

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 FINETUNE ONCE: DECOUPLING GENERAL & DOMAIN LEARNING WITH DYNAMIC ANCHOR ANNEALING

Anonymous authors

Paper under double-blind review

ABSTRACT

Large language models (LLMs) fine-tuning shows excellent implications. However, vanilla fine-tuning methods often require intricate data mixture and repeated experiments for optimal generalization. To address these challenges and streamline the training process, we propose an efficient and universal solution, Dynamic Anchor Annealing (DAA). We obtain a global gradient through zero-learning-rate training on general data, which is subsequently employed for gradient anchor and dynamic training step correction during domain training. In conjunction with annealing learning, we end up establishing a fine-tuning pipeline that relies solely on domain data without collapse. By evaluating both general and domain-specific performance across multiple tasks on several popular base models, DAA achieves an average improvement of 5.8% in joint performance over vanilla fine-tuning. Furthermore, since general data is no longer involved in annealing, repeated experiments led by data mixture are also eliminated. According to our tests, the DAA method can reduce GPU hours by 91.0% compared to the vanilla method.

1 INTRODUCTION

Large Language Models (LLMs) show significant promise in various applications due to their ability to understand and generate human-like text. Fine-Tuning (FT) LLMs on domain-specific tasks has become a common approach to enhance their performance in targeted applications Yang et al. (2023); Zhou et al. (2024); Chen et al. (2024); Huang et al. (2023). However, empirical evidence suggests that fine-tuned LLMs frequently demonstrate significant degradation of their original performance Chen et al. (2020); Luo et al. (2025); Lin et al. (2023); Korbak et al. (2022). Therefore, mitigating catastrophic forgetting in the fine-tuning process has emerged as a crucial research focus for LLMs (Table 1, row 1).

Data Mixture (DM) strategy was the basic and vanilla solution Wen et al. (2023); Wu et al. (2023); Zhang et al. (2024a); Wu et al. (2024); Held et al. (2025) to solve catastrophic problem. It combines general and domain-specific data in fine-tuning datasets to mitigate forgetting of general capabilities. Due to the coupling between data from different domains, each fine-tuning requires repeated experimentation to adjust the data mixture in order to achieve satisfactory performance (Table 1, row 2). As shown in Figure 1, the effectiveness of DM heavily depends on the mixing ratio, necessitating extensive empirical validation to determine optimal proportions for each domain. Alternative approaches, such as Low-Rank Adaptation (LoRA) Hu et al. (2021); Yang et al. (2023); Cui et al. (2023), have demonstrated some success in preserving general capabilities, yet they face inherent limitations in achieving peak domain-specific performance (Table 1, row 3). This ad-hoc process of data mixing is not only computationally prohibitive but also lacks scalability, as the optimal ratio for one domain rarely transfers to another. Consequently, an ideal fine-tuning framework must decouple domain adaptation from the costly cycle of data mixture experiments, while still effectively balancing specialization with the preservation of general knowledge.

To address the above challenges, we propose **Dynamic Anchor Annealing (DAA)**, a streamline fine-tuning framework that eliminates the requirements for data mixture and repeated experiments. First, to effectively isolating the contributions of general-domain and domain-specific data, we propose **Global Gradient Anchoring (GGA)**. Here, “**Anchoring**” refers to augmenting the domain-specific gradient with a pre-computed global one. This method initially estimates the global gradient in the general domain through zero-learning-rate learning. During fine-tuning, the global gradient is

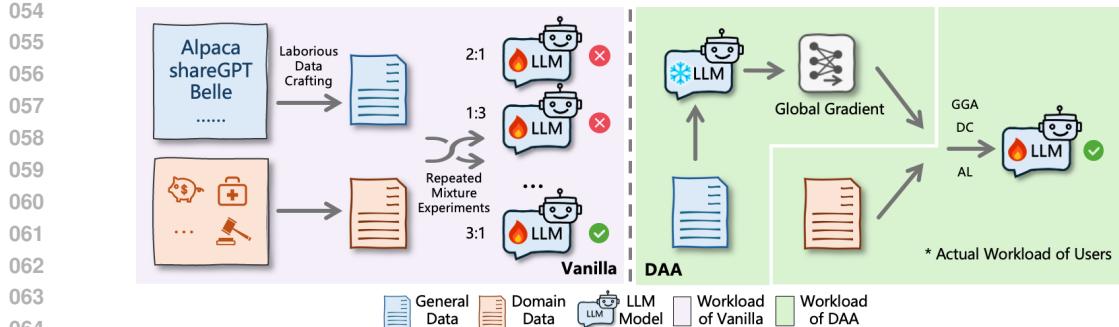


Figure 1: Comparison between vanilla and DAA. [*] is the part that users need to perform in SFT.

Table 1: Comparison of different methods. Repeated Exps indicates that the method requires hyper-parameter tuning or recipe adjustment for data mixture ratios to achieve optimal results. Collapse means losing generalization ability.

Method	No Data Mixture	No Repeated Exps	Reduce Cost	No Collapse	SOTA
Direct FT	✓	✓	✓	✗	✗
Vanilla FT	✗	✗	✗	✓	✗
LoRA-like FT	✓	✗	✓	✓	✗
DAA (Ours)	✓	✓	✓	✓	✓

used as guidance, combined with annealing learning, to mitigate catastrophic forgetting. Second, to achieve global optimal performance in specifics, we propose a domain similarity-guided **Dynamic Correction (DC)** strategy. This adaptive parameter update strategy modulates the optimization steps based on the gradient similarity between specific and general domains. It is important to emphasize that the fundamental goal of DAA extends beyond the reduction of fine-tuning expenses to specifically address the mitigation of catastrophic forgetting. The efficiency of this method emerges as a significant inherent benefit of its design, as it delivers comparable or superior performance while avoiding the substantial costs associated with data mixture search. As demonstrated in Table 1, DAA achieves superior performance compared to conventional fine-tuning approaches, while significantly reducing workload by eliminating the need for data mixing and repeated experiments. Our contributions are summarized as follows:

- We explore the impact of data mixture on both fine-tuning performance and workload, and propose new fine-tuning schemes.
- We propose DAA, a novel training framework designed to efficiently fine-tuning by gradient-based domain decoupling and similarity-guided adaptation.
- We conduct empirical evaluations across various tasks, demonstrating that our method effectively balances domain-specific performance while maintaining general capabilities with low workload.

2 MOTIVATION

2.1 RELATED WORK

Recent work on fine-tuning Xie et al. (2023); Zhang et al. (2023b); Bao et al. (2023); Yue et al. (2023); Chen et al. (2024); Zhou et al. (2024); Yang et al. (2023); Cui et al. (2023) including direct fine-tuning Xie et al. (2023); Zhu et al. (2024), vanilla fine-tuning Bao et al. (2023); Yue et al. (2023); Chen et al. (2024); Deng et al. (2023) and LoRA Yang et al. (2023); Chen et al. (2023); Cui et al. (2023) seeks to mitigate catastrophic forgetting while controlling cost. As shown in Table 1, direct fine-tuning is inexpensive yet the general performance of LLMs can collapse. Vanilla fine tuning employs data mixtures to suppress forgetting and often preserves general capability, although

108 the cost rises sharply. LoRA is effective in reducing both forgetting and cost, but performance in
 109 unfamiliar specific domains remains below that of full-parameter fine-tuning.
 110

111 Among these options, vanilla fine-tuning with data mixture offers the best balance between general
 112 and domain-specific performance, yet its experimental cost is substantial. Mixing ratios tend to be
 113 domain-dependent and therefore require repeated experimentation for each target domain Wen et al.
 114 (2023). In addition, when the ratio is swept from $1 : 1$ through $1 : N$, the total volume of processed
 115 data scales as $\sum_{n=1}^N (1 + n) = O(N^2)$ times the size of the specific domain set, which becomes
 116 prohibitive as the domain dataset grows. This computational inefficiency motivates more efficient
 117 and more generalizable fine-tuning methodology.
 118

Recent studies Wu et al. (2024); Dong et al. (2024) investigate the optimization of data mixtures.
 For instance, Mixture-of-Skills Wu et al. (2024) utilizes Reinforcement Learning (RL) to optimize
 data utilization ratios. While DAA eliminates the need for online mixture search, it is orthogonal to
 data selection strategies. Future work could integrate RL-based data curriculum to further enhance
 the quality of the Global Gradient Anchor. Additionally, Liang et al. (2025) employs a sample-level
 Evolving Interaction-guided Curriculum for multi-domain coordination, DAA adopts a domain-level
 strategy through the Global Gradient Anchor. This approach abstracts the general prior into a stable
 optimization direction to prioritize computational efficiency and scalability over complex sample
 dynamics, which facilitates a streamlined “Finetune Once” workflow. Furthermore, gradient surgery
 methods such as PCGrad Yu et al. (2020) and GradNorm Chen et al. (2018) manage gradient con-
 flicts in multi-task learning by storing individual task gradients, yet this approach incurs prohibitive
 memory costs for LLMs. In contrast, DAA employs a memory-efficient and pre-computed global
 anchor to specifically mitigate catastrophic forgetting.
 129

Another related area of research is Continual Learning (CL). CL is defined as a model learning from
 a dynamic data distribution Wen et al. (2023). Our setting can be viewed as single task continual
 learning in which, after adapting a pretrained model to one instruction task, we aim to mitigate
 degradation of its general capabilities.
 130

Besides, it’s important to differentiate the optimization challenges in DAA from those in Multi-
 Task Learning (MTL). Prior MTL works often focus on **imbalanced domains**, where high-resource
 tasks dominate the gradient direction at the expense of low-resource tasks Sener & Koltun (2018),
 and **cross-domain transfer**, which aims to leverage knowledge from a source domain to improve
 a distinct target domain Liu et al. (2019). In contrast, DAA targets Finetuning to Target Domain ,
 where the optimization landscape is defined by the tension between the General Prior and the Target
 Domain. 1) Handling Imbalance: Unlike MTL, where multiple specific domains compete for
 model capacity, the imbalance in our setting exists between the massive pre-training distribution and
 the smaller fine-tuning dataset. DAA explicitly regulates this imbalance through the Global Gra-
 dient Anchor and the decay coefficient γ_t , ensuring that the general prior is not overwhelmed by
 the specific domain, regardless of the domain’s data scale. 2) Nature of Transfer: DAA focuses on
 General-to-Specific retention rather than Specific-to-Specific transfer. The Global Gradient serves
 as a regularization term that enforces the preservation of general reasoning structures while allow-
 ing the model to adapt to the target distribution. Thus, the negative transfer or interference often
 seen between disjoint domains in MTL is mitigated by anchoring optimization to the robust general
 manifold.
 149

150

151 2.2 ROLE OF GRADIENT IN FINE-TUNING

152

153 During stochastic optimization, gradient variance strongly affects convergence and generalization
 154 Gurbuzbalaban et al. (2021). High variance slows convergence and complicates optimization Agar-
 155 wal et al. (2022); Xia et al. (2024), which can hinder domain adaptation. We provide qualitative and
 156 quantitative analyses of fine-tuning across general and specific domains and expose drawbacks of
 157 multi domain optimization.
 158

First, we qualitatively analyze the differences in convergence trajectory between general and spe-
 159 cific domain by visualizing the loss landscape. Following the methodology in Lucas et al. (2021),
 160 we interpolate between the weights θ_0 of Qwen3-1.7B Yang et al. (2025) base model and the fully
 161 fine-tuned weights θ_D , constructing a two-dimensional slice of the loss landscape. To ensure inde-
 162 pendence, we apply orthogonalization to the interpolation direction.
 163

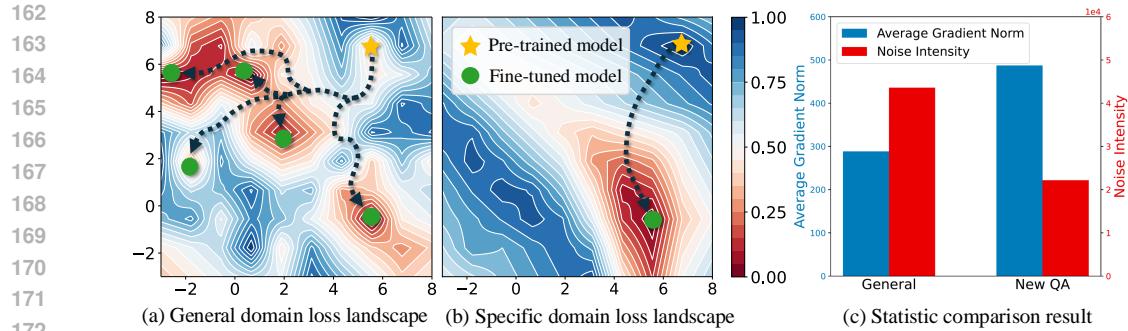


Figure 2: Qualitative and quantitative analysis result of general and specific domain.

Figure 2(a) and 2(b) show that the general domain has many local optima and a tortuous path details in appendix B, while the specific domain news QA details in Section 4.2 shows fewer optima and a more stable trajectory. General domain training can therefore constrain domain specific fine-tuning.

Second, we follow stochastic optimization methodology Ghadimi & Lan (2013) to compare average gradient norms and noise scale across the two domains. As shown in Figure 2(c), the general domain has nearly twice the noise scale of the specific domain. We randomly sample 1,000 instances from each domain. On the full general domain the gap may be larger. We therefore attempt to freeze general domain gradients to limit their impact on training.

In multi domain optimization, conflicts between domain gradients degrade efficiency Yu et al. (2020); Hadsell et al. (2020); Liu et al. (2021). Prior work reduces negative interactions by removing projection components between domain gradients Yu et al. (2020) or by automatic gradient balancing Liu et al. (2021). This motivates a balancing mechanism between specific and general domain gradients that preserves generalization while learning specific domain distributions.

3 METHOD

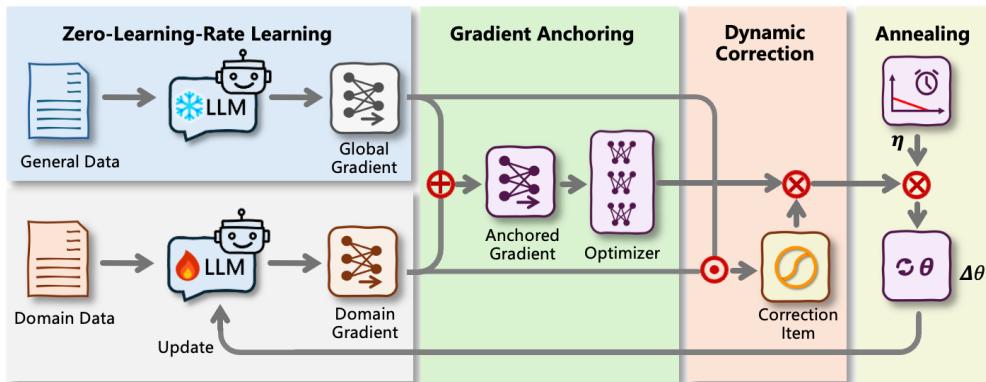


Figure 3: Overview of Dynamic Anchor Annealing. Our approach consists of two stages. In the first stage, global gradient is estimated in the general domain through zero-learning-rate learning, which serves as an independent preprocessing stage. In the second stage, the fine-tuning step, global gradient anchors the specific gradient to preserve general capability, while the similarity between global and specific gradients adaptively determines the parameter update magnitude. The learning rate with annealing strategy suppresses degradation.

In this section, we formally introduce the Dynamic Anchor Annealing (DAA) illustrated in Figure 3, which is based on annealing learning. It is worth to note that DAA is a full-parameter fine-tuning framework in which decoupling refers to the optimization-level separation of general and domain-specific signals within the gradient space rather than the division of model components. Initially,

216 the global gradient is independently estimated in the general domain through zero-learning-rate
 217 learning. During the fine-tuning stage, DAA anchors the gradient to preserve the general capability.
 218 Subsequently, the similarity between the global gradient and specific gradient adaptively selects the
 219 magnitude of parameter update. Finally, the learning rate with the annealing strategy suppresses
 220 degradation effectively.

222 3.1 GLOBAL GRADIENT ANCHORING

224 The global gradient serves as a stable optimization anchor. Our design is motivated by two key
 225 factors: 1) Variance Reduction: \hat{g}_G acts as a low-variance regularizer for the high-variance domain
 226 gradient $g_{D,t}$, preventing severe oscillations (see Eqn. 4). 2) Trajectory Smoothing: As shown in
 227 Figure 2(a), \hat{g}_G points towards regions of good generalization, helping the optimizer bypass poor
 228 local optima in the general domain landscape. The term “Anchoring” is metaphorical. We anchor
 229 or augment the domain-specific gradient at each step with this pre-computed stable anchor. This
 230 approach is distinct from traditional Gradient Boosting Machines.

231 In the joint learning of general and specific domains, the gradient is a weighted sum of the general
 232 gradient and the specific gradient, with weights determined by the data mixing ratio λ , that is

$$233 \quad g_{M,t} = \lambda g_{G,t} + (1 - \lambda) g_{D,t}. \quad (1)$$

235 We define \hat{g}_G as a fixed estimator of $g_{G,t}$ in joint training to diminish the volatility of the combined
 236 gradient.

$$237 \quad g_{B,t} = \gamma_t \hat{g}_G + (1 - \gamma_t) g_{D,t}, \quad (2)$$

238 where γ_t is the anchoring magnitude. The expectation and variance of $g_{B,t}$ are given by

$$240 \quad \mathbb{E}[g_{B,t}] = \gamma_t \hat{g}_G + (1 - \gamma_t) \mathbb{E}[g_{D,t}], \quad (3)$$

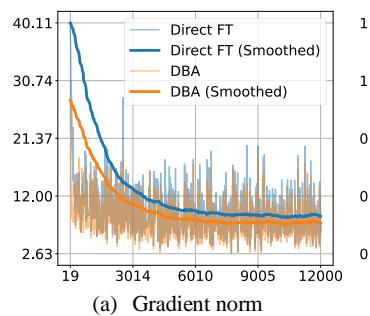
$$241 \quad \mathbb{E}[\|g_{B,t} - \mathbb{E}[g_{B,t}]\|^2] = (1 - \gamma_t)^2 \mathbb{E}[\|g_{D,t} - \mathbb{E}[g_{D,t}]\|^2]. \quad (4)$$

243 By fixed \hat{g}_G , we can significantly mitigate the randomness of parameter update while maintaining
 244 the regularization effect on optimization.

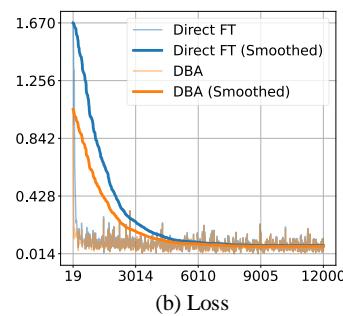
245 To estimate the global gradient \hat{g}_G , according to the derivation of Adam Kingma & Ba (2017),
 246 when the exponential decay rate for the 1st momentum estimates $\beta_1 \rightarrow 1$, the momentum $m_{G,t}$
 247 approximates the expectation of the gradient.

$$249 \quad \hat{g}_G = \mathbb{E}[g_{G,i}] = s^{-1} \sum_{i=1}^s g_{G,i} = \lim_{\beta_1 \rightarrow 1} m_{G,s}. \quad (5)$$

252 Therefore, we trained the LLM on the general domain with a learning rate of 0 and a decay rate
 253 $\beta_1 \rightarrow 1$, then stored the final momentum after s training steps. Notably, \hat{g}_G can be applied across
 254 all domains, rather than being obtained per domain. The acquisition of \hat{g}_G is consistent with the
 255 findings of LIMA Zhou et al. (2023) that prioritize diversity and quality over sheer quantity. As the
 256 general dataset attains sufficient diversity, the direction of \hat{g}_G stabilizes, which implies that massive
 257 datasets are not strictly required for robust anchoring.



268 (a) Gradient norm



268 (b) Loss

269 Figure 4: Comparative results of global gradient boosted learning.

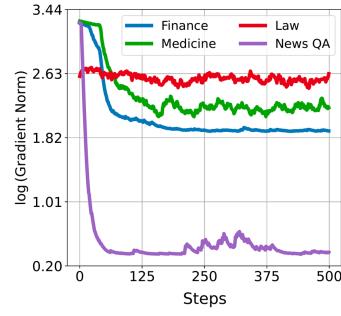


Figure 5: Norm of gradient projection.

We refer to this method as **Global Gradient Anchoring (GGA)**. The stored m_G from the general domain supplies a stable guidance signal during specific domain optimization and steers updates toward a joint optimum. As shown in Figure 4, **GGA** markedly reduces gradient norm and loss, especially early in fine-tuning, which indicates improved training stability on the specific domain with preservation of general capability.

The **anchoring** magnitude can be adapted over training. We set $\gamma_t = k_0 (1 - \frac{t}{T})$, so that the model increasingly emphasizes the specific domain near the end of fine tuning. Accounting for exponential averaging in Adam Kingma & Ba (2017), the exponential average of $\gamma_t m_G$ yields an effective coefficient α :

$$\alpha_t = k_0 \left(1 - \frac{t}{T}\right) + \frac{k_0 \beta_1 (1 - t\beta_1^{t-1} + (t-1)\beta_1^t)}{T(1 - \beta_1)(1 - \beta_1^t)}. \quad (6)$$

The nonlinear term is monotonically increasing and is bounded by $\frac{k_0 \beta_1 (1 - t\beta_1^{t-1} + (t-1)\beta_1^t)}{T(1 - \beta_1)(1 - \beta_1^t)}$. Since the cumulative contribution of the global gradient should be comparable across domains, it suffices to choose k_0 inversely proportional to T . With $T \gg k_0$, the nonlinear term becomes negligible and we use the approximation $\alpha_t \approx k_0 (1 - \frac{t}{T})$. This schedule is simple to implement and robust. Hyperparameter details and sensitivity analyses are provided in the Appendix A.

For deployment efficiency, storing m_G in 32 bit precision is memory intensive. We therefore apply singular value decomposition to m_G and retain a rank $r = 512$ approximation. During training, we reconstruct the low rank estimate of the global gradient and add it to each step. This saves memory and emphasizes the most informative components of the global signal.

3.2 DYNAMIC CORRECTION

Pre-trained language models have demonstrated remarkable capabilities by incorporating data from diverse domains during pre-training. However, these models often struggle with domains that are either private or temporally distinct from the pre-training distribution, necessitating extensive experiments for optimal performance. Applying uniform strategies across domains with varying degrees of familiarity can lead to suboptimal outcomes. To address this challenge, we propose a Dynamic Correction (DC) mechanism that modulates the magnitude of parameter update based on gradient similarity. **The core intuition is to adjust the learning step based on the interference strength s_t . When $s_t \approx 0$, gradients are orthogonal (low interference), allowing larger steps. When $s_t \approx 1$, gradients are correlated (high interference), requiring smaller steps to preserve general knowledge.**

To quantify the alignment between general domain and specific domain, we introduce a gradient similarity metric based on the L2 norm of the normalized projection of specific gradient $g_{D,t}$ onto the estimation of the global gradient m_G :

$$s_t = \frac{\|g_{D,t} \cdot \hat{g}_G\|}{\|g_{D,t}\| \cdot \|\hat{g}_G\|}, \quad (7)$$

where $g_{D,t}$ denotes gradients for the specific domain at time step t , and \hat{g}_G represents global gradient of the general domain. Our empirical analysis encompasses both familiar domains (finance Yang et al. (2023), medicine Wang et al. (2024), and law Fei et al. (2023)) and a temporally restrictive domain (news QA, details shown in section 4.2). As shown in Figure 5, each domain maintains a characteristic similarity range with the general domain. Notably, familiar domains exhibit similar magnitude, while the unfamiliar domain demonstrates significantly lower similarity values, differing by more than an order of magnitude.

Leveraging this similarity measurement, we introduce a dynamic correction coefficient:

$$c_t = s_t + c_0, \quad (8)$$

where c_0 represents a base coefficient that prevents excessive parameter updates and potential overfitting when similarity values are minimal. We set $c_0 = 0.01$ in practice. The resulting parameter update rule incorporating the dynamic correction is:

$$\Delta\theta = -\eta \frac{\hat{m}_t}{\sqrt{c_t \hat{v}_t} + \varepsilon}. \quad (9)$$

324 3.3 ANNEALING LEARNING
325

326 In the training of MiniCPM Hu et al. (2024) and Llama3 Grattafiori et al. (2024), Annealing Learning
327 (AL) is applied at the final stage of pre-training. Using learning rate with minimal initialization and
328 decay strategy, LLMs can learn downstream task knowledge from high-quality domain data without
329 forgetting. Suppose the conventional initialization of learning rate is η_0 , and η_0^a for annealing, the
330 parameter updates for both schemes are:

$$331 \Delta\theta_t = -\eta_0 \left(1 - \frac{t}{T}\right) \frac{\hat{m}_t}{\sqrt{\hat{v}_t + \varepsilon}}, \quad \Delta\theta_t^a = -\eta_0^a \left(1 - \frac{t}{T}\right) \frac{\hat{m}_t}{\sqrt{\hat{v}_t + \varepsilon}}. \quad (10)$$

334 Via comparative analysis of Eq. 10, we can measure the influence of annealing on the parameter
335 updates:

$$336 \Delta\theta_t - \Delta\theta_t^a = -(\eta_0 - \eta_0^a) \left(1 - \frac{t}{T}\right) \frac{\hat{m}_t}{\sqrt{\hat{v}_t + \varepsilon}}. \quad (11)$$

339 As shown in Eq. 11, annealing suppresses the learning of specific domains. Smaller parameter
340 updates thus reduce the risk of catastrophic forgetting. Therefore, we use the annealing learning
341 scheme in DAA. And we set $\eta_0^a = 1e^{-7}$ in the our experiment.

342 3.4 SUMMARY
343

344 After integrating the above learning strategies, we obtain the complete parameter update of DAA:

$$345 \Delta\theta_t^{\text{DAA}} = -\eta_0^a \left(1 - \frac{t}{T}\right) \frac{\hat{m}_{B,t}}{\sqrt{c_t \hat{v}_{B,t} + \varepsilon}}. \quad (12)$$

348 The integration of **GGA**, DC and AL facilitates the adaptation to specific domains while mitigating
349 forgetting.

351 4 EXPERIMENT
352353 4.1 EXPERIMENT SETTINGS
354

355 This study evaluates the effectiveness of DAA across diverse vertical domains in both English and
356 Chinese contexts, including finance, medicine, and law. The general-domain data used in our ex-
357 periments comprises Chinese and English corpora covering multiple tasks. Its detailed composition
358 can be found in Appendix B, Table 6. The evaluation utilizes multiple datasets: FinGPT Yang et al.
359 (2023), CMB Wang et al. (2024) and Fuzi-Mingcha Deng et al. (2023). To avoid the potential con-
360 tamination or overfitting of evaluation benchmarks during pre-training as new and improved LLMs
361 are developed Schaeffer (2023); Jain et al. (2024); Zhang et al. (2024b), we constructed a temporal
362 out-of-distribution (OOD) evaluation benchmark named News QA (details in Section 4.2).

363 For comparative analysis, we selected several representative fine-tuning methods. In addition to
364 direct fine-tuning and vanilla fine-tuning, we also compared the performance of LoRAHu et al.
365 (2021), DoRALiu et al. (2024), GaloreZhao et al. (2024), and our proposed DAA across diverse
366 vertical domains. Especially, for vanilla fine-tuning, we followed Wen et al. (2023) and combined
367 our vertical domain fine-tuning experience to choose three distinct data mixture ratios (specific data
368 : general data = 1:1, 1:3, 1:5). We ensured that the vanilla fine-tuning results presented in the
369 experimental tables all represent the optimal performance in the specific domain.

370 In addition, we have validated the effectiveness of the proposed method across multiple foundational
371 models, including Llama3.1-8B Grattafiori et al. (2024), Phi4-14B Abdin et al. (2024), and Qwen3-
372 8B Yang et al. (2025).

374 4.2 NEWS QA BENCHMARK
375

376 We constructed a benchmark comprising QA pairs extracted from news articles. As shown in Table
377 2, the dataset contains 30,613 news titles across three categories (Politics, Economics, and Culture),
with corresponding true/false questions designed to evaluate factual verification capabilities. The

task requires binary responses (“true” or “false”) for each statement. To ensure minimal overlap with foundational models’ pre-training corpus Yang et al. (2024); Grattafiori et al. (2024); Abdin et al. (2024), we specifically selected news articles published after December 2024. As shown in Table 2, row 3, all the foundational models exhibits limited factual verification capabilities , achieving only 31.06% average accuracy.

Table 2: Details of news QA benchmark. The first two rows show the data distribution. The third row presents performance (S_D) of Qwen3-8B Yang et al. (2025) on each category.

Split	Politics	Econ	Culture	Total
Train set	9823	10120	8670	28613
Test set	700	700	600	2000
S_D	29.05	30.70	33.43	31.06

4.3 METRICS

To evaluate the general performance of the models, we selected four benchmarks commonly used across all LLMs: MMLU Hendrycks et al. (2021a), MMLU-Pro, GSM8K Cobbe et al. (2021), MATH Hendrycks et al. (2021b) and M3Exam Zhang et al. (2023a). MMLU tests general knowledge across multiple subjects, CMMLU focus on Chinese-specific knowledge and reasoning, while GSM8K and MATH tests mathematical problem-solving skills. All raw scores for the evaluated benchmarks were within the interval of [0, 1]. Crucially, no external normalization procedures were applied during the computation of S_D , S_G , and S . To evaluate various vertical performance of the models, we selected suitable public benchmarks for evaluation. For the financial domain, we utilized the weighted F1 score average across the English FPB Malo et al. (2013), FiQA Maia et al. (2018), TFNS Zer (2024), and NWGI Yang (2024) financial sentiment analysis test sets as the metric for this domain. For the medical domain, the accuracy score from the Chinese CMB-Exam Wang et al. (2024) test set served as the domain-specific metric. For the legal domain, we employed Chinese LawBench Fei et al. (2023) for a comprehensive evaluation. Our overall metric design is structured as follows:

$$S_G = \text{Mean}(\{S_x \mid x \in \mathcal{X}\}), \quad (13)$$

$$S = \text{HarmonicMean}(S_D, S_G), \quad (14)$$

where $\mathcal{X} = \{\text{MMLU, MMLU-Pro, GSM8k, MATH, M3Exam}\}$. S_D is the score of the model’s vertical domain performance, S_G is the average score of the model’s general performance. S is the harmonic mean of S_D and S_G , meaning the model scores high only if both are large. If fine-tuning boosts domain performance but reduces general capability significantly, the score nears 0. Conversely, if it enhances domain performance while preserving general capability, the score approaches 1.

4.4 COST ANALYSIS

To demonstrate the efficiency of our approach, we compared GPU hours across fine-tuning methods using 16 Nvidia A100 GPUs. As shown in Table 3, DAA requires $T_{\text{DAA}} \approx 4.2$ GPU-hours per domain, similar to direct full-tuning ($T_{\text{Direct}} \approx 4.1$ GPU-hours) but without its drop in general-task performance. More importantly, DAA reduces training costs by over 90% compared to vanilla fine-tuning ($T_{\text{Vanilla}} \approx 46.7$ GPU-hours) while achieving notable gains in vertical (S_D) and general (S_G) scores.

While LoRA ($T_{\text{LoRA}} \approx 3.0$ GPU-hours) and DoRA ($T_{\text{DoRA}} \approx 3.1$ GPU-hours) are faster, DAA consistently outperforms them and Galore ($T_{\text{Galore}} \approx 5.1$ GPU-hours) in harmonic-mean score S , justifying the modest additional GPU time with superior task performance.

4.5 MAIN RESULTS

Table 3 demonstrates that our DAA method consistently achieves the best balance between vertical domain ability and general performance retention, as measured by the harmonic mean S , across four domains and three base models.

432

433
434
435
Table 3: Performance metrics across different domains and models. T is the GPU hours. S_D is the
score of the model’s vertical domain performance, S_G is the average score of the model’s general
performance change. S is the harmonic mean of S_D and S_G .

436

437 Domain	438 Method	439 Llama3.1				440 Phi4				441 Qwen3			
		$T \downarrow$	$S_D \uparrow$	$S_G \uparrow$	$S \uparrow$	$T \downarrow$	$S_D \uparrow$	$S_G \uparrow$	$S \uparrow$	$T \downarrow$	$S_D \uparrow$	$S_G \uparrow$	$S \uparrow$
442 Finance	Direct FT	3.40	80.01	54.69	64.97	6.22	89.72	77.92	83.41	3.38	85.83	63.74	73.15
	Vanilla FT	38.77	79.30	60.50	68.64	71.24	83.42	77.71	80.46	38.74	85.49	71.01	77.58
	LoRA	2.50	76.45	61.69	68.28	4.61	87.13	78.27	82.46	2.50	81.68	73.21	77.21
	DoRA	2.58	76.12	61.25	67.88	4.71	86.34	78.24	82.09	2.57	82.19	73.27	77.47
	Galore	5.09	77.31	60.78	68.06	9.34	86.23	78.55	82.21	5.09	84.37	74.59	79.18
	DAA (Ours)	3.45	79.84	61.75	69.64	6.35	87.73	78.50	82.86	3.41	85.32	76.49	80.66
443 Medicine	Direct FT	12.39	89.23	52.73	66.29	22.78	92.13	78.34	84.68	12.38	92.67	64.44	76.02
	Vanilla FT	141.81	87.32	59.47	70.75	260.58	91.24	79.26	84.83	141.81	81.81	68.74	74.71
	LoRA	9.18	81.76	59.97	69.19	16.87	90.30	79.02	84.28	8.21	84.00	69.51	76.07
	DoRA	9.31	81.21	59.07	68.40	17.05	90.31	78.74	84.13	8.37	84.74	69.67	76.47
	Galore	13.68	81.23	58.55	68.05	25.14	91.96	78.56	84.73	13.68	87.57	73.71	80.04
	DAA (Ours)	12.54	83.97	60.33	70.22	22.90	92.61	78.82	85.16	12.64	92.24	77.97	84.51
444 Law	Direct FT	4.39	56.81	48.94	52.58	8.04	42.63	71.06	53.29	4.39	55.28	70.13	61.83
	Vanilla FT	50.12	51.37	53.08	52.21	91.98	41.82	72.12	52.94	50.10	52.28	72.95	60.91
	LoRA	3.25	46.58	58.06	51.69	5.97	41.87	73.38	53.32	3.25	51.90	76.19	61.74
	DoRA	3.43	46.37	56.05	50.75	6.24	41.98	72.36	53.13	3.36	51.91	76.30	61.79
	Galore	6.13	47.80	56.13	51.63	11.26	40.12	71.53	51.41	6.13	52.76	77.53	62.79
	DAA (Ours)	4.53	49.93	56.68	53.09	8.17	41.95	73.12	53.31	4.60	52.79	79.38	63.41
445 News QA	Direct FT	1.27	82.38	36.34	50.44	2.32	87.38	26.40	40.55	1.27	79.37	3.83	7.31
	Vanilla FT	14.45	80.62	39.73	53.22	16.51	88.23	72.11	79.36	14.45	80.27	52.36	63.38
	LoRA	0.94	71.23	48.04	57.38	1.72	83.12	77.46	80.19	0.93	70.27	67.51	68.86
	DoRA	0.99	73.71	50.52	59.95	1.80	83.37	76.68	79.88	0.99	70.78	67.81	69.26
	Galore	2.02	79.13	53.05	63.52	3.70	85.02	76.27	80.41	2.03	80.88	68.43	74.14
	DAA (Ours)	1.32	82.37	52.00	63.76	2.42	89.19	77.91	83.17	1.28	80.82	68.52	74.16

456

457

458
459
460
461
462
463
Finance. On Llama3.1, DAA again leads with $S_D = 79.84\%$, $S_G = 61.75\%$ and $S = 69.64\%$,
surpassing all competitors. On Qwen3, direct fine-tuning and vanilla fine-tuning suffer considerable
general-performance drops ($S_G = 63.74\%$ and 71.01%), whereas DAA attains $S_G = 76.49\%$
(an improvement of 1.90 points over the next best) while maintaining a high domain score $S_D =$
 85.32% . This yields the highest overall score $S = 80.66\%$, outperforming direct fine-tuning ($S =$
 73.15%) and vanilla fine-tuning ($S = 77.58\%$).464
465
466
467
468
469
Medicine. On Llama3.1, DAA’s $S_G = 60.33\%$ and $S_D = 83.97\%$ produce $S = 70.22\%$, again the
best trade-off. On Phi4, DAA secures the highest domain accuracy ($S_D = 92.61\%$) and a strong
general score ($S_G = 78.82\%$), leading to an overall $S = 85.16\%$, which exceeds every baseline.
For Qwen3, direct fine-tuning and vanilla fine-tuning obtain only $S_G = 64.44\%$ and 68.74% , while
DAA achieves $S_G = 77.97\%$ coupled with $S_D = 92.24\%$, resulting in the top harmonic mean
 $S = 84.51\%$.470
471
472
473
474
Law. DAA attains the high S_G on Llama3.1 (56.68%), Phi4 (73.12%), and Qwen3 (79.38%), and
achieves harmonic means 53.09%, 53.31%, and $S = 63.41\%$ respectively. These results outperform
direct fine-tuning and vanilla fine-tuning, both of which incur larger general-performance regres-
sions. A more detailed discussion regarding the performance on the Law dataset is provided in the
Appendix C.475
476
477
478
479
News QA. In this strictly leak-free benchmark, direct fine-tuning collapses on general performance
($S_G = 3.83\%$ on Qwen3), while DAA preserves general knowledge ($S_G = 68.52\%$) while slightly
exceeding the domain score of direct fine-tuning ($S_D = 80.82\%$ vs. 79.37%), producing $S =$
 74.16% (versus 7.31%). On Phi4, DAA simultaneously achieves the highest $S_D = 89.19\%$ and
 $S_G = 77.91\%$, leading to $S = 83.17\%$, which outperforms the best baseline by 2.76 points.480
481
482
Overall, across all domains and models, DAA delivers the strongest joint performance S , validating
its effectiveness at vertical domain fine-tuning with minimal general-knowledge degradation.483
484
485
It is worth noting that the performance gains of DBA are more pronounced on the News QA bench-
mark compared to traditional domains like Finance or Medicine. We attribute this to the strong
pre-training coverage of the latter domains in modern base models, where fine-tuning primarily
serves as alignment. In contrast, News QA represents a Temporal Out-Of-Distribution (OOD) task

486

487

Table 4: Results of ablation on news QA benchmark. AL, **GGA** and DC are defined in Section 3.

488

489

490

491

492

493

494

495

496

497

498

499

500

501

containing knowledge strictly unseen during pre-training. The significant margin achieved by DBA on News QA demonstrates its capability, validating its effectiveness beyond simple style alignment.

4.6 ABLATION ANALYSIS

Furthermore, we conducted ablation studies on individual modules within the proposed DAA to quantify their contributions. The results are shown in Table 4. When solely applying annealing learning (row 1), the model shows decreased domain-specific performance and improved general domain performance, yet fails to match the overall effectiveness of DM. This indicates that while the annealing strategy helps mitigate catastrophic forgetting, its effectiveness is limited in isolation. The incorporation of **global gradient anchoring** with annealing learning (row 2) leads to enhanced performance in both domain-specific and general domains, demonstrating the significant impact of global gradient optimization.

Incorporating dynamic correction into annealing learning (row 3) leads to significant improvements in both domain-specific and general domain performance. This demonstrates that dynamic correction effectively optimizes the update step size, thereby enhancing the learning process. The combination of all three components (row 4) - annealing learning, **global gradient anchoring**, and dynamic correction - yields optimal performance across both domains, achieving highest joint performance S of 58.05%, 75.64%, and 74.16% on Llama3.1, Phi4, and Qwen3. These results validate the synergistic effects of DAA components in enhancing the model’s overall capabilities.

5 CONCLUSION

We present Dynamic Anchor Annealing, a fine-tuning method that mitigates catastrophic forgetting in LLMs. Using **global gradient anchoring** with similarity guided dynamic correction, DAA improves performance while reducing compute cost over prior methods.

Limitations. DAA is designed for dense models used in vertical domain tasks. Our experiments cover a few domains such as medical and finance. Robustness across vision, speech, reinforcement learning, continual fine-tuning, and large scale language modeling remains unverified. Although DAA scales linearly in theory, extremely deep or wide networks with billions of parameters and web scale datasets may reveal stability or convergence issues not seen in our mid scale benchmarks.

Applicability Analysis. DAA relies on gradient **anchoring** learning and magnitude adjustment of parameter updates, so it applies to other optimizers. In fine-tuning we focus on AdamW Loshchilov & Hutter (2019), which is widely used.

Future Work. Domain specific LLMs can equip workers with specialized AI in their fields. We will explore broader applications of DAA to inspire research on domain specific training. We will release code and associated global gradients, followed by additional global gradients matched to more base models for the community. **Future work will also explore integrating DAA with Parameter-Efficient Fine-Tuning (PEFT) methods**, such as LoRA Hu et al. (2022). By applying the Global Gradient Anchoring and Dynamic Correction mechanisms specifically to the gradients of low-rank adapters, we aim to synergize the computational efficiency of DAA, which eliminates the need for data mixture, with the memory efficiency of PEFT. This combination holds the promise of enabling robust, low-resource domain adaptation on consumer-grade hardware.

540 REFERENCES
541

542 Zeroshot/twitter-financial-news-sentiment . Datasets at Hugging Face.
543 2024. URL [https://huggingface.co/datasets/zeroshot/
544 twitter-financial-news-sentiment](https://huggingface.co/datasets/zeroshot/twitter-financial-news-sentiment).

545 Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar,
546 Michael Harrison, Russell J Hewett, Mojan Javaheripi, Piero Kauffmann, et al. Phi-4 techni-
547 cal report. *arXiv preprint arXiv:2412.08905*, 2024.

548 Chirag Agarwal, Daniel D’souza, and Sara Hooker. Estimating example difficulty using variance of
549 gradients. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recog-
550 nition*, pp. 10368–10378, 2022.

551 Zhijie Bao, Wei Chen, Shengze Xiao, Kuang Ren, Jiaao Wu, Cheng Zhong, Jiajie Peng, Xuanjing
552 Huang, and Zhongyu Wei. DISC-MedLLM: Bridging General Large Language Models and Real-
553 World Medical Consultation. 2023. doi: 10.48550/arXiv.2308.14346. URL <http://arxiv.org/abs/2308.14346>.

554 Junying Chen, Xidong Wang, Ke Ji, Anningzhe Gao, Feng Jiang, Shunian Chen, Hongbo Zhang,
555 Dingjie Song, Wenya Xie, Chuyi Kong, Jianquan Li, Xiang Wan, Haizhou Li, and Benyou Wang.
556 HuatuoGPT-II, One-stage Training for Medical Adaption of LLMs. 2024. doi: 10.48550/arXiv.
557 2311.09774. URL <http://arxiv.org/abs/2311.09774>.

558 Sanyuan Chen, Yutai Hou, Yiming Cui, Wanxiang Che, Ting Liu, and Xiangzhan Yu. Recall and
559 Learn: Fine-tuning Deep Pretrained Language Models with Less Forgetting. 2020. doi: 10.48550/
560 arXiv.2004.12651. URL <http://arxiv.org/abs/2004.12651>.

561 Wei Chen, Qiushi Wang, Zefei Long, Xianyin Zhang, Zhongtian Lu, Bingxuan Li, Siyuan Wang,
562 Jiarong Xu, Xiang Bai, Xuanjing Huang, and Zhongyu Wei. DISC-FinLLM: A Chinese Financial
563 Large Language Model based on Multiple Experts Fine-tuning. 2023. doi: 10.48550/arXiv.2310.
564 15205. URL <http://arxiv.org/abs/2310.15205>.

565 Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich. Gradnorm: Gradient
566 normalization for adaptive loss balancing in deep multitask networks. In *International conference
567 on machine learning*, pp. 794–803. PMLR, 2018.

568 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
569 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John
570 Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*,
571 2021.

572 Jiaxi Cui, Zongjian Li, Yang Yan, Bohua Chen, and Li Yuan. ChatLaw: Open-Source Legal Large
573 Language Model with Integrated External Knowledge Bases. 2023. doi: 10.48550/arXiv.2306.
574 16092. URL <http://arxiv.org/abs/2306.16092>.

575 Wentao Deng, Jiahuan Pei, Keyi Kong, Zhe Chen, Furu Wei, Yujun Li, Zhaochun Ren, Zhumin
576 Chen, and Pengjie Ren. Syllogistic Reasoning for Legal Judgment Analysis. In Houda Bouamor,
577 Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empirical Methods
578 in Natural Language Processing*, pp. 13997–14009. Association for Computational Linguistics,
579 2023. doi: 10.18653/v1/2023.emnlp-main.864. URL <https://aclanthology.org/2023.emnlp-main.864/>.

580 Guanting Dong, Hongyi Yuan, Keming Lu, Chengpeng Li, Mingfeng Xue, Dayiheng Liu, Wei
581 Wang, Zheng Yuan, Chang Zhou, and Jingren Zhou. How abilities in large language models are
582 affected by supervised fine-tuning data composition. In *Proceedings of the 62nd Annual Meeting
583 of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 177–198, 2024.

584 Zhiwei Fei, Xiaoyu Shen, Dawei Zhu, Fengzhe Zhou, Zhuo Han, Songyang Zhang, Kai Chen,
585 Zongwen Shen, and Jidong Ge. LawBench: Benchmarking Legal Knowledge of Large Language
586 Models. 2023. doi: 10.48550/arXiv.2309.16289. URL <http://arxiv.org/abs/2309.16289>.

594 Saeed Ghadimi and Guanghui Lan. Stochastic first-and zeroth-order methods for nonconvex stochas-
 595 tic programming. *SIAM journal on optimization*, 23(4):2341–2368, 2013.
 596

597 Aaron Grattafiori, Abhimanyu Dubey, and Jauhri et al. The Llama 3 Herd of Models. 2024. doi:
 598 10.48550/arXiv.2407.21783. URL <http://arxiv.org/abs/2407.21783>.

599 Mert Gurbuzbalaban, Umut Simsekli, and Lingjiong Zhu. The heavy-tail phenomenon in sgd. In
 600 *International Conference on Machine Learning*, pp. 3964–3975. PMLR, 2021.
 601

602 Raia Hadsell, Dushyant Rao, Andrei A Rusu, and Razvan Pascanu. Embracing change: Continual
 603 learning in deep neural networks. *Trends in cognitive sciences*, 24(12):1028–1040, 2020.

604 William Held, Bhargavi Paranjape, Punit Singh Koura, Mike Lewis, Frank Zhang, and Todor
 605 Mihaylov. Optimizing Pretraining Data Mixtures with LLM-Estimated Utility, 2025. URL
 606 <http://arxiv.org/abs/2501.11747>.
 607

608 Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob
 609 Steinhardt. Measuring massive multitask language understanding. *Proceedings of the Interna-*
 610 *tional Conference on Learning Representations (ICLR)*, 2021a.

611 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,
 612 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv*
 613 *preprint arXiv:2103.03874*, 2021b.
 614

615 Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,
 616 and Weizhu Chen. LoRA: Low-Rank Adaptation of Large Language Models. 2021. doi: 10.
 617 48550/arXiv.2106.09685. URL <http://arxiv.org/abs/2106.09685>.

618 Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,
 619 Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022.

620

621 Shengding Hu, Yuge Tu, Xu Han, Chaoqun He, Ganqu Cui, Xiang Long, Zhi Zheng, Yewei Fang,
 622 Yuxiang Huang, Weilin Zhao, Xinrong Zhang, Zheng Leng Thai, Kaihuo Zhang, Chongyi Wang,
 623 Yuan Yao, Chenyang Zhao, Jie Zhou, Jie Cai, Zhongwu Zhai, Ning Ding, Chao Jia, Guoyang
 624 Zeng, Dahai Li, Zhiyuan Liu, and Maosong Sun. MiniCPM: Unveiling the Potential of Small
 625 Language Models with Scalable Training Strategies. 2024. doi: 10.48550/arXiv.2404.06395.
 626 URL <http://arxiv.org/abs/2404.06395>.

627 Quzhe Huang, Mingxu Tao, Chen Zhang, Zhenwei An, Cong Jiang, Zhibin Chen, Zirui Wu, and
 628 Yansong Feng. Lawyer LLaMA Technical Report. 2023. doi: 10.48550/arXiv.2305.15062. URL
 629 <http://arxiv.org/abs/2305.15062>.
 630

631 Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando
 632 Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free
 633 evaluation of large language models for code. *arXiv preprint arXiv:2403.07974*, 2024.

634 Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. 2017. doi:
 635 10.48550/arXiv.1412.6980. URL <http://arxiv.org/abs/1412.6980>.
 636

637 Tomasz Korbak, Hady Elsahar, German Kruszewski, and Marc Dymetman. Controlling Condi-
 638 tional Language Models without Catastrophic Forgetting. In *Proceedings of the 39th Inter-*
 639 *national Conference on Machine Learning*, pp. 11499–11528. PMLR, 2022. URL <https://proceedings.mlr.press/v162/korbak22a.html>.
 640

641 Xize Liang, Lin Yang, Jie Wang, Yiyang Lu, Runyu Wu, Hanzhu Chen, and Jianye Hao. Boost-
 642 ing multi-domain fine-tuning of large language models through evolving interactions between
 643 samples. In *Forty-second International Conference on Machine Learning*, 2025.

644

645 Yong Lin, Lu Tan, Hangyu Lin, Zeming Zheng, Renjie Pi, Jipeng Zhang, Shizhe Diao, Haoxiang
 646 Wang, Han Zhao, Yuan Yao, and Tong Zhang. Speciality vs Generality: An Empirical Study on
 647 Catastrophic Forgetting in Fine-tuning Foundation Models, 2023. URL <http://arxiv.org/abs/2309.06256>.

648 Bo Liu, Xingchao Liu, Xiaojie Jin, Peter Stone, and Qiang Liu. Conflict-averse gradient descent
 649 for multi-task learning. *Advances in Neural Information Processing Systems*, 34:18878–18890,
 650 2021.

651

652 Hong Liu, Mingsheng Long, Jianmin Wang, and Michael Jordan. Transferable adversarial training:
 653 A general approach to adapting deep classifiers. In *International conference on machine learning*,
 654 pp. 4013–4022. PMLR, 2019.

655 Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Frank Wang, Kwang-
 656 Ting Cheng, and Min-Hung Chen. Dora: Weight-decomposed low-rank adaptation. In *Forty-first*
 657 *International Conference on Machine Learning*, 2024.

658

659 Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization, 2019. URL <http://arxiv.org/abs/1711.05101>.

660

661 James Lucas, Juhan Bae, Michael R Zhang, Stanislav Fort, Richard Zemel, and Roger Grosse.
 662 Analyzing monotonic linear interpolation in neural network loss landscapes. *arXiv preprint*
 663 *arXiv:2104.11044*, 1, 2021.

664

665 Yun Luo, Zhen Yang, Fandong Meng, Yafu Li, Jie Zhou, and Yue Zhang. An Empirical Study of
 666 Catastrophic Forgetting in Large Language Models During Continual Fine-tuning, 2025. URL
 667 <http://arxiv.org/abs/2308.08747>.

668

669 Macedo Maia, Siegfried Handschuh, André Freitas, Brian Davis, Ross McDermott, Manel Zarrouk,
 670 and Alexandra Balahur. WWW’18 Open Challenge: Financial Opinion Mining and Question
 671 Answering. In *Companion of the The Web Conference 2018 on The Web Conference 2018 -*
 672 *WWW ’18*, pp. 1941–1942. ACM Press, 2018. ISBN 978-1-4503-5640-4. doi: 10.1145/3184558.
 3192301. URL <http://dl.acm.org/citation.cfm?doid=3184558.3192301>.

673

674 Pekka Malo, Ankur Sinha, Pyry Takala, Pekka Korhonen, and Jyrki Wallenius. Good Debt or Bad
 675 Debt: Detecting Semantic Orientations in Economic Texts. 2013. doi: 10.48550/arXiv.1307.5336.
 676 URL <http://arxiv.org/abs/1307.5336>.

677

678 Rylan Schaeffer. Pretraining on the test set is all you need. *arXiv preprint arXiv:2309.08632*, 2023.

679

680 Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization. In *Advances*
 681 *in neural information processing systems*, volume 31, 2018.

682

683 Xidong Wang, Guiming Hardy Chen, Dingjie Song, Zhiyi Zhang, Zhihong Chen, Qingying Xiao,
 684 Feng Jiang, Jianquan Li, Xiang Wan, Benyou Wang, and Haizhou Li. CMB: A Comprehensive
 685 Medical Benchmark in Chinese. 2024. doi: 10.48550/arXiv.2308.08833. URL <http://arxiv.org/abs/2308.08833>.

686

687 Cheng Wen, Xianghui Sun, Shuaijiang Zhao, Xiaoquan Fang, Liangyu Chen, and Wei Zou. ChatH-
 688 ome: Development and Evaluation of a Domain-Specific Language Model for Home Renovation.
 689 2023. doi: 10.48550/arXiv.2307.15290. URL <http://arxiv.org/abs/2307.15290>.

690

691 Minghao Wu, Thuy Vu, Lizhen Qu, and Reza Haf. Mixture-of-skills: Learning to optimize data
 692 usage for fine-tuning large language models. In *Proceedings of the 2024 Conference on Empirical*
 693 *Methods in Natural Language Processing*, pp. 14226–14240, 2024.

694

695 Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabrowski, Mark Dredze, Sebastian Gehrmann, Prab-
 696 hanjan Kambadur, David Rosenberg, and Gideon Mann. BloombergGPT: A Large Language
 697 Model for Finance. 2023. doi: 10.48550/arXiv.2303.17564. URL <http://arxiv.org/abs/2303.17564>.

698

699 Mengzhou Xia, Sadhika Malladi, Suchin Gururangan, Sanjeev Arora, and Danqi Chen. Less: Se-
 700 lecting influential data for targeted instruction tuning. *arXiv preprint arXiv:2402.04333*, 2024.

701

702 Chaojun Xiao, Xueyu Hu, Zhiyuan Liu, Cunchao Tu, and Maosong Sun. Lawformer: A pre-
 703 trained language model for Chinese legal long documents. 2:79–84, 2021. ISSN 2666-
 704 6510. doi: 10.1016/j.aiopen.2021.06.003. URL <https://www.sciencedirect.com/science/article/pii/S2666651021000176>.

702 Qianqian Xie, Weiguang Han, Xiao Zhang, Yanzhao Lai, Min Peng, Alejandro Lopez-Lira, and
 703 Jimin Huang. PIXIU: A Comprehensive Benchmark, Instruction Dataset and Large Language
 704 Model for Finance. 36:33469–33484, 2023.

705 An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li,
 706 Dayiheng Liu, Fei Huang, Haoran Wei, et al. Qwen2. 5 technical report. *arXiv preprint*
 707 *arXiv:2412.15115*, 2024.

708 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu,
 709 Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint*
 710 *arXiv:2505.09388*, 2025.

711 Hongyang Yang. AI4Finance-Foundation/FinGPT. 2024. URL <https://github.com/AI4Finance-Foundation/FinGPT>.

712 Hongyang Yang, Xiao-Yang Liu, and Christina Dan Wang. FinGPT: Open-Source Financial Large
 713 Language Models. 2023. doi: 10.48550/arXiv.2306.06031. URL <http://arxiv.org/abs/2306.06031>.

714 Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn.
 715 Gradient surgery for multi-task learning. *Advances in neural information processing systems*, 33:
 716 5824–5836, 2020.

717 Shengbin Yue, Wei Chen, Siyuan Wang, Bingxuan Li, Chenchen Shen, Shujun Liu, Yuxuan Zhou,
 718 Yao Xiao, Song Yun, Xuanjing Huang, and Zhongyu Wei. DISC-LawLLM: Fine-tuning Large
 719 Language Models for Intelligent Legal Services. 2023. doi: 10.48550/arXiv.2309.11325. URL
 720 <http://arxiv.org/abs/2309.11325>.

721 Biao Zhang, Zhongtao Liu, Colin Cherry, and Orhan Firat. When Scaling Meets LLM Finetuning:
 722 The Effect of Data, Model and Finetuning Method, 2024a. URL <http://arxiv.org/abs/2402.17193>.

723 Hugh Zhang, Jeff Da, Dean Lee, Vaughn Robinson, Catherine Wu, Will Song, Tiffany Zhao, Pranav
 724 Raja, Dylan Slack, Qin Lyu, et al. A careful examination of large language model performance
 725 on grade school arithmetic. *arXiv preprint arXiv:2405.00332*, 2024b.

726 Wenzuan Zhang, Mahani Aljunied, Chang Gao, Yew Ken Chia, and Lidong Bing. M3exam: A
 727 multilingual, multimodal, multilevel benchmark for examining large language models. *Advances*
 728 *in Neural Information Processing Systems*, 36:5484–5505, 2023a.

729 Xuanyu Zhang, Qing Yang, and Dongliang Xu. XuanYuan 2.0: A Large Chinese Financial Chat
 730 Model with Hundreds of Billions Parameters. 2023b. doi: 10.48550/arXiv.2305.12002. URL
 731 <http://arxiv.org/abs/2305.12002>.

732 Jiawei Zhao, Zhenyu Zhang, Beidi Chen, Zhangyang Wang, Anima Anandkumar, and Yuandong
 733 Tian. Galore: Memory-efficient ILM training by gradient low-rank projection. *arXiv preprint*
 734 *arXiv:2403.03507*, 2024.

735 Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia
 736 Efrat, Ping Yu, Lili Yu, et al. Lima: Less is more for alignment. *Advances in Neural Information*
 737 *Processing Systems*, 36:55006–55021, 2023.

738 Zhi Zhou, Jiang-Xin Shi, Peng-Xiao Song, Xiao-Wen Yang, Yi-Xuan Jin, Lan-Zhe Guo, and Yu-
 739 Feng Li. LawGPT: A Chinese Legal Knowledge-Enhanced Large Language Model. 2024. doi:
 740 10.48550/arXiv.2406.04614. URL <http://arxiv.org/abs/2406.04614>.

741 Jingwei Zhu, Minghuan Tan, Min Yang, Ruixue Li, and Hamid Alinejad-Rokny. CollectiveSFT:
 742 Scaling Large Language Models for Chinese Medical Benchmark with Collective Instructions in
 743 Healthcare. 2024. doi: 10.48550/arXiv.2407.19705. URL <http://arxiv.org/abs/2407.19705>.

744

745

746

747

748

749

750

751

752

753

754

755

756 Finetune Once: Decoupling General & Domain Learning 757 758 with Dynamic Anchor Annealing 759

760 Appendix 761

762 This Appendix contains the following parts:
763

- 764 • **Hyper Parameters.** We delineate the specific hyperparameters for model training and eval-
765 uation, detailing the settings for gradient expectation estimation, momentum compression,
766 the AdamW optimizer, and the empirical justification for the anchoring learning coefficient
767 k_0 .
- 768 • **Dataset Details.** We provide a comprehensive description of the datasets utilized for both
769 general and vertical domain fine-tuning, detailing the specific sources, composition, and
770 quantities for the finance, medicine, law, and the constructed temporal out-of-distribution
771 News QA domains.
- 772 • **Performance on the Law Dataset.** We provide a contextual analysis of the performance on
773 the Law dataset, attributing the lower absolute scores to the domain’s complex and hetero-
774 geneous task mixture while underscoring the robustness of the DAA method in achieving
775 superior relative performance.
- 776 • **Practical Implementation Guide.** We outline a two-stage practical implementation guide
777 for DAA, involving a one-time, reusable pre-computation of the global gradient and its
778 subsequent integration into standard fine-tuning frameworks to ensure efficiency and ease
779 of adoption.
- 780 • **Significance Test.** We provide statistical validation of our performance gains.
- 781 • **Cost Analysis.** We provide a detailed cost comparison including memory usage and train-
782 ing time.
- 783 • **Additional Sensitivity Analysis.** We analyze the sensitivity of our method to data quantity
784 and anchor source.
- 785 • **Raw Experimental Results.** We report the raw scores for general benchmarks.

787 A HYPER PARAMETERS 788

789 This section will introduce the detailed process and hyperparameters involved in model training and
790 testing. In the main experiments and ablation experiments, we chose Qwen2-7B as our base model.
791 To obtain the gradient expectation estimation of the general domain, we set the learning rate of the
792 general domain training $\eta_G = 0$, meaning no parameter updates are performed in the general do-
793 main. Additionally, $\beta_1 = 0.999$, so the momentum approximates the gradient expectation. The
794 training batch size is 8, and only the gradient momentum is retained after training. Note that the
795 computation in the general domain only needs to be done **once**, and the same momentum is used for
796 different vertical domains subsequently. Since the original momentum is in F32 data format, loading
797 it directly into the GPU memory would occupy a large space. We performed singular value decom-
798 position on the momentum, retaining $r = 512$ dimensions of singular values and vectors. During
799 the **global gradient anchoring** in training, the low-rank approximation of the original momentum is
800 restored and then added to the gradient. In the fine-tuning phase, we set the initial learning rate
801 $\eta_D = 1e-7$, which is much lower than the usual fine-tuning learning rate. We used a linear decay
802 to zero learning rate schedule without warmup. The training batch size is 8, and we train for only
803 one epoch. We use the AdamW optimizer with $\beta_1 = 0.9$ and $\beta_2 = 0.95$. For the **global gradient**
804 **anchoring** coefficients defined in equations (7) and (8), we chose a linear decay scheme with
805 $k_0 = 200/T$, where T is the total number of steps in vertical domain fine-tuning. **We propose the**
806 **heuristic $k_0 \propto 1/T$ as a practical guideline.** This relationship maintains a consistent total contribu-
807 tion from global gradients irrespective of the training duration. Specifically, the total ‘regularization
808 influence’ exerted by \hat{g}_G can be approximated by the integral of the anchoring magnitude γ_t over
809 time: $\int_0^T [k_0(1 - t/T)]dt = k_0 \cdot T/2$. To ensure this total contribution remains a constant C inde-
810 pendent of the total steps T , k_0 must be proportional to $1/T$. This ensures that whether fine-tuning
811 for 1,000 or 5,000 steps, the total ‘pull’ from the global anchor remains consistent.

810
811 Table 5: Extended sensitivity analysis of hyperparameter k_0 on the NewsQA benchmark.
812

k_0	S
$50/T$	57.23
$100/T$	60.13
$150/T$	61.98
$175/T$	62.92
$200/T$	63.76
$225/T$	63.28
$250/T$	63.50
$275/T$	63.12
$300/T$	63.39

825 This is analogous to tuning LoRA, where practitioners often fix the dropout rate and primarily ex-
826periment with the rank (r) and scaling factor (α). In our case, the core tuning effort is simplified to a
827 single, well-behaved parameter governed by a clear rule. As shown in the Table 5, the performance
828 metric S of Llama3.1 on NewsQA improves as the hyperparameter k_0 increases. However, this
829 growth plateaus after k_0 reaches $200/T$. Since there is no significant performance gain beyond this
830 point, we select $k_0 = 200/T$ as the value for our experiments.

832 B DATASET DETAILS

834 We obtained validated our proposed method across a wide range of vertical domains, covering fi-
835nance, medicine, law and news QA.

837 **General Data:** Since the vertical domain tasks mainly cover Chinese and English languages and
838 include multiple-choice and conversational tasks, the general data needs to fully cover similar data
839 patterns. Therefore, we collected Chinese and English QA data, covering QA, conversations, and
840 multiple-choice questions. Specifically, the general data includes 54,042 Chinese QA conversation
841 pairs, 65,596 English QA conversation pairs, and 1,881 Chinese multiple-choice questions.

843 Table 6: Data sources and quantities

NAME	SOURCE	QUANTITY
CHINESE QA DATA	SELF-BUILT	54,042
ENGLISH QA DATA	SELF-BUILT	65,596
CHINESE MCQs	SELF-BUILT	1,881

850 **Finance:** We referred to the training data and testing methods of FinGPT Yang et al. (2023), se-
851lecting its sentiment analysis task as the financial vertical domain. This task requires the model to
852analyze the market sentiment of the input text as negative, neutral, or positive. According to Yang
853 et al. (2023), the training data was collected from FPB Malo et al. (2013), FiQA Maia et al. (2018),
854 TFNS Zer (2024), and NWGI Yang (2024). FinGPT designed three types of instructions for each
855 original data, resulting in a total of 76,772 training samples after filtering.

857 Table 7: Data sources and their quantities.

NAME	SOURCE	QUANTITY
ENGLISH SENTIMENT DATA	FPB	12,122
	FiQA	26,532
	TFNS	12,731
	NWGI	25,387

864 **Medicine:** We chose the CMB-Exam from the Chinese medicine Benchmark (CMB) Wang et al.
 865 (2024) as the medical domain. This dataset includes 280,839 medicine multiple-choice questions,
 866 covering 124,926 physician questions, 16,919 nursing questions, 27,004 medicine technician ques-
 867 tions, 33,354 pharmacist questions, 62,271 undergraduate exam questions, and 16,365 graduate en-
 868 trance exam questions. We randomly selected 11,200 questions from each category as the test set,
 869 with a total of 269,359 questions in the training set.

870
871 Table 8: Questions from various Chinese medicine exams.
872

NAME	SOURCE	QUANTITY
PHYSICIAN	PHYSICIAN EXAM	124,926
NURSING	NURSING EXAM	16,919
TECHNICIAN	TECHNICIAN EXAM	27,004
PHARMACIST	PHARMACIST EXAM	33,354
UNDERGRADUATE	MEDICINE EXAM	62,271
GRADUATE ENTRANCE	MEDICINE EXAM	16,365

880
881 Table 9: Law Data Statistics
882

NAME	SOURCE	QUANTITY
FACT RECALL	CAIL-LONG	4,200
CASE SUMMARIZATION	CAIL-LONG LAWGPT	5,750 35,000
LEGAL QA DATA	LAWYER LLAMA FUZI	11,000 32,050
SYLLOGISTIC REASONING	FUZI	11,237

891 **Law:** We referred to the data summarized by the Fuzi-Mingcha Deng et al. (2023) to filter suitable
 892 legal vertical fine-tuning data. The fine-tuning data composition is as follows: 4,200 recall data and
 893 5,750 summarization data from CAIL-Long Xiao et al. (2021), 35,000 legal QA data from LawGPT
 894 Zhou et al. (2024), 11,000 legal QA data from Lawyer Llama Huang et al. (2023), 32,050 legal QA
 895 data and 11,237 syllogistic reasoning judgment data independently constructed by Fuzi-Mingcha
 896 Deng et al. (2023). The total training data amounts to 99,237 samples.

897 **News QA:** To precisely evaluate the domain decoupling capabilities, we constructed a temporal
 898 out-of-distribution evaluation benchmark comprising QA pairs derived from news articles published
 899 after December 2024 for ablation study. We used Qwen2.5-72B Yang et al. (2024) to extract three
 900 factual QA questions for each headline. We ensured that there is no overlap between the vertical
 901 domain data and the general data.

902 The above datasets come from diverse sources, and the characteristics and distributions among the
 903 datasets vary significantly, providing ample and credible test scenarios for verifying the effectiveness
 904 of the Dynamic Anchor Annealing scheme.

906
907

C PERFORMANCE ON THE LAW DATASET

908

909 **Task Diversity and Complexity:** The Law fine-tuning data is a highly heterogeneous mixture of
 910 tasks, including not only multiple-choice questions but also complex generation tasks like **Case**
 911 **Summarization** and reasoning tasks like **Syllogistic Reasoning**. These generative and reason-
 912 ing tasks are fundamentally more challenging and diverge more significantly from the pre-training
 913 objectives than the classification-style tasks that dominate the Finance, Medicine, and News QA
 914 datasets.

915 **Performance Interpretation:** While the absolute score on Law is lower across all methods, it is im-
 916 portant to note that DAA still consistently achieves the best or second-best harmonic mean score (S)
 917 across all three base models (Table 3). For instance, on Qwen2.5, DAA achieves the highest S score
 (59.85), significantly outperforming Direct FT (58.27) and Vanilla FT (57.35) by better preserving

918 general capabilities (S_G). This demonstrates that even in this more complex, generation-heavy do-
 919 main, DAA’s regularization mechanism provides a tangible benefit over baselines by striking a better
 920 balance between domain specialization and knowledge retention.

921 **Conclusion on Generality:** The Law dataset does not necessarily indicate a weakness but rather
 922 highlights DAA’s robust performance on a more challenging and diverse task mixture. It showcases
 923 that DAA’s benefits are not confined to simple classification tasks but extend to complex, mixed-task
 924 scenarios.

927 D PRACTICAL IMPLEMENTATION GUIDE

929 While Dynamic Anchor Annealing (DAA) introduces steps beyond a standard fine-tuning script,
 930 it has been designed for high efficiency and straightforward integration. The methodology is in-
 931 tended to serve as a principal approach for domain specialization, analogous to the role of LoRA
 932 in parameter-efficient tuning. The practical implementation can be decomposed into two distinct
 933 stages.

934 The first stage is a one-time pre-computation of the global gradient \hat{g}_G , on general-domain data.
 935 This process is analogous to a standard training procedure but with the learning rate set to zero,
 936 representing a single, non-recurring computational cost. A critical feature of this approach is its
 937 reusability. The resulting gradient artifact is model-specific yet domain-agnostic, meaning that for
 938 a given foundation model like Llama3.1-8B, this computation is performed only once. The same
 939 global gradient can then be applied to fine-tuning tasks across any number of vertical domains, such
 940 as finance, law, or medicine. We propose that this pre-computation could become a standard practice,
 941 wherein foundation model developers release an official global gradient alongside model weights,
 942 leveraging their high-quality pre-training data. Such a community-driven effort would obviate this
 943 step entirely for downstream domain specialists.

944 The second stage is the integration of Global Gradient Anchoring (GGB) and Dynamic Correction
 945 (DC) into the fine-tuning loop. To facilitate seamless adoption, we have implemented our method
 946 within the LLaMA-Factory and DeepSpeed frameworks. We will release this implementation as
 947 open-source code and submit pull requests to these upstream projects, allowing practitioners to
 948 enable DAA via a simple command-line argument with minimal implementation overhead.

949 In summary, the initial setup cost of DAA is substantially offset by the elimination of repeated data
 950 mixing and extensive hyperparameter tuning. This modest, one-time investment yields significant
 951 and recurring savings in computational resources and engineering time during the iterative process
 952 of domain adaptation. A “Practical Implementation Guide” is provided in the Appendix to further
 953 detail these steps and emphasize the long-term efficiency benefits.

955 E SIGNIFICANCE TEST

957 We conducted a rigorous significance test using the News QA benchmark on Llama3.1-8B. We
 958 chose this setup as it represents a challenging temporal OOD task. We repeated the experi-
 959 ments for both Vanilla FT (at its optimal mixture ratio) and DAA using 5 different random seeds
 960 {42, 43, 44, 45, 46}.

963 Table 10: Significance test results on News QA (Llama3.1-8B).

Method	Seed 42	Seed 43	Seed 44	Seed 45	Seed 46	Mean (S)	Std Dev (σ)
Vanilla FT	53.22	52.15	54.05	52.80	53.75	53.19	0.76
DAA (Ours)	63.76	64.25	63.15	63.88	63.45	63.70	0.42

968 An independent t-test yields a p-value $\ll 0.001$, confirming that the performance improvement
 969 is significant and robust to random initialization. DAA also shows lower variance ($\sigma = 0.42$)
 970 compared to Vanilla FT ($\sigma = 0.76$), demonstrating that the Global Gradient Anchor effectively
 971 stabilizes the optimization process by reducing the impact of initialization randomness.

972 F COST ANALYSIS 973

974 We present the training overhead of DAA and Direct Fine-tuning (FT) for Llama3.1-8B on the News
975 QA dataset using 16 A100 GPUs. As shown in Table 11, DAA’s low-rank gradient anchors occupy
976 5.5 GB of memory. Relative to Direct Fine-tuning (FT), DAA exhibits a modest 0.7 GB increase in
977 Peak VRAM, due to the sharded gradient anchors and the computational cost for gradient correction.
978 Owing to the fact that the anchor gradients are loaded only once, the wall-clock time increased
979 slightly. Crucially, DAA does not introduce additional memory or I/O overhead for *storage* and
980 *repeated use* during training steps, as anchors are pre-loaded into VRAM. Crucially, DAA does not
981 introduce additional memory or I/O overhead for *storage* and *repeated use* during training steps, as
982 anchors are pre-loaded into VRAM.
983

984 Table 11: Training overhead comparison on Llama3.1-8B / News QA.
985

Method	Anchor Memory	Peak VRAM	Parameters	Wall-clock	GPU Hour	S
Direct FT	-	37.1GB	7B	3.91h	3.40h	50.44
DAA	5.5GB	37.8GB	7B	4.20h	3.45h	63.76

989 G ADDITIONAL SENSITIVITY ANALYSIS 990

991 G.1 SENSITIVITY TO GENERAL DATA 992

993 The general dataset used in our paper was filtered from a raw pool of approximately 2.4M samples
994 using the Platypus method. To evaluate sensitivity, we sampled various ratios from this 2.4M raw
995 pool. As shown in Table 12, performance (S) remains remarkably stable as data volume increases.
996 This confirms that \hat{g}_G is robust and does not require massive amounts of raw SFT data to be effective.
997 This aligns with findings from recent studies like LIMA Zhou et al. (2023), suggesting that data
998 quality and curation strategy outweigh sheer quantity. Once the general dataset achieves sufficient
999 diversity to cover broad capabilities, the direction of \hat{g}_G stabilizes, and the marginal utility of adding
1000 more data diminishes. Consequently, the cost of curating this data recipe is a one-time effort. We
1001 view \hat{g}_G as a standard “Model Artifact” as critical as model weights, and advocate for developers to
1002 release these gradients to eliminate data collection costs for end-users.
1003

1004 Table 12: Performance sensitivity to general data sample ratio.
1005

Sample Ratio	20% (Ours)	40%	60%	80%	100%
S	63.76	63.55	63.68	63.71	63.12

1006 G.2 ANCHOR SOURCE 1007

1008 Deriving \hat{g}_G from pre-training or post-pretraining data yields even better results, as these gradients
1009 align more closely with the base model’s intrinsic distribution. We conducted a preliminary exper-
1010 iment where we curated 100B tokens of post-pretraining data to compute \hat{g}_G , comparing it against
1011 the anchor derived from standard SFT data used in our main paper. As shown in Table 13, the anchor
1012 derived from post-pretraining data achieved a superior joint performance score (S) on the News QA
1013 task.
1014

1015 Table 13: Performance comparison of different anchor sources.
1016

Anchor Source	S_D	S_G	S
General SFT Data (Ours)	82.37	52.00	63.76
Post-Pretrain Data	82.51	53.85	65.17

1023 H RAW EXPERIMENTAL RESULTS 1024

1025 We provide the raw scores for experiments on general benchmarks here.

1026
 1027
 1028
 1029
 1030
 1031
 1032
 1033
 1034
 1035
 1036
 1037
 1038
 1039
 1040
 1041
 1042
 1043
 1044
 1045
 1046
 1047
 1048
 1049
 1050
 1051
 1052
 1053
 1054
 1055
 1056
 1057
 1058
 1059
 1060
 1061
 1062
 1063
 1064
 1065
 1066
 1067
 1068
 1069
 1070
 1071
 1072
 1073
 1074
 1075
 1076
 1077
 1078
 1079

Table 14: Raw scores for main comparison experiments.

Domain	Method	Phi4												Qwen3									
		$T \downarrow$	$S_D \uparrow$	S_{MATH}	$S_{MATH-Prev}$	$S_G \uparrow$	$S_D \downarrow$	S_{MATH}	$S_{MATH-Prev}$	$S_G \uparrow$	$T \downarrow$	$S_D \uparrow$	S_{MATH}	$S_{MATH-Prev}$	$S_G \uparrow$	$T \downarrow$	$S_D \uparrow$	S_{MATH}	$S_{MATH-Prev}$	$S_G \uparrow$	$S \uparrow$		
Finance	Direct FT	3.40	80.01	46.67	48.90	55.38	56.99	65.50	65.57	64.97	6.22	89.72	65.57	66.63	80.00	85.12	92.27	77.92	84.41	3.38	85.83	52.89	
	Vanilla FT	38.77	79.30	47.57	58.49	60.73	66.85	68.87	60.50	68.64	71.24	83.42	62.61	72.09	74.18	88.65	91.02	77.71	80.46	38.74	85.49	59.41	
	LoRA	2.50	76.45	50.25	54.01	61.99	71.10	71.10	61.69	68.28	4.61	87.13	67.83	67.28	75.52	84.52	96.20	78.27	82.46	2.50	81.68	58.42	62.18
	DoRA	2.58	76.12	47.83	54.85	58.60	67.60	77.38	61.25	67.88	4.71	86.34	56.99	67.30	82.48	84.69	100.00	78.24	2.57	82.19	62.19	63.02	
Galore	DoRA	5.09	77.31	49.93	56.04	55.89	68.25	73.79	60.78	68.06	9.34	86.23	68.14	73.86	73.35	83.67	93.73	78.55	82.21	5.69	84.37	67.88	74.68
	DAA (Ours)	3.45	79.84	52.13	51.00	63.42	64.97	77.23	61.75	69.64	6.35	87.73	61.89	76.79	74.85	85.24	93.73	78.50	82.86	3.41	85.32	63.33	72.89
	Direct FT	12.39	89.23	46.59	48.86	49.50	55.42	63.28	52.73	66.29	22.78	92.13	62.59	68.46	81.35	82.44	96.86	78.34	84.68	12.38	92.67	48.56	54.00
	Vanilla FT	141.81	87.32	45.02	54.23	58.12	68.25	71.73	59.47	70.75	266.58	91.24	59.32	71.71	86.06	85.68	93.53	79.26	84.83	141.81	81.81	51.22	
Medicine	DoRA	9.18	81.76	47.52	51.10	60.80	67.61	72.83	59.97	69.19	16.87	90.30	68.29	69.54	86.56	92.50	78.50	86.28	8.21	84.00	55.69	65.32	
	DoRA	9.31	81.21	45.40	54.15	60.67	66.61	73.29	59.07	68.40	17.05	90.31	60.76	72.38	82.27	84.30	82.69	82.09	87.74	84.13	8.37	84.74	
	Galore	13.68	81.23	44.21	53.48	57.66	65.68	71.72	91.96	91.96	25.14	91.96	58.08	65.48	77.96	91.53	99.77	78.56	84.73	13.68	87.57	63.91	63.44
	DAA (Ours)	12.54	83.97	48.85	54.15	57.38	66.91	74.37	60.33	70.22	23.90	92.61	60.59	74.66	76.43	88.61	94.01	79.26	85.16	12.64	92.24	62.67	75.23
Law	Direct FT	4.39	56.81	40.27	44.22	47.20	50.13	56.88	38.94	52.58	8.04	42.63	57.94	65.11	64.46	82.20	85.59	71.06	53.29	4.39	55.28	58.46	62.10
	Vanilla FT	50.12	51.37	43.89	48.69	52.03	57.33	63.46	53.08	52.21	91.98	41.82	60.67	60.57	68.52	84.29	86.55	72.12	52.94	50.10	52.28	56.62	64.67
	LoRA	3.25	46.58	46.71	54.68	58.21	65.29	65.42	58.06	51.69	5.97	41.87	61.18	63.69	71.58	83.53	73.38	53.32	3.25	51.90	58.00	62.69	
	DoRA	3.43	46.37	40.33	51.54	59.77	65.54	66.07	56.05	50.75	6.24	41.98	54.59	67.79	87.79	72.36	52.13	3.26	51.91	63.23	69.97	73.75	84.50
News QA	Galore	6.13	47.80	48.74	54.99	59.29	63.48	56.13	51.63	11.26	40.12	53.81	69.17	70.79	81.86	82.02	71.53	51.41	6.13	52.76	59.69	72.76	
	DAA (Ours)	4.53	49.93	42.56	52.79	53.99	64.28	69.78	36.68	53.09	8.17	41.95	61.46	66.30	69.39	78.05	90.41	73.12	53.31	4.60	52.79	66.39	68.63
	Direct FT	1.27	82.38	30.50	34.28	35.56	40.76	40.60	36.34	50.44	2.32	87.38	20.21	25.35	27.52	27.33	31.59	26.40	40.55	1.27	79.37	3.06	3.54
	Vanilla FT	14.45	80.62	34.62	37.22	39.19	42.13	45.48	39.73	53.22	16.51	58.23	59.90	59.66	81.24	84.03	92.40	77.46	80.19	93.93	70.27	55.48	
DAA (Ours)	LoRA	0.94	71.23	37.79	45.95	48.14	50.48	57.84	48.04	57.38	1.72	83.12	62.67	66.96	76.79	80.52	98.51	76.68	64.78	65.15	72.94	80.69	67.81
	DoRA	0.99	73.71	38.80	48.29	50.23	58.30	56.99	50.52	59.95	1.80	83.37	62.47	68.79	73.19	80.52	80.88	51.70	59.68	66.43	83.20	81.14	68.43
	Galore	2.02	79.13	42.95	53.13	53.13	54.54	64.45	53.05	63.52	3.70	85.02	60.55	66.79	74.46	89.89	89.66	76.27	80.41	2.03	80.88	51.70	54.14
	DAA (Ours)	1.32	82.37	40.84	44.16	54.54	56.67	63.80	52.00	63.76	2.42	89.19	59.08	71.97	72.04	85.05	96.41	77.91	83.17	1.28	80.82	51.32	65.66