to new tasks. E<sub>6</sub>-E<sub>8</sub> achieves an additional 2 improvements, demonstrating the effectiveness of the decomposed component weighting strategy in enhancing the task representations for each DATA.

#### 5 Conclusion

In this paper, we revisit existing methods for leveraging Large Language Models (LLMs) in continual learning (CL) and propose three ideal characteristics for such systems: *rehearsal-free*, *inference efficiency*, and *plasticity*. To address the challenges associated with these characteristics, we introduce Decomposed Attention-based Task Adaptation (DATA) as a solution to catastrophic forgetting and plasticity loss. Moreover, we propose a *decomposed component weighting* strategy to dynamically integrate the knowledge from both low-rank and high-rank DATA, thereby enhancing the unique task representation. Extensive experiments across multiple CL benchmarks and LLMs consistently validate the effectiveness of DATA.



## Limitations

**Method** The DATA method introduces a decomposed component weighting strategy and employs both high-rank and low-rank adapters, which increases the complexity of the model architecture. This complexity may lead to higher computational costs during training and inference, particularly when scaling to larger models or more tasks. Additionally, the need for dynamic weight adjustments based on task relevance and distinction may require more sophisticated optimization techniques, potentially limiting its applicability in resource-constrained environments. Furthermore, our current approach to stochastic recovery involves a step-level method, where a small proportion of parameters is recovered every 200 steps. There is significant potential to enhance this process by exploring dynamically adaptive methods that can more effectively select saturated or less important parameters for recovery. Additionally, establishing criteria to determine when recovery is necessary could optimize the process further, potentially improving model performance and efficiency.

**Task** Although DATA is designed to be rehearsal-free, it still relies on the availability of diverse and high-quality task-specific data for effective adaptation. In scenarios where task-specific data is scarce or of low quality, the method’s ability to adapt and generalize may be compromised. Additionally, the method’s performance on tasks with significant domain shifts or out-of-distribution data remains to be fully explored.

**Large Language Models** The effectiveness of DATA is highly dependent on the underlying LLM architecture. While the method shows promising results on models like LLaMA2, LLaMA3.1, and Qwen2.5, its performance may vary across different LLMs, especially those with significantly different architectures or pre-training objectives. Furthermore, we do not experiment with larger models like 13B and 72B due to computational or financial constraints.

## Ethical Considerations

Our approach does not introduce ethical concerns. The datasets we used are public, and there are no privacy issues.

## AI writing statement

This paper utilized AI assistance for language polishing of the manuscript, including vocabulary cor-

rection and spell checking.

## References

- BIG bench authors. 2023. [Beyond the imitation game: Quantifying and extrapolating the capabilities of language models](#). *Transactions on Machine Learning Research*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- OpenCompass Contributors. 2023. Opencompass: A universal evaluation platform for foundation models. <https://github.com/open-compass/opencompass>.
- Shibhansh Dohare, J. Fernando Hernandez-Garcia, Qingfeng Lan, Parash Rahman, A. Ruapm Mahmood, and Richard S. Sutton. 2024. Loss of plasticity in deep continual learning. *Nature*, 632:768–774.
- Shibhansh Dohare, Ashique Rupam Mahmood, and Richard S. Sutton. 2021. [Continual backprop: Stochastic gradient descent with persistent randomness](#). *ArXiv*, abs/2108.06325.
- Wenyu Du, Shuang Cheng, Tongxu Luo, Zihan Qiu, Zeyu Huang, Ka Chun Cheung, Reynold Cheng, and Jie Fu. 2024. [Unlocking continual learning abilities in language models](#). In *Conference on Empirical Methods in Natural Language Processing*.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony S. Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, and et al. 2024. [The llama 3 herd of models](#). *ArXiv*, abs/2407.21783.
- Yaroslav Ganin, E. Ustinova, Hana Ajakan, Pascal Germain, H. Larochelle, François Laviolette, Mario Marchand, and Victor S. Lempitsky. 2015. [Domain-adversarial training of neural networks](#). In *Journal of machine learning research*.
- Jinghan He, Haiyun Guo, Kuan Zhu, Zihan Zhao, Ming Tang, and Jinqiao Wang. 2024. [SEEKR: Selective attention-guided knowledge retention for continual learning of large language models](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 3254–3266, Miami, Florida, USA. Association for Computational Linguistics.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). In *International Conference on Learning Representations*.

705	Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski,	Colin Raffel, Noam M. Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. <a href="#">Exploring the limits of transfer learning with a unified text-to-text transformer</a> . <i>J. Mach. Learn. Res.</i> , 21:140:1–140:67.	760
706	Bruna Morrone, Quentin de Laroussilhe, Andrea Ges-		761
707	musundo, Mona Attariyan, and Sylvain Gelly. 2019.		762
708	<a href="#">Parameter-efficient transfer learning for nlp</a> . <i>ArXiv</i> ,		763
709	abs/1902.00751.		764
710	Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-	Anastasia Razdaibiedina, Yuning Mao, Rui Hou, Ma-	765
711	Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu	dian Khabsa, Mike Lewis, and Amjad Almahairi.	766
712	Chen. 2022. <a href="#">LoRA: Low-rank adaptation of large</a>	2023. Progressive prompts: Continual learning for	767
713	<a href="#">language models</a> . In <i>International Conference on</i>	language models. In <i>International Conference on</i>	768
714	<i>Learning Representations</i> .	<i>Learning Representations</i> .	769
715	Gangwei Jiang, Caigao JIANG, Zhaoyi Li, Siqiao Xue,	Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten	770
716	JUN ZHOU, Linqi Song, Defu Lian, and Ying Wei.	Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi,	771
717	2025. <a href="#">Unlocking the power of function vectors for</a>	Jingyu Liu, Romain Sauvestre, Tal Remez, et al. 2023.	772
718	<a href="#">characterizing and mitigating catastrophic forgetting</a>	Code llama: Open foundation models for code. <i>arXiv</i>	773
719	<a href="#">in continual instruction tuning</a> . In <i>The Thirteenth In-</i>	<i>preprint arXiv:2308.12950</i> .	774
720	<i>ternational Conference on Learning Representations</i> .		
721	Xiang Lisa Li and Percy Liang. 2021. <a href="#">Prefix-tuning:</a>	Thomas Scialom, Tuhin Chakrabarty, and Smaranda	775
722	<a href="#">Optimizing continuous prompts for generation</a> . In	Muresan. 2022. <a href="#">Fine-tuned language models are</a>	776
723	<i>Proceedings of the 59th Annual Meeting of the Asso-</i>	<a href="#">continual learners</a> . In <i>Proceedings of the 2022 Con-</i>	777
724	<i>ciation for Computational Linguistics and the 11th</i>	<i>ference on Empirical Methods in Natural Language</i>	778
725	<i>International Joint Conference on Natural Language</i>	<i>Processing</i> , pages 6107–6122, Abu Dhabi, United	779
726	<i>Processing (Volume 1: Long Papers)</i> , pages 4582–	Arab Emirates. Association for Computational Lin-	780
727	4597, Online. Association for Computational Lin-	guistics.	781
728	guistics.		
729	Jiaming Liu, Senqiao Yang, Peidong Jia, Ming Lu, Yan-	Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon	782
730	dong Guo, Wei Xue, and Shanghang Zhang. 2023.	Kim. 2017. <a href="#">Continual learning with deep generative</a>	783
731	<a href="#">Vida: Homeostatic visual domain adapter for contin-</a>	<a href="#">replay</a> . In <i>Neural Information Processing Systems</i> .	784
732	<a href="#">ual test time adaptation</a> . <i>ArXiv</i> , abs/2306.04344.		
733	Michael McCloskey and Neal J. Cohen. 1989. <a href="#">Catas-</a>	James Smith, Leonid Karlinsky, Vyshnavi Gutta, Paola	785
734	<a href="#">trophic interference in connectionist networks: The</a>	Cascante-Bonilla, Donghyun Kim, Assaf Arbel, Rameswar Panda, Rogério Schmidt Feris, and Zsolt	786
735	<a href="#">sequential learning problem</a> . volume 24 of <i>Psychol-</i>	Kira. 2022. <a href="#">Coda-prompt: Continual decomposed</a>	787
736	<i>ogy of Learning and Motivation</i> , pages 109–165. Aca-	<a href="#">attention-based prompting for rehearsal-free contin-</a>	788
737	ademic Press.	<a href="#">ual learning</a> . 2023 <i>IEEE/CVF Conference on Com-</i>	789
		<i>puter Vision and Pattern Recognition (CVPR)</i> , pages	790
		11909–11919.	791
			792
738	Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida,	Fan-Keng Sun, Cheng-Hao Ho, and Hung yi Lee. 2019.	793
739	Carroll Wainwright, Pamela Mishkin, Chong Zhang,	<a href="#">Lamol: Language modeling for lifelong language</a>	794
740	Sandhini Agarwal, Katarina Slama, Alex Ray, et al.	<a href="#">learning</a> . In <i>International Conference on Learning</i>	795
741	2022. Training language models to follow instruc-	<i>Representations</i> .	796
742	tions with human feedback. <i>Advances in neural in-</i>		
743	<i>formation processing systems</i> , 35:27730–27744.		
744	Adam Paszke, Sam Gross, Francisco Massa, Adam	Mirac Suzgun, Nathan Scales, Nathanael Scharli, Se-	797
745	Lerer, James Bradbury, Gregory Chanan, Trevor	bastian Gehrmann, Yi Tay, Hyung Won Chung,	798
746	Killeen, Zeming Lin, Natalia Gimelshein, Luca	Aakanksha Chowdhery, Quoc V. Le, Ed Huai hsin	799
747	Antiga, Alban Desmaison, Andreas Köpf, Edward	Chi, Denny Zhou, and Jason Wei. 2022. <a href="#">Challenging</a>	800
748	Yang, Zach DeVito, Martin Raison, Alykhan Tejani,	<a href="#">big-bench tasks and whether chain-of-thought can</a>	801
749	Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Jun-	<a href="#">solve them</a> . In <i>Annual Meeting of the Association for</i>	802
750	jie Bai, and Soumith Chintala. 2019. <i>PyTorch: an</i>	<i>Computational Linguistics</i> .	803
751	<i>imperative style, high-performance deep learning li-</i>		
752	<i>brary</i> . Curran Associates Inc., Red Hook, NY, USA.		
753	Chengwei Qin and Shafiq R. Joty. 2021. <a href="#">Lfmt5: A</a>	Rishabh Tiwari, Krishnateja Killamsetty, Rishabh K.	804
754	<a href="#">unified framework for lifelong few-shot language</a>	Iyer, and Pradeep Shenoy. 2021. <a href="#">Gcr: Gradient core-</a>	805
755	<a href="#">learning based on prompt tuning of t5</a> . <i>ArXiv</i> ,	<a href="#">set based replay buffer selection for continual learn-</a>	806
756	abs/2110.07298.	<a href="#">ing</a> . 2022 <i>IEEE/CVF Conference on Computer Vi-</i>	807
		<i>sion and Pattern Recognition (CVPR)</i> , pages 99–108.	808
757	Guanghui Qin and Jas’ Eisner. 2021. <a href="#">Learning how</a>	Hugo Touvron, Louis Martin, Kevin Stone, Peter Al-	809
758	<a href="#">to ask: Querying lms with mixtures of soft prompts</a> .	bert, Amjad Almahairi, Yasmine Babaei, Nikolay	810
759	<i>ArXiv</i> , abs/2104.06599.	Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti	811
		Bhosale, et al. 2023. Llama 2: Open founda-	812
		tion and fine-tuned chat models. <i>arXiv preprint</i>	813
		<i>arXiv:2307.09288</i> .	814

815	Laurens van der Maaten and Geoffrey E. Hinton. 2008.	<i>Methods in Natural Language Processing: System Demonstrations</i> , pages 38–45, Online. Association for Computational Linguistics.	872
816	<a href="#">Visualizing data using t-sne</a> . <i>Journal of Machine Learning Research</i> , 9:2579–2605.		873
817			874
818	Mingyang Wang, Heike Adel, Lukas Lange, Jannik	Tongtong Wu, Massimo Caccia, Zhuang Li, Yuan-Fang	875
819	Strotgen, and Hinrich Schütze. 2024a. <a href="#">Rehearsal-free modular and compositional continual learning for language models</a> . In <i>North American Chapter of the Association for Computational Linguistics</i> .	Li, Guilin Qi, and Gholamreza Haffari. 2022. <a href="#">Pre-trained language model in continual learning: A comparative study</a> . In <i>International Conference on Learning Representations</i> .	876
820			877
821			878
822			879
823	Qin Wang, Olga Fink, Luc Van Gool, and Dengxin	Tongtong Wu, Linhao Luo, Yuan-Fang Li, Shirui Pan,	880
824	Dai. 2022. <a href="#">Continual test-time domain adaptation</a> . <i>2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)</i> , pages 7191–7201.	Thuy-Trang Vu, and Gholamreza Haffari. 2024. Continual learning for large language models: A survey. <i>arXiv preprint arXiv:2402.01364</i> .	881
825			882
826			883
827	Xiao Wang, Tianze Chen, Qiming Ge, Han Xia, Rong	Yichen Wu, Hongming Piao, Long-Kai Huang, Renzhen	884
828	Bao, Rui Zheng, Qi Zhang, Tao Gui, and Xuanjing	Wang, Wanhua Li, Hanspeter Pfister, Deyu Meng,	885
829	Huang. 2023a. <a href="#">Orthogonal subspace learning for language model continual learning</a> . In <i>Findings of the Association for Computational Linguistics: EMNLP 2023</i> , pages 10658–10671, Singapore. Association for Computational Linguistics.	Kede Ma, and Ying Wei. 2025. <a href="#">S-loRA: Scalable low-rank adaptation for class incremental learning</a> . In <i>The Thirteenth International Conference on Learning Representations</i> .	886
830			887
831			888
832			889
833			
834	Xiao Wang, Tianze Chen, Qiming Ge, Han Xia, Rong	An Yang, Baosong Yang, Binyuan Hui, Bo Zheng,	890
835	Bao, Rui Zheng, Qi Zhang, Tao Gui, and Xuanjing	Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan	891
836	Huang. 2023b. <a href="#">Orthogonal subspace learning for language model continual learning</a> . In <i>Findings of the Association for Computational Linguistics: EMNLP 2023</i> , pages 10658–10671, Singapore. Association for Computational Linguistics.	Li, Dayiheng Liu, Fei Huang, Guanting Dong, Hao-	892
837		ran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian	893
838		Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jin	894
839		Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang	895
840		Lin, Kai Dang, Keming Lu, Ke-Yang Chen, Kexin	896
841	Xiao Wang, Yuansen Zhang, Tianze Chen, Songyang	Yang, Mei Li, Min Xue, Na Ni, Pei Zhang, Peng	897
842	Gao, Senjie Jin, Xianjun Yang, Zhiheng Xi, Rui	Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin,	898
843	Zheng, Yicheng Zou, Tao Gui, et al. 2023c. <a href="#">Trace: A comprehensive benchmark for continual learning in large language models</a> . <i>arXiv preprint arXiv:2310.06762</i> .	Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu,	899
844		Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng,	900
845		Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin	901
846		Wei, Xuancheng Ren, Yang Fan, Yang Yao, Yichang	902
847	Yifan Wang, Yafei Liu, Chufan Shi, Haoling Li, Chen	Zhang, Yunyang Wan, Yunfei Chu, Zeyu Cui, Zhenru	903
848	Chen, Haonan Lu, and Yujiu Yang. 2024b. <a href="#">InsCL: A data-efficient continual learning paradigm for fine-tuning large language models with instructions</a> . In <i>Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)</i> , pages 663–677, Mexico City, Mexico. Association for Computational Linguistics.	Zhang, and Zhi-Wei Fan. 2024. <a href="#">Qwen2 technical report</a> . <i>ArXiv</i> .	904
849			905
850		Yuexiang Zhai, Shengbang Tong, Xiao Li, Mu Cai, Qing	906
851		Qu, Yong Jae Lee, and Yi Ma. 2023. <a href="#">Investigating the catastrophic forgetting in multimodal large language model fine-tuning</a> . In <i>Conference on Parsimony and Learning (Proceedings Track)</i> .	907
852			908
853			909
854		Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015. <a href="#">Character-level convolutional networks for text classification</a> . In <i>Neural Information Processing Systems</i> .	910
855			911
856			912
857	Zifeng Wang, Zizhao Zhang, Chen-Yu Lee, Han Zhang,		913
858	Ruoxi Sun, Xiaoqi Ren, Guolong Su, Vincent Perot,	Weixiang Zhao, Shilong Wang, Yulin Hu, Yanyan Zhao,	914
859	Jennifer G. Dy, and Tomas Pfister. 2021. <a href="#">Learning to prompt for continual learning</a> . <i>2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)</i> , pages 139–149.	Bing Qin, Xuanyu Zhang, Qing Yang, Dongliang Xu,	915
860		and Wanxiang Che. 2024. <a href="#">Sapt: A shared attention framework for parameter-efficient continual learning of large language models</a> . In <i>Annual Meeting of the Association for Computational Linguistics</i> .	916
861			917
862			918
863	Thomas Wolf, Lysandre Debut, Victor Sanh, Julien		919
864	Chaumond, Clement Delangue, Anthony Moi, Pier-	Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan	920
865	ric Cistac, Tim Rault, Remi Louf, Morgan Funtow-	Ye, Zheyang Luo, Zhangchi Feng, and Yongqiang Ma.	921
866	icz, Joe Davison, Sam Shleifer, Patrick von Platen,	2024. <a href="#">Llamafactory: Unified efficient fine-tuning of 100+ language models</a> . In <i>Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)</i> , Bangkok, Thailand. Association for Computational Linguistics.	922
867	Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu,		923
868	Teven Le Scao, Sylvain Gugger, Mariama Drame,		924
869	Quentin Lhoest, and Alexander Rush. 2020. <a href="#">Trans-formers: State-of-the-art natural language processing</a> . In <i>Proceedings of the 2020 Conference on Empirical</i>		925
870			926
871			927



Wanjun Zhong, Ruixiang Cui, Yiduo Guo, Yaobo Liang, Shuai Lu, Yanlin Wang, Amin Saied, Weizhu Chen, and Nan Duan. 2023. Agieval: A human-centric benchmark for evaluating foundation models. *arXiv preprint arXiv:2304.06364*.

Da-Wei Zhou, Qi-Wei Wang, Zhi-Hong Qi, Han-Jia Ye, De-Chuan Zhan, and Ziwei Liu. 2024. *Class-Incremental Learning: A Survey*. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 46(12):9851–9873.

Didi Zhu, Zhongyi Sun, Zexi Li, Tao Shen, Ke Yan, Shouhong Ding, Kun Kuang, and Chao Wu. 2024. *Model tailor: Mitigating catastrophic forgetting in multi-modal large language models*. *ArXiv*, abs/2402.12048.

## A Overall Framework

Drawing from the insight that LoRA (Hu et al., 2022) has exhibited superior performance, we utilize a high- and low-rank framework to ensure stability during continual task adaptation. The overall framework and the details of our method DATA are shown in Figure 3.

LoRA (Low-Rank Adaptation), as proposed by (Hu et al., 2022), postulates that parameter changes ( $\Delta W$ ) during fine-tuning occur within a low-rank subspace. This is particularly applied to the layer weights  $W_0 \in \mathbb{R}^{m \times n}$  of a model  $f_\theta$  for a downstream task. The parameter update is formulated as  $\Delta W = A \times B$ , where  $A \in \mathbb{R}^{m \times r}$  and  $B \in \mathbb{R}^{r \times n}$  are two learnable matrices, and the rank  $r$  is significantly smaller than  $\min\{m, n\}$ .

For a specific layer in the model  $f_\theta$ , the LoRA update is expressed as:

$$h' = W_0 x + \Delta W x = (W_0 + AB)x$$

Here,  $h'$  represents the updated output, and  $x$  is the input to the layer. Importantly, the original weights  $W_0$  are kept frozen during the fine-tuning process, and only the matrices  $A$  and  $B$  are trainable.

## B Experimental Settings

### B.1 Datasets

**Train Tasks.** Tables 4 and 5 provide detailed information on the datasets utilized in our continual learning (CL) experiments. Table 4 presents the 15 datasets included in the Long Sequence Benchmark (Razdaibiedina et al., 2023), while Table 5 outlines the 8 datasets from TRACE (Wang et al., 2023c). Both tables include the corresponding evaluation metrics for each dataset.

**Generalization.** We select the 1) Multitask Language Understanding (MMLU) (Hendrycks et al., 2021), which includes multiple-choice questions across 57 subjects. 2) GSM8K (Cobbe et al., 2021), which is a high-quality linguistically diverse multi-step elementary math reasoning dataset. 3) BIG-Bench Hard (BBH) (Suzgun et al., 2022), which includes 27 challenging tasks spanning arithmetic, symbolic reasoning, and more, derived from BIG-Bench (BB) (bench authors, 2023). Most of the data consists of multiple-choice questions. 4) AGIEval (Zhong et al., 2023), which includes a wide range of high-quality official entrance exams, qualifying exams, and advanced competitions tailored to human participants.

### B.2 Task Sequence Orders

We report task orders used for our CL experiments in Table 6.

### B.3 Implementations

Our implementations are based on huggingface transformers v4.45.2 (Wolf et al., 2020) using PyTorch v2.3.1 (Paszke et al., 2019) and LlamaFactory (Zheng et al., 2024). All unseen tasks generalization evaluation conducted using the OpenCompass toolkit (Contributors, 2023), adopting its default configuration.

For Standard CL Benchmark and Long Sequence Benchmark (Order 1 - Order 6), We trained the models with 1 epoch, a constant learning rate of  $1e-4$ .

For TRACE Order 7 (C-STANCE, FOMC, MeetingBank, Py150, ScienceQA, NumGLUE-cm, NumGLUE-ds, 20Minuten), we trained with 5000 samples with a constant learning rate of  $1e-4$  for 5, 3, 7, 5, 3, 5, 5, 7 epochs respectively.

In a series of performance experiments, we configured various parameters as follows: the LoRA rank was set to 8 refer to Figure 5, and the proportion of past task data mixed in LoRAReplay was set to 2%. For the DATA model, the low-rank configuration was set to 2 and the high-rank configuration to 8. From Figure 7, it can be observed that the performance of 2 and 8 is optimal. Additionally, compared to LoRA, the increase in parameters is limited, with only an additional set of LoRA with a rank of 2, achieving a balance between resources and performance.

In terms of the decomposed component weighting strategy, we used a weight length ( $L_w$ ) of 8. The weight component for each task was set



Dataset Name	Category	Task	Domain	Metric
Yelp	CL Benchmark	Sentiment Analysis	Yelp Reviews	Accuracy
Amazon	CL Benchmark	Sentiment Analysis	Amazon Reviews	Accuracy
DBpedia	CL Benchmark	Topic Classification	Wikipedia	Accuracy
Yahoo	CL Benchmark	Topic Classification	Yahoo Q&A	Accuracy
AG News	CL Benchmark	Topic Classification	News	Accuracy
MNLI	GLUE	Natural Language Inference	Various	Accuracy
QQP	GLUE	Paragraph Detection	Quora	Accuracy
RTE	GLUE	Natural Language Inference	News, Wikipedia	Accuracy
SST-2	GLUE	Sentiment Analysis	Movie Reviews	Accuracy
WiC	SuperGLUE	Word Sense Disambiguation	Lexical Databases	Accuracy
CB	SuperGLUE	Natural Language Inference	Various	Accuracy
COPA	SuperGLUE	Question and Answering	Blogs, Encyclopedia	Accuracy
BoolQA	SuperGLUE	Boolean Question and Answering	Wikipedia	Accuracy
MultiRC	SuperGLUE	Question and Answering	Various	Accuracy
IMDB	SuperGLUE	Sentiment Analysis	Movie Reviews	Accuracy

Table 4: The details of 15 classification datasets in the Long Sequence Benchmark (Razdaibiedina et al., 2023). First five tasks correspond to the standard CL benchmark (Zhang et al., 2015).

to  $\frac{N_{\text{layer}}}{4}$ , leading to a total weight calculation of  $M = \frac{N \times N_{\text{layer}}}{4}$ , where  $N_{\text{layer}}$  is the number of model layers and  $N$  represents the number of tasks. The hyperparameter  $\beta$  was assigned a value of 10. For stochastic recovery, a simple strategy was applied where a small proportion of parameters was recovered every 200 training steps.

#### B.4 More Baselines

**IncLoRA:** Incremental learning of new LoRA parameters for a sequential series of tasks (without adding any regularization or replaying samples from previous tasks).

**LFPT5** (Qin and Joty, 2021): Continuously train a soft prompt that simultaneously learns to solve tasks and generate training samples, which are subsequently used in experience replay.

**ProgPrompt** (Razdaibiedina et al., 2023): Sequentially concatenates previously learned prompts to the current one during training and testing.

**SAPT** (Zhao et al., 2024): In the SAPT method, a Shared Attentive Learning and Selection Module (SALS) is used to guide training samples through optimal PET blocks for task-specific learning, using a unique instance-level attention mechanism. This process ensures efficient continual learning for large language models.

## C Extended Results

### C.1 Fine-grained Results for the Main Experiments

We report the results of each task order on the 3 benchmarks in Table 8. Overall, our proposed DATA demonstrates excellent capabilities in addressing CF and Loss of plasticity.

### C.2 Supplementary Motivation

In our analysis of LoRA’s performance with varying ranks on LLaMA2-7B for Order 1, as shown in Figure 5, we observed that increasing the rank enhances Adaptation Plasticity (AP) for new tasks. However, this improvement comes at the cost of greater forgetting, as it tends to overlook the shared knowledge between tasks. Conversely, lower ranks do not lead to significant changes in forgetting, indicating that they are more effective at capturing the shared knowledge and skills across tasks. This observation suggests the necessity of integrating both high- and low-rank adapters. By doing so, we can better balance the modeling of task-specific and shared knowledge, potentially optimizing both AP and reducing forgetting.

### C.3 Efficiency Analysis

In Table 9, we compare the FLOPs, trainable parameters, storage features, and average predict times of various CL methods. LoRA achieves the highest inference efficiency by bypassing module selection and expansion, allowing the learned

Dataset	Source	Avg len	Metric	Language	#Data
<i>Domain-specific</i>					
ScienceQA	Science	210	Accuracy	English	5,000
FOMC	Finance	51	Accuracy	English	5,000
MeetingBank	Meeting	2853	ROUGE-L	English	5,000
<i>Multi-lingual</i>					
C-STANCE	Social media	127	Accuracy	Chinese	5,000
20Minuten	News	382	SARI	German	5,000
<i>Code Completion</i>					
Py150	Github	422	Edim Similarity	Python	5,000
<i>Mathematical Reasoning</i>					
NumGLUE-cm	Math	32	Accuracy	English	5,000
NumGLUE-ds	Math	21	Accuracy	English	5,000

Table 5: The overview of dataset statistics in TRACE (Wang et al., 2023c). The ‘Source’ indicates the origin of the context. ‘Avg len’ denotes the average length, measured in word count for English, German, and code datasets, and in character count for Chinese datasets. ‘SARI’ is a score that is specific to evaluating simplification tasks.

Benchmark	Order	Task Sequence
Standard CL Benchmark	1	dbpedia → amazon → yahoo → ag
	2	dbpedia → amazon → ag → yahoo
	3	yahoo → amazon → ag → dbpedia
Long Sequence Benchmark	4	mnli → cb → wic → copa → qqp → boolqa → rte → imdb → yelp → amazon → sst-2 → dbpedia → ag → multirc → yahoo
	5	multirc → boolqa → wic → mnli → cb → copa → qqp → rte → imdb → sst-2 → dbpedia → ag → yelp → amazon → yahoo
	6	yelp → amazon → mnli → cb → copa → qqp → rte → imdb → sst-2 → dbpedia → ag → yahoo → multirc → boolqa → wic
TRACE	7	c-stance → fomc → meetingbank → py150 → scienceqa → numglue-cm → numglue-ds → 20minuten

Table 6: Seven distinct orders of task sequences were employed for the experiments in continual learning. Orders 1-3 align with the Standard CL Benchmarks, as adopted in previous studies (Zhang et al., 2015). Orders 4-6 pertain to the Long Sequence Benchmarks, which encompass a total of 15 tasks (Razdaibiedina et al., 2023). Order 7 refers to the TRACE benchmark, specifically designed for LLMs, and comprises eight datasets (Wang et al., 2023c).

Orders	2,16	2,8	4,8	4,16
1	79.4408	81.0921	80.8125	80.8387
2	78.5329	80.3684	80.4507	80.8585
3	78.9934	81.8882	80.2007	80.7500

Table 7: The performance of DATA using LLaMA3.1-8B with different high and low ranks on the standard CL benchmark.

ing both trainable parameters and storage features.

1089

LoRA weights to be directly integrated with the original model weights during testing. Furthermore, our method utilizes only two sets of high- and low-rank adapters with a small number of parameters for decomposed component weighting and eliminates the need to store sample features. This leads to excellent performance when consider-




Figure 5: Performance of different LoRA ranks.

Methods	Standard CL Benchmark (SC)				Long Sequence Benchmark (LS)				TRACE
	Order 1	Order 2	Order 3	Avg	Order 4	Order 5	Order 6	Avg	Order 7
# T5-Large based									
SeqLoRA	72.1	66.8	73.3	70.7	66.4	63.9	19.5	59.9	12.1
LoRAReplay	74.0	73.1	73.0	73.3	74.2	72.7	73.9	73.6	34.0
L2P* (Wang et al., 2021)	60.3	61.7	61.1	60.7	57.5	53.8	56.9	56.1	-
LFPT5* (Qin and Joty, 2021)	67.6	72.6	<b>77.9</b>	72.7	70.4	68.2	69.1	69.2	-
IncLoRA	66.5	64.6	66.1	65.7	59.1	60.7	59.4	59.7	-
ProgPrompt* (Razdaibiedina et al., 2023)	75.2	75.0	75.1	75.1	78.0	<b>77.7</b>	77.9	<b>77.9</b>	-
O-LoRA (Wang et al., 2023b)	73.2	72.4	70.4	72.0	69.9	68.5	65.3	67.9	-
+ MIGU (Du et al., 2024)	73.5	71.4	70.0	71.6	65.4	65.2	65.2	65.3	-
SAPT-LoRA* (Zhao et al., 2024)	-	-	-	-	<b>83.4</b>	-	<b>80.6</b>	-	-
<b>DATA (ours)</b>	73.7	70.5	73.8	72.7	71.5	70.5	68.0	70.0	16.7
+ Replay	<b>77.0</b>	<b>75.6</b>	75.2	<b>75.9</b>	75.6	73.2	74.1	74.3	<b>36.5</b>
# LLaMA2-7B based									
SeqLoRA	73.0	73.2	78.4	74.9	74.7	73.7	72.5	73.7	64.1
LoRAReplay	80.3	80.4	76.7	79.2	80.3	79.5	80.5	80.0	71.9
O-LoRA (Wang et al., 2023b)	76.2	76.3	76.8	76.4	68.5	67.8	67.5	67.9	35.0
<b>DATA (ours)</b>	79.5	79.9	<b>80.0</b>	79.8	76.6	77.0	77.2	76.9	66.0
+ Replay	<b>80.4</b>	<b>81.3</b>	78.4	<b>80.0</b>	<b>83.2</b>	<b>82.5</b>	<b>81.8</b>	<b>81.8</b>	<b>73.1</b>
# LLaMA3.1-8B based									
SeqLoRA	79.9	79.0	80.0	79.6	74.2	73.7	76.5	74.8	65.1
LoRAReplay	80.1	80.6	80.1	80.3	<b>83.2</b>	80.7	<b>82.2</b>	82.0	78.7
O-LoRA (Wang et al., 2023b)	71.8	72.2	72.8	72.3	73.1	69.4	71.6	71.4	36.7
<b>DATA (ours)</b>	<b>81.4</b>	80.7	80.5	<b>80.9</b>	80.7	77.7	81.5	80.0	72.7
+ Replay	80.6	<b>81.0</b>	<b>80.7</b>	80.8	83.1	<b>81.7</b>	81.8	<b>82.2</b>	<b>80.1</b>
# Qwen2.5-7B based									
SeqLoRA	80.0	77.9	78.4	78.8	79.5	79.1	81.1	79.9	65.1
LoRAReplay	<b>80.7</b>	<b>80.6</b>	<b>80.1</b>	<b>80.5</b>	83.3	<b>83.2</b>	82.7	83.1	75.7
<b>DATA (ours)</b>	79.8	79.1	79.4	79.4	79.8	80.2	81.5	80.5	70.4
+ Replay	80.3	<b>80.6</b>	79.9	80.3	<b>83.7</b>	82.9	<b>82.9</b>	<b>83.2</b>	<b>77.3</b>

Table 8: Summary of the results on 3 standard CL benchmarks with T5-Large, LLaMA2-7B, LLaMA3.1-8B and Qwen2.5-7B. Averaged accuracy after training on the last task (**FP**, Sec. 4.1) is reported.

Method	FLOPs ( $10^{16}$ ) ↓	Trainable Parameters (%) ↓	Stored Features (%) ↓	Predict Time (ms) ↓
SeqLoRA	2.6	0.30	0	89
LoRARepaly	4.8	0.30	2%	90
O-LoRA	8.8	0.46	0	196
<b>DATA</b>	4.6	0.38	0	141

Table 9: Comparison of the number of trainable parameters and FLOPs for Order 4 with LLaMA2-7B.