

ANCHORS: Adaptive Normalization for Continual High-Order Retention via Sigmoid-Guided Magnitude Control

Yuxuan Zhang^{§,*} Jiaqi Zhu^{§,*}

[§]South China Normal University

[§]University of Aberdeen

Avenum_Z@outlook.com

* Equal contribution. † Corresponding author.

Abstract

While Parameter-Efficient Fine-Tuning (PEFT) reduces the cost of adapting Large Language Models (LLMs), current methods like LoRA suffer from catastrophic forgetting due to unregulated updates on high-importance parameters. To address this stability-plasticity dilemma, we introduce **ANCHORS: Adaptive Normalization for Continual High-Order Retention via Sigmoid-Guided Magnitude Control**, a novel framework that ensures uninterrupted knowledge retention during aggressive adaptation. ANCHORS integrates CABR Decomposition to preserve informative subspaces and S-MagNorm, a bio-inspired normalization mechanism that utilizes Sigmoid-guided magnitude control to dynamically dampen updates on critical weights. Extensive experiments on reasoning and QA benchmarks demonstrate that ANCHORS consistently outperforms LoRA, DoRA, and CUR-LoRA. It achieves performance gains of up to **4.05%** in MCQ and **2.61%** in QA tasks while retaining **61.12%** of pre-trained knowledge—surpassing Experience Replay and EWC baselines in continuous learning scenarios.

1 Introduction

The rapid advancement of transformer-based generative models has led to a sharp increase in model parameter counts, supporting higher representational capacity and more complex multi-task behaviors. Notable examples in NLP, such as LLaMA 3.1 (Grattafiori et al., 2024) and Qwen 3 (Yang et al., 2025), demonstrate remarkable performance in tasks including text generation, logical reasoning, and multilingual processing. However, their massive scale (e.g., 70B+ parameters) creates severe deployment barriers: full fine-tuning requires over 1.5TB of GPU memory and weeks of computation. Even distilled versions remain prohibitively expensive for individual developers or small research labs.

Parameter-Efficient Fine-Tuning (PEFT) (Xu et al., 2023) methods mitigate these challenges by updating only a small portion of the parameters while freezing the backbone model. Among them, Low-Rank Adaptation (LoRA) (Hu et al., 2021) has emerged as a widely adopted approach, injecting trainable low-rank matrices to reduce memory costs. Despite its efficiency, standard LoRA remains an approximation of full fine-tuning that inherits a critical limitation: catastrophic forgetting (McCloskey and Cohen, 1989). While freezing the base model mitigates forgetting to some extent, the introduction of blind updates in LoRA matrices often reintroduces instability, eroding pre-trained knowledge before optimal convergence is reached.

To address this, recent variants have attempted to target specific improvements: **DoRA** (Liu et al., 2024) and **AdaLoRA** (Zhang et al., 2023) enhance accuracy through weight decomposition; **PiSSA** (Meng et al., 2024) accelerates convergence; and **CUR-LoRA** (Fawi, 2024) and **I-LoRA** (Ren et al., 2024) attempt to reduce information loss. However, these methods remain essentially **open-loop systems**: they determine *how* to update parameters but lack a mechanism to dynamically regulate the *intensity* of updates based on real-time parameter stress. Consequently, they are vulnerable to over-modifying high-importance weights, leaving the stability-plasticity dilemma unresolved.

To bridge this gap, we look to the self-stabilizing nature of biological neural networks, specifically the phenomenon of **Homeostatic Plasticity** (Turrigiano, 2008; Abbott and Nelson, 2000). Unlike the open-loop updates in PEFT, biological circuits employ intrinsic negative feedback mechanisms to normalize synaptic excitability, thereby acting as a safeguard against “runaway” potentiation that could dismantle existing memory structures. Guided by this principle and the insight from Pruning Theory that high-magnitude weights encode the model’s “skeletal” knowledge (Han et al.,

2015; Lv et al., 2024), we envisioned the **closed-loop control system** depicted in **Figure 1**. Instead of blind adaptation, we propose a framework that acts as a **dynamic dampener**: it continuously senses the magnitude of weight changes and applies a non-linear gate to suppress aggressive fluctuations in critical regions. This ensures that the model mimics biological stability—maintaining the structural integrity of vital knowledge while permitting plasticity in less sensitive subspaces.

Materializing this bio-inspired vision, we introduce **ANCHORS**: Adaptive Normalization for Continual High-Order Retention via Sigmoid-Guided Magnitude Control. ANCHORS translates the concept of homeostasis into two complementary innovations: **CABR Decomposition** and **S-MagNorm**. **CABR** ensures a “knowledge-rich” initialization by capturing critical subspaces often discarded in standard methods, while **S-MagNorm** acts as a dynamic regulatory valve, automatically scaling update intensity to balance plasticity and stability without external constraints.

Empirically, ANCHORS achieves up to **4.05%** gains in MCQ accuracy and **2.61%** on QA benchmarks over LoRA, rivaling even advanced variants such as DoRA. For knowledge retention, ANCHORS maintains over **61.12%** of base accuracy, outperforming state-of-the-art methods including Experience Replay (ER) (Fedus et al., 2020), Sequential Learning (SEQ) (McCloskey and Cohen, 1989; Ring, 1992), EWC (Kirkpatrick et al., 2017), I-LoRA, and CUR-LoRA.

Contributions. Our work makes three key contributions:

- We propose **ANCHORS**, a novel PEFT framework that simultaneously improves task performance and knowledge retention. By addressing the open-loop nature of previous methods, it effectively mitigates catastrophic forgetting in large LLMs.
- We introduce **CABR Decomposition** and **S-MagNorm Normalization**, complementary techniques that stabilize training by ensuring high-norm parameters are updated cautiously (stability) while low-norm parameters adapt aggressively (plasticity).
- We empirically validate ANCHORS on multiple LLMs and tasks, showing consistent improvements in MCQ and QA benchmarks while maintaining over 61.12% of pre-trained

knowledge, significantly surpassing existing LoRA variants and continuous learning baselines.

2 Background

2.1 Low-Rank Adaptation.

LoRA is a parameter-efficient fine-tuning method based on the hypothesis that only a small subset of parameters encodes most of the model’s knowledge. It introduces two low-rank matrices, A and B , to approximate updates without modifying the full weight matrix.

Formally, consider a weight matrix $W \in \mathbb{R}^{h \times d}$ within the original LLMs. LoRA introduces two low-rank matrices, $A \in \mathbb{R}^{h \times r}$ and $B \in \mathbb{R}^{r \times d}$, where $r \ll \min(h, d)$. Instead of directly updating the weight matrix, LoRA modifies the model’s forward pass according to the following equation:

$$W' = o + \Delta o = W + AB$$

By freezing the original weight and training only the low-rank matrices, LoRA achieves fine-tuning performance comparable to full adaptation while significantly reducing computational cost.

2.2 CUR-Decomposition and Its Application in CUR-LoRA.

CUR-Decomposition (Hamm and Huang, 2019) is a matrix factorization technique that approximates a matrix using selected subsets of its columns and rows. When integrated into fine-tuning, it can be expressed as:

$$W = CUR(x),$$

where $W \in \mathbb{R}^{h \times d}$ is the original parameter weight matrix, $C \in \mathbb{R}^{h \times r}$ and $R \in \mathbb{R}^{r \times d}$ are column and row subsets of W , and $U \in \mathbb{R}^{r \times r}$ is a trainable matrix, with $r \ll \min(h, d)$.

CUR-LoRA (Fawi, 2024) extends CUR-Decomposition by selecting rows and columns with the lowest norms as C and R . This strategy slows updates on less important parameters, which are more sensitive to fine-tuning, while leaving high-importance parameters largely unaffected. As a result, CUR-LoRA improves knowledge retention and training efficiency by focusing adaptation on the less critical parts of the model.

2.3 Dynamic Network Pruning.

Dynamic Network Pruning (Lin et al., 2020) adapts the network structure at inference time based on input characteristics. Unlike static pruning (Siciliano




Figure 1: S-MagNorm Update method: The figure illustrates the process flow of the S-MagNorm Normalization algorithm. Shows the steps of starting with the fusion of the former weight matrix with the CABR module and moving through the ratio loss matrix calculation, the normalization and sigmoid steps, followed by the final limiting of the ratio loss matrix.

et al., 2024), which permanently removes redundant weights, dynamic pruning selectively activates parameters depending on the input.

Formally, given a weight matrix $\mathbf{W} \in \mathbb{R}^{h \times d}$ and a binary mask matrix $\mathbf{M}(x) \in \{0, 1\}^{h \times d}$, where x represents the input data. Dynamic pruning modifies the forward pass as follows:

$$\mathbf{y} = \sigma((\mathbf{W} \odot \mathbf{M}(x))\mathbf{x})$$

where $\odot$ denotes element-wise multiplication, and σ represents the activation function.

The mask $\mathbf{M}(x)$ can be generated using input-dependent criteria such as weight importance, gradient magnitudes, or gating mechanisms. This ensures that only a subset of parameters is used for each input.

Dynamic pruning offers significant computational efficiency by reducing unnecessary operations, allowing the network to adjust its complexity based on input difficulty. However, generating $\mathbf{M}(x)$ and ensuring efficient real-time pruning remain active research challenges.

3 Methodology

Large language models (LLMs) excel at capturing broad knowledge, yet they remain vulnerable to catastrophic forgetting when fine-tuned on new tasks. Existing parameter-efficient fine-tuning (PEFT) methods, such as LoRA and its variants, partially address this issue by restricting updates to low-rank subspaces. However, these approaches often face two limitations: (1) they discard infor-

mation from important parameters during initialization, and (2) they lack mechanisms to control the dynamics of weight updates, leading to unstable adaptation.

To overcome these challenges, we introduce the **ANCHORS** framework, which is designed to balance *adaptation* to new tasks and *retention* of pre-trained knowledge. ANCHORS integrates three complementary ideas:

1. **CABR-LoRA inverse decomposition**, which carefully initializes low-rank adapters to preserve prior knowledge while creating capacity for new updates;
2. **S-MagNorm normalization**, a Sigmoid-based mechanism that dynamically scales weight updates to stabilize learning;
3. **ANCHORS merge strategies**, which control how fine-tuned adapters are integrated back into the base model, preventing destructive forgetting.

Together, these components form a coherent strategy: CABR-LoRA ensures the right initialization, S-MagNorm regulates the update process, and the merge strategies determine how the model evolves over time. In the following subsections, we explain each component in detail. The complete pseudocode is provided in Appendix D.




Figure 2: Gradient Analysis: The gradient variations during training under Qwen3-1.7B. Showing that LoRA exhibits higher fluctuations, indicating greater parameter reshaping, which may increase the risk of catastrophic forgetting. In contrast, ANCHORS (CABR-LoRA + S-MagNorm) and CABR-LoRA only show lower gradient changes, avoiding excessive drift and suggesting more stable parameter updates, thereby mitigating catastrophic forgetting and improving convergence stability. Additionally, the range and variance comparison highlights the essential role of S-MagNorm. Hyperparameter settings are detailed in Appendix C.

3.1 CABR-LoRA Inverse Decomposition Initialization

Basic CUR-LoRA employs CUR decomposition to reduce the number of trainable parameters by selecting representative rows and columns of the weight matrix. Although efficient, this process often discards parameters that encode critical knowledge, resulting in performance degradation.

CABR-LoRA addresses this limitation by introducing an *inverse low-rank decomposition* of the U matrix, illustrated in Figure 1. Instead of discarding information, CABR-LoRA preserves a higher-dimensional adaptation space while compressing less important rows and columns. Formally, the initialization is defined as:

$$\mathcal{W}_{\text{CABRWeight}} = C \cdot A_{\text{initial}} \cdot B_{\text{zeros}} \cdot R, \quad (1)$$

where C and R represent selected low-importance columns and rows, while A_{initial} and B_{zeros} denote the trainable matrices associated with weights $\mathcal{W}_A$ and $\mathcal{W}_B$.

The key novelty lies in how the U matrix is decomposed. Rather than removing important directions, CABR factorizes U into two complementary matrices:

$$\mathcal{W}_A \in \mathbb{R}^{r \times m}, \quad \mathcal{W}_B \in \mathbb{R}^{m \times r}, \quad \text{with } m > r.$$

Here, $\mathcal{W}_A$ is initialized using Singular Value Decomposition (SVD) to capture prior knowledge, while $\mathcal{W}_B$ is initialized with zeros, ensuring that adaptation starts smoothly without disturbing high-norm parameters.

The initialization process can be formally expressed as:

$$\mathcal{W}_{\text{base}} = \mathcal{U}_{\text{SVD}} \mathcal{S}_{\text{SVD}} \mathcal{V}_{\text{SVD}}^T, \quad (2)$$

where $\mathcal{U}_{\text{SVD}}$ and $\mathcal{V}_{\text{SVD}}$ are the left and right singular vectors of the base weight matrix, respectively, and $\mathcal{S}_{\text{SVD}}$ contains the corresponding singular values.

$$\mathcal{W}_A = \mathcal{U}_{\text{SVD}}^{(r,:)} \mathcal{S}_{\text{SVD}} \mathcal{V}_{\text{SVD}}^{(:,m)T}, \quad (3)$$

$$\mathcal{W}_B = \mathbf{0}_{m \times r}. \quad (4)$$

Through this design, CABR-LoRA achieves a balance: it retains critical knowledge in $\mathcal{W}_A$ while allowing $\mathcal{W}_B$ to absorb new task-specific information during fine-tuning.

3.2 S-MagNorm Normalization Algorithm

While CABR-LoRA creates capacity for adaptation, it does not fully prevent catastrophic forgetting. In particular, unconstrained updates to weight matrices can distort the relative importance of parameters. To address this, we propose the **S-MagNorm normalization algorithm**, which dynamically scales weight updates to ensure stability in Figure 2.

The central idea is to regulate parameter changes based on their relative magnitude, following the intuition that high-norm parameters encode essential knowledge and should be updated cautiously. Conversely, low-norm parameters can afford larger updates, making them natural candidates for task-specific adaptation.

Formally, we define the intermediate merged weight matrix $\mathcal{W}_{\text{ANCHORSMerged}}$, which combines the contribution of the CABR component with the weights of the base model $\mathcal{W}_{\text{base}}$. This serves as the foundation for subsequent normalization and scaling:

$$\mathcal{W}_{\text{ANCHORSMerged}} = \text{CABR} + \mathcal{W}_{\text{base}}. \quad (5)$$

We then compute the relative magnitude change between the merged and base weights:

$$\mathcal{M}_{\text{Mag}} = \left| \frac{\mathcal{W}_{\text{ANCHORSMerged}}}{\mathcal{W}_{\text{base}} + \epsilon} \right|, \quad (6)$$

where ϵ ensures numerical stability. To maintain a balanced update **Scale**, the values are normalized as:

$$\mathcal{M}_{\text{norm}} = \left[\frac{\mathcal{M}_{\text{Mag}}}{\max(\mathcal{M}_{\text{Mag}}) + \epsilon} - 0.5 \right] \cdot \text{Scale}. \quad (7)$$

Next, we apply a Sigmoid transformation σ to produce the rescaling matrix:

$$\mathcal{M}_{\text{Res}} = 2 - \sigma(\mathcal{M}_{\text{norm}}). \quad (8)$$

Finally, the updated weights are obtained as:

$$\mathcal{W}_{\text{updated}} = \frac{\mathcal{W}_{\text{ANCHORSMerged}}}{\mathcal{M}_{\text{Res}}}. \quad (9)$$

This dynamic process ensures that high-norm parameters are updated conservatively, while low-norm ones are adjusted more aggressively. As a result, S-MagNorm provides controlled pruning that both stabilizes fine-tuning and mitigates catastrophic forgetting in Figure 3.

3.3 Merge with Basic Model

Even with CABR and S-MagNorm, relying solely on adapters risks performance degradation, as base parameters are crucial for long-term retention. Freezing base weights can create a destructive cycle (as observed in our ablation study Exp 4.3): initially less important parameters accumulate updates, eventually growing large enough to accelerate the fine-tuning of critical parameters, thereby eroding retained knowledge and diminishing the model’s memory of prior tasks.

To break this cycle, ANCHORS introduces two complementary merging strategies:




Figure 3: Comparison of Weight SVD Norm across different methods: The "Basic" method represents the original LLM weights without fine-tuning. All experiments are conducted on the Qwen3-1.7B backbone. After 2000 training steps, DoRA shows a 1.309 decrease in norm, indicating some loss of information. ANCHORS preserves more knowledge with a minor 0.302 change, while LoRA shows a significant 3.083 increase, reflecting greater weight modification. Hyperparameter settings are detailed in Appendix C.

M1: Direct Weight Merge. Here, fine-tuned CABR-LoRA weights are directly merged with the base model:

$$\mathcal{W}_{\text{updated}} = \underbrace{\text{CABR-LoRA}}_{\text{CABR-LoRA}} + \mathcal{W}_{\text{base}}. \quad (10)$$

Here, A_{former} refers to the previously learned A matrix, and B_{train} is the newly trained B matrix. This approach enables stronger fine-tuning by directly modifying the base weights.

M2: Frozen Base Merge. Alternatively, the base weights remain frozen, while two ANCHORS instances manage updates: one accumulates historical knowledge and the other applies fresh updates:

$$\mathcal{W}_{\text{merged}} = \underbrace{\text{CABR-LoRA}}_{\text{Previous CABR-LoRA}} + \underbrace{\text{CABR-LoRA}}_{\text{New CABR-LoRA}} + \mathcal{W}_{\text{base}}. \quad (11)$$

Here, A_{frozen} denotes the frozen A matrix preserving historical knowledge, B_{accum} accumulates past updates, A_{train} is the freshly trained A matrix, and B_{reset} is reset for the current update.

The update rules are:

$$A_{\text{former}}, A_{\text{frozen}} \leftarrow A_{\text{train}}, \quad (12)$$

$$B_{\text{accum}} \leftarrow B_{\text{accum}} + B_{\text{train}}, \quad (13)$$

$$B_{\text{train}} \leftarrow \mathbf{0}. \quad (14)$$

M2 emphasizes retention by isolating base weights, progressively accumulating updates, and cyclically resetting parameters to avoid saturation.

Comparison. M1 excels in performance-critical domains (e.g., mathematics, programming), where aggressive fine-tuning is beneficial. M2, by contrast, is better suited for knowledge-sensitive domains (e.g., medicine, economics), where retention is paramount. Ablation studies (Table 4) confirm these trade-offs, and in practice, a hybrid of M1 and M2 often achieves the best balance between adaptability and knowledge preservation.

In summary, ANCHORS integrates careful initialization through CABR-LoRA, stable adaptation enabled by S-MagNorm, and knowledge-preserving merging strategies M1 and M2. Together, these components provide a principled framework for mitigating catastrophic forgetting while enabling effective task adaptation.

4 Experiments

4.1 Experiment Design

To comprehensively evaluate **ANCHORS**, we design three main experiments and one ablation study, following consistent setups across five backbone LLMs (Gemma2-2B (Team et al., 2024), Qwen3-1.7B (Yang et al., 2025), Qwen3-8B (Yang et al., 2025), LLaMA3.1-8B (Grattafiori et al., 2024), LLaMA3.2-3B (Grattafiori et al., 2024)). All experiments adopt the same training protocol unless otherwise stated, and hyperparameters are listed in Appendix C. Unless explicitly mentioned, ANCHORS uses the M1 merge method.

Experiment 1 (PEFT Comparison). This experiment compares LoRA-based PEFT variants, including DoRA, CUR-LoRA, and I-LoRA, using 4 multiple-choice (MCQ) tasks and 5 question-answering (QA) tasks. MCQ tasks are evaluated by **Accuracy (AC)**, while QA tasks are measured by **BLEU**, **ROUGE-1**, and **ROUGE-L**. The goal is to test whether ANCHORS can maintain or exceed LoRA’s fine-tuning capability while introducing better stability.

Experiment 2 (Catastrophic Forgetting with PEFT Comparison). To analyze whether ANCHORS truly provides a substantial solution to catastrophic forgetting, inspired by Kirkpatrick et al. (2017), we conduct a stress test with a learning rate of **1e-2** to further assess knowledge retention under extreme training conditions.

Experiment 3 (Catastrophic Forgetting with Retention Methods). We benchmark ANCHORS

Task	LoRA	DoRA	CUR-LoRA	I-LoRA	ANCHORS
MCQ Avg.					
AbsNarr	53.43	53.90	42.01	50.09	57.37 ^{+3.94}
ObjCount	56.67	57.76	46.69	53.67	60.18 ^{+3.51}
PlayDiag	53.74	55.33	50.33	50.39	57.89 ^{+4.15}
ColorReason	52.31	54.04	46.45	50.42	56.90 ^{+4.56}
QA Avg.					
GSM8K	75.51	78.28	72.97	74.88	77.84 ^{+2.33}
NewsDE	62.70	63.21	57.92	61.25	65.48 ^{+2.78}
NewsIT	65.23	65.59	60.18	64.50	69.18 ^{+3.95}
NewsES	64.80	66.11	59.51	63.68	68.37 ^{+3.57}
ALPACA	60.95	60.71	58.88	60.33	61.35 ^{+0.40}

Table 1: Performance comparison across 4 MCQ and 5 QA tasks averaged over 5 LLMs. Only ANCHORS shows relative increase vs. LoRA for clarity.

against classical continual-learning methods including EWC, ER, and SEQ, under a sequential fine-tuning regime across 16 tasks (each trained for 2K steps, totaling 32K steps) using a smaller learning rate of **2e-4**. Each PEFT variant merges its fine-tuned adapters back to the frozen base weights after each task (as in CUR-LoRA, I-LoRA, and ANCHORS). We then evaluate how much prior knowledge is retained using the fused **F-task set** (ARC-Challenge (Clark et al., 2018), GSM8K (Cobbe et al., 2021), Multipl-E (Cassano et al., 2022)) after fine-tuning each PEFT variant.

Ablation Study. We further investigate the role of merge frequency and merge strategy by testing different update cycles (1-step vs. 200-step) and comparing two merge formulations (M1 vs. M2). These studies validate the update cycle hypothesis and demonstrate how the merge depth impacts both fine-tuning and retention behavior.

4.2 Main Results

Performance on PEFT Comparison. Table 1 presents the results of Experiment 1 across 4 MCQ and 5 QA tasks. ANCHORS achieves consistent and significant gains compared to all PEFT variants, outperforming LoRA by an average of **4.05%** on MCQ tasks and **2.61%** on QA tasks. Notably, on the *ColorReason* task, ANCHORS demonstrates nearly a **5% improvement**, highlighting its exceptional fine-tuning ability even without modifying all model parameters. While DoRA offers slight improvement on specific tasks such as *GSM8K* (+0.44%), other variants designed for catastrophic forgetting (CUR-LoRA, I-LoRA) tend to trade performance for stability, resulting in overall lower fine-tuning accuracy.

Catastrophic Forgetting with PEFT Comparison. Table 2 shows the results of Experiment




Figure 4: The performance of ANCHORS compared to 16 tasks baselines sequentially trained using LLaMA3.1-8B (left) and Qwen3-8B (right) backbones, tested under F-task with learning rate $2e-4$. Detailed results are in Appendix F

Task	LoRA	DoRA	CUR-LoRA	I-LoRA	ANCHORS
AbsNarr	11.76	23.87	10.75	12.52	75.54
DisflQA	03.15	08.10	06.70	15.89	64.70
LangID	03.18	04.84	04.20	04.05	61.12
LogDeduc	06.37	07.89	08.77	06.93	64.59
ObjCount	06.06	07.90	10.24	08.68	63.90
PlayDiag	04.75	07.91	06.33	06.72	64.61
ColorReason	05.31	10.94	09.16	09.61	66.80
TrackObj	07.05	10.68	09.70	08.79	68.66

Table 2: Knowledge retention comparison (Experiment 2) under high learning rate $1e-2$. ANCHORS retains 61% base knowledge while other variants collapse below 24%.

2 with a high learning rate ($1e-2$). As previous studies suggest, a high learning rate amplifies catastrophic forgetting (Kirkpatrick et al., 2017). Consistent with these findings, most PEFT variants lose nearly all previously acquired knowledge. Even variants tailored for retention, such as CUR-LoRA and I-LoRA, only maintain around 10% to 20% of the base knowledge. In contrast, ANCHORS preserves an average of **66.24%** knowledge under the same stress conditions — a clear indication that the Sigmoid-based Magnitude Normalization effectively stabilizes updates without constraining learning flexibility.

Catastrophic Forgetting with Retention Methods. Experiment 3 extends the analysis to continual-learning baselines under sequential training. Figure 4 shows that while ER and I-LoRA partially recover prior knowledge through replay and re-integration, ANCHORS consistently maintains higher retention with comparable or better task performance. This balance between stability

Task	ANCHORS-1	ANCHORS-200	ANCHORS-1	ANCHORS-200
	Qwen-3 (8B)		LLaMA-3.1 (8B)	
AbsNarr	78.36	11.69	75.89	00.23
DisflQA	66.27	09.60	66.58	00.22
LangID	58.82	08.97	67.39	00.19
LogDeduc	65.36	09.36	69.95	00.21
ObjCount	60.71	09.85	70.95	00.22
PlayDiag	61.21	08.54	68.81	00.34
ColorReason	68.21	00.26	70.87	00.27
TrackObj	70.86	10.29	70.74	00.31

Table 3: Ablation on merge frequency (1-step vs 200-step) under high learning rate $1e-2$. Frequent merges (1-step) yield more stable retention across tasks.

and plasticity demonstrates that ANCHORS not only mitigates forgetting but also supports continual optimization across diverse domains.

4.3 Ablation Studies

Update Cycle Hypothesis. To validate the proposed dual-cycle hypothesis, we compare 1-step and 200-step merge frequencies. As shown in Table 3, the 1-step merge strategy consistently yields superior performance. We hypothesize that the degradation of 200-step merges occurs because the system begins entering a second update cycle without synchronization, causing drift accumulation. This finding underscores the critical role of frequent merges in maintaining stable learning dynamics.

Merge Method Comparison (M1 vs. M2). We further compare the two merge strategies proposed in Section 3. As summarized in Table 4, M1 (deep merge) modifies base weights more aggressively and thus excels in fine-tuning performance, outperforming M2 by **0.29%–2.63%** on accuracy metrics. In contrast, M2 (shallow merge) shows superior

Merge Under EXP1:	MergeMethod-1	MergeMethod-2
AbsNarr	91.14	88.63
ObjCount	95.28	92.65
PlayDiag	94.66	94.37
ColorReason	91.66	90.60
Merge Under EXP2:	MergeMethod-1	MergeMethod-2
AbsNarr	66.31 / 69.64 / 63.67	69.62 / 71.65 / 65.89
ObjCount	66.29 / 69.46 / 63.33	70.42 / 72.90 / 67.94
PlayDiag	66.04 / 68.74 / 62.57	70.07 / 71.82 / 67.00
ColorReason	64.47 / 67.24 / 61.20	69.23 / 71.05 / 65.90

Table 4: Comparison of fine-tuning efficiency and knowledge retention using MergeMethod-1 and MergeMethod-2 under EXP1 and EXP2, under the standard Qwen3-1.7B experimental setting. Values in EXP2 rows show BLEU / ROUGE-1 / ROUGE-L.

retention — improving BLEU by **4.76%**, ROUGE-1 by **3.81%**, and ROUGE-L by **4.70%** over M1. These results align with our theoretical expectations and demonstrate the flexibility of ANCHORS in balancing adaptation depth and stability.

4.4 Summary of Findings

Across all experiments, ANCHORS demonstrates:

- **Superior fine-tuning performance** compared to LoRA-based PEFTs (Table 1);
- **Strong resilience to catastrophic forgetting**, even under extreme learning conditions (Table 2);
- **Balanced optimization capability** between performance and retention through its merge strategy (Tables 3 and 4).

These results confirm that ANCHORS provides a practical and stable fine-tuning paradigm for large language models, achieving both efficiency and continual-learning robustness.

5 Discussion

5.1 Why Do We Update the Basic Parameters or Track ANCHORS Parameters But Never Change C and R ?

We assume parameters with higher norms carry more important information, while lower-norm ones are less critical for generalization. The column and row selectors C and R initialized in ANCHORS mark these less important positions. Our goal is to preserve essential knowledge encoded in high-norm regions, so we must avoid altering C and R , which define where updates occur.

Keeping C and R unchanged ensures critical knowledge remains stable while adaptation occurs only in controlled, less sensitive subspaces. If they changed during training, the system could lose its

reference to which weights are safe to modify, introducing instability and potentially accelerating forgetting. Additionally, fixing them reduces re-computation overhead and maintains consistent learning dynamics across tasks.

This design thus provides both stability and efficiency: it protects important parameter regions while focusing fine-tuning on adaptable, low-norm components.

5.2 What Are the Other Benefits of ANCHORS Compared to Other Catastrophic Forgetting Methods?

Having demonstrated ANCHORS’s performance experimentally, we further highlight its advantages over methods like EWC or Experience Replay, focusing on efficiency and simplicity. ANCHORS maintains strong performance while remaining lightweight and easily integrable into existing PEFT pipelines without architectural changes.

5.2.1 Smaller Weight

Traditional methods like EWC rely on the Fisher matrix to detect parameter changes, requiring storage of previous parameters and increasing computational and space costs. In contrast, ANCHORS achieves comparable or superior performance with fewer parameters than standard LoRA C. This efficiency enables better scalability across tasks and models while maintaining stable optimization.

5.2.2 No Need for Former Datasets and Calculation

Experience Replay prevents forgetting by storing and reusing past datasets, incurring significant time and computational overhead. ANCHORS requires no historical data yet performs effectively, eliminating dataset storage and replay computations. This makes it a more resource-efficient and practical solution for continual fine-tuning.

6 Conclusion

We propose **ANCHORS**, a bio-inspired PEFT framework integrating S-MagNorm and CABR decomposition to mitigate catastrophic forgetting. By regulating update intensity, our Sigmoid-guided control dynamically balances plasticity and stability, ensuring the preservation of critical knowledge. Extensive experiments across 16 datasets and 5 LLMs demonstrate that ANCHORS significantly outperforms LoRA, enabling scalable and sustainable continuous deployment.

7 Limitations

- **Lack of experiments on 70B+ LLMs:** Due to device limitations, further research on larger models (70B+ parameters) could not be conducted. Future work will explore scaling ANCHORS to more massive LLMs.
- **Computational Overhead:** S-MagNorm introduces additional matrix operations during training, resulting in higher per-step computational overhead compared to LoRA. Future work could focus on optimizing the normalization layer for real-time applications.

8 Acknowledgement

We would like to thank **Professor Biqing Zeng**, Dean of the School of Data Science and Artificial Intelligence at the Aberdeen Institute of South China Normal University, for the inspiration we drew from his lectures, which ignited our passion for research. We are also deeply grateful for his constant moral support and encouragement throughout our work. His kindness and belief in us have meant more than we can express.

9 Use of AI Assistants

The authors conducted all research and writing independently. AI tools were employed only for language refinement, and all outputs were manually checked for accuracy and originality.

References

- L. F. Abbott and Sacha B. Nelson. 2000. [Synaptic plasticity: taming the beast](#). *Nature Neuroscience*, 3(11):1178–1183.
- Federico Cassano, John Gouwar, Daniel Nguyen, Sydney Nguyen, Luna Phipps-Costin, Donald Pinckney, Ming-Ho Yee, Yangtian Zi, Carolyn Jane Anderson, Molly Q Feldman, Arjun Guha, Michael Greenberg, and Abhinav Jangda. 2022. [Multipl-e: A scalable and extensible approach to benchmarking neural code generation](#). *Preprint*, arXiv:2208.08227.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. [Think you have solved question answering? try arc, the ai2 reasoning challenge](#). *Preprint*, arXiv:1803.05457.
- 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](#). *Preprint*, arXiv:2110.14168.
- Muhammad Fawi. 2024. [Curlora: Stable llm continual fine-tuning and catastrophic forgetting mitigation](#).
- William Fedus, Prajit Ramachandran, Rishabh Agarwal, Yoshua Bengio, Hugo Larochelle, Mark Rowland, and Will Dabney. 2020. [Revisiting fundamentals of experience replay](#). *Preprint*, arXiv:2007.06700.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Keaton Hamm and Longxiu Huang. 2019. [Cur decompositions, approximations, and perturbations](#). *Preprint*, arXiv:1903.09698.
- Song Han, Jeff Pool, John Tran, and William J. Dally. 2015. [Learning both weights and connections for efficient neural networks](#). *Preprint*, arXiv:1506.02626.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. [Lora: Low-rank adaptation of large language models](#). *Preprint*, arXiv:2106.09685.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. 2017. [Overcoming catastrophic forgetting in neural networks](#). *Proceedings of the National Academy of Sciences*, 114(13):3521–3526.
- Tao Lin, Sebastian U. Stich, Luis Barba, Daniil Dmitriev, and Martin Jaggi. 2020. [Dynamic model pruning with feedback](#). *Preprint*, arXiv:2006.07253.
- Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Frank Wang, Kwang-Ting Cheng, and Min-Hung Chen. 2024. [Dora: Weight-decomposed low-rank adaptation](#). *Preprint*, arXiv:2402.09353.
- Bo Lv, Quan Zhou, Xuanang Ding, Yan Wang, and Zeming Ma. 2024. [Kvpruner: Structural pruning for faster and memory-efficient large language models](#). *Preprint*, arXiv:2409.11057.
- Michael McCloskey and Neal J. Cohen. 1989. [Catastrophic interference in connectionist networks: The sequential learning problem](#). In Gordon H. Bower, editor, *Psychology of Learning and Motivation*, volume 24 of *Psychology of Learning and Motivation*, pages 109–165. Academic Press.
- Fanxu Meng, Zhaohui Wang, and Muhan Zhang. 2024. [Pissa: Principal singular values and singular vectors adaptation of large language models](#). *Preprint*, arXiv:2404.02948.

Weijieying Ren, Xinlong Li, Lei Wang, Tianxiang Zhao, and Wei Qin. 2024. [Analyzing and reducing catastrophic forgetting in parameter efficient tuning](#). *Preprint*, arXiv:2402.18865.

Mark Ring. 1992. [Learning sequential tasks by incrementally adding higher orders](#). In *Advances in Neural Information Processing Systems*, volume 5. Morgan-Kaufmann.

Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). *Preprint*, arXiv:1704.04368.

Federico Siciliano, Francesca Pezzuti, Nicola Tonelotto, and Fabrizio Silvestri. 2024. [Static pruning in dense retrieval using matrix decomposition](#). *Preprint*, arXiv:2412.09983.

Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R. Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, Agnieszka Kluska, Aitor Lewkowycz, Akshat Agarwal, Alethea Power, Alex Ray, Alex Warstadt, Alexander W. Kocurek, Ali Safaya, Ali Tazarv, and 432 others. 2023. [Beyond the imitation game: Quantifying and extrapolating the capabilities of language models](#). *Preprint*, arXiv:2206.04615.

Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. [CommonsenseQA: A question answering challenge targeting commonsense knowledge](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.

Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, Johan Ferret, Peter Liu, Pouya Tafti, Abe Friesen, Michelle Casbon, Sabela Ramos, Ravin Kumar, Charline Le Lan, Sammy Jerome, and 179 others. 2024. [Gemma 2: Improving open language models at a practical size](#). *Preprint*, arXiv:2408.00118.

Gina G. Turrigiano. 2008. [The self-tuning neuron: Synaptic scaling of excitatory synapses](#). *Cell*, 135(3):422–435.

Jingwei Xu, Junyu Lai, and Yunpeng Huang. 2024. [Me-teora: Multiple-tasks embedded lora for large language models](#). *Preprint*, arXiv:2405.13053.

Lingling Xu, Haoran Xie, Si-Zhao Joe Qin, Xiaohui Tao, and Fu Lee Wang. 2023. [Parameter-efficient fine-tuning methods for pretrained language models: A critical review and assessment](#). *Preprint*, arXiv:2312.12148.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao,

Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.

Qingru Zhang, Minshuo Chen, Alexander Bukharin, Nikos Karampatziakis, Pengcheng He, Yu Cheng, Weizhu Chen, and Tuo Zhao. 2023. [Adalora: Adaptive budget allocation for parameter-efficient fine-tuning](#). *Preprint*, arXiv:2303.10512.

Algorithm 1: ANCHORS Methodology

Input: Base weights $\mathcal{W}_{\text{base}}$, hyperparameters r, m, ϵ , Scale

Output: Updated weights $\mathcal{W}_{\text{updated}}$

Step 1: CABR-LoRA Decomposition

1. *Initialize CABR matrices:*

$\mathcal{W}_{\text{frozen}} = \mathcal{U}_{\text{SVD}} \mathcal{S}_{\text{SVD}} \mathcal{V}_{\text{SVD}}^{\top}$ // SVD of base weights

$\mathcal{W}_A \leftarrow \mathcal{U}_{\text{SVD}}^{(r,:)} \mathcal{S}_{\text{SVD}} \mathcal{V}_{\text{SVD}}^{(:,m)\top}$ // Truncated SVD for A

$\mathcal{W}_B \leftarrow \text{Zeros}(m \times r)$ // Zero-initialized B

$\mathcal{W}_{\text{CABR}} \leftarrow C \cdot \mathcal{W}_A \cdot \mathcal{W}_B \cdot R$ // CABR reconstruction

Step 2: S-MagNorm Normalization

2. *Merge weights:*

$\mathcal{W}_{\text{ANCHORSmerged}} \leftarrow \mathcal{W}_{\text{CABR}} + \mathcal{W}_{\text{base}}$

3. *Compute magnitude ratio:*

$\mathcal{M}_{\text{Mag}} \leftarrow \left| \frac{\mathcal{W}_{\text{ANCHORSmerged}}}{\mathcal{W}_{\text{base}} + \epsilon} \right|$ // Avoid division by zero

4. *Normalize and scale:*

$\mathcal{M}_{\text{norm}} \leftarrow \left[\frac{\mathcal{M}_{\text{Mag}}}{\max(\mathcal{M}_{\text{Mag}}) + \epsilon} - 0.5 \right] \cdot \text{Scale}$

5. *Apply Sigmoid constraint:*

$\mathcal{M}_{\text{Res}} \leftarrow 2 - \sigma(\mathcal{M}_{\text{norm}})$ // σ : Sigmoid function

6. *Adjust updates:*

$\mathcal{W}_{\text{updated}} \leftarrow \mathcal{W}_{\text{ANCHORSmerged}} / \mathcal{M}_{\text{Res}}$
// Element-wise division

Step 3: Merge Strategy (M1/M2)

7. *if M1 (Deep Merge) then*

$\mathcal{W}_{\text{updated}} \leftarrow C A_{\text{former}} B_{\text{train}} R + \mathcal{W}_{\text{base}}$
// Update base weights

end

else if M2 (Shallow Merge) then

$\mathcal{W}_{\text{mergedBase}} \leftarrow C A_{\text{frozen}} B_{\text{accum}} R + \mathcal{W}_{\text{base}}$
 $\mathcal{W}_{\text{merged}} \leftarrow \mathcal{W}_{\text{mergedBase}} + C A_{\text{train}} B_{\text{reset}} R$
// Preserve $\mathcal{W}_{\text{basic}}$

end

return $\mathcal{W}_{\text{updated}}$

A Broader Impacts

- **Efficiency and Accessibility:** By eliminating the need for prior datasets and reducing training time for knowledge retention, ANCHORS lowers the computational barrier for fine-tuning LLMs, making advanced AI adaptation more accessible to researchers with limited resources.

- **Environmental Sustainability:** The reduced computational overhead compared to Experience Replay could decrease energy consumption in large-scale LLM deployments.
- **Ethical Risks:** While ANCHORS mitigates catastrophic forgetting, its ability to retain sensitive information from pre-training data (e.g., biased or private content) requires careful auditing before deployment. Future work should investigate selective forgetting mechanisms to align with ethical AI principles.
- **Generalization Beyond NLP:** The S-MagNorm algorithm, though initially designed for PEFT, could be extended to other continual learning scenarios (e.g., robotics, healthcare), potentially improving adaptive systems in dynamic environments.

B Details about 18 Tasks and Datasets

Table 5 includes detailed descriptions of each dataset’s name, keywords, main content and corresponding evaluation metrics. These 18 tasks, most of which are sourced from BIG-bench (Srivastava et al., 2023), include different topics, such as mathematical reasoning (Cobbe et al., 2021), logical deduction (Clark et al., 2018), abstractive summarization (See et al., 2017), programming challenges (Cassano et al., 2022), translation questions, and commonsense understanding (Talmor et al., 2019).

C Experimental Hyper-parameters on Three Main Experiments

All of our main experiments were conducted on a single machine equipped with two customized NVIDIA 2080 Ti GPUs (22GB each)¹ and two NVIDIA 3090 GPUs (24GB each).

For better clarity in gradient and knowledge retention visualizations (Figures 2 and 4), a small smoothing factor was applied: 5 for Figure 2 and 15 for Figure 4.

Tables 6, 7, 8, and 9 list all hyper-parameters for different LoRA baselines and our proposed ANCHORS across experiments.

¹These GPUs use standard architectures with extended memory for the 2080 Ti. The compute cores and driver stack remain unchanged. This setup demonstrates the efficiency of our method under constrained, non-datacenter hardware.




Figure 5: Pareto analysis of CABR decomposition under different (r, m) configurations. Each point represents a distinct decomposition setting evaluated on downstream adaptation and knowledge retention tasks. The results show that different decomposition dimensions exhibit different adaptation-retention tradeoff behaviors.

D Pseudo code deduction

Algorithm 1 shows whole deduction of ANCHORS:

E Further Discussion on the Effect of CABR Decomposition Dimensions

To further analyze the impact of CABR decomposition, we systematically evaluate different (r, m) configurations under the standard Qwen3-1.7B setting. This ablation study aims to elucidate how the rank and intermediate dimension jointly govern model behavior during continual learning.

Figure 5 illustrates the tradeoff between downstream adaptation and knowledge retention across varying decomposition dimensions. The results indicate that increasing decomposition capacity does not consistently improve performance; rather, different (r, m) settings exhibit distinctly non-monotonic adaptation-retention behaviors, suggesting the existence of an optimal operating region rather than a simple scaling law.

Specifically, smaller intermediate dimensions tend to limit downstream adaptation due to reduced representational flexibility, creating an information bottleneck that constrains the model’s ability to encode new task-specific patterns. Conversely, excessively large decomposition dimensions may reduce retention stability by introducing redundant parameters that are prone to overfitting on limited downstream data, thereby accelerating catastrophic forgetting of pre-trained knowledge.

These observations suggest that CABR decomposition dimensions critically influence the balance between adaptation flexibility and knowledge preservation. Consequently, selecting appropriate (r, m) values is essential for achieving robust continual learning performance, highlighting the necessity of dimension-aware configuration strategies in practical deployment.

F Experimental Results on All Datasets

Table 10, 11, 12, 13 and 14 present comprehensive results across different LoRA baselines and ANCHORS using Gemma2-2b, Qwen3-1.7B, Qwen3-8B, LLaMA3.2-3B and LLaMA3.1-8B backbones. To further discover the model’s ability to learn the specific dataset, we compute MCQ AC scores as the average of ROUGE-1 and BLEU-1 evaluated on fixed-length outputs, where ROUGE-1 reflects correctness of the answer and BLEU-1 further accounts for output length; in other words, $AC = (ROUGE-1 + BLEU-1)/2$, and the true answer accuracy considering only correctness will be higher. For clarity of visualization, the gradient trends in Figure 2 and the learning curves in Figure 4 were lightly smoothed to reduce noise while preserving overall trends.

Task Name	Keywords	Description	Evaluation Metrics	Met-
abstract_narrative_understanding (AbsNarr)	narrative understanding, multiple choice	Given a narrative, choose the most related proverb.	Accuracy	
cnn_dailymail (CNNDM)	summarization	Given news articles, write the summarization.	ROUGE	
disfl_qa (DisflQA)	contextual question-answering, reading comprehension	Pick the correct answer span from the context given the disfluent question.	Accuracy	
gsm8k (GSM8K)	mathematics	Solve the grade school math word problems.	Accuracy	
language_identification (LangID)	multilingual, multiple choice	Given a sentence, select the correct language.	Accuracy	
linguistics_puzzles (LingPuzz)	logical reasoning, linguistics	Solve Rosetta Stone-style linguistics puzzles.	BLEU, ROUGE	
logical_deduction (LogDeduc)	logical reasoning, multiple choice	Deduce the order of a sequence of objects.	Accuracy	
news_commentary_de (NewsDE)	multilingual, translation	Translate German sentences into English.	BLEU	
news_commentary_es (NewsES)	multilingual, translation	Translate Spanish sentences into English.	BLEU	
news_commentary_it (NewsIT)	multilingual, translation	Translate Italian sentences into English.	BLEU	
object_counting (ObjCount)	logical reasoning	Questions that involve enumerating objects and asking the model to count them.	Accuracy	
play_dialog_same_or_different (PlayDiag)	reading comprehension, multiple choice	Determine if nearby lines in a Shakespeare play were spoken by the same individual.	Accuracy	
question_selection (QuestSel)	reading comprehension, multiple choice	Given an answer along with its context, select the most appropriate question which has the given answer as its answer.	Accuracy	
reasoning_about_colored_objects (ColorReason)	reading comprehension, logical reasoning, multiple choice	Answer extremely simple questions about the colors of objects on a surface.	Accuracy	
strategyqa (StratQA)	logical reasoning, context-free question answering	Answer questions in which the required reasoning steps are implicit in the question.	BLEU, ROUGE, Accuracy	
tracking_shuffled_objects (TrackObj)	logical reasoning, multiple choice	Determine the final positions given initial positions and a description of a sequence of swaps.	Accuracy	
Abstraction and Reasoning Corpus - Challenge (ARC-C)	logical reasoning, multiple choice	Answer questions that require reasoning about complex, abstract scenarios, often requiring multi-step logic and understanding of concepts.	Accuracy	
Multiple-Choice Evaluations for Embodied Agents (Multiple-E)	reasoning, multiple choice, dynamic problems	Answer questions based on a sequence of logical events or transformations, with an emphasis on evolving scenarios and inference.	Accuracy	

Table 5: Details about the 18 selected tasks following (Xu et al., 2024).

Hyper Parameters	LoRA	DoRA	CUR-LoRA	I-LoRA	ANCHORS
Exp 1					
Rank	8	8	-	8	-
learning rates	1e-5	1e-4	1e-4	1e-3	1e-3
CABR/CUR r Rank	-	-	150	-	150
CABR m Rank	-	-	-	-	200
Fusion Step	-	-	-	-	1
Trainable parameters(each layer)	1.31e+5	1.48e+5	0.90e+5	2.62e+5	1.20e+5
Exp 2					
Rank	8	8	-	8	-
learning rates	1e-2	1e-2	1e-2	1e-2	1e-2
CABR/CUR r Rank	-	-	150	-	150
CABR m Rank	-	-	-	-	200
Fusion Step	-	-	-	-	1
Trainable parameters(each layer)	1.31e+5	1.48e+5	0.90e+5	2.62e+5	1.20e+5

Table 6: hyper-parameters of experiment 1&2

Exp 3	SEQ	EWC	ER	I-LoRA	CUR-LoRA	ANCHORS
Rank	-	-	-	8	-	-
Learning Rates	2e-4	2e-4	2e-4	2e-4	2e-4	2e-4
CABR/CUR r Rank	-	-	-	-	150	150
CABR m Rank	-	-	-	-	200	200
Fusion Step	2000	2000	2000	1	2000	1
Trainable Parameters (each layer)	-	-	-	2.62e+5	1.20e+5	1.00e+5

Table 7: Hyper-parameters of experiment 3

Gradient Exp	LoRA	CABR-LoRA	ANCHORS
Rank	8	-	-
Learning Rates	1e-3	1e-3	1e-3
CABR/CUR r Rank	-	150	150
CABR m Rank	-	200	200
Fusion Step	-	-	1
Trainable Parameters (each layer)	2.62e+5	1.20e+5	1.00e+5

Table 8: Hyper-parameters of Gradient comparison experiments

SVD Norm Exp	LoRA	DoRA	CABR-LoRA	ANCHORS
Rank	8	8	-	-
Learning Rates	1e-3	1e-3	1e-3	1e-3
CABR/CUR r Rank	-	-	150	150
CABR m Rank	-	-	200	200
Fusion Step	-	-	-	1
Trainable Parameters (each layer)	2.62e+5	1.48e+5	1.20e+5	1.00e+5

Table 9: Hyper-parameters of SVD Norm experiments

Task	LoRA	DoRA	CUR-LoRA	I-LoRA	ANCHORS	LoRA	DoRA	CUR-LoRA	I-LoRA	ANCHORS
Gemma2-2B						LLaMA3.2-3B				
AbsNarr	46.50	49.53	37.03	50.31	52.76	51.59	52.12	40.07	47.28	59.06
ObjCount	54.42	57.68	58.34	48.64	59.21	53.49	55.01	48.27	51.82	57.36
PlayDiag	55.14	52.07	53.39	47.53	60.00	54.17	55.96	53.28	48.09	58.42
ColorReason	55.63	60.82	39.27	49.90	62.63	53.01	52.67	51.38	47.59	54.13
Qwen3-1.7B						LLaMA3.1-8B				
AbsNarr	59.54	59.93	56.58	58.74	59.11	55.93	56.33	43.01	41.20	57.33
ObjCount	61.17	61.34	55.37	59.14	63.46	59.76	61.10	40.35	57.36	59.59
PlayDiag	61.25	61.53	56.14	60.01	62.61	53.13	53.27	54.29	54.19	54.20
ColorReason	63.97	64.51	55.29	56.51	61.94	44.07	47.46	50.15	43.40	52.31
Qwen3-8B						Avg.				
AbsNarr	53.60	51.59	33.35	52.93	58.60	53.43	53.90	42.01	50.09	57.37
ObjCount	54.53	53.65	31.12	51.40	61.30	56.67	57.76	46.69	53.67	60.18
PlayDiag	45.00	53.80	34.56	42.13	54.20	53.74	55.33	50.33	50.39	57.89
ColorReason	44.87	44.73	36.17	54.71	53.47	52.31	54.04	46.45	50.42	56.90

Table 10: EXP1-MCQ

Task	LoRA	DoRA	CUR-LoRA	I-LoRA	ANCHORS
Gemma2-2B					
GSM8K	73.00 / 78.51 / 75.07	73.15 / 78.65 / 75.25	69.28 / 72.65 / 66.98	71.30 / 76.74 / 73.27	77.10 / 79.22 / 74.71
NewsDE	56.66 / 60.02 / 58.31	56.54 / 59.76 / 58.05	63.59 / 61.46 / 59.97	54.62 / 57.96 / 56.56	69.79 / 69.72 / 68.88
NewsIT	58.08 / 61.38 / 59.48	58.23 / 61.67 / 60.15	65.67 / 62.65 / 60.82	56.74 / 60.32 / 58.70	72.80 / 73.45 / 72.07
NewsES	58.56 / 63.99 / 62.76	58.27 / 63.58 / 62.16	66.25 / 64.38 / 62.82	56.60 / 61.27 / 59.94	70.93 / 71.95 / 71.05
ALPACA	60.91 / 58.94 / 52.50	60.92 / 58.87 / 52.54	64.15 / 62.37 / 55.21	59.68 / 58.73 / 52.49	66.68 / 64.38 / 57.82
Qwen3-1.7B					
GSM8K	73.87 / 77.21 / 75.47	71.93 / 79.18 / 75.93	67.53 / 72.96 / 69.37	71.37 / 76.63 / 73.11	71.91 / 77.64 / 74.08
NewsDE	55.42 / 61.89 / 57.66	58.17 / 58.03 / 58.14	54.35 / 54.62 / 52.86	54.99 / 57.65 / 55.97	56.42 / 60.23 / 58.84
NewsIT	59.73 / 61.42 / 64.15	56.88 / 64.91 / 62.22	55.83 / 58.08 / 56.51	56.60 / 61.25 / 59.94	58.07 / 64.12 / 62.80
NewsES	59.27 / 60.97 / 59.79	59.64 / 60.12 / 60.28	58.01 / 58.36 / 56.84	57.08 / 60.30 / 58.65	58.88 / 62.54 / 60.93
ALPACA	62.18 / 58.03 / 52.14	59.27 / 59.85 / 53.20	58.82 / 57.26 / 50.87	59.92 / 58.27 / 52.17	60.73 / 58.49 / 51.93
Qwen3-8B					
GSM8K	68.53 / 70.00 / 69.42	80.39 / 82.68 / 79.97	76.44 / 80.95 / 77.72	68.50 / 69.21 / 68.43	80.03 / 82.25 / 79.45
NewsDE	66.39 / 66.15 / 64.72	65.85 / 65.96 / 64.63	62.83 / 61.43 / 59.94	64.98 / 64.88 / 63.55	66.00 / 67.13 / 66.16
NewsIT	68.05 / 68.47 / 67.52	69.61 / 69.02 / 67.87	65.09 / 64.57 / 63.10	69.36 / 68.96 / 67.48	69.71 / 70.34 / 69.23
NewsES	66.62 / 64.28 / 58.98	68.56 / 69.78 / 69.00	63.42 / 64.08 / 63.00	66.71 / 63.97 / 58.39	67.87 / 71.33 / 69.93
ALPACA	67.28 / 67.32 / 65.93	67.24 / 64.25 / 59.00	65.45 / 62.76 / 57.26	66.19 / 65.29 / 63.77	67.29 / 63.84 / 58.19
LLaMA3.2-3B					
GSM8K	78.36 / 81.39 / 76.77	78.71 / 82.10 / 77.74	73.22 / 77.92 / 74.37	78.70 / 81.45 / 77.15	79.38 / 81.35 / 78.26
NewsDE	65.07 / 64.25 / 63.11	66.13 / 64.80 / 63.35	58.50 / 54.63 / 53.27	65.43 / 64.18 / 63.03	66.72 / 67.08 / 65.68
NewsIT	68.16 / 67.21 / 65.79	67.29 / 67.31 / 65.84	61.70 / 59.04 / 57.31	68.98 / 68.27 / 66.85	70.75 / 71.46 / 70.32
NewsES	67.77 / 68.03 / 67.06	68.19 / 68.90 / 68.26	61.37 / 59.98 / 58.50	68.95 / 69.31 / 67.87	70.29 / 70.37 / 69.45
ALPACA	65.21 / 62.33 / 56.00	63.85 / 60.31 / 53.87	61.78 / 58.81 / 52.82	64.04 / 61.03 / 54.74	65.84 / 62.83 / 57.11
LLaMA3.1-8B					
GSM8K	78.02 / 80.69 / 76.34	78.71 / 82.10 / 77.74	67.76 / 72.47 / 64.86	78.70 / 81.45 / 77.15	76.85 / 79.73 / 75.67
NewsDE	67.55 / 67.16 / 66.12	69.83 / 69.96 / 68.93	56.20 / 53.67 / 51.44	66.14 / 65.03 / 63.85	66.73 / 67.08 / 65.68
NewsIT	69.65 / 70.18 / 69.14	70.77 / 71.73 / 70.44	59.69 / 57.55 / 55.02	68.98 / 68.27 / 66.85	70.75 / 71.46 / 70.32
NewsES	70.47 / 72.10 / 71.42	71.35 / 72.24 / 71.40	54.05 / 51.71 / 49.95	68.95 / 69.31 / 67.87	70.29 / 70.37 / 69.45
ALPACA	66.28 / 62.55 / 56.66	67.10 / 63.03 / 57.35	62.86 / 60.13 / 52.66	67.24 / 63.74 / 57.61	65.89 / 62.74 / 56.40
Avg.					
GSM8K	74.36 / 77.56 / 74.61	76.58 / 80.94 / 77.33	70.85 / 75.39 / 70.66	73.71 / 77.10 / 73.82	77.05 / 80.04 / 76.43
NewsDE	62.22 / 63.89 / 61.98	63.30 / 63.70 / 62.62	59.09 / 57.16 / 55.50	61.23 / 61.94 / 60.59	65.13 / 66.25 / 65.05
NewsIT	64.73 / 65.73 / 65.22	64.56 / 66.93 / 65.30	61.60 / 60.38 / 58.55	64.13 / 65.41 / 63.96	68.42 / 70.17 / 68.95
NewsES	64.54 / 65.87 / 64.00	65.20 / 66.92 / 66.22	60.62 / 59.70 / 58.22	63.66 / 64.83 / 62.54	67.65 / 69.31 / 68.16
ALPACA	64.37 / 61.83 / 56.65	63.68 / 61.26 / 55.19	62.61 / 60.27 / 53.76	63.41 / 61.41 / 56.16	65.29 / 62.46 / 56.29

Table 11: EXP1-QA BLEU/Rouge1/RougeL

Task	LoRA	DoRA	CUR-LoRA	I-LoRA	ANCHORS
Gemma2-2B					
AbsNarr	11.76 / 18.21 / 14.46	44.61 / 43.37 / 39.92	45.26 / 48.11 / 42.62	23.70 / 28.04 / 22.38	79.10 / 80.59 / 76.18
DisflQA	07.96 / 16.23 / 12.88	17.37 / 13.87 / 09.79	27.70 / 23.99 / 16.72	13.55 / 16.00 / 12.56	65.99 / 67.22 / 60.65
LangID	08.84 / 09.89 / 09.11	11.48 / 07.69 / 07.01	19.57 / 16.18 / 13.77	11.15 / 06.57 / 06.40	71.78 / 73.19 / 66.56
LogDeduc	10.77 / 15.24 / 12.37	11.74 / 16.52 / 12.73	28.37 / 25.65 / 17.07	17.63 / 20.47 / 15.41	70.34 / 71.87 / 65.43
ObjCount	12.41 / 15.56 / 12.37	14.44 / 17.46 / 13.25	32.09 / 31.86 / 20.38	21.88 / 22.54 / 16.43	70.13 / 71.51 / 65.07
PlayDiag	11.01 / 14.50 / 12.84	13.38 / 18.26 / 14.09	24.22 / 25.98 / 17.40	12.41 / 17.05 / 13.70	69.10 / 71.65 / 64.76
ColorReason	07.57 / 18.08 / 14.16	15.75 / 17.47 / 13.33	28.71 / 28.44 / 19.90	28.36 / 28.29 / 21.19	70.14 / 72.00 / 65.32
TrackObj	09.91 / 19.02 / 15.17	16.50 / 22.88 / 17.51	29.41 / 30.99 / 22.57	26.74 / 28.52 / 21.29	70.85 / 72.62 / 66.66
Qwen3-1.7B					
AbsNarr	16.21 / 14.51 / 11.53	21.86 / 17.85 / 14.53	00.50 / 00.22 / 00.22	05.51 / 09.56 / 07.92	68.75 / 71.00 / 63.86
DisflQA	01.67 / 04.40 / 03.92	01.25 / 04.76 / 04.25	01.14 / 02.79 / 02.62	47.10 / 47.26 / 36.14	59.14 / 59.37 / 50.10
LangID	06.04 / 03.33 / 03.21	01.44 / 01.50 / 01.45	00.20 / 00.00 / 00.00	04.51 / 02.23 / 02.14	56.11 / 55.19 / 44.05
LogDeduc	06.50 / 05.92 / 05.71	06.30 / 07.43 / 07.11	01.96 / 05.86 / 05.85	05.06 / 05.81 / 05.50	55.12 / 55.28 / 45.67
ObjCount	05.48 / 04.87 / 04.71	06.05 / 04.59 / 04.56	05.87 / 06.33 / 06.00	05.19 / 03.99 / 03.99	53.16 / 53.64 / 43.93
PlayDiag	01.95 / 01.84 / 01.82	06.63 / 06.04 / 05.89	00.20 / 00.00 / 00.00	03.80 / 05.54 / 05.21	57.59 / 58.00 / 48.16
ColorReason	07.94 / 10.78 / 09.59	10.20 / 13.56 / 10.59	06.66 / 12.81 / 10.84	01.14 / 04.77 / 04.13	59.86 / 60.58 / 51.41
TrackObj	04.41 / 03.90 / 03.87	07.29 / 09.92 / 09.45	06.04 / 04.13 / 03.85	04.43 / 04.56 / 04.55	63.58 / 63.32 / 55.46
Qwen3-8B					
AbsNarr	29.36 / 27.63 / 23.33	19.50 / 21.65 / 17.02	04.75 / 08.34 / 06.90	14.82 / 14.99 / 12.91	79.14 / 80.54 / 75.40
DisflQA	00.05 / 00.00 / 00.00	08.03 / 09.92 / 09.22	01.89 / 05.18 / 04.72	08.24 / 12.59 / 10.08	68.08 / 68.94 / 61.78
LangID	00.08 / 00.00 / 00.00	04.98 / 03.63 / 03.48	04.46 / 03.47 / 03.47	06.30 / 09.19 / 07.79	61.55 / 62.51 / 52.40
LogDeduc	05.28 / 07.87 / 07.71	05.76 / 08.25 / 08.09	03.23 / 07.72 / 07.52	04.26 / 05.40 / 05.39	68.02 / 68.05 / 60.02
ObjCount	03.64 / 07.24 / 06.55	04.51 / 06.38 / 06.36	02.76 / 06.12 / 06.12	07.28 / 08.68 / 07.85	63.12 / 64.03 / 54.98
PlayDiag	04.22 / 03.32 / 03.32	03.32 / 06.82 / 06.74	00.45 / 00.30 / 00.30	04.54 / 04.32 / 04.32	63.07 / 65.68 / 56.88
ColorReason	00.07 / 00.26 / 00.26	04.59 / 10.46 / 09.38	01.69 / 08.20 / 06.24	03.75 / 01.11 / 09.36	68.91 / 71.29 / 64.42
TrackObj	05.66 / 07.72 / 07.59	05.85 / 08.80 / 08.55	04.61 / 08.84 / 07.93	04.80 / 03.44 / 03.43	71.83 / 73.22 / 67.53
LLaMA3.2-3B					
AbsNarr	01.53 / 00.64 / 00.64	11.74 / 16.61 / 13.06	00.37 / 00.00 / 00.00	05.87 / 10.07 / 08.86	79.13 / 78.88 / 72.92
DisflQA	00.04 / 00.00 / 00.00	09.27 / 07.56 / 07.19	00.04 / 00.00 / 00.00	04.65 / 07.55 / 07.33	71.53 / 72.54 / 65.45
LangID	02.45 / 02.33 / 02.32	16.30 / 00.90 / 00.87	00.22 / 00.00 / 00.00	01.38 / 01.31 / 01.30	60.60 / 61.56 / 49.06
LogDeduc	05.84 / 06.27 / 06.05	06.09 / 04.73 / 04.64	04.43 / 03.43 / 03.43	01.25 / 01.16 / 01.16	68.94 / 69.13 / 61.18
ObjCount	00.44 / 01.20 / 01.20	07.43 / 06.63 / 06.42	05.43 / 04.11 / 04.11	05.48 / 04.69 / 04.60	71.04 / 72.19 / 64.88
PlayDiag	04.67 / 03.43 / 03.43	04.29 / 04.73 / 04.63	04.53 / 04.48 / 04.47	05.87 / 04.41 / 04.41	71.41 / 71.90 / 64.50
ColorReason	03.13 / 03.89 / 03.88	07.54 / 11.14 / 09.61	00.04 / 00.00 / 00.00	04.62 / 04.63 / 04.63	71.09 / 71.94 / 64.40
TrackObj	04.09 / 05.27 / 05.26	07.44 / 09.43 / 08.73	00.26 / 01.61 / 01.60	05.33 / 04.32 / 04.32	72.59 / 73.08 / 66.85
LLaMA3.1-8B					
AbsNarr	01.06 / 02.86 / 02.71	26.51 / 26.79 / 23.03	00.90 / 01.50 / 01.50	05.49 / 08.93 / 08.68	77.38 / 78.04 / 72.25
DisflQA	00.04 / 00.00 / 00.00	03.95 / 07.75 / 07.38	04.00 / 04.87 / 04.87	02.01 / 06.83 / 06.40	68.57 / 70.12 / 61.04
LangID	00.07 / 00.00 / 00.00	06.17 / 02.93 / 02.78	01.04 / 00.28 / 00.28	00.10 / 00.18 / 00.18	69.68 / 70.93 / 61.56
LogDeduc	00.07 / 00.00 / 00.00	05.66 / 06.75 / 06.60	05.92 / 05.59 / 05.56	04.40 / 05.54 / 05.54	71.49 / 73.35 / 65.00
ObjCount	04.85 / 05.24 / 05.24	05.79 / 07.37 / 07.20	06.21 / 08.17 / 07.98	05.87 / 05.86 / 05.85	71.73 / 73.79 / 65.33
PlayDiag	00.25 / 02.35 / 02.35	08.82 / 07.82 / 07.24	04.87 / 03.87 / 03.87	01.80 / 06.72 / 06.70	70.35 / 72.43 / 63.66
ColorReason	00.04 / 00.00 / 00.00	07.29 / 12.35 / 10.90	01.43 / 06.22 / 06.21	07.80 / 10.64 / 09.77	72.22 / 73.05 / 65.34
TrackObj	04.80 / 04.57 / 04.57	03.51 / 12.62 / 11.66	02.36 / 11.08 / 10.23	04.94 / 05.63 / 05.60	71.63 / 73.64 / 66.96
Avg.					
AbsNarr	11.98 / 12.77 / 10.53	24.84 / 25.25 / 21.51	10.36 / 11.63 / 10.25	11.08 / 14.32 / 12.15	76.70 / 77.81 / 72.12
DisflQA	01.95 / 04.13 / 03.36	07.97 / 08.77 / 07.57	06.95 / 07.37 / 05.79	15.11 / 18.05 / 14.50	66.66 / 67.64 / 59.80
LangID	03.50 / 03.11 / 02.93	08.07 / 03.33 / 03.12	05.10 / 03.99 / 03.50	04.69 / 03.90 / 03.56	63.94 / 64.68 / 54.73
LogDeduc	05.69 / 07.06 / 06.37	07.11 / 08.74 / 07.83	08.78 / 09.65 / 07.89	06.52 / 07.68 / 06.60	66.78 / 67.54 / 59.46
ObjCount	05.36 / 06.82 / 06.01	07.64 / 08.49 / 07.56	10.47 / 11.32 / 08.92	09.14 / 09.15 / 07.74	65.84 / 67.03 / 58.84
PlayDiag	04.42 / 05.09 / 04.75	07.29 / 08.73 / 07.72	06.85 / 06.93 / 05.21	05.68 / 07.61 / 06.87	66.30 / 67.93 / 59.59
ColorReason	03.75 / 06.60 / 05.58	09.07 / 13.00 / 10.76	07.71 / 11.13 / 08.64	09.13 / 09.89 / 09.82	68.44 / 69.77 / 62.18
TrackObj	05.77 / 08.10 / 07.29	08.12 / 12.73 / 11.18	08.54 / 11.33 / 09.24	09.25 / 09.29 / 07.84	70.10 / 71.18 / 64.69

Table 12: EXP2-Forgetting BLEU/Rouge1/RougeL

Datasets	SEQ	EWC	ER	I-LoRA	ANCHORS
AbsNarr-1k	74.70 / 76.90 / 72.46	70.23 / 71.72 / 67.12	74.25 / 76.66 / 72.22	75.26 / 76.69 / 72.51	78.10 / 78.67 / 74.36
AbsNarr-2k	71.87 / 73.78 / 69.63	69.78 / 71.59 / 66.90	74.04 / 76.20 / 71.85	74.98 / 76.49 / 72.31	77.81 / 78.43 / 74.32
DisflQA-1k	73.65 / 75.03 / 70.69	71.26 / 72.75 / 67.52	75.55 / 77.63 / 73.26	74.82 / 76.73 / 72.24	78.13 / 78.53 / 74.17
DisflQA-2k	74.19 / 76.10 / 71.64	70.73 / 72.81 / 67.72	74.14 / 76.87 / 72.22	74.16 / 76.94 / 72.31	78.07 / 78.72 / 74.52
CNNNDM-1k	74.82 / 76.44 / 71.42	71.43 / 71.69 / 66.45	75.57 / 76.83 / 72.29	76.35 / 77.97 / 73.39	78.64 / 78.85 / 74.51
CNNNDM-2k	73.78 / 75.17 / 70.35	69.30 / 70.63 / 64.91	75.67 / 77.07 / 72.47	74.55 / 76.31 / 71.51	78.63 / 78.84 / 74.50
LangID-1k	72.11 / 73.72 / 69.03	68.15 / 69.80 / 64.59	74.46 / 76.59 / 71.71	72.40 / 74.56 / 69.80	77.86 / 78.26 / 73.98
LangID-2k	73.26 / 75.12 / 70.54	67.61 / 69.15 / 63.94	74.56 / 76.15 / 71.49	72.58 / 74.20 / 69.56	77.62 / 77.93 / 73.77
LingPuzz-1k	73.15 / 75.27 / 70.66	71.21 / 72.27 / 66.47	74.69 / 76.87 / 72.38	71.18 / 74.23 / 69.66	78.07 / 78.27 / 73.89
LingPuzz-2k	71.42 / 75.96 / 71.32	70.70 / 72.06 / 66.29	74.60 / 77.02 / 72.36	72.22 / 74.89 / 70.22	78.09 / 78.87 / 74.66
LogDeduc-1k	73.01 / 75.52 / 71.10	70.36 / 72.54 / 66.95	75.01 / 77.29 / 72.94	70.98 / 74.13 / 69.49	78.42 / 78.85 / 74.36
LogDeduc-2k	74.13 / 75.04 / 70.17	70.25 / 71.18 / 65.23	74.46 / 76.36 / 71.85	72.13 / 74.27 / 69.38	78.12 / 78.53 / 74.39
NewsDE-1k	72.63 / 75.13 / 70.47	70.37 / 71.91 / 66.67	75.21 / 77.17 / 72.78	73.49 / 75.60 / 70.99	78.01 / 78.43 / 74.13
NewsDE-2k	74.32 / 76.13 / 71.22	71.03 / 71.95 / 66.19	74.46 / 76.29 / 71.77	74.27 / 76.15 / 71.51	77.72 / 78.28 / 74.11
NewsES-1k	74.56 / 76.75 / 72.04	70.86 / 72.55 / 66.97	73.97 / 76.52 / 71.67	74.88 / 76.83 / 72.35	78.46 / 78.92 / 74.66
NewsES-2k	73.99 / 76.40 / 71.74	70.48 / 71.30 / 65.83	72.87 / 74.96 / 70.17	74.13 / 76.45 / 71.99	77.24 / 77.82 / 73.54
NewsIT-1k	72.66 / 75.47 / 70.25	69.45 / 70.32 / 64.82	74.95 / 76.87 / 71.90	72.70 / 74.78 / 69.97	77.54 / 77.80 / 73.49
NewsIT-2k	72.21 / 74.95 / 70.27	70.00 / 70.85 / 64.88	73.88 / 75.70 / 70.98	72.60 / 74.73 / 70.14	77.13 / 77.96 / 73.75
ObjCount-1k	70.69 / 73.38 / 68.69	67.76 / 69.20 / 63.74	73.10 / 76.85 / 70.17	72.96 / 75.07 / 70.35	77.67 / 77.51 / 73.14
ObjCount-2k	70.22 / 72.59 / 67.67	70.25 / 70.70 / 65.00	73.49 / 75.70 / 70.92	72.59 / 74.46 / 69.76	77.29 / 77.61 / 73.36
PlayDiag-1k	68.18 / 72.19 / 67.28	61.19 / 63.55 / 58.32	73.65 / 75.39 / 70.57	71.93 / 73.94 / 68.97	78.33 / 77.90 / 73.37
PlayDiag-2k	69.71 / 73.18 / 68.59	65.02 / 66.27 / 60.71	72.50 / 74.91 / 70.30	72.72 / 74.57 / 69.74	77.56 / 77.97 / 73.58
QuesSel-1k	69.15 / 71.89 / 67.24	64.51 / 66.07 / 60.80	73.08 / 75.02 / 70.00	71.83 / 73.41 / 68.53	77.27 / 77.85 / 73.43
QuesSel-2k	68.59 / 71.27 / 66.30	62.79 / 64.80 / 59.60	72.10 / 74.53 / 69.55	71.63 / 73.35 / 68.54	77.16 / 77.68 / 73.49
ColorReason-1k	00.69 / 04.12 / 03.61	66.78 / 68.06 / 62.29	73.00 / 75.08 / 70.51	71.88 / 73.82 / 68.70	77.06 / 77.23 / 72.88
ColorReason-2k	00.50 / 00.12 / 00.11	66.97 / 67.63 / 61.64	73.02 / 74.78 / 69.82	71.65 / 73.63 / 68.65	77.29 / 77.33 / 72.85
StratQA-1k	71.99 / 74.41 / 68.99	70.30 / 71.13 / 65.20	74.27 / 76.35 / 71.60	74.06 / 76.45 / 71.73	77.59 / 77.62 / 73.01
StratQA-2k	71.97 / 74.63 / 68.82	69.80 / 70.48 / 63.99	72.71 / 75.17 / 70.29	73.78 / 76.19 / 71.27	78.05 / 78.07 / 73.52
TrackObj-1k	68.24 / 71.24 / 65.49	67.80 / 68.65 / 62.75	73.70 / 75.33 / 70.39	73.60 / 75.03 / 70.24	77.45 / 77.61 / 73.29
TrackObj-2k	67.77 / 70.67 / 64.72	67.28 / 68.14 / 62.18	73.19 / 74.73 / 69.69	72.63 / 74.60 / 69.44	77.48 / 77.41 / 72.99
ALPACA-1k	73.21 / 74.79 / 70.25	72.44 / 73.32 / 68.25	74.93 / 76.41 / 71.39	71.92 / 73.11 / 70.06	77.86 / 77.67 / 73.39
ALPACA-2k	75.07 / 76.18 / 71.50	72.43 / 72.82 / 67.28	75.92 / 77.50 / 72.74	70.16 / 72.04 / 69.87	77.22 / 77.50 / 73.42

Table 13: EXP3 Qwen3-8B

Datasets	SEQ	EWC	ER	I-LoRA	ANCHORS
AbsNarr-1k	66.43 / 69.08 / 60.61	74.94 / 75.45 / 70.48	32.34 / 45.39 / 35.28	70.48 / 71.89 / 64.78	74.09 / 74.32 / 66.71
AbsNarr-2k	66.63 / 68.30 / 60.47	73.95 / 74.59 / 69.67	60.59 / 66.00 / 58.00	70.90 / 72.22 / 64.95	73.77 / 74.26 / 66.88
DisflQA-1k	67.92 / 70.91 / 63.14	73.44 / 74.47 / 69.42	67.53 / 72.21 / 63.35	71.35 / 74.16 / 66.50	74.04 / 74.36 / 66.66
DisflQA-2k	64.96 / 69.89 / 61.70	73.79 / 74.91 / 69.56	67.09 / 70.48 / 59.69	70.09 / 73.37 / 65.68	73.38 / 74.10 / 66.46
CNNNDM-1k	68.72 / 70.46 / 62.42	76.45 / 76.50 / 71.10	68.35 / 71.13 / 63.48	69.91 / 72.35 / 64.02	74.72 / 74.73 / 67.23
CNNNDM-2k	68.59 / 70.30 / 62.68	75.11 / 75.48 / 70.17	66.40 / 69.16 / 60.80	68.02 / 71.00 / 62.93	74.15 / 74.52 / 67.27
LangID-1k	03.07 / 08.98 / 07.41	68.40 / 68.08 / 63.03	00.33 / 01.35 / 01.29	65.29 / 71.45 / 62.65	73.68 / 74.09 / 66.76
LangID-2k	15.74 / 19.36 / 15.14	70.39 / 69.98 / 65.16	59.98 / 63.09 / 50.47	65.83 / 70.77 / 61.78	73.11 / 73.58 / 66.07
LingPuzz-1k	48.95 / 48.94 / 36.97	75.24 / 76.05 / 70.19	67.29 / 69.87 / 60.44	67.95 / 70.43 / 61.33	74.07 / 74.45 / 67.01
LingPuzz-2k	53.57 / 54.63 / 41.63	75.08 / 75.18 / 69.74	67.95 / 70.02 / 61.06	68.21 / 70.36 / 61.20	73.89 / 74.90 / 67.47
LogDeduc-1k	01.13 / 01.46 / 01.39	64.59 / 67.35 / 61.50	67.13 / 68.67 / 59.66	66.99 / 71.44 / 62.27	74.62 / 75.00 / 67.46
LogDeduc-2k	00.08 / 00.14 / 00.14	57.12 / 59.51 / 53.00	64.57 / 67.13 / 57.02	65.24 / 69.81 / 60.40	73.75 / 74.04 / 66.56
NewsDE-1k	04.25 / 05.36 / 05.34	74.30 / 75.21 / 69.41	68.79 / 71.09 / 63.37	66.80 / 71.05 / 62.65	74.17 / 74.47 / 67.07
NewsDE-2k	09.90 / 08.28 / 08.17	74.23 / 74.91 / 68.99	67.84 / 70.82 / 61.94	68.19 / 72.45 / 64.00	74.17 / 74.79 / 67.44
NewsES-1k	07.93 / 08.84 / 08.11	74.10 / 74.49 / 68.26	69.06 / 71.34 / 62.63	67.46 / 71.54 / 62.82	74.07 / 74.29 / 66.92
NewsES-2k	04.52 / 08.01 / 07.56	74.45 / 74.67 / 69.09	66.20 / 69.59 / 60.93	68.09 / 71.47 / 62.99	73.57 / 74.03 / 66.58
NewsIT-1k	06.32 / 10.47 / 09.71	73.40 / 74.17 / 68.46	67.96 / 69.61 / 60.05	67.43 / 71.15 / 62.89	74.02 / 74.40 / 67.12
NewsIT-2k	07.03 / 07.82 / 07.32	74.28 / 74.56 / 68.39	65.59 / 67.50 / 58.50	66.86 / 70.69 / 62.42	73.76 / 74.62 / 67.37
ObjCount-1k	01.24 / 00.08 / 00.08	60.70 / 64.46 / 61.00	63.00 / 65.83 / 54.97	67.26 / 71.33 / 63.56	74.95 / 74.81 / 67.38
ObjCount-2k	01.82 / 00.04 / 00.04	52.97 / 58.52 / 53.36	65.46 / 67.53 / 58.35	66.73 / 71.07 / 63.09	73.78 / 74.19 / 66.91
PlayDiag-1k	00.88 / 01.65 / 01.62	58.33 / 61.73 / 56.09	62.17 / 64.27 / 54.82	66.60 / 71.05 / 62.74	74.90 / 74.73 / 66.85
PlayDiag-2k	01.93 / 01.51 / 01.50	57.43 / 60.60 / 59.28	57.02 / 61.87 / 49.83	66.59 / 71.31 / 62.78	74.31 / 74.61 / 67.23
QuesSel-1k	00.88 / 01.29 / 01.26	63.78 / 65.67 / 65.01	58.69 / 61.48 / 50.89	66.59 / 70.42 / 62.38	73.70 / 74.38 / 66.91
QuesSel-2k	00.64 / 01.05 / 01.05	63.99 / 65.81 / 65.12	60.85 / 63.44 / 51.63	64.38 / 68.86 / 60.82	73.78 / 74.32 / 66.99
ColorReason-1k	00.14 / 00.86 / 00.82	60.27 / 60.56 / 60.53	00.12 / 00.69 / 00.69	64.26 / 68.75 / 58.34	74.14 / 74.10 / 66.80
ColorReason-2k	00.07 / 00.18 / 00.17	61.36 / 62.38 / 62.24	00.07 / 00.21 / 00.21	63.73 / 69.12 / 59.66	74.01 / 74.08 / 66.63
StratQA-1k	00.87 / 07.93 / 07.03	70.37 / 72.07 / 64.64	00.58 / 04.17 / 04.13	68.47 / 72.56 / 64.53	73.96 / 74.18 / 66.64
StratQA-2k	03.50 / 10.43 / 09.05	70.97 / 71.19 / 64.15	05.39 / 08.21 / 07.49	67.52 / 71.44 / 63.01	74.84 / 75.00 / 67.42
TrackObj-1k	00.09 / 00.50 / 00.50	59.42 / 62.32 / 54.04	05.76 / 05.42 / 05.32	66.38 / 70.51 / 61.74	74.47 / 74.61 / 67.09
TrackObj-2k	00.06 / 01.97 / 01.94	55.30 / 56.99 / 48.89	06.69 / 04.55 / 04.50	65.61 / 70.34 / 60.98	74.10 / 73.87 / 66.66
ALPACA-1k	08.43 / 11.87 / 10.45	72.34 / 72.70 / 66.75	08.37 / 10.13 / 9.40	69.55 / 72.41 / 65.85	74.91 / 74.73 / 67.46
ALPACA-2k	05.74 / 13.15 / 12.27	70.13 / 71.02 / 69.84	09.37 / 12.27 / 10.82	71.14 / 72.53 / 66.50	73.32 / 73.99 / 66.74

Table 14: EXP3 LLaMA3.1-8B