

Towards Understanding Pause Token Fine-Tuning Dynamics: A Mode Retention Perspective

Jaehyeon Kim^{1,*} Suhwan Kim^{2,*} Nakyung Lee² Yeongoon Kim²
Jimin Seo² Giho Lee¹ Jungwoo Lee^{1,2,†}

¹HodooAI Lab, Seoul, South Korea

²Seoul National University, Seoul, South Korea

*Equal contribution. †Corresponding author.

Correspondence: junglee@snu.ac.kr

Abstract

Pause-token methods improve LLM reasoning by inserting special tokens into sequences. Prior work explains these gains through computational expressivity. However, there is relatively little investigation into the training dynamics of pause tokens. We explore how pause tokens reshape the training dynamics of fine-tuning. Two controlled pilots expose distinct asymmetries. On a synthetic continual-learning task, masked pauses overwrite a previously-learned distribution roughly 4× less at matched final adaptation (H1, **mode retention**); on a synthetic math-reasoning probe, the boundary-adjacent token comes to encode substantially more downstream-step information (H2, **non-myopic compression**). We formalize a training rule consistent with both — Masked Boundary Pause (MBP), pause tokens placed at reasoning-step boundaries with their loss masked. Across 1B-8B Qwen and Llama models, MBP consistently improves reasoning, achieving gains of up to 6 points on math and 2.5 points on code, while preserving general language understanding abilities. We further demonstrate that these gains extend to GRPO. These results recast pause tokens as a training-dynamics intervention on the retention–adaptation trade-off, rather than merely an inference-time computation device.

1 Introduction

Recent progress in large language models (LLMs) has produced a wide range of training and test-time methods for improving reasoning, including chain-of-thought prompting, test-time scaling, and reinforcement learning for reasoning (Wei et al., 2022b; Snell et al., 2025; Wang et al., 2023). These methods achieve strong performance on complex tasks but incur substantial costs in generation length, inference latency, sampling budget, or training compute. *Pause-* and *filler-token* methods instead improve performance through a small, architecture-

preserving change: inserting special tokens without modifying the model or optimizer (Goyal et al., 2024; Burtsev et al., 2021).

Pause-token gains have been attributed to forward-pass computation and expressivity (Pfau et al., 2024; London and Kanade, 2025), focal-loss amplification at hard tokens (Kim et al., 2025), and structural anchoring at sentence boundaries (Liu et al., 2026). These explanations center on effects at the inserted positions themselves. Yet much less is known about how pause training shapes the pre-trained distribution from which fine-tuning starts, even though this distribution is central to modern adaptation pipelines that increasingly warm-start from existing checkpoints (Gururangan et al., 2020; Roziere et al., 2023; Azerbayev et al., 2024).

We argue that this is precisely where an under-examined role of pause tokens is most consequential. Pause tokens also participate in the forward pass during fine-tuning, where they alter the context under which target tokens are predicted and reshape gradient flow through the surrounding positions (Wu et al., 2024; Pfau et al., 2024). These training-dynamics effects are present even when pause tokens are absent at inference and never appear as supervised targets, and they directly modulate the retention–adaptation trade-off central to fine-tuning. In a controlled continual-learning setting, we observe that loss-masked pauses reduce overwriting of a previously learned distribution by roughly 4× at matched final adaptation, revealing a separation that neither inference-time nor low-confidence framings predict.

We investigate this through two controlled pilots that yield distinct empirical phenomena: a *retention asymmetry* on synthetic continual learning and an *information-packing asymmetry* on iGSM (Ye et al., 2025) probing, which we elevate to working hypotheses H1 (mode retention) and H2 (non-myopic compression). We then formalize a training rule consistent with both hypotheses, **Masked Bound-**




Figure 1: Comparison of standard fine-tuning and masked boundary pause fine-tuning. Standard SFT applies next-token prediction (NTP) loss at every task-token position, allowing gradients from the target token to flow directly into the preceding context representation. This can strongly adapt the model to the task-data distribution, but may also erode the pretrained predictive distribution. Masked boundary pause fine-tuning (MBP) instead inserts loss-masked pause tokens at selected boundaries, blocking the direct gradient path into the boundary context state. Neighboring token predictions still provide indirect updates, enabling task adaptation while better preserving the pretrained predictive distribution.

ary Pause (MBP)¹, and test the hypotheses’ scaled predictions across 1B–8B fine-tuning and GRPO. The headline finding is that *boundary placement and loss masking are jointly necessary* for the gains, with specific results reported in §5–§7.3.

In summary, our contributions are:

- Through controlled pilots, we identify two hypotheses: (H1) mode retention, and (H2) non-myopic compression.
- Motivated by these, we propose **Masked Boundary Pause (MBP)** fine-tuning, where pause tokens are inserted at sentence or reasoning-step boundaries while masking their cross-entropy loss.
- Across 1B–8B models, MBP improves math and code reasoning while preserving general capability; gains extend to GRPO. A four-way placement ablation shows that boundary placement and loss masking are jointly necessary.

2 Preliminaries

We formalize masked-pause fine-tuning under a single notation that subsequent sections re-use, and

¹Code is available at <https://github.com/jkpizza/emnlp2026-pause-token-dynamics>.

locate prior pause-token methods within the same framework.

2.1 Pause-token training

Formulation. Let (x, y) be a (prompt, response) pair drawn from a fine-tuning dataset $\mathcal{D}$, with response $y = (y_1, \dots, y_N)$. Standard supervised fine-tuning (SFT) minimizes the per-token conditional cross-entropy of y given x ; we recover it as the $g(x, y) = \emptyset$ case of the formulation below.

A pause-token training method is parameterized by a *position selector* g that returns insertion indices $g(x, y) = \{i_1 < i_2 < \dots < i_M\} \subseteq \{0, \dots, N\}$. With the conventions $i_0 = 0$ and $i_{M+1} = N$, these indices partition y into $M + 1$ contiguous segments

$$u_n = (y_{i_{n-1}+1}, \dots, y_{i_n}), \quad n = 1, \dots, M + 1, \quad (1)$$

with $y = u_1 \parallel u_2 \parallel \dots \parallel u_{M+1}$. A distinguished pause token τ is inserted between adjacent segments, yielding the pause-augmented sequence

$$\tilde{y} = u_1 \parallel \tau \parallel u_2 \parallel \tau \parallel \dots \parallel \tau \parallel u_{M+1} \quad (2)$$

of length $N + M$. Let $S_\tau = \{k : \tilde{y}_{k+1} = \tau\}$ be the set of prediction positions whose target is a pause token. Following prior work (Goyal et al., 2024;

Kim et al., 2025), we exclude S_τ from the loss: pause tokens appear in the autoregressive context but never as supervised outputs. The per-example training loss is

$$\mathcal{L}^g(\theta) = -\frac{1}{N} \sum_{k \notin S_\tau} \log p_\theta(\tilde{y}_{k+1} \mid x, \tilde{y}_{1:k}), \quad (3)$$

averaged over $(x, y) \sim \mathcal{D}$. Setting $g(x, y) = \emptyset$ recovers standard SFT.

Relation to existing work. Prior work differs in the position selector g : APPEND (Goyal et al., 2024; Pfau et al., 2024), DIT (Kim et al., 2025), and BOUNDARY at sentence endings (Liu et al., 2026). We additionally test RANDOM placement as an unstructured control. Of these, Goyal et al. (2024) and Kim et al. (2025) exclude pause positions from the loss as in Eq. 3, whereas Liu et al. (2026) trains pauses as predicted targets. Our method, Masked Boundary Pause Training (§5), takes $g = \text{BOUNDARY}$ with S_τ masked, and Section 6.1 compares the four placements.

2.2 GRPO

Formulation. Group Relative Policy Optimization (GRPO) (Shao et al., 2024) optimizes a policy π_θ on prompts $q \in \mathcal{Q}$ by sampling G responses $\{y_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot \mid q)$ from the previous policy, scoring each with a reward $r_i = R(q, y_i)$, and computing the group-relative advantage $A_i = (r_i - \mu_r) / (\sigma_r + \epsilon)$ where μ_r, σ_r are the group’s mean and standard deviation of rewards. With per-token policy ratio $\rho_{i,t}(\theta) = \pi_\theta(y_{i,t} \mid q, y_{i,<t}) / \pi_{\theta_{\text{old}}}(y_{i,t} \mid q, y_{i,<t})$, GRPO maximizes

$$\mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E} \left[\frac{1}{G} \sum_{i,t} c_{i,t}(\theta) - \beta D_{\text{KL}} \right], \quad (4)$$

$$c_{i,t}(\theta) = \min(\rho_{i,t} A_i, \text{clip}(\rho_{i,t}, 1-\epsilon, 1+\epsilon) A_i).$$

Here, D_{KL} is a per-token KL penalty against a reference policy.

3 Pilot Experiments

The introduction posed a mechanistic question: what is masked-pause fine-tuning actually doing to a pretrained model’s representations? We approach this in two controlled studies. The first asks whether masked pauses change *what gets overwritten* during adaptation (§3.1). The second asks *what*

gets encoded at the boundary-adjacent token (§3.2). Each pilot yields a distinct empirical asymmetry that we crystallize into a working hypothesis.

3.1 Modular continual learning: a retention asymmetry

Our first pilot asks whether masked pauses change what gets overwritten during adaptation. We use the simplest setting that can support a clean before/after comparison, a synthetic continual-learning task with two disjoint regimes. Detailed experimental settings are provided in Appendix A; Appendix A.2 extends the pilot to a four-regime chain.

Setup. Sequences are generated by a *Modular Linear Chain* (MLC),

$$x_{t+1} = (a_t x_t + b_t) \bmod p, \quad (5)$$

where a *regime* is defined by a length- L coefficient sequence $\{(a_t, b_t)\}_{t=1}^L$. We sample two disjoint coefficient sets to define regimes D_1 and D_2 , and train a small GPT-style Transformer sequentially: Phase 1 trains on D_1 until accuracy exceeds 0.99; Phase 2 forks from a single Phase-1 checkpoint and continues on D_2 under two matched conditions sharing architecture, optimizer, batch size, learning rate, and update budget:

- NO PAUSE: standard next-token prediction on raw D_2 sequences.
- PAUSE: one pause token inserted after every content token; pause-target positions are excluded from the loss.

Results. Figure 2 shows the Phase-2 loss curves on both regimes. Both conditions adapt to D_2 to the same $\sim 10^{-7}$ floor, with NO PAUSE marginally faster in descent but indistinguishable at convergence. The conditions diverge sharply in how much they overwrite D_1 : NO PAUSE plateaus near 6 nats of held-out D_1 loss while PAUSE plateaus near 1.5 — a $\sim 4\times$ retention gap, consistent across seeds. The gap occurs at *matched* D_2 fit and is not explained by the number of tokens processed per step.

Hypothesis. Masked-pause training matches the new regime’s final fit while preserving the previous regime substantially better:

H1 (mode retention). *Masked-pause training reduces overwriting*




Figure 2: Phase-2 dynamics from a D_1 -pretrained backbone ($n=8$ seeds, $\pm 1\sigma$). **Top:** D_2 adaptation (held-out loss, log scale). **Bottom:** D_1 retention (held-out loss, linear scale). Both conditions reach the same D_2 floor; PAUSE retains $D_1 \sim 4\times$ better than NO PAUSE.

of previously-learned distributions during adaptation, relative to standard fine-tuning at matched final performance.

3.2 iGSM step probing: an information-packing asymmetry

The MLC pilot shows masked pauses change adaptation dynamics, but not what the representation absorbs. To localize the mechanism, we probe what the *boundary-adjacent token* — the period immediately preceding an inserted [PAUSE] — encodes after training. We use iGSM (Ye et al., 2025), a synthetic mathematical reasoning dataset where each reasoning step concludes with a deterministic integer answer, providing the clean per-step probe targets this analysis requires.

Setup. We fine-tune Qwen3-0.6B-Base (Yang et al., 2025) on 3K iGSM problems under two conditions — NO PAUSE and PAUSE (one [PAUSE] inserted at each step boundary, loss masked) — and compare against the pretrained BASE checkpoint ($n=3$ seeds each). For each held-out problem we extract the final-layer hidden state at the boundary-adjacent period of every reasoning step and fit two linear probes: one for the integer answer of the *current* step, one for the *next* step. Both fine-tuned conditions reach comparable held-out task accuracy, so the probe compares representations at matched task fit.




Figure 3: Linear-probe accuracy on the boundary-adjacent token’s final hidden state ($n=3$ seeds, $\pm 1\sigma$). **Left:** current step’s integer answer. **Right:** next step’s integer answer. PAUSE (UNMASKED) trains the same boundary pauses as prediction targets.

Results. Figure 3 shows masked-pause SFT yielding the highest decodability on both targets: current-step accuracy is roughly twice that under NO PAUSE (0.282 vs. 0.133, and clearly above the base 0.182), and next-step accuracy is highest as well (0.104 vs. 0.062 for no pause and 0.075 for base). Training the same boundary pauses as prediction targets preserves most of the effect (0.270 and 0.094), indicating that the compression follows primarily from pause placement, with masking a small additional increment. The pattern is consistent across seeds.

Hypothesis. Under ordinary SFT, the period token must predict the first token of the next reasoning step; inserting a pause replaces that target with [PAUSE], and masking additionally removes its loss. Either way, the period’s representation is freed from immediate next-token specialization and can instead absorb gradient from downstream positions:

H2 (non-myopic compression). *Boundary pauses redirect gradient flow such that boundary-adjacent tokens encode more downstream information than under ordinary SFT.*

3.3 From hypotheses to predictions

Motivated by H1 and H2, the rest of the paper tests three predictions that follow if they hold at scale on real reasoning tasks:

- **RQ1 (scaling).** Does masked boundary pause fine-tuning at 1B–8B improve math and code reasoning while preserving general capability (MMLU (Hendrycks et al., 2021a), GPQA (Rein et al., 2024), BBH (Suzgun et al., 2023)) better than standard fine-tuning? H1 predicts this trade-off.

- **RQ2 (placement).** Does pause placement affect performance at matched pause-token budget? If so, which strategy is optimal? H1 and H2 both act locally at the inserted positions, predicting that placement matters and that boundary insertion outperforms random, prompt-append (Goyal et al., 2024), and low-confidence (Kim et al., 2025) alternatives.
- **RQ3 (RL exploration).** Do GRPO rollouts under MBP retain wider behavioral support than standard GRPO — more per-prompt reward variance, slower distribution sharpening, and improved convergence? H1 predicts so.

Sections 5–8 formalize MBP and answer these questions via 1B–8B math/code SFT with OOD-retention measurement, and GRPO experiments.

4 A retention certificate for masked-pause SFT

H1 admits a first-order retention certificate. The intuition: old D_1 predictions on no-pause prefixes survive D_2 fine-tuning when their top-token margins outlast the logit movement induced by adaptation. Bounding this movement uniformly over a protected prefix set converts retention into a comparison between the certified drifts of two procedures. Let $\Phi(r)$ denote the fraction of protected D_1 prefixes whose original margin exceeds r .

Theorem 4.1 (Retention certificate, informal). *In the local linearized regime around the D_1 -trained parameters, with both procedures matched at the same D_2 fitting level, let δ_O be the certified protected drift of ordinary SFT on no-pause D_1 logits. Masked-pause SFT has certified protected drift at most $\beta_{\text{ctx}}\delta_O$, where β_{ctx} measures how separated the pause-perturbed prefix is from the ordinary prefix in the Jacobian sense. Consequently,*

$$\begin{aligned}\text{Retain}(W_{\text{SFT}}) &\geq \Phi(2\delta_O), \\ \text{Retain}(W_{\text{pause}}) &\geq \Phi(2\beta_{\text{ctx}}\delta_O).\end{aligned}$$

Since Φ is non-increasing, whenever $\beta_{\text{ctx}} < 1$, masked-pause SFT inherits a stronger retention guarantee than ordinary SFT.

The condition $\beta_{\text{ctx}} < 1$ has a geometric reading: masked-pause SFT fits D_2 with less drift on D_1 no-pause logits whenever the pause-perturbed state is sufficiently separated from the ordinary state. This separation is structural rather than architectural — a property of *where* pauses are inserted,

which the boundary selector of MBP encodes by design. The full setup, lemma statements (margin stability, protected drift bounds), and proofs are in Appendices B and C.

5 Method

Motivated by H1 and H2 in §3, we formalize the masked-pause training strategy predicted to perform optimally.

5.1 Masked Boundary Pause Training

MBP specializes the position selector g from §2 to a *boundary selector* g_{Boundary} that produces the segments u_n of Eq. 1 that are *semantic units*: sentences, reasoning steps, or code lines. Further details can be found in Appendix E.2.

Boundary placement is the advantageous choice for both hypotheses:

- **H1 (mode retention).** Step boundaries are by construction where semantic transitions occur, and we expect fine-tuning’s structural learning to concentrate there. A pause embedding at these positions absorbs that learning, preserving the pretrained distribution that would otherwise drift (§3.1).
- **H2 (non-myopic compression).** Inserted pause positions encourage implicit summarization at the preceding token (§3.2), so semantic boundaries are the natural placement: there the compression encapsulates a complete semantic unit, rather than being forced to predict the next token mid-sentence.

The MBP training loss is Eq. 3 with $g = g_{\text{Boundary}}$, and all placement baselines in §6.1 use the same pause-token budget M so that comparisons isolate the role of g .

5.2 Extension to GRPO

MBP extends to GRPO without changing rollouts or rewards: responses are sampled and scored on the clean policy, while pauses appear only when computing the policy update. Given a sampled response y_i , we form $\tilde{y}_i$ (Eq. 2) and apply the GRPO objective at non-pause positions, with policy ratio and KL penalty computed under the pause-perturbed prefix $(q, \tilde{y}_{i, < j})$ rather than the clean prefix $(q, y_{i, < j})$. Pauses thus remain outside the action space, perturbing only the policy update’s context — the mechanism through which we test RQ3 (§3.3): whether retention of the pretrained distribution improves GRPO’s exploration.

6 Experiment Setup

We test the predictions of §3.3 in two settings: downstream math/code supervised fine-tuning (§6.1, RQ1 and RQ2) and GRPO-based reasoning training (§6.2, RQ3). Full hyperparameters and data details are in Appendix D.

6.1 Downstream Math/Code SFT

Setup. We fine-tune Qwen3 (1.7B, 4B, 8B Base)(Yang et al., 2025) and Llama-3.1-8B-Instruct (Grattafiori et al., 2024) on math data (83K curated DeepMath-103K (He et al., 2025) subset) and code data (100K OpenCodeInstruct subset (Ahmad et al., 2025)), inserting pauses only on the response side. We apply LoRA (Hu et al., 2022) for 8B models. We evaluate reasoning on math (GSM8K, MATH, AMC23, AIME24, Minerva) (Cobbe et al., 2021; Hendrycks et al., 2021b; Lewkowycz et al., 2022) and code (HumanEval(+), MBPP(+)) (Chen et al., 2021; Austin et al., 2021; Liu et al., 2023) benchmarks, and general-capability retention on MMLU, GPQA, HellaSwag, BBH, PIQA (Hendrycks et al., 2021a; Rein et al., 2024; Zellers et al., 2019; Suzgun et al., 2023; Bisk et al., 2020). We report pass@1 for math and code benchmarks and pass@4 for general-capability benchmarks.

Baselines (matched pause-token budget). SFT (no pauses); APPEND: pauses at the prompt-response boundary (Goyal et al., 2024); RANDOM: pauses at random response positions; DIT: pauses before low-likelihood target tokens (Kim et al., 2025); MBP (ours): pauses at reasoning-step or code-line boundaries. All pause variants are loss-masked except MBP w/o MASKING, which supervises pause tokens as targets (§8.1). Matching the budget isolates placement from token count.

6.2 GRPO Setup

We compare STANDARD GRPO against MBP GRPO on Qwen3-1.7B and 4B, training on 5K-problem subsets of DeepMath-103K (He et al., 2025) and SimpleRLZoo (Zeng et al., 2025). We report pass@1 on the math benchmarks above, and track two training-dynamics metrics: convergence AUC (area under the training reward curve, measuring how quickly reward accumulates), and mixed-group ratio (proportion of rollout groups containing both correct and incorrect rollouts).

Model	Method	GSM8K	MATH500	MATH	AMC23	Minerva	Avg.
Qwen3-1.7B Base	Base	65.13	54.20	54.82	35.00	17.65	45.36
	SFT	82.41	63.40	62.98	27.50	20.59	51.38
	MBP	83.02	61.00	64.56	37.50	23.16	53.85
Qwen3-4B Base	Base	82.56	66.00	67.64	37.50	28.68	56.48
	SFT	91.36	69.20	71.24	45.00	29.78	61.32
	MBP	90.07	73.40	72.32	50.00	30.51	63.26
Qwen3-8B Base	Base	88.40	65.40	65.86	42.50	26.84	57.80
	SFT	92.95	71.60	72.54	52.50	32.72	64.46
	MBP	92.57	75.00	72.84	55.00	35.29	66.14
Llama3.2-1B Instruct	Base	38.97	23.80	25.70	15.00	7.72	22.24
	SFT	50.72	28.00	38.48	10.00	8.09	27.06
	MBP	51.63	31.80	41.66	27.50	14.34	33.39
Llama3.2-3B Instruct	Base	70.43	40.00	43.22	10.00	17.65	36.26
	SFT	76.80	46.40	55.02	27.50	13.97	43.94
	MBP	77.03	48.20	55.92	30.00	19.85	46.20
Llama3.1-8B Instruct	Base	77.48	46.40	48.62	<u>26.50</u>	22.43	44.29
	SFT	<u>84.84</u>	<u>53.20</u>	<u>56.28</u>	17.50	<u>21.69</u>	<u>46.70</u>
	MBP	86.28	58.00	61.02	27.50	22.43	51.05

Table 1: Math-domain results. Average denotes the macro-average over GSM8K, MATH500, MATH, AMC23, and Minerva.

Model	Method	MBPP	MBPP+	HumanEval	HumanEval+	Avg.
Qwen3-1.7B Base	Base	69.05	59.26	58.54	53.05	59.98
	SFT	69.58	62.43	67.07	61.59	65.17
	MBP	73.02	62.43	70.12	65.24	67.70
Qwen3-4B Base	Base	74.34	64.81	76.83	72.56	72.14
	SFT	78.04	67.99	79.27	73.78	74.77
	MBP	79.10	65.08	82.93	78.05	76.29
Qwen3-8B Base	Base	82.80	72.49	88.41	84.15	81.96
	SFT	86.51	75.13	85.98	81.71	82.33
	MBP	87.30	75.66	86.59	82.32	82.97
Llama3.1-8B Instruct	Base	67.46	59.52	67.68	62.20	64.22
	SFT	68.25	60.05	68.29	64.02	65.15
	MBP	68.52	60.05	69.51	64.02	65.53

Table 2: Code-domain results. Average denotes the macro-average over MBPP, MBPP+, HumanEval, and HumanEval+.

7 Results

We answer RQ1 (§7.1), RQ2 (§7.2), and RQ3 (§7.3) in order. Across these results, the findings show the patterns predicted by H1 and H2.

7.1 Scaling and Retention (RQ1)

Reasoning gains scale across model sizes and families. Tables 1 and 2 show that MBP outperforms standard SFT on every model variant we tested. Math gains over SFT range from +1.7 to +6.3 points across Qwen3 (1.7B-8B-Base) and Llama (3.2-1B/3B-Instruct, 3.1-8B-Instruct); code gains range from +0.6 to +2.5 across Qwen3. The improvement is consistent in sign across every model variant, showing that pause training is not a

Trained Domain	Method	MMLU	GPQA	HellaSwag	BBH	PIQA	Avg.
Math	Base	62.95	27.46	38.69	38.84	62.13	46.01
	SFT	62.32	24.78	40.39	42.06	66.10	47.13
	MBP	63.93	27.23	48.32	44.13	67.63	50.25
Code	Base	62.95	29.91	38.00	38.84	62.13	46.37
	SFT	60.19	23.88	36.44	39.15	60.66	44.06
	MBP	63.78	25.67	44.01	42.88	68.01	48.87

Table 3: General-domain results on Qwen3-1.7B-Base. Average denotes the macro-average over MMLU, GPQA, HellaSwag, BBH, and PIQA.

model-specific phenomenon.

General capability is preserved. Table 3 reports MMLU, GPQA, HellaSwag, BBH, and PIQA on Qwen3-1.7B-Base after math and code SFT. Standard SFT modestly improves general capability after math SFT but degrades it after code SFT. By contrast, MBP improves the general-benchmark average over base in both settings (+4.2 after math, +2.5 after code), beating SFT by +3.1 and +4.8, respectively. The MBP-vs-SFT gap on general benchmarks is larger than the gap on the trained-domain benchmarks themselves. This suggests that MBP does not trade reasoning gains for retention loss, but improves both simultaneously.

These findings support H1 at scale: MBP preserves the pretrained distribution while still acquiring new-task reasoning.

7.2 Placement and Masking (RQ2)

H1 and H2 predict that boundary placement activates both mechanisms: placements without alignment to semantic transitions (append, random) or correlated with hard tokens rather than transitions (DIT) should fail to do so. We test this on Qwen3-1.7B-Base at matched pause-token budget.

Boundary placement is the only consistent winner. Table 4 reveals a 7.5-point spread across the four placements on math (from DIT’s -5.0 to MBP’s $+2.5$ relative to SFT), despite identical pause-token budgets. Of the four, only MBP is reliably positive on both math and code domains. The other placements fluctuate: Append and DIT modestly help code ($+1.4$, $+1.7$) but hurt math (-0.9 , -5.0). Random is near-flat on code ($+0.2$) and clearly negative on math (-1.7). The inconsistency suggests that random, append, and low-confidence positions activate H1 and H2’s mechanisms only when an inserted pause happens to coincide with a meaningful boundary, whereas BOUNDARY placement targets these positions by construction and

Method	Math					Avg.
	GSM8K	MATH500	MATH	AMC23	Minerva	
Base	65.13	54.20	54.82	<u>35.00</u>	17.65	45.36
SFT	<u>82.41</u>	63.40	62.98	27.50	20.59	<u>51.38</u>
Append	79.68	59.20	60.96	30.00	<u>22.79</u>	50.53
Random	81.35	60.80	<u>63.26</u>	20.00	23.16	49.71
DIT	68.01	57.20	56.38	30.00	20.22	46.36
MBP (ours)	83.02	61.00	64.56	37.50	23.16	53.85
MBP w/o masking	80.97	<u>62.20</u>	62.64	27.50	23.16	51.29

Method	Code				Avg.
	MBPP	MBPP+	HumanEval	HumanEval+	
Base	69.05	59.26	58.54	53.05	59.98
SFT	69.58	62.43	67.07	61.59	65.17
Append	69.84	60.58	<u>70.73</u>	65.24	66.60
Random	<u>70.37</u>	<u>61.11</u>	67.07	<u>62.80</u>	65.34
DIT	69.84	60.58	71.95	65.24	<u>66.90</u>
MBP (ours)	73.02	62.43	70.12	65.24	67.70
MBP w/o masking	70.63	58.99	67.68	62.20	64.88

Table 4: Method ablation results on Qwen3-1.7B-Base. The upper block reports math benchmarks, and the lower block reports code benchmarks.

delivers reliable gains across domains.

7.3 Extension to GRPO exploration (RQ3)

If MBP preserves the pretrained distribution (H1), GRPO rollouts should retain wider behavioral support, yielding a richer gradient signal. We test this on Qwen3-1.7B-Base and Qwen3-4B-Base.

Faster convergence, higher final accuracy. Table 5 shows MBP-GRPO improves final four-benchmark average accuracy by $+1.31$ points (Qwen3-1.7B) and $+0.40$ points (Qwen3-4B) over standard GRPO, while normalized convergence AUC increases by $+0.048$ and $+0.045$, respectively.

Early rollout diversity drives the gain. Figure 4 shows that during the earliest steps, MBP-GRPO’s mixed-group ratio spikes substantially above standard GRPO, generating more useful gradient updates when the advantage signal is most valuable. This can be observed in reward and accuracy curves. MBP-GRPO pulls ahead from the first steps and stays ahead throughout training. Later, MBP’s mixed-group ratio falls below standard’s, consistent with earlier convergence on solved prompts. The pattern matches H1 in that a retained pretrained distribution diversifies rollouts when they most help, and the early advantage compounds.




Figure 4: **Training dynamics of MBP-GRPO.** Columns correspond to Qwen3-1.7B-Base (left) and Qwen3-4B-Base (right), while rows show GSM8K validation pass@1 (top), training reward (middle), and mixed-group ratio (bottom). Red denotes MBP-GRPO and blue denotes standard GRPO. MBP-GRPO exhibits a higher mixed-group ratio in the early stage of training, providing more useful update signals and leading to faster convergence than standard GRPO.

8 Analysis

8.1 Retention Analyses

We examine H1’s retention mechanism through a decomposition of post-training correctness, an ablation on loss-masking, and next-token entropy at step boundaries.

MBP preserves base-correct examples while solving new ones. If H1 holds at scale, MBP should preserve pretrained capabilities while fitting the downstream distribution. We measure both directly: BCP (Base-Correct Preservation) is the fraction of base-correct examples still correct after fine-tuning, and NSR (New-Solve Rate) is the fraction of base-wrong examples that become correct. Table 6 shows MBP improves both BCP and NSR. This shows that MBP achieves sufficient adaptation while better preserving the pretrained distribution.

Masking is essential at boundary placement. We hypothesize that masking pause-token targets is advantageous for retaining the pretrained distribution. To test this, we perform ablation experiments

Method	GSM8K	MATH	AMC23	Minerva	AUC
Qwen3-1.7B-Base					
STANDARD	80.52	59.68	37.50	20.96	0.366
MBP	81.35	61.24	40.00	21.32	0.414
Qwen3-4B-Base					
STANDARD	92.19	74.30	50.00	30.15	0.505
MBP	91.58	75.26	57.50	23.90	0.550

Table 5: Final GRPO pass@1 accuracy (%) on GSM8K, MATH, AMC23, and Minerva, together with normalized training-reward AUC, for Qwen3-1.7B-Base and Qwen3-4B-Base. Full results are reported in Appendix Table 19.

Metric	Method	GSM8K	MATH	Overall
BCP $\uparrow$	STANDARD	92.08	88.33	89.18
	MBP	93.36	89.09	90.00
NSR $\uparrow$	STANDARD	63.04	33.42	38.20
	MBP	66.74	36.83	41.11
Acc. $\uparrow$	STANDARD	81.96	63.52	67.14
	MBP	84.08	65.48	68.87

Table 6: BCP (Base-Correct Preservation), NSR (New-Solve Rate), and final accuracy of Standard-SFT and MBP on Qwen3-1.7B-Base.

in math and code training where pauses are predicted as targets (Liu et al., 2026). Table 4 (MBP w/o MASKING) shows that predicting rather than masking pause tokens reduces performance by 2.6 points on math and 2.8 points on code; on code, predicted boundary placement even falls below standard SFT, confirming that masking is load-bearing.

Entropy is retained at step boundaries. H1 predicts that pause training selectively preserves diversity at boundaries. We therefore measure entropy both at step-boundary positions and across all token positions on out-of-distribution (OOD) general language understanding benchmarks, with the in-distribution GSM8K and MATH benchmarks included for comparison (Table 7). At step boundaries, MBP maintains substantially higher entropy than SFT across both in-domain and OOD tasks, suggesting that MBP selectively preserves the base model’s distributional diversity at these positions. In contrast, the average entropy over all tokens remains almost identical between SFT and MBP. Taken together, these results indicate that MBP selectively preserves distributional diversity at step boundaries while committing to a distribution within steps.

Model	OOD		In-domain	
	MMLU	GPQA	GSM8K	MATH
<i>Step boundaries</i>				
Base	0.707	0.607	0.500	0.579
SFT	0.403	0.380	0.278	0.339
MBP	0.463	0.440	0.365	0.380
<i>All tokens</i>				
Base	0.587	0.489	0.358	0.258
SFT	0.361	0.321	0.189	0.144
MBP	0.368	0.332	0.192	0.145

Table 7: Next-token vocabulary entropy on Qwen3-1.7B-Base rollouts under the Base, math-SFT, and MBP checkpoints, measured at step boundaries and across all response tokens. MMLU and GPQA are out of domain; GSM8K and MATH are in domain.

Dataset	Base	SFT	MBP
GSM8K	0.00	1.43	6.31
MATH	0.00	2.64	4.84
DeepMath Eval	0.00	4.67	8.29

Table 8: Base-referenced normalized MI proxy between context-conditioned step-boundary tokens and final answers. Results use macro aggregation over the full evaluation sets.

8.2 Non-Myopic Compression

An intuitive way to empirically quantify the extent to which H2 holds is to measure mutual information (MI) between step-boundary tokens and the final answer. We propose the following MI proxy:

$$I(X; Y) = \mathbb{E}_{(x,y) \sim p_{\theta}(\cdot, \cdot | q)} \left[\log \frac{p_{\theta}(x, y | q)}{p_{\theta}(x | q) p_{\theta}(y | q)} \right]$$

Here, x denotes the intermediate reasoning steps, y the final answer, and q the input prompt.

This measure evaluates the predictability of the final answer from the intermediate step boundary, which reflects the extent to which boundary tokens compress final answer-related information. Table 8 compares base-referenced normalized MI proxy values (Appendix E.4) for MBP and standard SFT across mathematical reasoning benchmarks. Results show that MBP compresses substantially more final-answer-related information at step boundaries than standard SFT does, providing empirical evidence consistent with H2.

9 Related Work

Pause and Dummy Tokens. Prior work has used pause tokens to enhance LLM reasoning by delaying generation (Goyal et al., 2024; Herel and

Mikolov, 2024); London and Kanade (2025) prove that pause tokens strictly increase the expressivity of constant-depth transformers. Building on this, recent approaches have explored dynamic token insertion at strategic points, like low-confidence positions or sentence boundaries, to better guide cognitive processing (Kim et al., 2025; Liu et al., 2026). Unlike these prior works that primarily view token insertion as a mechanism for increasing computational bandwidth or providing direct generation cues, our work uniquely isolates the role of pause tokens as training-time context perturbations.

Post-Training for Reasoning and Alignment.

Contemporary post-training paradigms for LLMs aim to align behavior and enhance complex problem-solving, typically progressing from SFT to reinforcement learning frameworks such as GRPO (Wei et al., 2022a; Chung et al., 2024; Shao et al., 2024). A critical challenge across these stages is the trade-off between reasoning adaptation and mode retention, as fine-tuning on a new reasoning task often degrades the model’s previously established language fluency and factual knowledge (Luo et al., 2025; Zheng et al., 2025; Lin et al., 2024). Our work is complementary to these post-training paradigms. We show that a lightweight token-insertion strategy can improve both SFT and GRPO when applied at semantic boundaries and masked from direct prediction, while mitigating the adaptation-retention trade-off.

10 Conclusions

In this work, we extend the study of pause tokens to their effect on fine-tuning dynamics. Through controlled pilot experiments, we identify two hypotheses, mode retention (H1) and non-myopic compression (H2), which naturally lead to Masked Boundary Pause Training. Through this lens, we show improvements in math and code reasoning while preserving general capability on unrelated benchmarks. We further show that boundary placement reliably outperforms SFT across model scales, as our hypotheses predict, and that retention also helps exploration in RL training. We hope this training-dynamics perspective, a shift from computational expressivity to train-time mode retention and non-myopic compression, supports further research across domains and adaptation regimes.

Limitations

Our current study focuses on full fine-tuning and evaluates retention primarily through per-example correctness, using BCP and NSR to measure whether base-correct examples remain correct and base-wrong examples become correct. Additionally, initial LoRA experiments show more moderate gains, suggesting that MBP’s interaction with parameter-efficient fine-tuning requires further study. Future work should examine retention at finer token- or distribution-level granularity and systematically analyze MBP under parameter-efficient adaptation.

Acknowledgments

This work is in part supported by the National Research Foundation of Korea (NRF, RS-2024-00451435(15%), RS-2024-00413957(15%)), Institute of Information & Communications Technology Planning & Evaluation (IITP, RS-2025-02305453(10%), RS-2025-02273157(10%), RS-2025-25442149(20%), RS-2021-II211343(10%)) grant funded by the Ministry of Science and ICT (MSIT), Institute of New Media and Communications (INMAC), the BK21 FOUR program of the Education, Artificial Intelligence Graduate School Program (Seoul National University), Research Program for Future ICT Pioneers, Seoul National University in 2026, and the KT (Korea Telecom)–(Seoul National University/Korea University) AICT R&D Center (20%).

References

- Wasi Uddin Ahmad, Aleksander Ficek, Mehrzad Samadi, Jocelyn Huang, Vahid Noroozi, Somshubra Majumdar, and Boris Ginsburg. 2025. [OpenCode-Instruct: A large-scale instruction tuning dataset for code LLMs](#). *Preprint*, arXiv:2504.04030.
- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, and Charles Sutton. 2021. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*.
- Zhangir Azerbayev, Hailey Schoelkopf, Keiran Paster, Marco Dos Santos, Stephen Marcus McAleer, Albert Q. Jiang, Jia Deng, Stella Biderman, and Sean Welleck. 2024. [Llemma: An open language model for mathematics](#). In *The Twelfth International Conference on Learning Representations*.
- Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi. 2020. PIQA: Reasoning about physical commonsense in natural language. In *AAAI Conference on Artificial Intelligence*.
- Mikhail S. Burtsev, Yuri Kuratov, Anton Peganov, and Grigory V. Sapunov. 2021. [Memory transformer](#). *Preprint*, arXiv:2006.11527.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, and 16 others. 2024. [Scaling instruction-finetuned language models](#). *Journal of Machine Learning Research*, 25(70):1–53.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Sachin Goyal, Ziwei Ji, Ankit Singh Rawat, Aditya Krishna Menon, Sanjiv Kumar, and Vaishnavh Nagarajan. 2024. [Think before you speak: Training language models with pause tokens](#). In *The Twelfth International Conference on Learning Representations*.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A Smith. 2020. Don’t stop pretraining: Adapt language models to domains and tasks. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pages 8342–8360.
- Zhiwei He, Tian Liang, Jiahao Xu, Qiuzhi Liu, Xingyu Chen, Yue Wang, Linfeng Song, Dian Yu, Zhenwen Liang, Wenxuan Wang, Zhuosheng Zhang, Rui Wang, Zhaopeng Tu, Haitao Mi, and Dong Yu. 2025. [DeepMath-103K: A large-scale, challenging, decontaminated, and verifiable mathematical dataset for advancing reasoning](#). *Preprint*, arXiv:2504.11456.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021a. [Measuring massive multitask language understanding](#). In *International Conference on Learning Representations*.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and

- Jacob Steinhardt. 2021b. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*.
- David Herel and Tomas Mikolov. 2024. [Thinking tokens for language modeling](#). *Preprint*, arXiv:2405.08644.
- Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Jian Hu, Xibin Wu, Wei Shen, Jason Klein Liu, Weixun Wang, Songlin Jiang, Haoran Wang, Hao Chen, Bin Chen, Wenkai Fang, Xianyu, Yu Cao, Haotian Xu, and Yiming Liu. 2025. [OpenRLHF: A ray-based easy-to-use, scalable and high-performance RLHF framework](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 656–666, Suzhou, China. Association for Computational Linguistics.
- Eunki Kim, Sangryul Kim, and James Thorne. 2025. [Learning to insert \[pause\] tokens for better reasoning](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, page 23760–23777. Association for Computational Linguistics.
- Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, et al. 2022. Solving quantitative reasoning problems with language models. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Yong Lin, Hangyu Lin, Wei Xiong, Shizhe Diao, Jianmeng Liu, Jipeng Zhang, Rui Pan, Haoxiang Wang, Wenbin Hu, Hanning Zhang, Hanze Dong, Renjie Pi, Han Zhao, Nan Jiang, Heng Ji, Yuan Yao, and Tong Zhang. 2024. [Mitigating the alignment tax of RLHF](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 580–606, Miami, Florida, USA. Association for Computational Linguistics.
- Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. 2023. Is your code generated by ChatGPT really correct? rigorous evaluation of large language models for code generation. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Zhichen Liu, Yongyuan Li, and Yang Xu. 2026. [Think in sentences: Explicit sentence boundaries enhance language model’s capabilities](#). In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2275–2288, San Diego, California, United States. Association for Computational Linguistics.
- Charles London and Varun Kanade. 2025. [Pause tokens strictly increase the expressivity of constant-depth transformers](#). In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Yun Luo, Zhen Yang, Fandong Meng, Yafu Li, Jie Zhou, and Yue Zhang. 2025. [An empirical study of catastrophic forgetting in large language models during continual fine-tuning](#). *Preprint*, arXiv:2308.08747.
- Jacob Pfau, William Merrill, and Samuel R. Bowman. 2024. [Let’s think dot by dot: Hidden computation in transformer language models](#). In *First Conference on Language Modeling*.
- 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*.
- Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Romain Sauvestre, Tal Remez, et al. 2023. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950*.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024. [DeepSeekMath: Pushing the limits of mathematical reasoning in open language models](#). *Preprint*, arXiv:2402.03300.
- Charlie Victor Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. 2025. [Scaling LLM test-time compute optimally can be more effective than scaling parameters for reasoning](#). In *The Thirteenth International Conference on Learning Representations*.
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc Le, Ed Chi, Denny Zhou, and Jason Wei. 2023. [Challenging BIG-bench tasks and whether chain-of-thought can solve them](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 13003–13051, Toronto, Canada. Association for Computational Linguistics.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V. Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. [Self-consistency improves chain of thought reasoning in language models](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V Le. 2022a. [Finetuned language models are zero-shot learners](#). In *International Conference on Learning Representations*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi, Quoc V Le, and Denny Zhou. 2022b. [Chain of thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems*.
- Wilson Wu, John Xavier Morris, and Lionel Levine. 2024. [Do language models plan ahead for future tokens?](#) In *First Conference on Language Modeling*.

- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Tian Ye, Zicheng Xu, Yuanzhi Li, and Zeyuan Allen-Zhu. 2025. [Physics of language models: Part 2.1, grade-school math and the hidden reasoning process](#). In *The Thirteenth International Conference on Learning Representations*.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. HellaSwag: Can a machine really finish your sentence? In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Weihao Zeng, Yuzhen Huang, Qian Liu, Wei Liu, Keqing He, Zejun Ma, and Junxian He. 2025. [SimpleRL-Zoo: Investigating and taming zero reinforcement learning for open base models in the wild](#). *Preprint*, arXiv:2503.18892.
- Junhao Zheng, Xidi Cai, Shengjie Qiu, and Qianli Ma. 2025. [Spurious forgetting in continual learning of language models](#). In *The Thirteenth International Conference on Learning Representations*.

A Pilot Experiment Details

A.1 MLC continual learning pilot

Task and data. Sequences are generated by the Modular Linear Chain (MLC) recurrence $x_{t+1} = (a_t x_t + b_t) \bmod p$ with prime $p = 37$ and sequence length $L = 8$. A *regime* is a fixed length- L coefficient sequence $\{(a_t, b_t)\}_{t=1}^L$ with each a_t, b_t drawn uniformly from $\{0, \dots, p-1\}$. We sample two disjoint coefficient sets to define regimes D_1 and D_2 . Vocabulary size is $p+2$ (digits, BOS, and one pause token).

Model and optimizer. We use a 2-layer, 4-head GPT-style Transformer with hidden size $d = 128$ and learned positional embeddings. We train it with AdamW (learning rate 3×10^{-4} , weight decay 0.01) using a batch size of 64.

Two-phase training. *Phase 1:* train on D_1 with standard next-token prediction until D_1 accuracy exceeds 0.99 (approximately 2,000 optimizer steps). *Phase 2:* fork from a single Phase-1 checkpoint and continue training on D_2 under two matched conditions (NO PAUSE and PAUSE), each run for an additional 2,000 steps. The PAUSE condition inserts one pause token after every content token, doubling sequence length; pause-target positions are excluded from the cross-entropy loss. Both conditions share architecture, optimizer, batch size, learning rate, and update budget. We report $n = 8$ seeds with shaded bands at $\pm 1\sigma$.

Evaluation. We track held-out loss on D_2 (adaptation) and D_1 (retention) every 5 Phase-2 steps. D_2 loss is computed on pause-augmented sequences for PAUSE and bare sequences for NO PAUSE; D_1 loss is bare for both conditions (matching the no-pause inference distribution).

A.2 Multi-stage continual learning

The main pilot (§3.1) uses a single $D_1 \rightarrow D_2$ transition. To test whether the retention advantage is specific to one adaptation or persists as adaptations accumulate, we extend the same setup to a four-regime chain.

Setup. We sample four disjoint coefficient sets defining regimes $D_1, \dots, D_4$ ($p = 37, L = 8$); all task, model, and optimizer settings follow Appendix A.1. After Phase-1 training on D_1 , each condition adapts sequentially to D_2, D_3 , and D_4 for 2,000 steps per stage under matched budget and optimizer settings: NO PAUSE runs standard SFT

After	NO PAUSE	PAUSE	Gap
D_2	5.76 ± 0.56	2.74 ± 0.81	$+3.01 \pm 0.75$
D_3	8.60 ± 1.44	4.62 ± 0.46	$+3.98 \pm 1.59$
D_4	8.83 ± 0.16	6.04 ± 0.84	$+2.80 \pm 0.97$

Table 9: Held-out D_1 loss (nats, mean $\pm 1\sigma$, $n = 8$ seeds) after each stage of the $D_1 \rightarrow D_2 \rightarrow D_3 \rightarrow D_4$ chain. Lower is better; **Gap** is NO PAUSE – PAUSE, so positive values favor masked pauses.

on each bare D_k , and PAUSE runs masked-pause SFT. We report $n = 8$ seeds and track held-out bare-format D_1 loss.

Results. Table 9 reports held-out D_1 loss at the end of each stage. The retention gap is positive in 8/8 seeds at every stage and never inverts across the 6,000 steps, holding at roughly 3–4 nats throughout the chain. The absolute advantage therefore neither accumulates nor decays as adaptations pile up. Both conditions reach the same $\sim 10^{-7}$ adaptation floor on every D_k , so the retained D_1 behavior does not come at the cost of adaptation quality.

A.3 iGSM probing pilot

Base model and data. We fine-tune Qwen3-0.6B-Base on 3,000 problems sampled from the iGSM generator (Ye et al., 2025). We use the max_op=15, max_edge=20, perm_level=5 difficulty setting.

SFT. We train three conditions for two epochs with learning rate 1×10^{-5} , batch size 32, AdamW optimizer, and bf16 precision: NO PAUSE (standard next-token prediction on bare iGSM solutions), PAUSE (one [PAUSE] token inserted at every reasoning-step boundary, pause-target positions excluded from the loss), and PAUSE (UNMASKED), identical but with pause positions included in the loss. The BASE condition is the pretrained checkpoint with no fine-tuning. We report $n = 3$ seeds.

Probe training. For each held-out problem, we extract the final-layer hidden state at the period token immediately preceding each reasoning step boundary (the same character position across all four protocols). We train two separate linear probes on these representations: one predicting the integer answer of the *current* step and one predicting the integer answer of the *next* step. Probes are trained with a by-problem 70/30 split, with each problem’s boundary positions assigned to either train or test based on the problem split.

B Theoretical analysis for masked-pause SFT

B.1 A first-order view of masked-pause mode retention

Old D_1 behaviors survive D_2 fine-tuning as long as their logits do not move enough to erase their old margins. We make this precise on a protected set of D_1 no-pause prefixes and then show that masked-pause SFT yields a smaller certified protected logit drift than ordinary SFT. The result is a first-order sufficient condition within the local linearized model around W_1 ; it certifies preservation of the D_1 top token on selected no-pause prefixes, not preservation of the full D_1 predictive distribution.

Margin stability. Whether an old top token survives depends only on whether its margin outlasts the logit movement at that prefix; bounding the movement uniformly over $\mathcal{C}_1$ converts this into a retention guarantee.

Lemma B.1 (Margin stability). *Let W_1 be the shared parameters after D_1 training and $\mathcal{C}_1$ a set of D_1 no-pause prefixes whose top-token behavior we protect. For $c \in \mathcal{C}_1$, let $a(c) = \arg \max_v z_{W_1}(c)_v$ be the D_1 top token and $\gamma(c)$ its top-vs-runner-up margin at W_1 — the amount a competing logit must close to flip the prediction at c . Define the margin survival function, the retention rate, and the protected logit drift*

$$\begin{aligned}\Phi(r) &= \frac{1}{|\mathcal{C}_1|} |\{c \in \mathcal{C}_1 : \gamma(c) > r\}|, \\ \text{Retain}(W) &= \frac{1}{|\mathcal{C}_1|} |\{c \in \mathcal{C}_1 : \\ &\quad \arg \max_v z_W(c)_v = a(c)\}|, \\ \Delta_1(W) &= \max_{c \in \mathcal{C}_1} \|z_W(c) - z_{W_1}(c)\|_\infty,\end{aligned}$$

with Φ non-increasing in r . Then

$$\Delta_1(W) \leq \delta \implies \text{Retain}(W) \geq \Phi(2\delta).$$

The factor of two appears because the old top logit may move down by δ while a competitor may move up by δ . Comparing the certified retention of two adaptation procedures therefore reduces to comparing their certified protected logit drifts on $\mathcal{C}_1$.

What masked-pause SFT changes. Ordinary SFT predicts the next D_2 target y_{i+1} from the no-pause prefix b_i ; masked-pause SFT predicts the same target from the pause-perturbed prefix $b_i \circ p$

and masks the loss for the pause token p itself. Masking is essential: were p a training target, b_i would again receive a direct supervised gradient, restoring the interference channel masked-pause SFT removes. Four scalars control the resulting first-order drifts on $\mathcal{C}_1$. The interference exposure I_O is large when ordinary prefixes b_i push the parameters in directions that strongly move protected D_1 logits. Decomposing the Jacobian of $b_i \circ p$ into a part inherited from the ordinary boundary Jacobian and a pause-local residual yields a worst-case carry coefficient λ and a worst-case pause-local overlap ν with $\mathcal{C}_1$. Finally, $\kappa = M_P/M_O$ is the ratio of cumulative normalized shared-update masses the two procedures need to reach the same D_2 fitting level (matched D_2 validation loss in our experiments).

Lemma B.2 (Protected drift bounds). *In the local linearized regime around W_1 with $I_O > 0$, the protected logit drifts of the ordinary and masked-pause endpoints satisfy*

$$\begin{aligned}\Delta_1(\widehat{W}_O) &\leq M_O I_O, \\ \Delta_1(\widehat{W}_P) &\leq M_P (\lambda I_O + \nu).\end{aligned}$$

Equivalently, defining $\beta_{\text{ctx}} := \kappa(\lambda + \nu/I_O)$, the masked-pause bound can be written as

$$\Delta_1(\widehat{W}_P) \leq \beta_{\text{ctx}} M_O I_O.$$

Both bounds share the ordinary scale $M_O I_O$, and in this certificate β_{ctx} is the factor multiplying the masked-pause drift bound. Feeding these bounds through Lemma B.1 turns the comparison into a retention guarantee.

Theorem B.3 (First-order context-perturbation retention certificate). *Assume the local small-step regime around W_1 and $I_O > 0$, with ordinary and masked-pause SFT using the same normal D_2 targets and reaching the same D_2 fitting level. Combining the drift bounds of Lemma B.2 with the margin-stability bound of Lemma B.1 yields*

$$\begin{aligned}\text{Retain}(\widehat{W}_O) &\geq \Phi(2 M_O I_O), \\ \text{Retain}(\widehat{W}_P) &\geq \Phi(2 \beta_{\text{ctx}} M_O I_O).\end{aligned}$$

Masked-pause SFT yields the stronger certificate whenever $\beta_{\text{ctx}} < 1$.

The condition $\beta_{\text{ctx}} < 1$ admits a geometric reading: masked-pause SFT fits D_2 with less first-order drift on D_1 no-pause logits whenever the pause state $b_i \circ p$ is sufficiently separated from the ordinary boundary b_i in the Jacobian sense captured

by λ and ν . This separation is structural rather than architectural — a property of the pause insertion policy, which our empirical preference for reasoning-step boundaries effectively encodes. The full setup, lemmas, proofs, and the role of masking and pause-insertion position are in Appendix C.

C Theoretical analysis for masked-pause SFT - Proofs

This appendix contains the full setup, lemmas, and proofs supporting Theorem B.3. The analysis is a local first-order sufficient condition on a protected no-pause prefix set $\mathcal{C}_1$. It does *not* claim preservation of the full next-token distribution, of D_1 multimodality, or of the realized retained fraction; it compares certified lower bounds within the local linearized model around W_1 .

C.1 Setup and notation

Let $\mathcal{V}$ be the vocabulary, W the ordinary shared model parameters, and e_p the pause embedding. We write $\circ$ for token-sequence concatenation, so $b_i \circ p$ denotes the prefix obtained by appending a single pause token p to b_i . For a no-pause prefix c , let $z_W(c) \in \mathbb{R}^{|\mathcal{V}|}$ be the next-token logit vector; for a prefix containing p , write $z_{W,e_p}(s)$ when the dependence on e_p should be explicit.

Let W_1 be the shared parameters after D_1 training and let $\mathcal{C}_1$ be the protected set of D_1 no-pause prefixes. For each $c \in \mathcal{C}_1$, the D_1 top token, its margin, and the margin survival function are

$$\begin{aligned} a(c) &= \arg \max_{v \in \mathcal{V}} z_{W_1}(c)_v, \\ \gamma(c) &= z_{W_1}(c)_{a(c)} - \max_{v \neq a(c)} z_{W_1}(c)_v, \\ \Phi(r) &= \frac{1}{|\mathcal{C}_1|} |\{c \in \mathcal{C}_1 : \gamma(c) > r\}|. \end{aligned}$$

Old-mode retention and protected no-pause logit drift are

$$\begin{aligned} \text{Retain}(W) &= \frac{1}{|\mathcal{C}_1|} \sum_{c \in \mathcal{C}_1} \mathbf{1}[\arg \max_v z_W(c)_v = a(c)], \\ \Delta_1(W) &= \max_{c \in \mathcal{C}_1} \|z_W(c) - z_{W_1}(c)\|_\infty. \end{aligned}$$

Because every prefix in $\mathcal{C}_1$ is a no-pause prefix, both quantities depend only on the shared parameters W ; in particular, $\Delta_1(\widehat{W}_P)$ and $\text{Retain}(\widehat{W}_P)$ do not directly involve the learned pause embedding $\widehat{e}_P$.

For a matrix B , the operator norm induced by ℓ_∞ is

$$\begin{aligned} \|B\|_{\infty \rightarrow \infty} &= \sup_{u \neq 0} \frac{\|Bu\|_\infty}{\|u\|_\infty} \\ &= \max_j \sum_k |B_{jk}|, \end{aligned}$$

and satisfies $\|XY\|_{\infty \rightarrow \infty} \leq \|X\|_{\infty \rightarrow \infty} \|Y\|_{\infty \rightarrow \infty}$ whenever XY is defined.

C.2 Direct supervised-prefix relocation

Let $b_i = (x, y_1, \dots, y_i)$ be a no-pause D_2 reasoning boundary and $\ell(z, y)$ the softmax cross-entropy loss for target y . The two methods give the per-step direct losses

$$\begin{aligned} \ell_i^O(W) &= \ell(z_W(b_i), y_{i+1}), \\ \ell_i^P(W, e_p) &= \ell(z_{W,e_p}(b_i \circ p), y_{i+1}), \end{aligned}$$

with the pause-token prediction loss masked out in the second case. Thus the direct supervised prefix for the normal target y_{i+1} moves from b_i to $b_i \circ p$. Earlier tokens of b_i can still receive indirect gradient through the pause position; those contributions are folded into the pause-state Jacobian below.

We denote the resulting (multi)sets of supervised prefixes by

$$\begin{aligned} \mathcal{S}_O &= \{b_i : \\ &\quad \text{loss predicts a normal } D_2 \text{ token} \\ &\quad \text{from } b_i\}, \\ \mathcal{S}_P &= \{b_i \circ p : \\ &\quad \text{loss predicts the same normal token} \\ &\quad \text{from } b_i \circ p\}, \end{aligned}$$

read as multisets if a prefix is supervised more than once.

C.3 First-order interference and Jacobians

For a prefix s , define the start-of-adaptation shared-parameter logit Jacobian

$$J_1(s) = \begin{cases} \left. \frac{\partial z_W(s)}{\partial W} \right|_{W_1} & s \text{ contains no } p, \\ \left. \frac{\partial z_{W,e_p^{(0)}}(s)}{\partial W} \right|_{W_1} & s \text{ contains } p, \end{cases}$$

where $e_p^{(0)}$ is the pause embedding at the start of D_2 adaptation; in the pause case, $e_p^{(0)}$ is held fixed and the derivative is taken with respect to W only.

For $c \in \mathcal{C}_1$ and a supervised prefix s , the first-order interference matrix and exposure are

$$K_1(c, s) = J_1(c)J_1(s)^\top, \\ I_1(\mathcal{C}_1, \mathcal{S}) = \max_{c \in \mathcal{C}_1} \sum_{s \in \mathcal{S}} \|K_1(c, s)\|_{\infty \rightarrow \infty}.$$

The operator-norm choice matches the ℓ_∞ logit-drift quantity Δ_1 . We write $I_O = I_1(\mathcal{C}_1, \mathcal{S}_O)$. Operationally, I_O is large when ordinary D_2 supervised prefixes share gradient directions with protected D_1 prefixes.

C.4 Pause-state Jacobian decomposition

For each paired boundary, write $J_i^O = J_1(b_i)$ and $J_i^P = J_1(b_i \circ p)$. Let P_i^O be the orthogonal projection in parameter space onto the row space of J_i^O , acting by right-multiplication, so that $J_i^P P_i^O$ projects each row of J_i^P onto $\text{row}(J_i^O)$. Define

$$J_i^{\text{carry}} = J_i^P P_i^O, \\ J_i^{\text{loc}} = J_i^P (I - P_i^O),$$

so $J_i^P = J_i^{\text{carry}} + J_i^{\text{loc}}$. Because every row of J_i^{carry} lies in $\text{row}(J_i^O)$, there exists A_i with $J_i^{\text{carry}} = A_i J_i^O$, and we take the canonical minimum-norm choice

$$A_i = J_i^P (J_i^O)^+,$$

where $(\cdot)^+$ is the Moore–Penrose pseudoinverse. If $J_i^O = 0$, set $A_i = 0$. The carry and pause-local interference summaries are

$$\lambda = \max_i \|A_i^\top\|_{\infty \rightarrow \infty}, \\ \nu = \max_{c \in \mathcal{C}_1} \sum_i \|J_1(c)(J_i^{\text{loc}})^\top\|_{\infty \rightarrow \infty}.$$

Operationally, λ is small when the pause state does not reuse the same gradient-sensitive directions as the ordinary boundary, and ν is small when the pause-local residual does not align with the gradient-sensitive directions of protected D_1 prefixes. Whether λ and ν are small for a given architecture and insertion policy is a separate structural question; see Section C.9.

C.5 Training-regime quantities

Remark C.1 (Local small-step regime). The D_2 adaptation trajectory is analyzed in a neighborhood of W_1 in which the linearization $z_W(s) \approx z_{W_1}(s) + J_1(s)(W - W_1)$ is accurate; for the full nonlinear model, the same statements hold up to a standard local linearization error in Δ_1 .

For $A \in \{O, P\}$, let $w_t^A(s)$ be the normalized weight of supervised prefix s in the minibatch loss at step t and η_t the learning rate. Define

$$m_A(s) = \sum_t \eta_t w_t^A(s), \\ M_A = \max_s m_A(s), \quad \kappa = \frac{M_P}{M_O}.$$

M_A is the largest cumulative normalized update mass assigned to any single supervised prefix along the trajectory. Throughout this appendix, “the same D_2 fitting level” refers to matched D_2 validation loss (or task accuracy, when that is the empirical comparison metric used in the main paper); $\kappa > 1$ then quantifies the additional shared-parameter update mass masked-pause SFT requires, if any, to reach that level.

C.6 Additional lemmas

(Lemmas B.1 and B.2 are stated in the main text; their proofs are included below.)

Lemma C.2 (Direct-loss relocation). *Write $J_W(s) = \partial z_W(s)/\partial W$ and $J_{W|e_p}(s) = \partial z_{W,e_p}(s)/\partial W$ (with e_p held fixed). Then*

$$\nabla_W \ell_i^O(W) = J_W(b_i)^\top r_i^O, \\ \nabla_W \ell_i^P(W, e_p) = J_{W|e_p}(b_i \circ p)^\top r_i^P,$$

with $r_i^O = \nabla_z \ell(z_W(b_i), y_{i+1})$ and $r_i^P = \nabla_z \ell(z_{W,e_p}(b_i \circ p), y_{i+1})$. *There is no direct loss term of the form $\ell(z_W(b_i), p)$, because the pause-token prediction loss is masked. Under the local linearized convention, J_W and $J_{W|e_p}$ reduce to J_1 at the start of D_2 adaptation.*

Lemma C.3 (Interference separation). *Under the canonical decomposition of Section C.4,*

$$I_1(\mathcal{C}_1, \mathcal{S}_P) \leq \lambda I_O + \nu.$$

If $I_O > 0$, this is equivalently $I_1(\mathcal{C}_1, \mathcal{S}_P) \leq (\lambda + \nu/I_O) I_O = \alpha_{\text{ctx}} I_O$.

Lemma C.4 (Interference controls drift). *In the local linearized regime, any D_2 adaptation procedure with supervised multiset $\mathcal{S}$ and normalized update masses $m(s)$ satisfies*

$$\Delta_1(W) \leq \max_{c \in \mathcal{C}_1} \sum_{s \in \mathcal{S}} m(s) \|K_1(c, s)\|_{\infty \rightarrow \infty} \\ \leq M I_1(\mathcal{C}_1, \mathcal{S}),$$

where $M = \max_s m(s)$.

C.7 Proofs

Proof of Lemma B.1 (margin stability). Fix $c \in \mathcal{C}_1$ with $\gamma(c) > 2\delta$ and any $v \neq a(c)$. Since each logit moves by at most δ ,

$$\begin{aligned} z_W(c)_{a(c)} - z_W(c)_v &\geq (z_{W_1}(c)_{a(c)} - z_{W_1}(c)_v) - 2\delta \\ &\geq \gamma(c) - 2\delta > 0. \end{aligned}$$

Thus $a(c)$ remains the top token under W , and the fraction of prefixes c with $\gamma(c) > 2\delta$ lower-bounds $\text{Retain}(W)$. $\square$

Proof of Lemma C.2. Ordinary SFT predicts y_{i+1} from b_i , so its CE gradient is attached to the output state at b_i ; the chain rule gives the stated form. Masked-pause SFT inserts p after b_i , masks the would-be loss for predicting p from b_i , and predicts y_{i+1} from the pause-state output position $b_i \circ p$; again the chain rule gives the stated form. Indirect gradient on earlier tokens flows through $J_{W|e_p}(b_i \circ p)$ and is handled by the carry/local decomposition. $\square$

Proof of Lemma C.3. Fix $c \in \mathcal{C}_1$. Using $J_i^P = A_i J_i^O + J_i^{\text{loc}}$,

$$\begin{aligned} K_1(c, b_i \circ p) &= J_1(c)(J_i^P)^\top \\ &= K_1(c, b_i) A_i^\top + J_1(c)(J_i^{\text{loc}})^\top. \end{aligned}$$

By submultiplicativity of $\|\cdot\|_{\infty \rightarrow \infty}$,

$$\begin{aligned} \|K_1(c, b_i \circ p)\|_{\infty \rightarrow \infty} &\leq \lambda \|K_1(c, b_i)\|_{\infty \rightarrow \infty} \\ &\quad + \|J_1(c)(J_i^{\text{loc}})^\top\|_{\infty \rightarrow \infty}. \end{aligned}$$

Summing over i and taking $\max_{c \in \mathcal{C}_1}$ on both sides, with $\max(A(c) + B(c)) \leq \max A(c) + \max B(c)$, gives $I_1(\mathcal{C}_1, \mathcal{S}_P) \leq \lambda I_O + \nu$. $\square$

Proof of Lemma C.4. In the linearized regime, the step is $\Delta W_t = -\eta_t \sum_s w_t(s) J_1(s)^\top r_t(s)$, where the softmax CE residual $r_t(s) = q_t(s) - e_{y_s}$ satisfies $\|r_t(s)\|_\infty \leq 1$ uniformly along the trajectory. For $c \in \mathcal{C}_1$, summing first-order logit changes,

$$\begin{aligned} \|\Delta z(c)\|_\infty &\leq \sum_t \eta_t \sum_s w_t(s) \\ &\quad \|K_1(c, s)\|_{\infty \rightarrow \infty} \|r_t(s)\|_\infty \\ &\leq \sum_s m(s) \|K_1(c, s)\|_{\infty \rightarrow \infty}. \end{aligned}$$

Taking $\max_{c \in \mathcal{C}_1}$ and then $M = \max_s m(s)$ gives both inequalities. $\square$

Proof of Lemma B.2. Applying Lemma C.4 to the ordinary supervised multiset $\mathcal{S}_O$ gives $\Delta_1(\widehat{W}_O) \leq M_O I_1(\mathcal{C}_1, \mathcal{S}_O) = M_O I_O$. For masked-pause SFT, Lemma C.4 on $\mathcal{S}_P$ followed by Lemma C.3 gives

$$\begin{aligned} \Delta_1(\widehat{W}_P) &\leq M_P I_1(\mathcal{C}_1, \mathcal{S}_P) \\ &\leq M_P(\lambda I_O + \nu). \end{aligned}$$

When $I_O > 0$, factoring $M_O I_O$ out of the last bound gives

$$\begin{aligned} M_P(\lambda I_O + \nu) &= \frac{M_P}{M_O} \left(\lambda + \frac{\nu}{I_O} \right) M_O I_O \\ &= \beta_{\text{ctx}} M_O I_O. \quad \square \end{aligned}$$

Proof of Theorem B.3. By Lemma B.2, $\Delta_1(\widehat{W}_O) \leq M_O I_O$ and $\Delta_1(\widehat{W}_P) \leq \beta_{\text{ctx}} M_O I_O$. Applying Lemma B.1 with $\delta = M_O I_O$ and $\delta = \beta_{\text{ctx}} M_O I_O$, respectively, yields

$$\begin{aligned} \text{Retain}(\widehat{W}_O) &\geq \Phi(2 M_O I_O), \\ \text{Retain}(\widehat{W}_P) &\geq \Phi(2 \beta_{\text{ctx}} M_O I_O). \end{aligned}$$

Since Φ is non-increasing and $I_O > 0$, the condition $\beta_{\text{ctx}} < 1$ gives $\beta_{\text{ctx}} M_O I_O < M_O I_O$, so the masked-pause certificate dominates the ordinary-SFT one. The dominance is strict whenever a nonzero fraction of $\mathcal{C}_1$ has $\gamma(c) \in (2\beta_{\text{ctx}} M_O I_O, 2 M_O I_O]$. $\square$

C.8 Why masking the pause-token loss matters

If the pause-token prediction loss is not masked, each inserted pause introduces an extra direct loss $\ell_i^{\text{pause}}(W) = \ell(z_W(b_i), p)$ whose gradient is attached to the original no-pause boundary b_i . This re-introduces a supervised multiset $\mathcal{S}_{\text{pause}} = \{b_i : \text{loss predicts } p \text{ from } b_i\}$ and changes the drift certificate to

$$\begin{aligned} \Delta_1(\widehat{W}_P^{\text{unmask}}) &\leq M_{\text{norm}}(\lambda I_O + \nu) \\ &\quad + M_{\text{pause}} I_{\text{pause}}, \end{aligned}$$

where $I_{\text{pause}} = I_1(\mathcal{C}_1, \mathcal{S}_{\text{pause}})$. Since $\mathcal{S}_{\text{pause}}$ consists of the original no-pause boundaries, this added term can erase the benefit of context perturbation. Masking removes this loss and leaves the certificate in Theorem B.3 unchanged.

C.9 Pause insertion position

An insertion policy A determines $\mathcal{S}_P^A$ and therefore changes $\lambda(A)$, $\nu(A)$, and $\kappa(A) = M_P^A/M_O$,

giving $\beta_{\text{ctx}}(A) = \kappa(A)(\lambda(A) + \nu(A)/I_O)$. The theorem applies as soon as $\beta_{\text{ctx}}(A) < 1$. The remark below collects the qualitative considerations; it is not part of the formal theorem.

Remark C.5 (Why reasoning-step boundaries are natural insertion points). (i) At a reasoning-step boundary, the previous step has ended, so the pause state $b_i \circ p$ is under less pressure to copy a mid-phrase token representation, which tends to reduce the carry coefficient $\lambda(A)$. (ii) The pause-local residual at a complete boundary need not preserve ordinary syntactic-continuation features tied to D_1 , which tends to reduce $\nu(A)$. (iii) Predicting the next reasoning step from a boundary-plus-pause prefix is still a natural next-token task, so the fitting-budget ratio $\kappa(A)$ stays moderate. Insertion that is too sparse leaves most supervised prefixes close to $\mathcal{S}_O$; insertion at unnaturally dense positions can force the pause state to copy the immediately preceding prefix, increasing $\lambda(A)$, $\nu(A)$, or $\kappa(A)$.

D Experimental Settings

D.1 Hyperparameter settings

We report the hyperparameter settings used for each training stage. Across experiments, we use FusedAdam as the optimizer, CosineAnnealingLR as the learning-rate scheduler, a warmup ratio of 0.03, and bf16 precision for training, while adapting the learning rate and batch configuration to each training objective and model family.

For SFT on mathematical reasoning data, we train each model for 3 epochs with a batch size of 32. The learning rate is selected depending on the model family: 5×10^{-6} for Qwen models, 1×10^{-5} for Llama3.2, and 5×10^{-5} for Llama3.1. The full set of hyperparameters for math SFT is provided in Table 10.

For SFT on code generation data, we also train for 3 epochs with a batch size of 32, using the same optimizer, scheduler, warmup ratio, and training precision as in math SFT. The learning rate is model-specific: 5×10^{-7} for Qwen3-1.7B-Base, 5×10^{-6} for Qwen3-4B-Base, 5×10^{-5} for Qwen3-8B-Base, and 1×10^{-8} for Llama3.1-8B-Instruct. The complete configuration for code SFT is summarized in Table 11.

For GRPO training, we use a learning rate of 5×10^{-7} , train for 1 episode, and use a batch size of 128. For each prompt, we sample 8 rollouts and apply a KL coefficient of 0.001. The optimizer, scheduler, warmup ratio, and training dtype are

Hyperparameter	SFT on math
Learning rate	Qwen: 5×10^{-6} ; Llama3.2: 1×10^{-5} ; Llama3.1: 5×10^{-5}
Epochs	3
Batch size	32
Optimizer	FusedAdam
Scheduler	CosineAnnealingLR
Warmup ratio	0.03
Training dtype	bf16

Table 10: SFT hyperparameters for mathematical reasoning training.

Hyperparameter	SFT on code
Learning rate	Qwen3-1.7B-Base: 5×10^{-7} ; Qwen3-4B-Base: 5×10^{-6} ; Qwen3-8B-Base: 5×10^{-5} ; Llama3.1-8B-Instruct: 1×10^{-8}
Epochs	3
Batch size	32
Optimizer	FusedAdam
Scheduler	CosineAnnealingLR
Warmup ratio	0.03
Training dtype	bf16

Table 11: SFT hyperparameters for code generation training.

kept consistent with the SFT settings. The GRPO hyperparameters are listed in Table 12.

D.2 Computation resources

All training experiments in this paper were conducted with OpenRLHF (Hu et al., 2025) on 2× NVIDIA RTX PRO 6000 and 2× NVIDIA RTX A100 GPUs.

D.3 Dataset Statistics

We use a subset of DeepMath-103K (He et al., 2025) data for mathematical reasoning. Table 13 summarizes the statistics of the mathematical datasets used for training and evaluation.

D.4 Evaluation setup

For mathematical reasoning benchmarks, we report pass@1 accuracy under a deterministic decoding setting with temperature 0.0. For code generation benchmarks, we report pass@1 using the EvalPlus package, and a prediction is counted as correct only when it passes all test cases. For general benchmarks, we report pass@4 accuracy with temperature 0.7 using the LightEval package.

Hyperparameter	GRPO
Learning rate	5×10^{-7}
Episodes	1
Batch size	128
Rollouts per prompt	8
Optimizer	FusedAdam
Scheduler	CosineAnnealingLR
Warmup ratio	0.03
Training dtype	bf16
KL coefficient	0.001

Table 12: GRPO hyperparameters.

Statistic	Train	Eval	Total
Examples	83,438	4,391	87,829
Full SFT tokens	44.14M	2.36M	46.51M
Pause tokens	1.34M	71.72K	1.41M
Pause-token ratio	3.02%	3.03%	3.03%
Pauses/sample	16.003	16.332	16.020

Table 13: Statistics of the step-aligned pause dataset.

E Implementation Details

We implemented most components using OpenRLHF (Hu et al., 2025) and used the default OpenRLHF settings for hyperparameters not explicitly specified.

E.1 Objective function

Let x denote the prompt and $y = (y_1, \dots, y_T)$ denote the target response tokens. For a prefix s , write $p_\theta(\cdot | s)$ for the model distribution and $\ell_\theta(s, a) = -\log p_\theta(a | s)$ for the next-token cross-entropy loss. Standard SFT optimizes

$$\mathcal{L}_{\text{SFT}} = \frac{1}{T} \sum_{t=1}^T \ell_\theta(x \circ y_{<t}, y_t).$$

For pause-based training, let $u = (u_1, \dots, u_L)$ be the sequence obtained after inserting pause tokens into y , where each u_j is either a normal target token or the pause token p . Masked-pause SFT uses the same next-token CE loss, but removes the loss terms whose target is p :

$$m_j = \mathbf{1}[u_j \neq p],$$

$$\mathcal{L}_{\text{mask}} = \frac{1}{\sum_j m_j} \sum_{j=1}^L m_j \ell_\theta(x \circ u_{<j}, u_j).$$

Thus the model is trained to predict normal answer tokens from prefixes that may contain pauses, while it is not directly trained to emit the pause token itself. The pause-predicted variant keeps the

same inserted sequence u , but does not mask pause targets:

$$\mathcal{L}_{\text{pred}} = \frac{1}{L} \sum_{j=1}^L \ell_\theta(x \circ u_{<j}, u_j).$$

The distinction is important: masked-pause SFT changes the context used for predicting ordinary tokens, whereas pause-predicted SFT additionally adds direct supervision for predicting p .

E.2 Step-Boundary Extraction

We identify reasoning-step boundaries with deterministic preprocessing. For math, line breaks are the primary boundary candidates, while wrapped continuations and line breaks inside inline or display-math expressions are preserved. Displayed equations written with $\backslash[\dots \backslash]$, $\$ \$ \dots \$ \$$, or standard equation and alignment environments remain intact and are attached to immediately preceding explanatory text when present. Outside math environments, we split at transitions that form an independent reasoning unit, including bullets, final-answer headings, discourse cues, and completed calculations, and merge short non-semantic fragments with the preceding step. For code, we retain the extracted code region, discard surrounding natural-language text and blank lines, and treat each remaining source line as one step. One pause token is inserted at each resulting boundary, and its target position is masked from the loss.

E.3 Training Overhead

MBP inserts approximately 16 pause tokens per math example and 27 per code example on average. Pause tokens account for approximately 3.0% and 9.8% of tokens in the corresponding augmented math and code corpora, respectively. Training is otherwise identical to standard SFT: pause positions are simply excluded from the loss, and insertion is a deterministic rule-based preprocessing step.

E.4 MI-Proxy Computation

For each prompt q_i , we take y_i to be the final balanced expression in the generated solution and let $x_{i,k}$ denote the cumulative solution prefix ending at the k -th step boundary before the answer. For model m , we compute

$$r_{i,k}^{(m)} = \log \frac{p_m(x_{i,k}, y_i | q_i)}{p_m(x_{i,k} | q_i) p_m(y_i | q_i)}.$$

We weight each observed pair by its joint sequence probability, $w_{i,k}^{(m)} = p_m(x_{i,k}, y_i \mid q_i) r_{i,k}^{(m)}$, and macro-average first over the K_i boundaries of each example and then over the N examples:

$$\hat{I}_w(m) = \frac{1}{N} \sum_{i=1}^N \frac{1}{K_i} \sum_{k=1}^{K_i} w_{i,k}^{(m)}.$$

We report the signed, Base-referenced normalized score

$$\begin{aligned} \Delta_B(m) &= \hat{I}_w(m) - \hat{I}_w(B), \\ G_B(m) &= \text{sign}(\Delta_B(m)) \\ &\quad \cdot \log_{10} \left(1 + \frac{|\Delta_B(m)|}{|\hat{I}_w(B)|} \right), \end{aligned}$$

with $G_B(B) = 0$, where B denotes the Base model. All results use macro aggregation over the full evaluation set (`filter = all`).

E.5 Examples of Methods

The following examples illustrate the actual data formats used for different pause-token placement strategies, including MBP (Masked Boundary Pause), Random, Append, and DIT. Examples for math data are shown in Table 14, while examples for code data are provided in Table 15 and Table 16.

F Additional Results

This section reports the full evaluation results that are not shown in the main text. We include mathematical reasoning, code generation, and general benchmark results for Qwen3-1.7B-Base, Qwen3-4B-Base, Qwen3-8B-Base, Llama3.2-1B-Instruct, Llama3.2-3B-Instruct, and Llama3.1-8B-Instruct. The tables compare the base models, standard SFT, and the pause-insertion variants Append, Random, DIT, MBP, and MBP w/o masking where available. The full results for math and general language understanding tasks are shown in Table 17, while the full results for code and general language understanding tasks are presented in Table 18. Table 19 reports the full six-benchmark GRPO results.

Method	Math example
MBP	Question. Round each number to the nearest tenth and add the results: 543.55, 25.23, and 299.98. Answer. < PAUSE >First, let's round each number to the nearest tenth: < PAUSE >1. 543.55 rounds to 543.6 (since the hundredths place is 5). < PAUSE >2. 25.23 rounds to 25.2 (since the hundredths place is less than 5). < PAUSE >3. 299.98 rounds to 300.0 (since the hundredths place is 8). < PAUSE >Now, let's add the rounded numbers: < PAUSE >543.6 + 25.2 + 300.0 = 868.8. < PAUSE >So, the final answer is 868.8.
append	Question. Round each number to the nearest tenth and add the results: 543.55, 25.23, and 299.98.< PAUSE >< PAUSE >< PAUSE > < PAUSE >< PAUSE >< PAUSE >< PAUSE > Answer. First, let's round each number to the nearest tenth: 1. 543.55 rounds to 543.6 (since the hundredths place is 5). 2. 25.23 rounds to 25.2 (since the hundredths place is less than 5). 3. 299.98 rounds to 300.0 (since the hundredths place is 8). Now, let's add the rounded numbers: 543.6 + 25.2 + 300.0 = 868.8. So, the final answer is 868.8.
random	Question. Round each number to the nearest tenth and add the results: 543.55, 25.23, and 299.98. Answer. First, let's round each number to the nearest tenth: 1. 543.55 rounds to 5< PAUSE >4< PAUSE >3.6 (since the hundredths place is 5). 2. 25.23 rounds to 25.< PAUSE >2 (since the hundred< PAUSE >ths place is less than 5). 3. 299.98 rounds to 3< PAUSE >00.0 (since the hundredths place is 8). Now, let's add the rounded numbers: 543.6 + 25.2 + 300.0 = 868.8. So, the final< PAUSE > answer is 868.8.
dit	Question. Round each number to the nearest tenth and add the results: 543.55, 25.23, and 299.98. Answer. < PAUSE >First,< PAUSE > let's round each number to the nearest tenth: < PAUSE >1. 543.55 rounds to 543.6 (since the hundredths place is 5)< PAUSE >. 2. 25.23 rounds to 25.2 (since the hundredths place is< PAUSE > less than 5). 3. 299.98 rounds to 300.0 (since the hundredths place is 8). Now, let's add the rounded numbers: 543.6 + 25.2 + 300.0 = 868.8. So, the< PAUSE > final answer is< PAUSE > 868.8.

Table 14: Examples of pause-token placement for the math task.

Method	Code example
MBP	Question. Ninja Atsushi stands guard atop the Ninja Building, which consists of two adjacent buildings of equal height. Each building has a series of walls of varying heights. Atsushi must jump between buildings to patrol the rooftops. Given the heights of the walls for both buildings, determine the minimum height Atsushi must jump for each wall position. Input: two lists of integers, heights_building_1 and heights_building_2, with the same length. Output: a list of integers where each element is the minimum height Atsushi must jump for each wall position. Constraints: the common list length is between 1 and 10⁵, and each wall height is between 1 and 10⁹. Sample input: heights_building_1 = [3, 6, 2], heights_building_2 = [5, 4, 8]. Sample output: [5, 6, 8]. Answer. <pre>< PAUSE >def min_jumps(heights_building_1, heights_building_2): < PAUSE >Calculate the minimum height Atsushi must jump for each wall position between two buildings. < PAUSE >:param heights_building_1: list of integers representing the heights of walls in the first building. < PAUSE >:param heights_building_2: list of integers representing the heights of walls in the second building. < PAUSE >:return: list of integers representing the minimum jump heights for each wall position. < PAUSE >if len(heights_building_1) != len(heights_building_2): < PAUSE >raise ValueError("Both buildings must have the same number of walls."). < PAUSE >return [max(heights_building_1[i], heights_building_2[i]) for i in range(len(heights_building_1))].</pre>
append	Question. Ninja Atsushi stands guard atop the Ninja Building, which consists of two adjacent buildings of equal height. Each building has a series of walls of varying heights. Atsushi must jump between buildings to patrol the rooftops. Given the heights of the walls for both buildings, determine the minimum height Atsushi must jump for each wall position. Input: two lists of integers, heights_building_1 and heights_building_2, with the same length. Output: a list of integers where each element is the minimum height Atsushi must jump for each wall position. Constraints: the common list length is between 1 and 10⁵, and each wall height is between 1 and 10⁹. Sample input: heights_building_1 = [3, 6, 2], heights_building_2 = [5, 4, 8]. Sample output: [5, 6, 8].< PAUSE >< PAUSE >< PAUSE >< PAUSE >< PAUSE >< PAUSE >< PAUSE >< PAUSE >< PAUSE > < PAUSE >< PAUSE > Answer. <pre>def min_jumps(heights_building_1, heights_building_2): Calculate the minimum height Atsushi must jump for each wall position between two buildings. :param heights_building_1: list of integers representing the heights of walls in the first building. :param heights_building_2: list of integers representing the heights of walls in the second building. :return: list of integers representing the minimum jump heights for each wall position. if len(heights_building_1) != len(heights_building_2): raise ValueError("Both buildings must have the same number of walls."). return [max(heights_building_1[i], heights_building_2[i]) for i in range(len(heights_building_1))].</pre>

Table 15: Examples of pause-token placement for the code task.

Method	Code example
random	Question. Ninja Atsushi stands guard atop the Ninja Building, which consists of two adjacent buildings of equal height. Each building has a series of walls of varying heights. Atsushi must jump between buildings to patrol the rooftops. Given the heights of the walls for both buildings, determine the minimum height Atsushi must jump for each wall position. Input: two lists of integers, heights_building_1 and heights_building_2, with the same length. Output: a list of integers where each element is the minimum height Atsushi must jump for each wall position. Constraints: the common list length is between 1 and 10^5, and each wall height is between 1 and 10^9. Sample input: heights_building_1 = [3, 6, 2], heights_building_2 = [5, 4, 8]. Sample output: [5, 6, 8]. Answer. <pre>def min_jumps(heights_building_1, heights_building_2): Calculate the minimum height Atsushi must jump for each wall position between two buildings. :param heights_building_1: list of integers representing the heights of walls in the first building. :param heights_building_2: list of integers representing the heights of walls in the second building. :return: list of integers representing the minimum jump heights for each wall position. if len(heights_building_1) != len(heights_building_2): raise ValueError("Both buildings must have the same number of walls.") return [max(heights_building_1[i], heights_building_2[i]) for i in range(len(heights_building_1))]</pre>
dit	Question. Ninja Atsushi stands guard atop the Ninja Building, which consists of two adjacent buildings of equal height. Each building has a series of walls of varying heights. Atsushi must jump between buildings to patrol the rooftops. Given the heights of the walls for both buildings, determine the minimum height Atsushi must jump for each wall position. Input: two lists of integers, heights_building_1 and heights_building_2, with the same length. Output: a list of integers where each element is the minimum height Atsushi must jump for each wall position. Constraints: the common list length is between 1 and 10^5, and each wall height is between 1 and 10^9. Sample input: heights_building_1 = [3, 6, 2], heights_building_2 = [5, 4, 8]. Sample output: [5, 6, 8]. Answer. <pre><def min_jumps(heights_building_1, heights_building_2): Calculate the minimum height Atsushi must jump for each wall position between two buildings. :param heights_building_1: list of integers representing the heights of walls in the first building. :param heights_building_2: list of integers representing the heights of walls in the second building. :return: list of integers representing the minimum jump heights for each wall position. if len(heights_building_1) != len(heights_building_2): raise ValueError("Both buildings must have the same number of walls.") <return [max(heights_building_1[i], heights_building_2[i]) for i in range(len(heights_building_1))]></pre>

Table 16: Examples of pause-token placement for the code task, continued.

Method	GSM8K	MATH500	MATH	AIME24	AMC23	Minerva	MMLU	GPQA	HellaSwag	BBH	PIQA	Average
Qwen3-1.7B-Base												
Base	65.13	54.20	54.82	0.00	<u>35.00</u>	17.65	62.95	27.46	38.69	38.84	62.13	41.53
SFT	<u>82.41</u>	63.40	62.98	0.00	27.50	20.59	62.32	24.78	40.39	42.06	66.10	44.78
Append	79.68	59.20	60.96	0.00	30.00	<u>22.79</u>	62.93	25.67	41.53	42.62	66.59	44.72
Random	81.35	60.80	<u>63.26</u>	0.00	20.00	23.16	<u>63.64</u>	23.66	48.73	43.96	69.64	45.29
DIT	68.01	57.20	56.38	6.67	30.00	20.22	62.90	25.67	44.75	42.66	<u>69.31</u>	43.98
MBP	83.02	61.00	64.56	<u>3.33</u>	37.50	23.16	63.93	<u>27.23</u>	<u>48.32</u>	<u>44.13</u>	67.63	47.62
MBP w/o masking	80.97	<u>62.20</u>	62.64	6.67	27.50	23.16	63.53	24.55	45.86	44.17	66.70	<u>46.18</u>
Qwen3-4B-Base												
Base	82.56	66.00	67.64	<u>6.67</u>	37.50	28.68	73.05	30.80	63.69	56.64	73.29	53.32
SFT	<u>91.36</u>	69.20	<u>71.24</u>	<u>6.67</u>	<u>45.00</u>	<u>29.78</u>	73.33	29.91	75.53	54.64	77.91	56.78
Append	89.92	67.00	68.62	10.00	42.50	28.68	73.35	29.02	76.47	54.94	78.45	56.27
Random	91.43	<u>70.40</u>	71.06	10.00	42.50	29.04	<u>73.57</u>	28.79	<u>76.14</u>	54.89	77.64	<u>56.86</u>
DIT	87.19	63.00	64.00	<u>6.67</u>	40.00	29.04	73.60	<u>30.36</u>	70.35	54.49	<u>79.43</u>	54.38
MBP	90.07	73.40	72.32	10.00	50.00	30.51	73.54	26.79	76.04	54.55	79.65	57.90
MBP w/o masking	90.98	69.20	70.14	<u>6.67</u>	40.00	26.84	73.03	26.12	76.07	54.44	75.08	55.32
Qwen3-8B-Base												
Base	88.40	65.40	65.86	13.33	42.50	26.84	77.79	<u>29.46</u>	70.82	58.48	78.24	56.10
SFT	92.95	71.60	72.54	<u>6.67</u>	52.50	32.72	78.21	27.90	74.80	<u>58.56</u>	85.09	59.41
Append	<u>93.25</u>	75.40	<u>73.00</u>	<u>10.00</u>	60.00	33.82	78.20	35.04	76.45	58.49	84.17	61.62
Random	93.03	73.80	72.26	<u>6.67</u>	47.50	34.56	78.53	22.32	74.40	58.82	84.98	58.81
DIT	93.93	74.40	73.54	<u>6.67</u>	52.50	35.66	78.00	26.12	73.64	58.14	85.53	<u>59.83</u>
MBP	92.57	<u>75.00</u>	72.84	3.33	<u>55.00</u>	<u>35.29</u>	<u>78.25</u>	23.66	<u>75.52</u>	57.28	<u>85.36</u>	59.46
Llama3.2-1B-Instruct												
Base	38.97	23.80	25.70	<u>0.00</u>	15.00	7.72	42.34	27.90	30.63	<u>32.43</u>	61.26	27.80
SFT	50.72	28.00	38.48	<u>0.00</u>	10.00	<u>8.09</u>	<u>42.09</u>	28.79	28.54	31.77	<u>58.92</u>	<u>29.58</u>
Append	45.87	28.20	33.62	<u>0.00</u>	15.00	6.62	41.34	25.00	<u>29.03</u>	32.33	58.11	28.65
Random	<u>51.93</u>	28.60	38.94	3.33	10.00	7.72	39.70	25.45	26.91	31.49	58.65	29.34
DIT	43.75	21.80	29.12	<u>0.00</u>	<u>17.50</u>	6.25	39.74	24.33	27.62	33.07	56.86	27.28
MBP	51.63	31.80	41.66	<u>0.00</u>	27.50	14.34	40.81	27.68	28.43	32.16	55.82	31.98
MBP w/o masking	52.16	<u>30.60</u>	<u>39.64</u>	<u>0.00</u>	2.50	6.99	40.90	26.12	27.35	32.38	55.11	28.52
Llama3.2-3B-Instruct												
Base	70.43	40.00	43.22	0.00	10.00	<u>17.65</u>	59.33	27.90	56.28	43.06	76.77	40.42
SFT	<u>76.80</u>	46.40	<u>55.02</u>	<u>3.33</u>	<u>27.50</u>	13.97	55.99	28.57	49.78	41.09	73.50	42.90
Append	73.77	44.60	52.26	6.67	25.00	15.44	56.13	<u>29.24</u>	50.76	40.78	74.05	42.61
Random	75.97	<u>47.80</u>	54.16	<u>3.33</u>	<u>27.50</u>	13.24	56.62	29.69	51.42	41.15	<u>74.10</u>	<u>43.18</u>
DIT	70.36	38.00	43.00	<u>3.33</u>	<u>17.50</u>	12.13	<u>57.60</u>	28.12	<u>52.39</u>	<u>42.41</u>	74.05	39.90
MBP	77.03	48.20	55.92	<u>3.33</u>	30.00	19.85	56.46	25.89	50.04	41.56	73.72	43.82
MBP w/o masking	76.19	47.40	54.58	<u>3.33</u>	20.00	15.81	55.64	25.22	49.93	40.73	73.78	42.06
Llama3.1-8B-Instruct												
Base	77.48	46.40	48.62	6.67	<u>26.50</u>	<u>22.43</u>	67.56	28.79	73.30	49.42	82.05	48.11
SFT	84.84	<u>53.20</u>	56.28	<u>3.33</u>	17.50	21.69	59.49	24.33	49.52	46.72	74.65	44.69
Append	84.46	49.80	53.08	<u>3.33</u>	20.00	20.59	60.84	26.56	50.24	<u>48.26</u>	73.67	44.62
Random	<u>85.90</u>	52.80	<u>56.44</u>	<u>3.33</u>	20.00	20.59	59.79	<u>29.91</u>	49.93	46.44	73.72	45.35
DIT	77.41	44.00	47.14	<u>3.33</u>	20.00	22.79	58.51	31.92	<u>52.53</u>	45.89	<u>75.46</u>	43.54
MBP	86.28	58.00	61.02	6.67	27.50	<u>22.43</u>	<u>61.30</u>	<u>29.91</u>	50.78	48.13	72.25	<u>47.66</u>

Table 17: Additional SFT results on mathematical reasoning and general benchmarks. All entries report accuracy. Average denotes the macro-average over all reported benchmark columns.

Method	MBPP	MBPP+	HumanEval	HumanEval+	MMLU	GPQA	HellaSwag	BBH	PIQA	Average
Qwen3-1.7B-Base										
Base	69.05	59.26	58.54	53.05	62.95	<u>27.46</u>	38.69	38.84	62.13	52.22
SFT	69.58	62.43	67.07	61.59	60.19	23.88	36.44	39.15	60.66	53.44
Append	69.84	60.58	<u>70.73</u>	65.24	61.53	27.90	36.58	40.65	64.09	55.24
Random	<u>70.37</u>	<u>61.11</u>	67.07	62.80	64.03	24.78	<u>48.36</u>	42.50	68.77	56.64
DIT	69.84	60.58	71.95	65.24	63.71	25.45	50.40	43.09	67.41	57.52
MBP	73.02	62.43	70.12	65.24	<u>63.78</u>	25.67	44.01	<u>42.88</u>	<u>68.01</u>	<u>57.24</u>
Qwen3-4B-Base										
Base	74.34	64.81	76.83	72.56	73.05	<u>30.80</u>	63.69	<u>56.64</u>	73.29	65.11
SFT	78.04	<u>67.99</u>	79.27	73.78	74.19	28.12	75.35	56.85	<u>78.45</u>	68.00
Append	80.42	68.25	80.49	74.39	<u>74.43</u>	34.15	77.15	53.38	75.63	68.70
Random	78.84	67.20	<u>82.32</u>	78.66	74.37	28.12	76.08	53.22	78.29	68.57
DIT	75.93	66.93	<u>81.10</u>	76.22	74.36	28.12	<u>76.38</u>	53.43	76.77	67.69
MBP	<u>79.10</u>	65.08	82.93	<u>78.05</u>	74.56	29.24	74.72	54.94	78.89	<u>68.61</u>
Qwen3-8B-Base										
Base	82.80	72.49	88.41	84.15	77.79	<u>29.46</u>	70.82	58.48	78.24	71.40
SFT	<u>86.51</u>	<u>75.13</u>	85.98	81.71	77.96	29.24	75.33	61.25	83.30	72.93
Append	85.19	72.75	83.54	78.05	77.89	23.88	<u>79.01</u>	61.39	82.97	71.63
Random	86.24	<u>75.13</u>	84.76	81.10	78.09	28.12	79.63	59.85	85.75	<u>73.19</u>
DIT	84.92	74.07	<u>87.20</u>	81.10	77.75	32.14	76.98	60.58	<u>83.62</u>	73.15
MBP	87.30	75.66	86.59	<u>82.32</u>	<u>78.02</u>	29.24	78.16	<u>61.31</u>	83.03	73.51
Llama3.1-8B-Instruct										
Base	67.46	59.52	67.68	62.20	67.56	28.79	73.30	49.42	82.05	62.00
SFT	68.25	60.05	68.29	64.02	<u>67.47</u>	26.56	73.63	49.55	81.88	62.19
Append	<u>69.58</u>	60.85	67.68	62.20	67.44	25.45	<u>73.87</u>	49.47	82.15	62.08
Random	69.31	60.05	<u>68.90</u>	63.41	67.37	25.45	73.60	49.76	<u>82.10</u>	62.22
DIT	70.11	<u>60.58</u>	68.29	64.02	67.42	29.24	73.75	<u>49.58</u>	<u>82.10</u>	62.79
MBP	68.52	60.05	69.51	64.02	67.41	28.35	73.93	49.36	<u>82.10</u>	<u>62.58</u>

Table 18: Additional SFT results on code generation and general benchmarks. Code benchmark entries report pass@1, while general benchmark entries report accuracy. Average denotes the macro-average over all reported benchmark columns.

Method	GSM8K	MATH500	MATH	AIME24	AMC23	Minerva	Avg.	AUC
Qwen3-1.7B-Base								
STANDARD	80.52	61.20	59.68	3.33	37.50	20.96	43.87	0.366
MBP	81.35	61.80	61.24	3.33	40.00	21.32	44.84	0.414
Qwen3-4B-Base								
STANDARD	92.19	74.40	74.30	13.33	50.00	30.15	55.73	0.505
MBP	91.58	76.00	75.26	16.67	57.50	23.90	56.82	0.550

Table 19: Full GRPO pass@1 accuracy (%) on six mathematical reasoning benchmarks and normalized training-reward AUC for Qwen3-1.7B-Base and Qwen3-4B-Base. Average denotes the macro-average over the six benchmarks.