

ProFit: Leveraging High-Value Signals in SFT via Probability-Guided Token Selection

Tao Liu^{1,*} Taiqiang Wu^{2,*},† Runming Yang² Shaoning Sun¹ Junjie Wang^{1,‡} Yujiu Yang^{1,‡}

¹Tsinghua University ²The University of Hong Kong

*Equal contribution †Project Leader ‡Corresponding authors

🔗 <https://github.com/Utaotao/ProFit>

Abstract

Supervised fine-tuning (SFT) is a fundamental post-training strategy to align Large Language Models (LLMs) with human intent. However, traditional SFT often ignores the *one-to-many* nature of language by forcing alignment with a single reference answer, leading to the model overfitting to non-core expressions. Although our empirical analysis suggests that introducing multiple reference answers can mitigate this issue, the prohibitive data and computational costs necessitate a strategic shift: prioritizing the mitigation of single-reference overfitting over the costly pursuit of answer diversity. To achieve this, we reveal the intrinsic connection between token probability and semantic importance: high-probability tokens carry the core logical framework, while low-probability tokens are mostly replaceable expressions. Based on this insight, we propose **ProFit**, which selectively masks low-probability tokens to prevent surface-level overfitting. Extensive experiments confirm that ProFit consistently outperforms traditional SFT baselines on general reasoning and mathematical benchmarks.

1 Introduction

Large Language Models (LLMs) have demonstrated remarkable general capabilities (Jaech et al., 2024; Guo et al., 2025; Yang et al., 2025a). To adapt them to specific downstream tasks, Supervised Fine-Tuning (SFT) has become the prevailing paradigm (Chung et al., 2024). Traditional SFT is based on an autoregressive objective, forcing the model to strictly align with a single reference answer at the token level. However, this rigid objective neglects the *one-to-many* nature of language (Li et al., 2016; Yang et al., 2025b), where diverse expressions can convey the same intent. Therefore, this strategy of forcibly fitting a single reference is often suboptimal and can easily lead to the model simply memorizing specific samples. (Gudibandé et al., 2023; Chu et al., 2025;

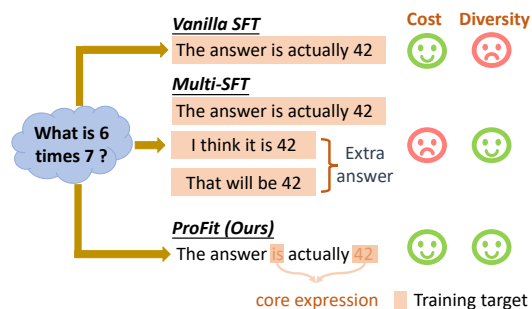


Figure 1: **Breaking the trade-off between training cost and semantic diversity.** While Multi-reference SFT offers semantic richness at prohibitive data and computational costs, standard SFT is efficient but semantically limited. ProFit achieves the best of both: by focusing supervision on high-value tokens, it captures core semantic integrity without sacrificing the efficiency of single-reference training.

Wang et al., 2025a, 2026).

While introducing multiple reference answers can alleviate this problem (Yuan et al., 2023; Li et al., 2024b; Shi and Shen, 2025), it faces the dual challenges of expensive data construction and difficulties in training convergence (please refer to Section 3.2 for details). To solve this problem while maintaining the existing low-cost *single instruction-single response* data configuration, we propose a more efficient strategy (as illustrated in Figure 1): **instead of striving for comprehensive coverage with multiple answers, it’s better to avoid overfitting to a single answer.** To achieve this, we need a mechanism to filter out truly high-value training signals from a single reference answer.

To accurately identify these high-value training signals, we conducted a semantic analysis. Specifically, we comparatively analyze multiple answers to a given question, aiming to identify those tokens defined as decisive for the answer’s correctness, with their importance judged by Gemini-3-Pro. For-

unately, we found that the predicted probability of the token can serve as an efficient and accurate proxy metric (Kadavath et al., 2022; Huang et al., 2025; Bentegeac et al., 2025). Further hypothesis testing confirmed this significant pattern: high-probability tokens tend to carry core reasoning logic or key semantics, while low-probability tokens correspond more to non-core expressions.

Inspired by this, we propose the **ProFit** method, which leverages the online probabilities predicted by the model currently being trained as the core clue to locate high-value signals. Specifically, **ProFit** employs a strategic masking mechanism: it selectively retains and trains high-probability tokens that carry crucial semantic information, while masking low-probability, non-essential tokens (Lin et al., 2024; Ruan et al., 2025). We further provided theoretical derivations demonstrating that the gradients of low-probability tokens can overshadow the optimization direction of crucial tokens.

We conducted extensive evaluations on general reasoning (GPQA-Diamond (Rein et al., 2024)), mathematics (MATH-500 (Lightman et al., 2023), AIME'24 (American Institute of Mathematics, 2024), GSM8K (Cobbe et al., 2021)), and instruction following (IFEval (Zhou et al., 2023b)). The results consistently show that **ProFit** outperforms the traditional SFT baseline. Notably, on the Qwen3 family, ProFit surpasses SFT by a significant margin of 3.0% to 10.9% in average accuracy, validating the effectiveness of our strategy.

Our contributions can be summarized as follows:

- We identify a positive correlation between prediction probability and semantic importance, revealing that low-probability tokens typically represent non-essential expression.
- We propose **ProFit**, a probability-guided masking strategy. We theoretically prove that masking low-probability tokens prevents their large gradients from overshadowing key semantic signals.
- Extensive experiments on general reasoning and math benchmarks demonstrate that **ProFit** consistently outperforms standard SFT baselines.

2 Related Work

2.1 Data-Efficient Instruction Tuning

Recent SFT adheres to the *less is more* principle (Zhou et al., 2023a; Zhang et al., 2025;

Li et al., 2025c; Ji et al., 2026), evolving from complexity-based selection (Cao et al., 2023; Li et al., 2024a) to 2025’s importance-aware metrics like MIWV (Jiang et al., 2025) and ICL-based filtering (Wang et al., 2025b; Jiang et al., 2025). However, these *coarse-grained sample-level* methods treat pairs as atomic units, overlooking intra-sample low-information segments or stylistic noise (Pang et al., 2025; Qin et al., 2025), thus limiting the model’s focus on dense logical signals.

2.2 Token-Level Training Objectives

To address granularity limitations, research has pivoted to *token-level* optimization (Kong et al., 2025). While classical methods like Focal Loss (Lin et al., 2017) and Unlikelihood Training (Welleck et al., 2019) targeted hard tokens, recent LLM approaches like Rho-1 (Lin et al., 2024) and TIS-DPO (Liu et al., 2024) rely on costly external reference models. More efficient intrinsic methods like DFT (Wu et al., 2025b) employ probability-driven *soft reweighting*, and CFT (Ruan et al., 2025) further validates the need for supervising critical regions. Recently, Li et al. (2025a) theoretically validated via a model-capability continuum that suppressing low-probability tokens benefits strong-prior domains like math, supporting our masking strategy. While sharing the probability-driven motivation with DFT, **ProFit** diverges by adopting a strict *hard masking* strategy to efficiently filter out non-core expressions in a single step.

3 Preliminaries and Motivation

3.1 Preliminaries

Supervised Fine-Tuning. Given a dataset \mathcal{D} containing pairs of inputs x and reference responses $y^* = (y_1^*, \dots, y_T^*)$, SFT optimizes the policy π_θ by minimizing the negative log-likelihood:

$$\mathcal{L}_{\text{SFT}}(\theta) = \mathbb{E}_{(x, y^*) \sim \mathcal{D}} \left[- \sum_{t=1}^{|y^*|} \log \pi_\theta(y_t^* | x, y_{<t}^*) \right]. \quad (1)$$

Let z_t be the logits at step t . The probability is $p_t = \text{softmax}(z_t)$. The gradient of the per-token loss $\ell_t = -\log p_{t, y_t^*}$ with respect to the logits z_t is:

$$\frac{\partial \ell_t}{\partial z_{t,v}} = p_{t,v} - \mathbb{I}[v = y_t^*]. \quad (2)$$

Equation 2 drives $p_{t, y_t^*} \rightarrow 1$ while suppressing alternatives ($p_{t,v} \rightarrow 0$). This mechanism indis-




Figure 2: **Performance comparison on diverse benchmarks.** While multi-reference training (SFT w/ 3 ans) offers sporadic gains, it suffers from optimization instability and stagnation on complex tasks. In contrast, **ProFit** achieves superior and robust performance across all metrics by selectively extracting high-value signals from a single reference.

criminally suppresses all non-reference tokens ($v \neq y_t^*$), including valid paraphrases, thereby penalizing semantic flexibility and driving the model toward *surface-form overfitting*.

Low-Rank Adaptation (LoRA). LoRA is a parameter-efficient fine-tuning technique based on the premise that weight updates ΔW possess a low intrinsic rank (Hu et al., 2022; Wu et al., 2024). For a frozen weight matrix $W_0 \in \mathbb{R}^{d \times k}$, LoRA approximates the update via a low-rank decomposition BA , where $B \in \mathbb{R}^{d \times r}$ and $A \in \mathbb{R}^{r \times k}$ with $r \ll \min(d, k)$. The forward pass is modified as:

$$h = W_0 x + \frac{\alpha}{r} \Delta W x = W_0 x + \frac{\alpha}{r} B A x \quad (3)$$

where α is a scaling hyperparameter for adaptation stability.

3.2 Motivation

Traditional SFT relies on strict single-reference token-level alignment, which inherently over-penalizes paraphrastic variants by treating valid semantic equivalents as incorrect predictions. Intuitively, introducing multiple reference answers could theoretically bridge this gap. However, our pilot investigation reveals that this approach faces substantial practical barriers rather than offering a straightforward solution.

First, constructing a diverse, high-quality response set entails prohibitive costs: collecting K distinct references per instruction scales the annotation burden linearly, and ensuring high quality often necessitates expert annotators, particularly

for complex reasoning or mathematical tasks. Second, and more critically, **simply expanding the reference set often introduces distributional conflicts, leading to optimization instability.** As illustrated in Figure 2, we compared SFT trained with single versus multiple (3) reference answers. While multiple references yield marginal improvements on specific reasoning tasks like MATH-500, they fail to generalize consistently. Surprisingly, on complex benchmarks like GPQA-Diamond, the performance stagnates or even slightly degrades (dropping from 34.1% to 33.5%) compared to the single-answer baseline. Furthermore, both standard and multi-answer SFT struggle to maintain base capabilities on instruction-following tasks like IFEval, where performance trails behind the Vanilla model. This suggests that blindly fitting diverse distributions can confuse the model, causing it to struggle with convergence due to conflicting gradient directions.

To address this dilemma, we propose relaxing the strict alignment objectives. Instead of performing indiscriminate full-scale fitting on all tokens or relying on expensive multi-reference datasets, we implement a selective alignment strategy. **We need a mechanism that can accurately filter out the high-value training signals that truly carry the core reasoning logic from a single reference answer,** thereby achieving robust performance across diverse benchmarks.

4 Methodology

4.1 Semantic Analysis

To accurately extract high-quality training signals, we performed a joint analysis of semantic importance and prediction probability for the tokens in multiple reference answers. Specifically, we utilize Gemini-3-pro (Google DeepMind, 2025) as a semantic evaluator to annotate the tokens in the reference answers, classifying them into *trivial tokens*, which represent interchangeable stylistic variations, and *core tokens*, which encapsulate the essential reasoning logic. Subsequently, we used Qwen3-4B-Base (Yang et al., 2025a) to perform forward propagation and calculate the predicted probability for each token. This choice provides a computation-efficient yet capable proxy for language modeling, striking a balance between estimation quality and inference cost.

As shown in Figure 3, the two token types exhibit distinct probability distributions: core tokens are




Figure 3: Probability density estimation of semantic tokens. We categorize tokens into semantically *Core* and *Trivial* groups. While *core tokens* are heavily concentrated in high-confidence zones, *trivial tokens* exhibit a significant **long-tail distribution**, disproportionately dominating the low-probability spectrum. A hypothesis test confirms this significant distributional difference ($p = 1 \times 10^{-6}$).

highly concentrated in the high-probability region, showing strong determinism, whereas trivial tokens display a long-tail distribution. Although some trivial tokens appear in the high-confidence interval, their density is much higher in the low-probability region than that of core tokens, making them the dominant component there.

To rigorously verify this observation, we conducted a statistical hypothesis test. We formally defined the null hypothesis as follows:

- **Null Hypothesis (H_0):** The probability distributions of *trivial tokens* and *core tokens* are statistically identical (i.e., drawn from the same population).
- **Alternative Hypothesis (H_1):** The probability distributions of *trivial tokens* and *core tokens* are statistically distinct (i.e., drawn from different populations).

The test yielded a p -value of 1×10^{-6} , leading to a significant rejection of H_0 . This empirical result strongly supports our hypothesis: low prediction probability is a strong indicator of semantic non-essentiality, as low-probability regions are primarily dominated by *trivial tokens*.

4.2 ProFit

Based on the motivation and semantic analysis, we propose **ProFit**. The core intuition of this method is to utilize the model’s own prediction probabilities as a dynamic indicator to locate the core tokens. During training, by implementing

a threshold-based masking operator, ProFit selectively backpropagates gradients only from high-probability tokens, effectively isolating parameter updates from the interference of low-probability *trivial tokens*.

To operationalize this strategy, we employ a **stop-gradient** mechanism to decouple the masking criterion from the gradient computation. We define the binary validity mask \mathcal{M}_t using the detached probability:

$$\mathcal{M}_t = \mathbb{I}[\text{sg}(\pi_\theta(y_t^* | x, y_{<t}^*)) > \tau], \quad (4)$$

where $\text{sg}(\cdot)$ denotes the stop-gradient operator, $\tau \in [0, 1]$ is a static threshold, and $\mathbb{I}[\cdot]$ is the indicator function. By strictly enforcing $p > \tau$, we ensure the optimization is driven solely by high-value semantic signals. Crucially, the stop-gradient ensures that \mathcal{M}_t acts as a fixed gate during backpropagation, avoiding the differentiability issues of the step function.

Formally, the optimization objective of ProFit is defined as:

$$\mathcal{L}_{\text{ProFit}}(\theta) = \mathbb{E}_{\mathcal{D}} \left[-\frac{1}{T} \sum_{t=1}^T \mathcal{M}_t \log \pi_\theta(y_t^* | x, y_{<t}^*) \right], \quad (5)$$

where $T = |y^*|$ denotes the sequence length.

4.3 Deeper Insights

Equation 2 shows that low-probability tokens induce significantly larger logit gradients. To quantify how this amplification propagates to the parameter space, we derive the following theorem, establishing a lower bound for the parameter gradient (proof in Appendix A):

Theorem 1 (Token-Wise Gradient Norm Lower Bound). Consider the prediction of a single target token y_t^* at step t , given the instruction x and the preceding ground-truth tokens $y_{<t}^*$. Let $z \in \mathbb{R}^{|\mathcal{V}|}$ be the output logits and $\ell(\theta) = -\log \pi_\theta(y_t^* | x, y_{<t}^*)$ be the loss for this step. Let $J_\theta(z) = \nabla_{\theta} z \in \mathbb{R}^{|\mathcal{V}| \times |\theta|}$ denote the Jacobian of the logits with respect to parameters $\theta \in \mathbb{R}^{|\theta|}$. Under the local non-degeneracy assumption that the Jacobian is full row-rank (i.e., the model is locally surjective, satisfying $\sigma_{\min}(J_\theta(z)) \geq \gamma > 0$), the gradient norm satisfies:

$$\|\nabla_{\theta} \ell\|_2 \geq \gamma \cdot (1 - \pi_\theta(y_t^* | x, y_{<t}^*)). \quad (6)$$

This lower bound theoretically guarantees that tokens with lower prediction probabilities inevitably induce larger parameter gradients.

Model	Method	GPQA-Diamond	GSM8K	MATH-500	AIME'24	IFEval	Avg.	Δ
Qwen3-0.6B-Base	Vanilla	4.36	45.99	7.28	0.10	15.6	14.67	-
	SFT	17.93	54.53	45.38	1.56	23.15	28.51	+13.84
	Entropy	20.58	53.67	45.83	2.19	22.02	28.86	+14.19
	DFT	17.68	62.42	47.92	2.29	20.63	30.20	+15.53
	ProFit	22.85	59.78	49.90	2.19	22.71	31.49	+16.82
Qwen3-4B-Base	Vanilla	23.17	82.50	57.67	8.23	33.04	40.92	-
	SFT	34.15	55.43	74.80	11.25	31.33	41.39	+0.47
	Entropy	34.91	46.54	74.75	10.94	30.71	39.57	-1.35
	DFT	31.69	87.83	77.50	10.83	43.67	50.30	+9.38
	ProFit	34.34	87.55	77.85	13.02	48.87	52.33	+11.41
Qwen3-14B-Base	Vanilla	37.69	83.44	78.32	13.33	52.61	53.08	-
	SFT	46.02	78.00	79.22	16.04	36.69	51.20	-1.88
	Entropy	47.85	78.23	80.12	14.79	38.01	51.80	-1.28
	DFT	43.81	88.98	81.97	17.29	49.86	56.38	+3.30
	ProFit	46.53	89.62	82.85	16.56	58.02	58.72	+5.64
OLMo-2-7B	Vanilla	21.72	68.08	11.03	0.10	14.44	23.07	-
	SFT	13.64	78.43	24.95	0.62	22.50	28.03	+4.96
	Entropy	12.94	78.22	24.77	0.52	22.64	27.82	+4.75
	DFT	12.31	76.09	23.57	0.31	23.24	27.10	+4.03
	ProFit	14.71	78.25	25.45	0.31	23.98	28.54	+5.47
Llama-3.1-8B	Vanilla	8.71	51.88	3.35	0.00	21.42	17.07	-
	SFT	23.30	60.42	24.75	0.31	24.72	26.70	+9.63
	Entropy	23.61	61.08	25.70	0.21	24.31	26.98	+9.91
	DFT	8.40	60.07	17.65	0.62	25.14	22.38	+5.31
	ProFit	21.40	62.11	24.90	0.62	26.16	27.04	+9.97

Table 1: Main results across five benchmarks. We report the accuracy of the Vanilla baseline, standard SFT, and varying strategies (Entropy, DFT, ProFit) on multiple model families. Regarding the evaluation settings, results on **AIME'24** are averaged over 32 samples, while results on the other datasets are averaged over 8 samples. The values in parentheses indicate the performance difference relative to the Vanilla baseline.

5 Experiments

5.1 Experimental Setup

Training. For the training data, we curated a subset of 2,000 samples from the BAAI-InfinityInstruct Dataset (Zhou et al., 2023a; Li et al., 2025b; Muennighoff et al., 2025), prioritizing high reward scores as done in Shadow-FT (Wu et al., 2025a). To comprehensively verify the effectiveness and generalization of our method, we conducted evaluations across a diverse set of LLMs, including the Qwen3 series (Yang et al., 2025a), Llama 3 series (AI@Meta, 2024), and OLMo 2 series (OLMo et al., 2024). We employed LLaMA-Factory (Zheng et al., 2024) as our primary training framework. All experiments are conducted on 8 H20 GPUs. For detailed hyperparameter settings, please refer to the Appendix B.

Baselines. To evaluate the effectiveness of ProFit, we compare it against several representative fine-tuning paradigms:

- **Supervised Fine-tuning:** The vanilla baseline that minimizes the cross-entropy loss across all tokens indiscriminately.
- **Dynamic Fine-tuning (Wu et al., 2025b):** A probability-aware method that assigns a dynamic scale to the loss of each token based on its confidence.
- **Entropy-based tuning (Wang et al., 2025c):** A selective strategy that updates parameters only on tokens with high entropy.

Evaluation. To comprehensively evaluate the downstream performance of our fine-tuned models, we conducted extensive assessments across a diverse set of benchmarks, including GPQA-Diamond (Rein et al., 2024), MATH-500 (Lightman et al., 2023), GSM8K (Cobbe et al., 2021), AIME'24 (American Institute of Mathematics, 2024), and IFEval (Zhou et al., 2023b). For the evaluation pipeline, we utilized OpenCompass (Contributors, 2023b) as the primary inference




Figure 4: Ablation study on the probability threshold τ . The **dashed line** represents the performance of the standard SFT baseline. (a) Training exclusively on low-probability tokens ($p(y_t^*) < \tau$) results in performance consistently below the baseline, indicating that non-core expressions are insufficient for constructing effective reasoning chains. (b) Conversely, the proposed strategy ($p(y_t^*) > \tau$), which masks low-probability noise, consistently outperforms the baseline across all tasks, validating the effectiveness of focusing on core logic.

framework, integrated with lmdeploy (Contributors, 2023a) and vllm (Kwon et al., 2023) as the acceleration backend. For the decoding strategy, we performed sampling and reported the average accuracy across 32 generations for AIME’24 and 8 for other benchmarks. Detailed inference hyperparameters and configuration settings are provided in the Appendix B.

5.2 Main Results

Table 1 presents the comparative results of our proposed method, ProFit, against the Vanilla baseline, standard SFT, and other strategies (Entropy and DFT) across five diverse benchmarks. The experimental results demonstrate that ProFit consistently achieves superior performance across all evaluated model families and scales.

Compared to standard SFT, ProFit delivers substantial improvements in average accuracy. For instance, on the Qwen3-4B-Base model, ProFit achieves an average accuracy of 52.33%, surpassing standard SFT (41.39%) by a significant margin of 10.94%. Similarly, on Llama-3.1-8B, ProFit improves the average score to 27.04%, outperforming SFT (26.70%) and showing a +9.97% gain over the Vanilla baseline. Notably, on Qwen3-14B-Base, standard SFT experiences a performance drop of 1.88% relative to Vanilla, likely stemming from the interference of non-core expressions or superficial stylistic patterns. ProFit successfully mitigates this issue, reversing the decline to achieve a distinct gain of +5.64%. This underscores the stability

of our probability-guided filtering mechanism in mitigating negative transfer while enhancing downstream performance.

ProFit also consistently outperforms other baselines. As shown in Table 1, while DFT and Entropy strategies generally offer improvements over the Vanilla baseline, they often fall short of the gains achieved by ProFit. For example, on Qwen3-0.6B, ProFit achieves the highest average accuracy of 31.49%, exceeding DFT (30.20%) and Entropy (28.86%). This trend holds across architectures like OLMo-2 and Llama-3.1, showing our probability-based identification of non-core expressions outperforms entropy-based filtering and dynamic reweighting methods.

In summary, ProFit shows strong scalability and universality, consistently boosting performance across model sizes and architectures without the high data costs of multi-reference fine-tuning.

6 Extensive Analysis

6.1 Analysis of τ parameter

To investigate the impact of the probability threshold τ , we compared the experimental results of retaining low-probability tokens ($p(y_t^*) < \tau$) versus retaining high-probability tokens ($p(y_t^*) > \tau$), as illustrated in Figure 4.

The Figure 4(b) reveals that when masking low-probability tokens—which typically represent semantic diversity—the model’s performance consistently exceeds the full-token fine-tuning baseline (dashed line) across all threshold settings. This sug-

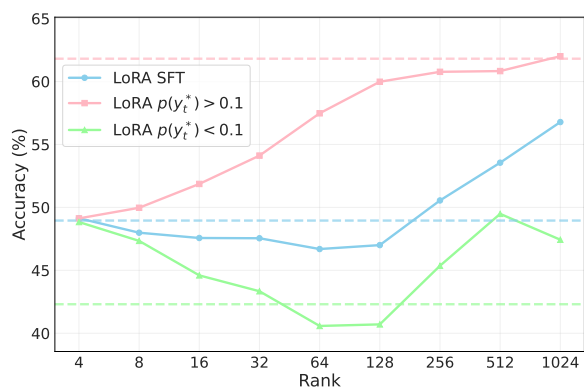


Figure 5: Average performance variation across different LoRA ranks ($r \in \{4, \dots, 1024\}$). The dashed lines represent the baseline performance of full-parameter fine-tuning for each corresponding setting. While core tokens ($p(y_t^*) > 0.1$) exhibit monotonic improvement driven by capacity, non-core tokens ($p(y_t^*) < 0.1$) and standard SFT show a U-shaped trend, revealing optimization interference at medium ranks.

gests that in traditional SFT, forcing the model to fit these surface-level stylistic variations distracts it from learning the underlying reasoning patterns. By alleviating this unnecessary learning burden, the model can focus more effectively on the invariant logical core. Although knowledge-intensive tasks like GPQA-Diamond show a slight performance decline as the threshold increases, likely due to specific long-tail entities falling into the low-probability range, their absolute performance remains significantly above the baseline.

Conversely, the Figure 4(a) highlights the irreplaceable skeletal role of high-probability tokens. When the model is restricted to learning only low-probability *diverse expressions* without the support of high-probability structural tokens, performance suffers a catastrophic decline. Crucially, even as the threshold τ increases (introducing more high-frequency tokens), while performance recovers slightly, it never reaches the level of the full-token fine-tuning baseline and remains far inferior to the strategy shown in the 4(b). This persistent performance gap indicates that low-probability tokens are essentially auxiliary components contingent upon the logical skeleton. Without the support of high-confidence core logic, merely increasing the fitting of these non-core expressions fails to establish effective reasoning links and ultimately limits the model’s generalization potential.




Figure 6: Average performance trajectory across training epochs. ProFit ($p > \tau$) demonstrates rapid convergence and a superior performance ceiling, whereas focusing on low-probability tokens ($p < \tau$) results in training instability and limited capacity.

6.2 Impact of LoRA Rank

We explore the influence of trainable parameter volume by scaling the LoRA rank from 4 to 1024, as shown in Figure 5 (refer to Appendix C for details). The results reveal a distinct divergence: core tokens ($p > 0.1$) benefit monotonically from increased rank, indicating they strictly require model capacity. In contrast, trivial tokens ($p < 0.1$) and standard SFT display a U-shaped trend, where medium ranks struggle with optimization interference. Interestingly, for trivial tokens, LoRA (Rank 1024) outperforms full fine-tuning, demonstrating that low-rank constraints serve as effective regularization, preventing the model from overfitting to non-essential statements. The fact that global SFT follows the trivial token trend underscores that non-core expressions act as the primary bottleneck in standard training.

6.3 Performance Evolution across Epochs

Figure 6 illustrates the training trajectory over 5 epochs. ProFit ($p > \tau$) demonstrates superior efficiency, converging immediately to 60.1% accuracy in the first epoch—already surpassing the Baseline’s peak performance of 54.9%. In stark contrast, training exclusively on low-probability tokens ($p < \tau$) leads to stagnation in a suboptimal range (40% – 50%) and training instability. These results confirm that high-probability tokens contain the essential gradient signals for alignment, whereas low-probability regions offer negligible or even detrimental supervision.




Figure 7: Performance trajectories on MATH-500, OlympiadBench, and Minerva datasets. ProFit combines high initial performance with continuous learning capability, ultimately achieving the best Pass@4 and Avg@4 scores across all benchmarks.

6.4 Superior RL Initialization

We conduct experiments based on the Qwen3-0.6B-Base model. Specifically, we employ GRPO (Group Relative Policy Optimization (Shao et al., 2024)) to optimize the models initialized with three different strategies: Base, DFT, and our proposed ProFit. We employ the DeepScaleR dataset as training prompts following Liu et al. (2025). All experiments are conducted on 32 H20 GPUs. For evaluation, in addition to the previously introduced MATH-500, we employ Minerva (Hendrycks et al., 2021) and OlympiadBench (He et al., 2024), which serve as comprehensive benchmarks covering tasks from general mathematical reasoning to diverse competition-level challenges.

We evaluated the potential of ProFit as an initialization for subsequent reinforcement learning by analyzing its training dynamics across three mathematical benchmarks, as illustrated in Figure 7. The experimental results demonstrate that ProFit consistently outperforms comparison methods, offering not only a superior starting point but also more robust training stability for the RL stage. Please refer to Appendix G for detailed training dynamics, such as KL divergence, entropy, and response length.

Specifically, on the MATH-500 dataset, ProFit achieves an Avg@4 of **57.3%** and a Pass@4 of **76.4%** at the final training stage, significantly surpassing the Baseline which reaches 53.1% and 70.6% respectively. Crucially, the analysis of training curves reveals distinct convergence behaviors: while the Baseline exhibits a pronounced *cold start* phenomenon requiring more steps to learn effective patterns, and DFT shows earlier plateauing

in Avg@4 under the same budget, ProFit maintains continuous performance growth throughout the training process.

This advantage is particularly evident on the challenging OlympiadBench, where ProFit attains an Avg@4 of **24.3%**, distinctively outperforming the competing methods which hover around the 21.1% level. Similar consistent improvements are also observed on Minerva, further validating the generalization capability of our method.

6.5 Analysis of Fine-Tuning Instruct Models

Further fine-tuning instruction-aligned models typically incurs an alignment tax, degrading general capabilities (Wu et al., 2025a,b; Diao et al., 2026). We evaluate ProFit on the Qwen2.5-Instruct series to test its robustness against this issue. As shown in Table 2, while standard SFT and DFT consistently suffer performance drops (e.g., -1.83 average degradation for SFT on 7B), ProFit effectively mitigates catastrophic forgetting. It not only prevents typical declines but actually improves the average scores for the 0.5B (+1.24) and 7B (+0.17) models, with notable gains in reasoning tasks like GSM8K and AIME’24, proving its superiority in refining aligned models without sacrificing foundational instruction-following skills.

6.6 Analysis of Performance on General Tasks

To evaluate whether ProFit preserves comprehensive abilities while optimizing for specific domains, we tested it on general benchmarks across Qwen3 scales. As shown in Table 3, ProFit demonstrates strong generalization without severe catastrophic

Model	Method	GPQA-Diamond	GSM8K	MATH-500	IFEval	AIME'24	Avg.	Δ
Qwen2.5-0.5B-Instruct	Vanilla	17.42	37.38	27.98	25.28	0.21	21.65	-
	SFT	17.55	37.96	24.47	23.66	0.21	20.77	-0.88
	DFT	10.67	40.48	27.65	25.25	0.31	20.87	-0.78
	ProFit	16.60	42.97	29.55	24.72	0.62	22.89	+1.24
Qwen2.5-3B-Instruct	Vanilla	32.01	74.73	65.78	58.99	4.38	47.18	-
	SFT	30.30	75.29	63.52	57.44	5.31	46.37	-0.81
	DFT	26.58	77.27	65.77	50.72	7.60	45.59	-1.59
	ProFit	29.10	77.90	65.83	52.22	7.19	46.45	-0.73
Qwen2.5-7B-Instruct	Vanilla	35.48	81.00	75.28	69.29	12.50	54.71	-
	SFT	33.27	83.82	74.65	61.95	10.73	52.88	-1.83
	DFT	32.77	85.60	74.65	60.28	13.02	53.26	-1.45
	ProFit	35.35	85.61	75.15	63.79	14.48	54.88	+0.17

Table 2: Performance evaluation on previously instruction-tuned models (Qwen2.5-Instruct series). We report the accuracy of the original Instruct model (Vanilla), standard SFT, DFT, and ProFit.

Model	Method	HellaSwag	ARC-c	HumanEval	MMLU-Pro History	MMLU-Pro Chemistry	Avg.	Δ
Qwen3-0.6B-Base	Vanilla	24.54	1.69	1.45	1.31	3.18	6.43	-
	SFT	33.37	45.76	32.09	15.22	15.72	28.43	+22.00
	Entropy	33.16	50.51	31.78	11.02	15.90	28.47	+22.04
	DFT	32.72	53.22	37.80	10.76	14.93	29.89	+23.46
	ProFit	33.80	55.59	38.64	17.32	19.61	32.99	+26.56
Qwen3-4B-Base	Vanilla	65.55	34.92	75.38	27.03	16.96	43.97	-
	SFT	74.83	92.20	82.85	41.99	62.46	70.87	+26.90
	Entropy	74.77	90.85	81.78	41.73	61.31	70.09	+26.12
	DFT	62.12	90.85	81.40	42.26	59.28	67.18	+23.21
	ProFit	72.21	91.19	80.79	44.62	61.04	69.97	+26.00
Qwen3-14B-Base	Vanilla	84.02	94.92	83.69	54.59	57.95	75.03	-
	SFT	82.18	94.58	86.97	53.81	70.05	77.52	+2.49
	Entropy	81.43	93.90	87.04	55.38	70.41	77.63	+2.60
	DFT	86.96	94.58	87.65	56.43	69.70	79.06	+4.03
	ProFit	87.30	94.92	89.18	59.84	70.32	80.31	+5.28

Table 3: Performance comparison on General Tasks across different Qwen3 model scales (0.6B, 4B, 14B). We report the accuracy of the Vanilla baseline, standard SFT, and varying fine-tuning strategies (Entropy, DFT, ProFit).

forgetting. It achieves the highest average improvements on the 0.6B (+26.56) and 14B scales over the Vanilla model. Even on the 4B scale, where standard SFT slightly edges out in the overall average, ProFit remains highly competitive and excels in specific reasoning-heavy subsets, proving it effectively balances broad knowledge retention with its primary fine-tuning objectives.

7 Conclusion

In this work, we propose ProFit, a novel method to fine-tune LLMs by leveraging token prediction probabilities as a proxy for semantic importance. Inspired by our hypothesis testing results revealing that high-probability tokens carry core semantics while low-probability ones correspond to non-core expressions, we propose ProFit to selectively mask the latter, aiming to alleviate the overfitting to surface-level phrasing and capture the underlying

logic. Extensive experiments across multiple LLM series, including Qwen, Llama, and OLMo2, demonstrate that ProFit consistently outperforms conventional full-parameter SFT and other data selection methods on diverse benchmarks covering general reasoning, mathematics, and instruction following. We further conduct comprehensive ablation studies on training epochs, probability thresholds, and LoRA ranks. ProFit also serves as a superior initialization for reinforcement learning, promoting stable convergence and deeper reasoning exploration. These analyses confirm that low-probability tokens serve as a primary source of optimization interference, validating that our approach offers a robust and effective solution for improving model generalization.

Acknowledgements

This work was partly supported by the National Natural Science Foundation of China (Grant No. 62576191), the Shenzhen Science and Technology Program (ZDCY20250901103533010) and Tsinghua SIGS KA Cooperation Fund.

Limitation

Despite the promising results, this work has limitations.

First, our core assumption that low-probability tokens represent non-core expressions primarily holds for logic-intensive tasks (e.g., reasoning, mathematics). For creative generation tasks, such tokens may contribute to stylistic diversity, which requires further investigation.

Second, ProFit currently employs a static probability threshold across all samples. While this design choice was made to ensure implementation simplicity and training stability and has yielded consistent empirical gains, we acknowledge that future iterations could further enhance performance by exploring adaptive mechanisms that dynamically adjust the threshold based on instance-specific difficulty.

References

- AI@Meta. 2024. [Llama 3 model card](#).
- American Institute of Mathematics. 2024. Aime 2024 competition mathematical problems.
- Raphaël Bentegeac, Bastien Le Guellec, Grégory Kuchcinski, Philippe Amouyel, and Aghiles Hamroun. 2025. Token probabilities to mitigate large language models overconfidence in answering medical questions: Quantitative study. *Journal of medical Internet research*, 27:e64348.
- Yihan Cao, Yanbin Kang, Chi Wang, and Lichao Sun. 2023. Instruction mining: Instruction data selection for tuning large language models. *arXiv preprint arXiv:2307.06290*.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, and 1 others. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.
- Tianzhe Chu, Yuexiang Zhai, Jihan Yang, Shengbang Tong, Saining Xie, Dale Schuurmans, Quoc V Le, Sergey Levine, and Yi Ma. 2025. Sft memorizes, rl generalizes: A comparative study of foundation model post-training. *arXiv preprint arXiv:2501.17161*.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, and 1 others. 2024. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70):1–53.
- 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. *arXiv preprint arXiv:1803.05457*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, and 1 others. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- LMDeploy Contributors. 2023a. Lmdeploy: A toolkit for compressing, deploying, and serving llm. <https://github.com/InternLM/lmdeploy>.
- OpenCompass Contributors. 2023b. Opencompass: A universal evaluation platform for foundation models. <https://github.com/open-compass/opencompass>.
- Muxi Diao, Lele Yang, Wuxuan Gong, Yutong Zhang, Zhonghao Yan, Yufei Han, Kongming Liang, Weiran Xu, and Zhanyu Ma. 2026. Entropy-adaptive fine-tuning: Resolving confident conflicts to mitigate forgetting. *arXiv preprint arXiv:2601.02151*.
- Google DeepMind. 2025. [Gemini 3 pro model card](#). Technical report, Google.
- Arnav Gudibande, Eric Wallace, Charlie Snell, Xinyang Geng, Hao Liu, Pieter Abbeel, Sergey Levine, and Dawn Song. 2023. The false promise of imitating proprietary llms. *arXiv preprint arXiv:2305.15717*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shitong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, and 1 others. 2024. Olympiadbench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal scientific problems. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3828–3850.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*.

- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, and 1 others. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3.
- Guanhua Huang, Tingqiang Xu, Mingze Wang, Qi Yi, Xue Gong, Siheng Li, Ruibin Xiong, Kejiao Li, Yuhao Jiang, and Bo Zhou. 2025. Low-probability tokens sustain exploration in reinforcement learning with verifiable reward. *arXiv preprint arXiv:2510.03222*.
- Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, and 1 others. 2024. Openai o1 system card. *arXiv preprint arXiv:2412.16720*.
- Hongru Ji, Yuyin Fan, Meng Zhao, Xianghua Li, Lianwei Wu, and Chao Gao. 2026. [Stride-ed: A strategy-grounded stepwise reasoning framework for empathetic dialogue systems](#). *Preprint*, arXiv:2604.07100.
- Tingyu Jiang, Shen Li, Yiyao Song, Lan Zhang, Hualei Zhu, Yuan Zhao, Xiaohang Xu, Kenjiro Taura, and Hao Henry Wang. 2025. Importance-aware data selection for efficient llm instruction tuning. *arXiv preprint arXiv:2511.07074*.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, and 1 others. 2022. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221*.
- Zhenglun Kong, Yize Li, Fanhu Zeng, Lei Xin, Shvat Messica, Xue Lin, Pu Zhao, Manolis Kellis, Hao Tang, and Marinka Zitnik. 2025. Token reduction should go beyond efficiency in generative models—from vision, language to multimodality. *arXiv preprint arXiv:2505.18227*.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.
- Gaotang Li, Ruizhong Qiu, Xiushi Chen, Heng Ji, and Hanghang Tong. 2025a. Beyond log likelihood: Probability-based objectives for supervised fine-tuning across the model capability continuum. *arXiv preprint arXiv:2510.00526*.
- Jijie Li, Li Du, Hanyu Zhao, Bo-wen Zhang, Liangdong Wang, Boyan Gao, Guang Liu, and Yonghua Lin. 2025b. Infinity instruct: Scaling instruction selection and synthesis to enhance language models. *arXiv preprint arXiv:2506.11116*.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and William B Dolan. 2016. A diversity-promoting objective function for neural conversation models. In *Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: human language technologies*, pages 110–119.
- Mengqi Li, Lei Zhao, Anthony Man-Cho So, Ruoyu Sun, and Xiao Li. 2025c. Online sft for llm reasoning: Surprising effectiveness of self-tuning without rewards. *arXiv preprint arXiv:2510.18814*.
- Ming Li, Yong Zhang, Shwai He, Zhitao Li, Hongyu Zhao, Jianzong Wang, Ning Cheng, and Tianyi Zhou. 2024a. Superfiltering: Weak-to-strong data filtering for fast instruction-tuning. *arXiv preprint arXiv:2402.00530*.
- Ziniu Li, Congliang Chen, Tian Xu, Zeyu Qin, Jiancong Xiao, Zhi-Quan Luo, and Ruoyu Sun. 2024b. Preserving diversity in supervised fine-tuning of large language models. *arXiv preprint arXiv:2408.16673*.
- Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023. Let’s verify step by step. In *The Twelfth International Conference on Learning Representations*.
- Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988.
- Zhenghao Lin, Zhibin Gou, Yeyun Gong, Xiao Liu, Yelong Shen, Ruochen Xu, Chen Lin, Yujiu Yang, Jian Jiao, Nan Duan, and 1 others. 2024. Rho-1: Not all tokens are what you need. *arXiv preprint arXiv:2404.07965*.
- Aiwei Liu, Haoping Bai, Zhiyun Lu, Yanchao Sun, Xiang Kong, Simon Wang, Jiulong Shan, Albin Madappally Jose, Xiaojiang Liu, Lijie Wen, and 1 others. 2024. Tis-dpo: Token-level importance sampling for direct preference optimization with estimated weights. *arXiv preprint arXiv:2410.04350*.
- Wei Liu, Ruochen Zhou, Yiyun Deng, Yuzhen Huang, Junteng Liu, Yuntian Deng, Yizhe Zhang, and Junxian He. 2025. [Learn to reason efficiently with adaptive length-based reward shaping](#). *Preprint*, arXiv:2505.15612.
- Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori B Hashimoto. 2025. s1: Simple test-time scaling. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 20286–20332.
- Team OLMo, Pete Walsh, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Shane Arora, Akshita Bhagia, Yuling Gu, Shengyi Huang, Matt Jordan, and 1 others. 2024. 2 olmo 2 furious. *arXiv preprint arXiv:2501.00656*.

- Jinlong Pang, Na Di, Zhaowei Zhu, Jiaheng Wei, Hao Cheng, Chen Qian, and Yang Liu. 2025. Token cleaning: Fine-grained data selection for llm supervised fine-tuning. *arXiv preprint arXiv:2502.01968*.
- Xiaohan Qin, Xiaoxing Wang, Ning Liao, Cancheng Zhang, Xiangdong Zhang, Mingquan Feng, Jingzhi Wang, and Junchi Yan. 2025. sstoken: Self-modulated and semantic-aware token selection for llm fine-tuning. *arXiv preprint arXiv:2510.18250*.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. 2024. Gpqa: A graduate-level google-proof q&a benchmark. In *First Conference on Language Modeling*.
- Zhiwen Ruan, Yixia Li, He Zhu, Yun Chen, Peng Li, Yang Liu, and Guanhua Chen. 2025. Enhancing large language model reasoning via selective critical token fine-tuning. *arXiv preprint arXiv:2510.10974*.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, and 1 others. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Wentao Shi and Yiqing Shen. 2025. Reinforcement fine-tuning for reasoning towards multi-step multi-source search in large language models. *arXiv preprint arXiv:2506.08352*.
- Chengbing Wang, Yang Zhang, Wenjie Wang, Xiaoyan Zhao, Fuli Feng, Xiangnan He, and Tat-Seng Chua. 2025a. Think-while-generating: On-the-fly reasoning for personalized long-form generation. *arXiv preprint arXiv:2512.06690*.
- Chengbing Wang, Wuqiang Zheng, Yang Zhang, Fengbin Zhu, Junyi Cheng, Yi Xie, Wenjie Wang, and Fuli Feng. 2026. Perm: Psychology-grounded empathetic reward modeling for large language models. *arXiv preprint arXiv:2601.10532*.
- Shaobo Wang, Xiangqi Jin, Ziming Wang, Jize Wang, Jiajun Zhang, Kaixin Li, Zichen Wen, Zhong Li, Conghui He, Xuming Hu, and 1 others. 2025b. Data whispherer: Efficient data selection for task-specific llm fine-tuning via few-shot in-context learning. *arXiv preprint arXiv:2505.12212*.
- Shenzhi Wang, Le Yu, Chang Gao, Chujie Zheng, Shixuan Liu, Rui Lu, Kai Dang, Xionghui Chen, Jianxin Yang, Zhenru Zhang, and 1 others. 2025c. Beyond the 80/20 rule: High-entropy minority tokens drive effective reinforcement learning for llm reasoning. *arXiv preprint arXiv:2506.01939*.
- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, and 1 others. 2024. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark. *Advances in Neural Information Processing Systems*, 37:95266–95290.
- Sean Welleck, Ilia Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. 2019. Neural text generation with unlikelihood training. *arXiv preprint arXiv:1908.04319*.
- Taiqiang Wu, Jiahao Wang, Zhe Zhao, and Ngai Wong. 2024. Mixture-of-subspaces in low-rank adaptation. *arXiv preprint arXiv:2406.11909*.
- Taiqiang Wu, Runming Yang, Jiayi Li, Pengfei Hu, Yik-Chung Wu, Ngai Wong, and Yujiu Yang. 2025a. [Shadow-ft: Tuning instruct model via training on paired base model](#). *Preprint*, arXiv:2505.12716.
- Yongliang Wu, Yizhou Zhou, Zhou Ziheng, Yingzhe Peng, Xinyu Ye, Xinting Hu, Wenbo Zhu, Lu Qi, Ming-Hsuan Yang, and Xu Yang. 2025b. On the generalization of sft: A reinforcement learning perspective with reward rectification. *arXiv preprint arXiv:2508.05629*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025a. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Chun-Hao Yang, Bo-Han Feng, Tzu-Yuan Lai, Yan Yu Chen, Yin-Kai Dean Huang, and Shou-De Lin. 2025b. Training llms beyond next token prediction—filling the mutual information gap. *arXiv preprint arXiv:2511.00198*.
- Zheng Yuan, Hongyi Yuan, Chengpeng Li, Guanting Dong, Keming Lu, Chuanqi Tan, Chang Zhou, and Jingren Zhou. 2023. [Scaling relationship on learning mathematical reasoning with large language models](#). *Preprint*, arXiv:2308.01825.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. Hellaswag: Can a machine really finish your sentence? In *Proceedings of the 57th annual meeting of the association for computational linguistics*, pages 4791–4800.
- Bolin Zhang, Jiahao Wang, Qianlong Du, Jiajun Zhang, Zhiying Tu, and Dianhui Chu. 2025. A survey on data selection for llm instruction tuning. *Journal of Artificial Intelligence Research*, 83.
- Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyang Luo, Zhangchi Feng, and Yongqiang Ma. 2024. Llamafactory: Unified efficient fine-tuning of 100+ language models. *arXiv preprint arXiv:2403.13372*.
- Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, and 1 others. 2023a. Lima: Less is more for alignment. *Advances in Neural Information Processing Systems*, 36:55006–55021.
- Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. 2023b. Instruction-following evaluation for large language models. *arXiv preprint arXiv:2311.07911*.

Appendix

A Proof of Theorem 1

Proof A.1 Consider the loss $\ell(\theta)$ for the target token y_t^* conditioned on the input x and history $y_{<t}^*$. Let $z \in \mathbb{R}^{|\mathcal{V}|}$ be the logit vector and $p = \text{softmax}(z)$ be the probability distribution. Recall from Eq. 2 that the gradient of the loss with respect to the logits is $\nabla_z \ell = p - \mathbf{e}_{y_t^*}$, where $\mathbf{e}_{y_t^*}$ is the one-hot vector corresponding to the target.

First, we analyze the ℓ_2 -norm of the logit gradient $\nabla_z \ell$. Let $p_{y_t^*} = p_\theta(y_t^* | x, y_{<t}^*)$ denote the predicted probability of the ground-truth token. We derive:

$$\begin{aligned} \|\nabla_z \ell\|_2 &= \sqrt{(p_{y_t^*} - 1)^2 + \sum_{v \neq y_t^*} p_v^2} \\ &= \sqrt{(1 - p_{y_t^*})^2 + \sum_{v \neq y_t^*} p_v^2} \\ &\geq 1 - p_{y_t^*}. \end{aligned} \quad (7)$$

Next, applying the chain rule, the gradient with respect to the model parameters θ is given by:

$$\nabla_\theta \ell = \left(\frac{\partial z}{\partial \theta} \right)^\top \nabla_z \ell = J_\theta(z)^\top \nabla_z \ell, \quad (8)$$

where $J_\theta(z) \in \mathbb{R}^{|\mathcal{V}| \times |\theta|}$ is the Jacobian matrix. We utilize the spectral inequality for matrix-vector products: for any matrix \mathbf{A} and vector \mathbf{v} , $\|\mathbf{A}\mathbf{v}\|_2 \geq \sigma_{\min}(\mathbf{A})\|\mathbf{v}\|_2$. Applying this to our context:

$$\begin{aligned} \|\nabla_\theta \ell\|_2 &= \|J_\theta(z)^\top \nabla_z \ell\|_2 \\ &\geq \sigma_{\min}(J_\theta(z)^\top) \cdot \|\nabla_z \ell\|_2 \\ &= \sigma_{\min}(J_\theta(z)) \cdot \|\nabla_z \ell\|_2. \end{aligned} \quad (9)$$

Here, we use the property that a matrix and its transpose share the same singular values. Finally, substituting Eq. 7 and the non-degeneracy assumption $\sigma_{\min}(J_\theta(z)) \geq \gamma$, we obtain:

$$\|\nabla_\theta \ell\|_2 \geq \gamma \cdot (1 - p_\theta(y_t^* | x, y_{<t}^*)). \quad (10)$$

This concludes the proof.

B Hyper parameters

Regarding the hyperparameter configuration, we set the per-device batch size to 1 and employed a gradient accumulation strategy with 4 steps. The input sequences were truncated to a maximum length

of 8,192 tokens. All models were fine-tuned for a single epoch.

During inference, we enable stochastic sampling across all models to ensure generation diversity. For Qwen3 and OLMo-2, we configure the hyperparameters with : do_sample=True, temperature=0.7, top_p=0.8, top_k=20. For the Llama series, we adopt a configuration : do_sample=True, temperature=0.6, top_p=0.9. Regarding generation length, we set the maximum output tokens to 32,768 for the AIME'24 benchmark to accommodate extensive reasoning, while limiting it to 8,192 for all other evaluations.

C Detailed Analysis of LoRA Rank across Datasets

We present the detailed performance trajectories for each dataset (GSM8K, MATH-500, GPQA-Diamond, and IFEval) across varying LoRA ranks in Figure 8. The results reveal consistent optimization patterns across diverse tasks:

Universal Monotonicity for Core Tokens ($p > 0.1$). A striking commonality across **all four datasets** is the monotonic performance growth observed when training on core tokens (pink lines). Whether for reasoning (GSM8K, MATH-500), knowledge recall (GPQA), or instruction following (IFEval), the model's performance steadily improves as the rank increases from 4 to 1024. This universality confirms a fundamental hypothesis: **core task semantics possess a high intrinsic dimensionality**. Consequently, providing larger parameter capacities allows the model to better resolve these critical patterns without saturation, regardless of the task type.

Interference from Non-Core Expressions ($p < 0.1$). In contrast, focusing on non-core tokens (green lines) consistently leads to suboptimal outcomes, though the manifestation varies slightly by task:

- **Reasoning Tasks (GSM8K, MATH-500):** These tasks are particularly sensitive to noise. At high ranks (e.g., $r = 1024$), the excess capacity leads to severe overfitting on non-core expressions, causing a sharp performance drop.
- **Knowledge & Instruction Tasks (GPQA, IFEval):** The standard SFT baseline (blue




Figure 8: Detailed performance comparison across different LoRA ranks for individual datasets. The trends corroborate our hypothesis: core tokens ($p > 0.1$) benefit from increased capacity, while non-core tokens ($p < 0.1$) induce optimization interference, particularly at high ranks.

lines) and non-core token training often exhibit fluctuations or stagnation at medium-to-high ranks, struggling to maintain the upward trajectory seen in the core token setting. This further highlights that filtering out non-core expressions is essential for stable scaling.

D Detailed Analysis of Training Dynamics across Epochs

We further examine the training stability and convergence speed on individual datasets, as illustrated in Figure 9. The trajectories provide granular evidence for the efficiency of our method:

Rapid Convergence and Stability. Across all benchmarks, ProFit settings (solid lines, $p > \tau$) demonstrate superior convergence efficiency, typically reaching near-optimal performance within just 2 epochs and maintaining stability throughout the training process. In contrast, the SFT baseline

often requires more steps to plateau or exhibits fluctuations.

Overfitting to Non-Core Expressions. The risks of training on low-probability tokens are clearly visible in reasoning tasks. For instance, in GSM8K and MATH-500, the performance of the $p < 0.1$ setting (yellow line) peaks early but subsequently degrades as training progresses (Epoch 3-5). This *inverted-U* pattern strongly suggests that prolonged exposure to low-probability tokens leads the model to overfit to non-core expressions, thereby impairing its underlying reasoning logic.

The Gap in Instruction Following. In IFEval, a distinct performance chasm is observed: models trained on high-probability tokens stabilize around 50% accuracy, whereas those focused on low-probability tokens stagnate below 30%. This indicates that the core semantics required for instruction following are almost exclusively encoded

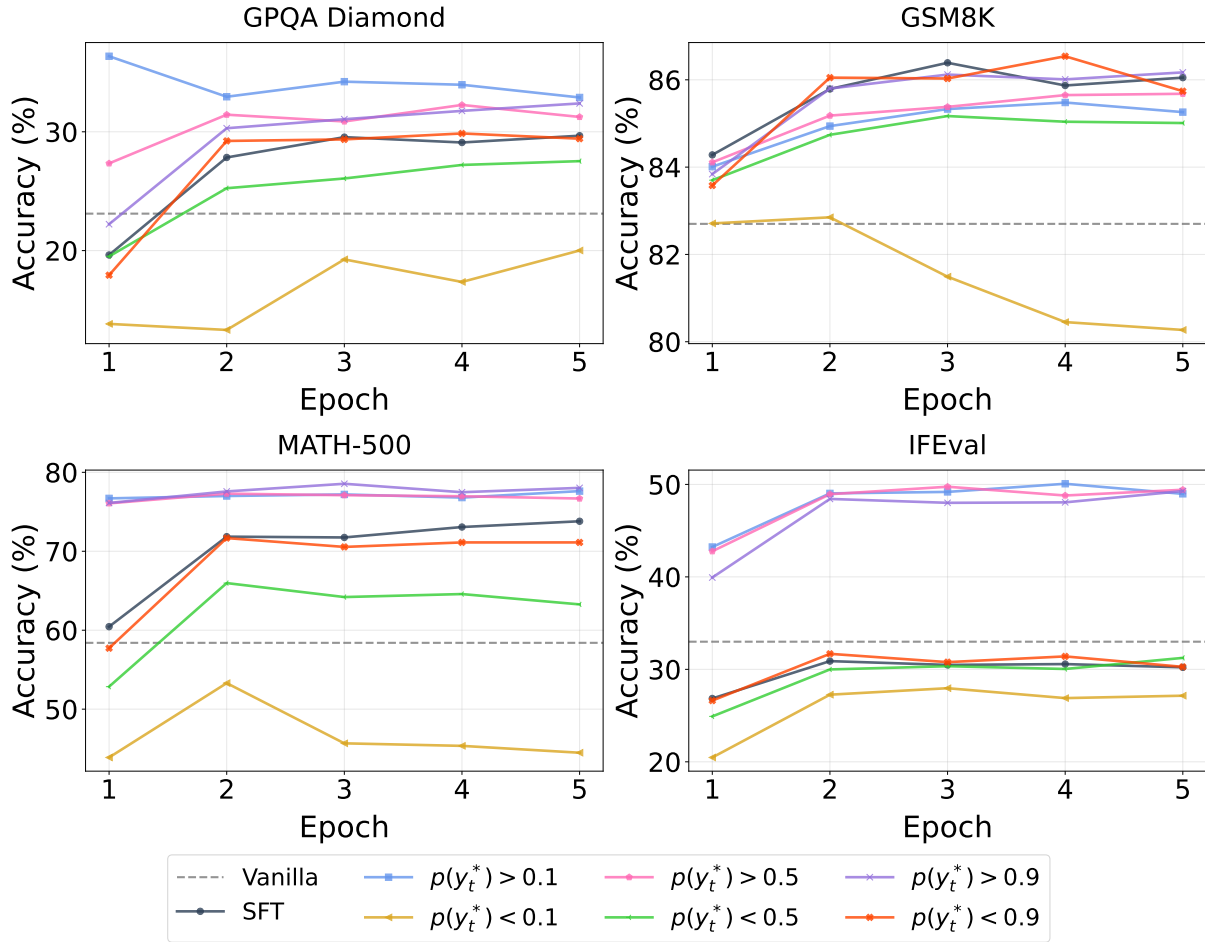


Figure 9: Performance evolution across training epochs for individual datasets. ProFit ($p > \tau$) exhibits rapid convergence and stability, whereas training on low-probability tokens ($p < \tau$) suffers from instability and overfitting to non-core expressions.

in high-probability tokens, while low-probability tokens contribute little to this capability.

E Case Study

As shown in Table 4, SFT succumbs to logical hallucinations by ignoring key interaction terms, whereas ProFit maintains a coherent reasoning chain to derive the correct solution. This demonstrates that masking non-core expressions effectively safeguards the model’s core logic against superficial errors.

F Benchmark Details

We evaluate our method on diverse benchmarks covering reasoning and instruction following capabilities:

- **GPQA-Diamond** (Rein et al., 2024): A dataset of 198 expert-level, *Google-proof* science questions testing advanced reasoning.

- **IFEval** (Zhou et al., 2023b): Evaluates the model’s ability to follow verifiable formatting instructions and constraints.
- **GSM8K** (Cobbe et al., 2021): A classic benchmark consisting of grade-school math word problems requiring multi-step reasoning.
- **MATH-500** (Lightman et al., 2023): A challenging subset of 500 competition-level mathematics problems from the Minerva dataset.
- **AIME’24** (American Institute of Mathematics, 2024): A set of high-difficulty problems from the 2024 American Invitational Mathematics Examination, testing frontier mathematical capabilities.
- **Minerva** (Hendrycks et al., 2021): A comprehensive dataset of 12,500 challenging competition mathematics problems ranging from




Figure 10: Training Dynamics in RL Stage. We compare the KL Divergence, Entropy, and Response Length of models initialized with Base, DFT, and ProFit strategies. ProFit demonstrates superior stability (low KL), confident convergence (low Entropy), and evolves deeper reasoning capabilities (highest Response Length).

pre-algebra to calculus, serving as a standard benchmark for mathematical reasoning.

- **OlympiadBench** (He et al., 2024): A large-scale, bilingual, and multimodal benchmark featuring Olympiad-level problems in mathematics and physics, designed to evaluate AGI capabilities in complex scientific reasoning.
- **HellaSwag** (Zellers et al., 2019): A benchmark for evaluating commonsense natural language inference, requiring models to choose the most logical and natural continuation of a given text snippet.
- **ARC-c** (Clark et al., 2018): The Challenge set of the AI2 Reasoning Challenge, consisting of difficult grade-school science questions that require advanced logic and reasoning beyond simple information retrieval.
- **HumanEval** (Chen et al., 2021): A benchmark designed to evaluate code generation capabilities, featuring programming problems (primarily in Python) that test algorithmic logic and basic coding skills.
- **MMLU-Pro** (Wang et al., 2024): An enhanced, more challenging version of the Massive Multitask Language Understanding (MMLU) benchmark. In our evaluation, we specifically report performance on the **History** and **Chemistry** subsets to assess advanced, domain-specific knowledge and reasoning.

G Training Dynamics Analysis in RL Stage

To further investigate the impact of different initialization strategies on the reinforcement learning (RL) process, we visualized the training dynamics of three key metrics: KL Divergence, Entropy, and Response Length. The comparative results are presented in Figure 10.

KL Divergence Stability. As shown in the left panel of Figure 10, the Base model (blue line) exhibits a rapid and uncontrolled increase in KL divergence, reaching approximately 0.17 by the end of training. This sharp rise suggests that without a robust SFT warm-up, the policy drifts significantly from the reference model, potentially leading to reward hacking or language degeneration. In contrast, **ProFit** (red line) and **DFT** (green line) maintain a remarkably low and stable KL divergence (staying below 0.05). This indicates that ProFit effectively constrains the policy update within a safe trust region, ensuring that the model improves its mathematical reasoning capabilities while preserving its general linguistic coherence.

Entropy. The entropy curves reveal distinct convergence behaviors. The Base model demonstrates chaotic behavior, where entropy initially collapses and then drastically rebounds to extremely high values (> 6.0), indicating a failure to converge to a stable policy and likely degenerating into generating high-randomness noise. Conversely, **ProFit** maintains a consistently low entropy (ending around 0.09), similar to DFT. This low entropy signifies high confidence in the generated reasoning paths.

Notably, unlike DFT which saturates near 0.03, ProFit maintains a slightly higher entropy margin, suggesting it retains a healthy level of exploration potential while remaining focused on high-reward solutions.

Response Length. The most significant differentiator lies in the response length, which serves as a proxy for the depth of Chain-of-Thought (CoT) reasoning. While all models show an increasing trend in response length, **ProFit** demonstrates the most robust growth pattern. Starting from $\sim 1,400$ tokens, ProFit rapidly learns to expand its reasoning steps, surpassing the Base model around step 150 and eventually achieving the highest average response length of over **8,000 tokens** at the end. Compared to DFT (ending at 7,100 tokens) and Base (6,400 tokens), ProFit's superior length indicates that the probability-guided initialization encourages the RL algorithm to explore deeper and more complex reasoning chains. This aligns perfectly with our main results, where ProFit excels in complex tasks like OlympiadBench, which require extensive multi-step deductions.

Question	The polynomial $x^3 - 3x^2 + 4x - 1$ is a factor of $x^9 + px^6 + qx^3 + r$. Enter the ordered triple (p, q, r) . Please reason step by step, and put your final answer within \square .
Reference	Let α be a root of $x^3 - 3x^2 + 4x - 1 = 0$, so $\alpha^3 = 3\alpha^2 - 4\alpha + 1$. Substituting this relation repeatedly into higher powers: $\alpha^9 = -63\alpha^2 + 190\alpha - 54$. Then $\alpha^9 + p\alpha^6 + q\alpha^3 + r = (-5p + 3q - 63)\alpha^2 + (-11p - 4q + 190)\alpha + (4p + q + r - 54)$. For this to be 0 for all roots, we set coefficients to 0: $-5p + 3q = 63$, $11p + 4q = 190$, $4p + q + r = 54$. Solving this system yields $(p, q, r) = \square(6, 31, -1)$.
SFT	To determine the ordered triple (p, q, r) , we assume $x^9 + px^6 + qx^3 + r = (x^3 - 3x^2 + 4x - 1) \cdot g(x)$, where $g(x) = ax^6 + bx^5 + cx^4 + dx^3 + ex^2 + fx + g$. Step 4: Match coefficients Let's expand the product step-by-step: <ul style="list-style-type: none"> • The x^9 term: $x^3 \cdot ax^6 = ax^9 \implies a = 1$. • The x^8 term: $x^3 \cdot bx^5 = bx^8 \implies b = 0$ (Incorrectly ignores interaction terms). • The x^7 term: $x^3 \cdot cx^4 = cx^7 \implies c = 0$. • The x^6 term: Combining terms yields $(d - 3)x^6 - cx^4$. We match this to p: $d - 3 = p$ and $c = 0$. • The x^5 term: Combining terms yields $ex^5 - dx^3$. We need this to match 0: $e = 0$ and $d = 0$. • The x^4 term: Combining terms yields fx^4. We need this to match 0: $f = 0$. • The constant term: $-g = r \implies g = r$. Step 5: Combine the results From the above steps, we have: $a = 1, b = 0, c = 0, d = 0, e = 0, f = 0, g = 0, r = 0$. Thus: $p = d - 3 = 0 - 3 = -3, q = g = 0, r = 0$. Final Answer: $\square(-3, 0, 0)$
ProFit	To determine the ordered triple (p, q, r) , we denote $f(x) = x^3 - 3x^2 + 4x - 1$. We assume $g(x) = x^6 + ax^5 + bx^4 + cx^3 + dx^2 + ex + f$. Then: $(x^3 - 3x^2 + 4x - 1)(x^6 + ax^5 + \dots + f) = x^9 + px^6 + qx^3 + r$. Step-by-step Coefficient Comparison: Expanding the LHS and comparing with $x^9 + px^6 + qx^3 + r$: <ol style="list-style-type: none"> 1. x^8 coeff: $a - 3 = 0 \implies a = 3$. 2. x^7 coeff: $b - 3a + 4 = 0 \implies b - 9 + 4 = 0 \implies b = 5$. 3. x^6 coeff: $c - 3b + 4a - 1 = p \implies c - 15 + 12 - 1 = p \implies p = c - 4$. 4. x^5 coeff: $d - 3c + 4b - 1a = 0 \implies d - 3c + 20 - 3 = 0 \implies d = 3c - 17$. 5. x^4 coeff: $e - 3d + 4c - 1b = 0 \implies e - 3(3c - 17) + 4c - 5 = 0 \implies e = 5c - 46$. 6. x^3 coeff: $f - 3e + 4d - 1c = q \implies f - 3(5c - 46) + 4(3c - 17) - c = q \implies q = f - 4c + 70$. 7. x^2 coeff: $-3f + 4e - 1d = 0 \implies \dots \implies f = \frac{17c - 167}{3}$. 8. x^1 coeff: $-f + 4e - 1d = 0 \implies \dots \implies 34c - 340 = 0 \implies c = 10$. 9. Constant: $-f = r \implies r = -1$. Final Calculation: Now, substituting $c = 10$ back into the expressions: $p = c - 4 = 10 - 4 = 6$. $q = f - 4c + 70 = 1 - 40 + 70 = 31$. $r = -1$. Thus, the ordered triple is: $\square(6, 31, -1)$

Table 4: **Comparison of model outputs on the polynomial factorization task.** SFT hallucinates constraints (e.g., setting $b = 0$ and $c = 0$ independently), leading to an incorrect solution. ProFit maintains a coherent logical chain throughout the coefficient comparison, deriving the correct ordered triple.

We are given: A worker deposits $\$20,000$ at the end of each year. He makes 3 deposits: at the end of Year 1, Year 2, and Year 3. By the end of the third year, he wants to have at least $\$66,200$ in the account. The account earns compound interest annually, and we are to find the minimum annual interest rate r (as a percentage) that makes this possible.

We are to find the minimal compound interest rate such that the future value of the three deposits is at least $\$66,200$.

--- Step 1: Future Value of Annuity (End of Year Deposits)

Since the deposits are made at the end of each year, this is an ordinary annuity. The future value FV of an ordinary annuity with annual payment P , interest rate r , and n years is:

$$FV = P \cdot \frac{(1+r)^n - 1}{r}$$

Here: $P = 20,000$, $n = 3$

$$FV \geq 66,200$$

So:

$$20,000 \cdot \frac{(1+r)^3 - 1}{r} \geq 66,200$$

Divide both sides by 20,000:

$$\frac{(1+r)^3 - 1}{r} \geq \frac{66,200}{20,000} = 3.31$$

So we need to solve:

$$\frac{(1+r)^3 - 1}{r} = 3.31$$

We need to find the smallest r (as a decimal) such that this inequality holds.

Let's try some values of r .

--- Try $r = 0.10$ (10%)

$$(1+0.10)^3 = 1.331$$

$$\frac{1.331 - 1}{0.10} = \frac{0.331}{0.10} = 3.31$$

Exactly 3.31! So at $r = 10\%$, the future value is exactly $\$20,000 \times 3.31 = \$66,200$.

--- Check if a lower rate would work?

Try $r = 0.09$ (9%)

$$(1.09)^3 = 1.295029$$

$$\frac{1.295029 - 1}{0.09} = \frac{0.295029}{0.09} \approx 3.2781$$

$$3.2781 < 3.31 \rightarrow \text{Not enough.}$$

So 9% is too low.

At 10%, we get exactly 3.31 \rightarrow meets the requirement.

Thus, 10% is the minimal rate that satisfies the condition.

--- Final Answer: $\boxed{10}$

0.0 (Low Prob)



1.0 (High Prob)

Figure 11: Visualization of token probability