arXiv:2611.33951v1 [cs.LG] 25 Nov 2026

BReD: Block Replay Dithering for Stable Low-Bit EMA Optimizer States

Heshen Zhan[1,4], **Youhan Huang**[1], **Yunke Peng**[2], **Yao Wang**[2], **Linghui Kong**[2]
Ziwei Zhu[1], **Qingyu Han**[1,4], **Yaoyuan Wang**[2], **Congliang Chen**[3,4,†], **Ruoyu Sun**[1,3,4,†]
[1]The Chinese University of Hong Kong, Shenzhen
[2]Huawei
[3]Shenzhen Loop Area Institute
[4]SICIAM, Shenzhen Research Institute of Big Data

Abstract

Low-bit storage of optimizer states can substantially reduce the memory footprint of large-scale training. In practice, however, these states are not quantized once: at every training step, they are dequantized, updated, requantized, and written back. This iterative quantization pipeline can destabilize training. We identify a failure mode behind this instability: for states updated via exponential moving averages (EMAs), quantization errors are recursively fed into subsequent updates. As a result, biased errors accumulate over time while high-variance errors are amplified, particularly when the EMA decay factor is close to 1. To address this error accumulation, we propose BReD (**B**lock **Re**play **D**ithering), a method that reduces rounding bias and controls rounding variance during quantization. BReD adapts classical subtractive dithering to optimizer-state quantization by combining deterministic seed replay with block-wise shared dither values, avoiding auxiliary random-tensor storage and per-element dither values generation. Across the default pretraining settings (five optimizers at 120M; AdamW and Muon at 1.1B and 3.4B model) and supervised fine-tuning (AdamW and Muon at 7B model), 4-bit BReD closely matches training with full-precision optimizer states, with PPL shifts ≤ 0.5 in pretraining and ≤ 0.2 in fine-tuning and average downstream performance changes within 0.7 points. Moreover, 3-bit BReD preserves stable convergence in the evaluated settings, offering a more aggressive alternative. [1]

1 Introduction

In large-scale model training, many optimizers maintain auxiliary states, such as moment estimates or preconditioner states. For example, AdamW [Kingma and Ba, 2014, Loshchilov and Hutter, 2017] keeps first- and second-moment estimates, Muon [Jordan et al.] uses a momentum buffer, and SOAP [Vyas et al., 2024] maintains matrix preconditioner states. These auxiliary states are used to compute the optimizer update, but they also introduce additional memory cost because they must be kept throughout training. Since this cost increases with model size, storing optimizer states in low-bit formats is a direct way to lower the memory footprint of training.

A growing line of work has studied low-bit optimizer-state storage as a way to reduce this cost, including low-bit Adam variants [Li et al., 2023, Xu et al., 2025], low-bit Muon methods [Wu et al., Gupta et al., 2025], and low-bit Shampoo methods [Wang et al., 2024, Li et al., 2025]. These methods show that optimizer-state memory can be substantially reduced. However, at aggressive low bitwidths,

[†]Corresponding authors: Ruoyu Sun (sunruoyu@cuhk.edu.cn) ; Congliang Chen (chencongliang@slai.edu.cn).
[1]Our implementation of BReD is available at: https://github.com/ZhanHeshen/BReD.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

stability becomes harder to maintain, and successful recipes are often tied to particular optimizer families, update modifications, or retuned hyperparameters. This raises a complementary question: can low-bit optimizer-state storage be stabilized across different optimizers while preserving the original optimizer update and hyperparameters?

We address this question by shifting focus from optimizer update rules to a shared mechanism in many optimizers: exponential moving average (EMA) recursions. For EMA optimizer states, low-bit storage is not a one-time quantization problem. At each training step, the low-bit state is dequantized for the optimizer update, then requantized and written back for future use. Let e_t denote the difference between the dequantized stored state and the corresponding full-precision state, and let ϵ_{t+1} be the error injected when the updated state is quantized. Then

$$e_{t+1} = \beta e_t + \epsilon_{t+1},$$

where β is the EMA decay factor (Sec. 3.1). Thus, the effect of quantization is governed not only by the error injected at each step, but also by how this error accumulates through the EMA recursion.

The rounding rule is a controllable component of quantization. Ignoring clipping or saturation, quantization error is determined by the rounding rule for a given scale and quantization grid. The error accumulation points to three design goals for the rule. First, rounding error should be unbiased, as even small systematic biases can accumulate into persistent drift. Second, unbiasedness alone is not enough: the error variance should also be controlled, because the EMA recursion amplifies the variance term as well. Third, the rule must be practical at training scale: optimizer states may contain billions of entries, so it should not require storing large random tensors or generating independent random numbers for every element.

Guided by these design goals, we propose **BReD** (**B**lock **Re**play **D**ithering), a rounding-based scheme for low-bit storage of EMA optimizer states. BReD applies subtractive dithering [Roberts, 1962] in quantization and dequantization, producing unbiased reconstruction error with controlled variance. While subtractive dithering is classical, applying it to optimizer-state storage is nontrivial: identical dither values must be available at the dequantization time, which typically requires storing large random tensors or regenerating dither values for every element. BReD addresses this challenge using deterministic seed replay and block-wise sharing of dither values. It enables practical subtractive dithering in large-scale training without modifying optimizer updates or hyperparameters.

We evaluate BReD on pretraining and supervised fine-tuning. The pretraining experiments cover five optimizers at 120M model, AdamW/Muon at 1.1B and 3.4B model and a long-horizon AdamW run, while the fine-tuning experiments evaluate AdamW and Muon at 7B model. Across these settings, 4-bit BReD incurs small changes relative to the corresponding full-precision baselines, with perplexity shifts within 0.5 in pretraining and 0.2 in fine-tuning, and average downstream performance changes within 0.7 points. This stability is not simply a consequence of the low-bit format: under the same scaling rule and quantization grid, replacing BReD with nearest or stochastic rounding can diverge or downgrade during pretraining. Moreover, 3-bit BReD remains stable in the evaluated settings, offering an aggressive alternative.

Our contributions are:

- **Mechanism.** We identify error accumulation as a critical failure mode behind low-bit storage of EMA optimizer states, and we characterize its underlying mechanism: the bias and variance of quantization errors continuously propagate through EMA recursive updates.
- **Method.** We propose BReD, a quantization scheme for low-bit storage of EMA optimizer states. It adapts classical subtractive dithering with two designs: deterministic seed replay and block-wise shared dither values. Importantly, BReD leaves the original optimizer update rules and hyperparameter configurations unchanged.
- **Evaluation.** We evaluate BReD across various optimizers, model scales, and training setups. Our experiments include rounding baselines, low-bit optimizer methods, a longer-horizon test, and memory-perplexity Pareto analysis. BReD enables robust 4-bit optimizer state compression in all evaluated settings, and remains convergent at 3-bit precision in the evaluated settings.

2 Related Work

Low-bit optimizer states. Optimizer states are a major memory bottleneck in large-scale training. Existing approaches reduce this cost structurally, as in Adafactor [Shazeer and Stern, 2018] and Adam-mini [Zhang et al., 2024] or at the system level, as in ZeRO-style sharding/offloading [Rajbhandari et al., 2020, Ren et al., 2021, Rajbhandari et al., 2021]. We study a complementary axis: low-bit storage of EMA optimizer states. Prior work has shown strong results for block-wise 8-bit optimizers [Dettmers et al., 2021] and 4-bit AdamW [Li et al., 2023]; we refer to the latter as Adam-LPMM. Recent methods further push low-bit AdamW states by focusing on EMA dynamics, as in SOLO [Xu et al., 2025], or target Muon states [Wu et al., Gupta et al., 2025] and preconditioner states, as in 4-bit Shampoo [Wang et al., 2024] and 4-bit Preconditioned Optimization [Li et al., 2025]. In contrast, BReD targets the shared quantize–write-back path of EMA states. It changes the rounding and reconstruction rule used for state storage, while keeping the optimizer update and hyperparameters unchanged.

Unbiased rounding and dithering. Unbiased rounding reduces systematic rounding bias. Stochastic rounding is unbiased in expectation [Von Neumann and Goldstine, 1947, Alistarh et al., 2017]. Subtractive dithering is a classical technique that gives unbiased reconstruction [Roberts, 1962, Jayant and Rabiner, 1972, Schuchman, 1964, Sherwood, 1985]; under the local non-clipped interval model used in 4.3, it has smaller worst-case conditional variance than stochastic rounding. However, subtractive dithering requires the same dither value at quantization and dequantization time, which is impractical at LLM scale if per-element dither values must be stored or generated. BReD makes subtractive dithering practical for persistent EMA optimizer states through deterministic seed replay and block-wise shared dither values.

3 Understanding Instability: Quantized EMA Error Accumulation

Many optimizer states follow EMA recursions, such as the first and second moments in AdamW and the momentum buffer in Muon. When these states are quantized and written back every step, the injected quantization error is not a one-time perturbation; it can be recursively fed into later EMA updates. We formulate this storage-level error accumulation mechanism below. The goal is to isolate one instability mechanism in low-bit optimizer-state storage, not to provide a full convergence proof of the optimizer.

3.1 Quantized-EMA Error Model

Consider a full-precision EMA state and its low-bit stored counterpart:

$$\begin{aligned} s_{t+1} &= \beta s_t + (1-\beta)u_{t+1}, \qquad 0 < \beta < 1, \\ s'_{t+1} &= \beta s_t^{dq} + (1-\beta)u_{t+1}, \\ \epsilon_{t+1} &:= s_{t+1}^{dq} - s'_{t+1}, \qquad e_t := s_t^{dq} - s_t. \end{aligned} \tag{1}$$

Here u_{t+1} is the per-step update term, s_t^{dq} is the dequantized stored state, ϵ_{t+1} is the error injected when the updated state is quantized and written back, and e_t is the accumulated state error relative to the full-precision EMA. A direct subtraction gives

$$e_{t+1} = \beta e_t + \epsilon_{t+1}. \tag{2}$$

Thus, injected errors are fed back through the EMA state rather than appearing as independent one-step perturbations.

3.2 Theoretical Analysis: Mean and Noise Amplification in EMA

We now characterize how the EMA recursion propagates the bias and variance of the injected error.

Lemma 3.1 (Bias induces persistent error under EMA recursion)**.** *Assume the injected errors satisfy* $\mathbb{E}[\epsilon_t] = b$*. For the recursion* $e_{t+1} = \beta e_t + \epsilon_{t+1}$ *with* $0 < \beta < 1$*, we have*

$$\mathbb{E}[e_t] = \beta^t \mathbb{E}[e_0] + \frac{1-\beta^t}{1-\beta} b \quad \longrightarrow \quad \frac{b}{1-\beta}. \tag{3}$$

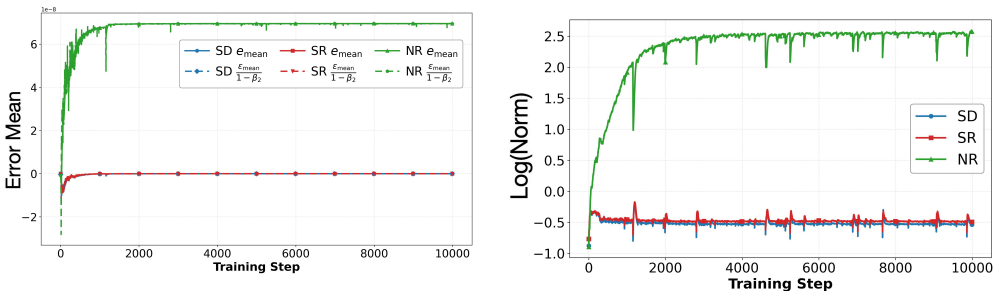

(a) Mean drift matches $\mathbb{E}[\epsilon_t]/(1-\beta_2)$.

(b) Relative error $\log_{10}(\|e_t\|_F/\|s_t\|_F)$.

Figure 1: In-training diagnostics for the quantized-EMA error model. During LLaMA-120M AdamW pretraining, we track the accumulated state error e_t and injected write-back error ϵ_t for the second-moment state. (a) The measured mean error follows $\mathbb{E}[\epsilon_t]/(1-\beta_2)$ after the transient, consistent with EMA amplification of rounding bias. (b) NR accumulates large relative error, whereas SR and SD maintain bounded relative errors in this diagnostic, with SD slightly smaller.

In particular, even a small nonzero bias b yields a persistent error whose magnitude scales as $O(\|b\|/(1-\beta))$ when $\beta \to 1$.

Theorem 3.2 (Steady-state second moment under quantized EMA)**.** *Assume $\epsilon_t = b + \xi_t$, where $\{\xi_t\}_{t\geq 1}$ are independent with $\mathbb{E}[\xi_t] = 0$ and $\mathbb{E}\|\xi_t\|_F^2 = \sigma^2$. Assume also that e_0 has finite second moment and is independent of $\{\xi_t\}_{t\geq 1}$. For $0 < \beta < 1$, the state error satisfies*

$$\lim_{t\to\infty} \mathbb{E}\|e_t\|_F^2 = \underbrace{\frac{\|b\|_F^2}{(1-\beta)^2}}_{\text{mean (bias) term}} + \underbrace{\frac{\sigma^2}{1-\beta^2}}_{\text{variance (noise) term}}. \tag{4}$$

Remark 3.3 (Mean dominates when $\beta \to 1$). The ratio between the mean term and the variance term in Eq. (4) equals

$$\frac{\|b\|_F^2/(1-\beta)^2}{\sigma^2/(1-\beta^2)} = \frac{(1+\beta)\|b\|_F^2}{(1-\beta)\sigma^2}, \tag{5}$$

which grows as $1/(1-\beta)$ when $\beta \approx 1$. Thus, even a tiny nonzero mean in rounding error ($b \neq 0$) can dominate the long-horizon error budget for large-β EMA states, such as AdamW's second moment.

Proofs of Lemma 3.1 and Theorem 3.2 are given in Appendix E. Eq. (4)–(5) motivates a two-step design rule for low-bit EMA states: first reduce systematic rounding bias, then control the variance.

3.3 Real Setting: In-training Evidence

We verify that the same mechanism appears in real LLM training. During LLaMA-120M pretraining with AdamW, we focus on the second-moment tensor v in layer `11.attn.proj` and maintain measurement-only quantized EMA copies under NR, SR, and SD. These copies are recursively updated and requantized, but do not affect model updates or subsequent gradients. Using the definitions above, Eq. (2) predicts the steady-state relation $\mathbb{E}[e_\infty] \approx \mathbb{E}[\epsilon_t]/(1-\beta_2)$ after the initial transient. As shown in Fig. 1, the measured $\mathbb{E}[e_t]$ closely tracks this prediction: NR exhibits persistent drift, while SR and SD keep the mean near zero. The relative-error plot further shows that NR accumulates larger state error, whereas SR and SD maintain bounded relative errors in this diagnostic, with SD slightly smaller than SR. Additional layer diagnostics are provided in Appendix F.

4 Method: BReD for EMA Optimizer States

4.1 Overview

Section 3 shows that, when an EMA state is quantized and written back every step, the mean component of the injected error can be amplified by the EMA recursion, and the variance is also amplified. This motivates a simple design rule: remove systematic rounding bias first, then control the remaining variance.

We implement this rule as BReD, a rounding-centric framework built from three components: subtractive dithering (SD), deterministic seed replay (DSR), and block-wise randomization (BWR). SD gives unbiased reconstruction with controlled variance under the local non-clipped interval model. DSR deterministically reproduces the same dither value at dequantization time without storing random tensors. BWR amortizes random-number generation by sharing one random scalar within each block. Combining these components gives a practical subtractive-dithering module for EMA optimizer states. The same scheme can be applied to different EMA optimizer states by dequantizing for computation and requantizing for storage each step, without changing optimizer updates or hyperparameters.

4.2 Element-wise View: Interval Parameterization

Given a scale and a quantization grid, and ignoring clipping, the local quantization error is determined by the rounding and reconstruction rule. We therefore adopt an element-wise perspective on rounding a scalar to a quantization grid. Let $p_0(x) \le x \le p_1(x)$ be the two neighboring grid points that bracket x, and define the interval width and coordinate

$$\Delta(x) := p_1(x) - p_0(x), \qquad \alpha(x) := \frac{x - p_0(x)}{\Delta(x)} \in [0, 1]. \tag{6}$$

Hence, rounding can be written as choosing a binary endpoint indicator $c \in \{0, 1\}$, followed by reconstruction of $\hat{\alpha}$ and $\hat{x}$:

$$c \in \{0, 1\}, \qquad \hat{x} = p_0(x) + \Delta(x)\hat{\alpha}. \tag{7}$$

The rounding–reconstruction gap is

$$\epsilon_{\text{elem}}(x) := \hat{x} - x = \Delta(x)\big(\hat{\alpha} - \alpha(x)\big). \tag{8}$$

A rounding rule is locally unbiased when $\mathbb{E}[\hat{\alpha} \mid \alpha] = \alpha$, which gives $\mathbb{E}[\epsilon_{\text{elem}} \mid x] = 0$ by Eq.(8). Clipping or saturation can introduce additional error; in our comparisons, the scale, grid, and clipping behavior are kept identical across NR, SR, and BReD.

4.3 Subtractive Dithering as an Unbiased then Low Variance Primitive

Section 3 shows that EMA error feedback amplifies the mean component of quantization error. We therefore first enforce unbiased rounding, then reduce the remaining variance.

Nearest rounding (NR). NR deterministically selects the nearest endpoint:

$$c_{\text{nr}}(\alpha) := \mathbb{1}\{\alpha \ge 1/2\}, \qquad \hat{\alpha} = c_{\text{nr}}(\alpha). \tag{9}$$

NR is computationally efficient, but it can introduce biased rounding errors for asymmetric data distributions.

Stochastic rounding (SR). SR draws $u \sim \text{Unif}(0, 1)$ and selects the upper grid point with probability α:

$$c_{\text{sr}}(\alpha; u) := \mathbb{1}\{u < \alpha\}, \qquad \hat{\alpha} = c_{\text{sr}}(\alpha; u). \tag{10}$$

Then $\mathbb{E}[\hat{\alpha} \mid \alpha] = \alpha$, hence $\mathbb{E}[\epsilon_{\text{elem}} \mid x] = 0$, and

$$\text{Var}(\epsilon_{\text{elem}} \mid x) = \Delta(x)^2\alpha(1 - \alpha) \le \Delta(x)^2/4. \tag{11}$$

Although SR achieves unbiasedness, it can still leave substantial residual noise, which is also accumulated by the EMA recursion.

Subtractive dithering (SD). SD draws $u \sim \text{Unif}(0, 1)$, rounds a dithered coordinate to a grid point, and reconstructs a continuous value using the same u:

$$c_{\text{sd}}(\alpha; u) := \mathbb{1}\{\alpha + u \ge 1\} \in \{0, 1\}, \tag{12}$$

$$\hat{\alpha} := c_{\text{sd}}(\alpha; u) - u + \frac{1}{2}, \qquad \hat{x} := p_0(x) + \Delta(x)\hat{\alpha}. \tag{13}$$

SD is unbiased, i.e., $\mathbb{E}[\epsilon_{\text{elem}} \mid x] = 0$, and has

$$\text{Var}(\epsilon_{\text{elem}} \mid x) = \Delta(x)^2/12, \tag{14}$$

which is independent of α. Therefore, SD does not dominate SR pointwise in conditional variance, but it provides α-independent noise and a smaller worst-case conditional variance ($\Delta(x)^2/12$ versus $\Delta(x)^2/4$ for SR). A short proof is deferred to Appendix E.

Under the local non-clipped interval model, SR errors can approach one bin width, whereas SD errors are bounded by half a bin (Appendix E.1). This difference offers a possible explanation for the observed SR failures, but is not a divergence criterion derived from Theorem 3.2.

4.4 Deterministic Seed Replay (DSR): Stateless Reconstruction of Dither Values

Subtractive dithering requires using the same random dither value u at quantization and dequantization; otherwise the subtractive correction does not cancel. Storing u per element would introduce an extra tensor comparable to the state itself, defeating compression. We therefore regenerate dither values on the fly via a stateless keyed pseudo-random number generator (PRNG).

For each EMA state tensor, we assign a fixed identifier `state_id`, such as a tuple containing the optimizer-state type and tensor index. For AdamW, for example, the state type distinguishes the first and second moments; the tensor index distinguishes different parameter tensors. When the state is written at training step t, for each block index j, we form a deterministic key

$$\kappa(j,t) := (\texttt{global_seed}, \texttt{state_id}, t, j), \tag{15}$$

and draw the dither as

$$u \leftarrow \mathrm{PRNG}(\kappa(j,t)) \in (0,1). \tag{16}$$

When decoding the stored state written at step t, we reuse the same key $\kappa(j,t)$ to replay the identical u, thus enabling subtractive dithering without storing random tensors.

4.5 Block-wise Randomization (BWR): Amortizing Randomness

Randomized rounding, including SR and SD, requires generating random numbers at scale. To reduce random-number generation overhead, BWR shares one random scalar within each block.

Let $\{B_k\}_{k=1}^K$ be a partition of indices into blocks of size B. Instead of sampling i.i.d. randomness per coordinate, BWR samples one $u^{(k)} \sim \mathrm{Unif}(0,1)$ per block and reuses it for all $i \in B_k$. For example, block-wise SR uses

$$c_i := \mathbb{1}\{u^{(k)} < \alpha(x_i)\}, \qquad i \in B_k. \tag{17}$$

For SD, the same block dither $u^{(k)}$ is used in Eq.12–13 for all elements in the block and is replayed during dequantization. If the underlying element-wise rule is unbiased, BWR preserves coordinate-wise marginal unbiasedness after taking expectation over $u^{(k)}$. BWR reduces PRNG calls from $O(d)$ to $O(d/B)$ while introducing within-block correlation, whose impact we study via block-size ablations.

4.6 BReD: Block Replay Dithering (SD + DSR + BWR)

We now introduce **BReD** (**B**lock **Re**play **D**ithering), a practical method for EMA optimizer-state quantization. BReD combines: (i) SD to achieve unbiased reconstruction error with controlled variance; (ii) DSR to replay the same dither values at dequantization time without storing random tensors; and (iii) BWR to amortize PRNG calls by sharing one dither value per block.

At write time, for each block written at training step t, BReD generates a single dither value via the keyed PRNG in Eq.16, applies subtractive dithering to the elements in the block, and stores the resulting low-bit codes together with the usual block scale. At read time, the same $u^{(k)}$ is replayed to decode the stored values and reconstruct the state. BReD can be combined with other scaling methods; in this work, we primarily use Microscaling data formats [Rouhani et al., 2023b,a].

5 Experiments

We evaluate BReD for low-bit storage of EMA optimizer states in pretraining and supervised fine-tuning. The pretraining experiments follow four questions: whether the rounding rule matters under

a fixed scaling rule and quantization grid; whether BReD can be combined with optimizer-specific low-bit methods; whether stability persists at larger scales and longer token horizons; and how BReD changes the memory–perplexity trade-off. We then evaluate supervised fine-tuning and ablate the block-wise randomization design.

Unless otherwise stated, all low-bit variants use the same optimizer update and hyperparameters as the corresponding full-precision optimizer-state baseline. The only change is how the selected EMA optimizer states are stored between training steps. Standard pretraining and fine-tuning results report mean $\pm$ standard deviation over three runs. The 100B-token experiment uses one run due to compute cost. The additional high-β check uses one paired run, and each SOLO and Shampoo configuration is evaluated once.

5.1 Experimental setup

Training setups. For pretraining, we evaluate AdamW [Kingma and Ba, 2014, Loshchilov and Hutter, 2017], Muon [Jordan et al.], NAdamW [Dozat, 2016], SOAP [Vyas et al., 2024], and MARS [Yuan et al., 2024] at the 120M scale, and AdamW and Muon at the 1.1B and 3.4B scales. The 120M and 1.1B models use TinyLlama-style configurations, while the 3.4B model follows OpenLLaMA. Models are trained on SlimPajama for 10.5B tokens at 120M and 31.4B tokens at 1.1B/3.4B. We report validation perplexity (PPL) and zero-shot accuracies on HellaSwag [Zellers et al., 2019], OpenBookQA [Mihaylov et al., 2018], WinoGrande [Sakaguchi et al., 2021], ARC [Clark et al., 2018], BoolQ [Clark et al., 2019], and PIQA [Bisk et al., 2020]. For supervised fine-tuning, we fine-tune Qwen2.5-7B [Qwen et al., 2025] on UltraFeedback [Cui et al., 2023] and evaluate HumanEval [Chen et al., 2021], MBPP [Austin et al., 2021], GSM8K [Cobbe et al., 2021], TruthfulQA [Lin et al., 2022], and validation PPL. Our fine-tuning recipe follows prior work on instruction-tuning [Yu et al., 2023, Liu et al., 2023, Cui et al., 2023]. Optimizer hyperparameters are provided in Appendices A and B; task-wise pretraining results are provided in Appendix C.

Quantized states and baselines. BReD quantizes EMA optimizer states, such as the first and second moments in AdamW and the momentum states in Muon-style optimizers. For SOAP, we additionally quantize the preconditioner states. Following prior low-bit optimizer-state work [Dettmers et al., 2021], we keep embedding-layer optimizer states in full precision during pretraining. Unless otherwise stated, NR, SR, and BReD use the same block scaling rule and quantization grid. For 4-bit storage, we use MXFP4 by default; for 3-bit storage, we use the grid $\{0, \pm 1, \pm 2, \pm 3\}$ with Microscaling. An additional MXINT4 check with 120M AdamW also converges, reaching a validation PPL of 18.9 (Appendix C.4). We compare rounding baselines (NR/SR), optimizer-specific low-bit methods (AdamW-LPMM [Li et al., 2023] and Muon-Grasp[Wu et al.]).

5.2 Pretraining: rounding controls and low-bit methods

We first isolate the rounding rule under a fixed scaling rule, quantization grid, optimizer update, and hyperparameters. As shown in Table 1(a), NR and SR fail or substantially degrade in all five 120M optimizer settings: four optimizers produce NaN or divergence, and Muon completes training only with much worse PPL. In contrast, BReD remains stable across all five optimizers, with validation PPL between 18.0 and 18.9. This controlled comparison shows that stability in aggressive 4-bit EMA-state storage is not a consequence of the low-bit format alone. An additional paired run with AdamW at $\beta_2 = 0.999$ also remains stable, with a PPL increase of 0.7 relative to the matched FP32-state baseline (Appendix C.3). An additional single-run check for each configuration with 120M Shampoo also finds that BReD completes training with 4-bit preconditioner states, whereas NR and SR diverge (Appendix C.6).

We next test whether BReD can be combined with optimizer-specific low-bit methods. Table 1(b) shows that adding BReD improves AdamW-LPMM from 18.7 to 18.4 PPL, and matches the best observed Muon-Grasp PPL of 17.8 while increasing average zero-shot accuracy from 40.6 to 41.2. Thus, BReD can serve as a rounding component for EMA-state write-back rather than only as a standalone MXFP4 state wrapper. We also evaluate SOLO [Xu et al., 2025] in our 120M pretraining setup (Appendix C.5). We report these checks separately from the controlled rounding comparison because SOLO also changes the quantization grids and first-moment dynamics.

Table 1: Pretraining controls at the 120M scale. (a) Controlled 4-bit rounding comparison. NR, SR, and BReD use the same scaling rule, quantization grid, optimizer update, and hyperparameters. Numbers in parentheses are validation PPL. (b) Combination with optimizer-specific low-bit methods.

(a) Controlled 4-bit rounding comparison

Optimizer	FP32-State	NR	SR	BReD
AdamW	18.3	Diverged	Diverged	**Stable (18.7)**
Muon	17.8	Degraded (31.0)	Degraded (21.7)	**Stable (18.0)**
NAdamW	18.9	Diverged	Diverged	**Stable (18.9)**
SOAP	17.6	Diverged	Diverged	**Stable (18.1)**
MARS	18.7	Diverged	Diverged	**Stable (18.9)**

(b) Combination with optimizer-specific low-bit methods

Method	Avg. acc. ↑		PPL ↓	
	Base	+BReD	Base	+BReD
AdamW-LPMM	40.3	40.5 (+0.2)	18.7	18.4 (-0.3)
Muon-Grasp	40.6	41.2 (+0.6)	17.8	17.8 (+0.0)

Table 2: Scale check for 4-bit BReD in pretraining. Parentheses denote changes relative to the corresponding full-precision optimizer-state baseline.

Scale	Optimizer	Avg. acc. ↑		PPL ↓	
		FP32-state	4-bit BReD	FP32-state	4-bit BReD
1.1B	AdamW	46.5	46.4 (-0.1)	10.8	10.9 (+0.1)
1.1B	Muon	47.6	47.1 (-0.5)	10.3	10.6 (+0.3)
3.4B	AdamW	47.9	47.6 (-0.3)	9.7	9.8 (+0.1)
3.4B	Muon	48.2	47.7 (-0.5)	9.6	9.8 (+0.2)

5.3 Pretraining: scale, horizon, and memory trade-offs

We next test whether 4-bit BReD remains stable beyond the 120M pretraining runs. Table 2 reports larger-scale results for AdamW and Muon at 1.1B and 3.4B. Across these settings, BReD remains close to the corresponding full-precision optimizer-state baselines, with PPL shifts no larger than 0.3 and average zero-shot accuracy changes within 0.5 points. Details can be found in Table 8.

We also run a longer-horizon AdamW pretraining test for 100B tokens at the 120M scale. As shown in Figure 2(a), 4-bit BReD remains close to the full-precision optimizer-state loss curve throughout training. This experiment is intentionally scoped to AdamW, but uses a much longer token horizon than the standard 120M runs.

Finally, we report a memory-perplexity case study for 120M AdamW in Figure 2(b). Moving from the 8-bit state-storage setting to 4-bit BReD reduces optimizer-state memory from 0.28GB to 0.19GB while keeping validation PPL close. Moving to 3-bit further reduces memory to 0.17GB but increases validation PPL to 19.0, while the 2-bit setting fails. This 2-bit failure is specific to the evaluated recipe and does not establish a fundamental lower bound for subtractive dithering. Across the evaluated optimizers and scales, packed 4-bit BReD reduces optimizer-state memory by 35.9%–86.0%; the full memory table is provided in Appendix C.7.

End-to-end efficiency. Using the default training configurations reported in the paper, Table 3 reports matched FP32-state and 4-bit BReD measurements for AdamW. BReD reduces persistent optimizer-state storage by 78.9%–86.0%, with throughput reductions of 5.1%–16.0%. Timing includes forward, backward, and all optimizer-state processing, including decoding and requantization, but excludes compilation, evaluation, and checkpointing. For gradient accumulation, micro-steps are aggregated into optimizer steps. The relative overhead is smaller in the larger configurations, but model size and gradient accumulation change together, so this trend cannot be attributed solely to model scale.

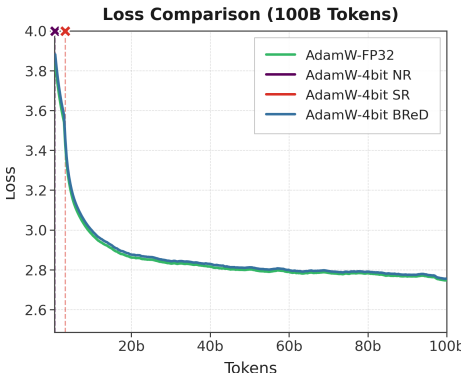

(a) 100B-token AdamW loss

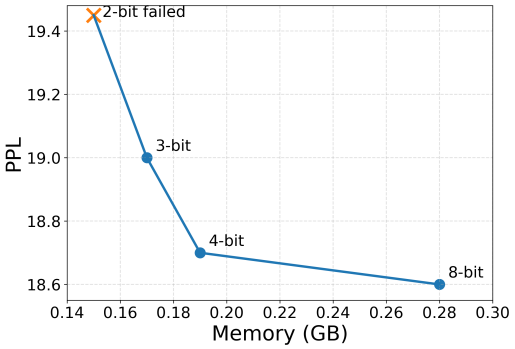

(b) 120M AdamW Pareto

Figure 2: Pretraining horizon and memory-quality checks. (a) Longer-horizon AdamW pretraining at the 120M scale. (b) Memory-perplexity trade-off for 120M AdamW pretraining. The Pareto plot is a case study; broader stability evidence is provided by the controlled and scale experiments.

Table 3: End-to-end AdamW efficiency. FP32 and BReD denote full-precision and 4-bit optimizer-state storage, respectively. Memory reports persistent optimizer-state storage, not peak device memory.

	Step time (s)		Throughput (M tok/s)		State memory (GB)	
Scale	FP32	BReD	FP32	BReD	FP32	BReD
120M	0.547	0.651	1.917	1.611	0.90	0.19
1.1B	2.881	3.178	0.364	0.330	8.20	1.18
3.4B	10.386	10.945	0.101	0.096	25.53	3.58

5.4 Supervised fine-tuning

We next test whether the same optimizer-state storage rule transfers beyond pretraining. We fine-tune Qwen2.5-7B on UltraFeedback and evaluate HumanEval, MBPP, GSM8K, TruthfulQA, and validation PPL. As shown in Table 4, 4-bit BReD remains close to the corresponding full-precision optimizer-state baselines for both AdamW and Muon. For AdamW, the average downstream score changes from 54.2 to 54.9 with no change in validation PPL. For Muon, the average downstream score changes from 55.2 to 54.9, again with no change in validation PPL. These results suggest that BReD can be applied in a downstream fine-tuning recipe without changing optimizer hyperparameters.

5.5 Aggressive 3-bit compression and ablations

The main results use 4-bit BReD as the default operating point. We also evaluate 3-bit BReD as a more aggressive compression setting. As summarized in Table 5(a), 3-bit BReD remains convergent in all evaluated pretraining and fine-tuning settings, but the quality shifts are larger and less uniform than in the 4-bit case. We therefore treat 3-bit BReD as an aggressive memory-quality trade-off rather than the default setting; results for each evaluated setting are provided in Appendix C.2.

We further ablate the block-wise randomization component of BReD. BWR shares one dither value within each block, reducing PRNG calls from $O(d)$ to $O(d/B)$ while introducing within-block correlation. Table 5(b) shows that the default $B = 32$ remains close to element-wise dithering in quality, while reducing the number of generated random values by $32\times$ and lowering generation time from 486.32 ± 29 ms to 340.88 ± 41 ms.

6 Conclusion

We studied a storage-level failure mode in low-bit EMA optimizer states: stepwise quantization and write-back can accumulate biased rounding error and amplify residual noise when the EMA coefficient is large. We proposed **BReD** (**B**lock **Re**play **D**ithering), which makes subtractive dithering

Table 4: Supervised fine-tuning with 4-bit BReD-quantized optimizer states. We fine-tune Qwen2.5-7B on UltraFeedback and report average downstream score and validation PPL. Parentheses denote changes relative to the corresponding full-precision optimizer-state baseline.

Optimizer	Avg. score ↑		PPL ↓	
	FP32-state	4-bit BReD	FP32-state	4-bit BReD
AdamW	54.2	54.9 (+0.7)	4.6	4.6 (+0.0)
Muon	55.2	54.9 (-0.3)	4.6	4.6 (+0.0)

Table 5: Aggressive compression and BWR ablation. (a) Summary of 3-bit BReD results. ΔAvg. denotes the range of changes relative to the corresponding full-precision optimizer-state baseline. (b) Block-wise randomization ablation for 4-bit AdamW at the 120M scale.

(a) 3-bit BReD summary

Setting	Architecture / size	Optimizers	ΔAvg.	ΔPPL
Pretrain	LLaMA-style / 120M	5 opt.	$[+0.0, +1.0]$	$[+0.1, +1.2]$
Pretrain	LLaMA-style / 1.1B	AdamW, Muon	$[-0.9, -0.7]$	$[+0.1, +0.6]$
SFT	Qwen2.5 / 7B	AdamW, Muon	$[-0.8, +0.8]$	$[-0.1, +0.1]$

(b) BWR ablation

Block size B	Avg. acc. ↑	PPL ↓	Random ratio	Generation time
1	40.9	18.8	1	486.32 ± 29 ms
32	40.7	18.7	$1/32$	340.88 ± 41 ms
256	40.6	18.8	$1/256$	335.47 ± 35 ms

practical with deterministic seed replay and block-wise shared dither values, without changing optimizer updates. Across pretraining and fine-tuning, BReD provides robust 4-bit optimizer-state storage and stable 3-bit convergence in the evaluated settings.

7 Limitations

Due to limited computational resources, our largest pretraining experiments use a 3.4B model and cover only AdamW and Muon. Broader validation at larger scales (e.g., tens of billions of parameters), longer token budgets, and additional optimizer families would further strengthen the conclusions.

Cross-optimizer breadth is evaluated at 120M, whereas larger-scale pretraining covers only AdamW and Muon. We do not provide matched comparisons with Adafactor [Shazeer and Stern, 2018], Adam-mini [Zhang et al., 2024], 4-bit Shampoo [Wang et al., 2024], or 4-bit Preconditioned Optimization [Li et al., 2025]. Joint use with gradient or activation quantization remains unevaluated.

Our storage-error analysis does not characterize how quantization errors affect eigendecomposition or other spectral operations on preconditioner states.

Acknowledgment

Heshen Zhan would like to thank Qihao Cheng and Xuefei Wang for the discussions of experimental observations.

This paper is supported by NSFC (No. 12326608); Hetao Shenzhen-Hong Kong Science and Technology Innovation Cooperation Zone Project (No.HZQSWS-KCCYB-2024016); Guangdong Provincial Key Laboratory of Mathematical Foundations for Artificial Intelligence (2023B1212010001).

References

Dan Alistarh, Demjan Grubic, Jerry Li, Ryota Tomioka, and Milan Vojnovic. Qsgd: Communication-efficient sgd via gradient quantization and encoding. *Advances in neural information processing*

systems, 30, 2017.

Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*, 2021.

Yonatan Bisk, Rowan Zellers, Jianfeng Gao, Yejin Choi, et al. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 7432–7439, 2020.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.

Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. Boolq: Exploring the surprising difficulty of natural yes/no questions. *arXiv preprint arXiv:1905.10044*, 2019.

Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge. *ArXiv*, abs/1803.05457, 2018.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.

Ganqu Cui, Lifan Yuan, Ning Ding, Guanming Yao, Bingxiang He, Wei Zhu, Yuan Ni, Guotong Xie, Ruobing Xie, Yankai Lin, et al. Ultrafeedback: Boosting language models with scaled ai feedback. *arXiv preprint arXiv:2310.01377*, 2023.

Tim Dettmers, Mike Lewis, Sam Shleifer, and Luke Zettlemoyer. 8-bit optimizers via block-wise quantization. *arXiv preprint arXiv:2110.02861*, 2021.

Timothy Dozat. Incorporating nesterov momentum into adam. 2016.

Aman Gupta, Rafael Celente, Abhishek Shivanna, DT Braithwaite, Gregory Dexter, Shao Tang, Hiroto Udagawa, Daniel Silva, Rohan Ramanath, and S Sathiya Keerthi. Effective quantization of muon optimizer states. *arXiv preprint arXiv:2509.23106*, 2025.

Shengding Hu, Yuge Tu, Xu Han, Chaoqun He, Ganqu Cui, Xiang Long, Zhi Zheng, Yewei Fang, Yuxiang Huang, Weilin Zhao, et al. Minicpm: Unveiling the potential of small language models with scalable training strategies. *arXiv preprint arXiv:2404.06395*, 2024.

NS Jayant and LR Rabiner. The application of dither to the quantization of speech signals. *Bell System Technical Journal*, 51(6):1293–1304, 1972.

Keller Jordan, Yuchen Jin, Vlado Boza, You Jiacheng, Franz Cecista, Laker Newhouse, and Jeremy Bernstein. Muon: An optimizer for hidden layers in neural networks, 2024. *URL https://kellerjordan. github. io/posts/muon*, 6.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Bingrui Li, Jianfei Chen, and Jun Zhu. Memory efficient optimizers with 4-bit states. *Advances in Neural Information Processing Systems*, 36:15136–15171, 2023.

Jingyang Li, Kuangyu Ding, Kim-Chuan Toh, and Pan Zhou. Memory-efficient 4-bit preconditioned stochastic optimization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22633–22643, 2025.

Ziniu Li, Congliang Chen, Tian Xu, Zeyu Qin, Jiancong Xiao, Zhi-Quan Luo, and Ruoyu Sun. Preserving diversity in supervised fine-tuning of large language models. *arXiv preprint arXiv:2408.16673*, 2024.

Stephanie Lin, Jacob Hilton, and Owain Evans. Truthfulqa: Measuring how models mimic human falsehoods. In *Proceedings of the 60th annual meeting of the association for computational linguistics (volume 1: long papers)*, pages 3214–3252, 2022.

Jingyuan Liu, Jianlin Su, Xingcheng Yao, Zhejun Jiang, Guokun Lai, Yulun Du, Yidao Qin, Weixin Xu, Enzhe Lu, Junjie Yan, et al. Muon is scalable for llm training. *arXiv preprint arXiv:2502.16982*, 2025.

Wei Liu, Weihao Zeng, Keqing He, Yong Jiang, and Junxian He. What makes good data for alignment? a comprehensive study of automatic data selection in instruction tuning. *arXiv preprint arXiv:2312.15685*, 2023.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.

Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. Can a suit of armor conduct electricity? a new dataset for open book question answering. *arXiv preprint arXiv:1809.02789*, 2018.

Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report, 2025. URL `https://arxiv.org/abs/2412.15115`.

Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. Zero: Memory optimizations toward training trillion parameter models. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–16. IEEE, 2020.

Samyam Rajbhandari, Olatunji Ruwase, Jeff Rasley, Shaden Smith, and Yuxiong He. Zero-infinity: Breaking the gpu memory wall for extreme scale deep learning. In *Proceedings of the international conference for high performance computing, networking, storage and analysis*, pages 1–14, 2021.

Jie Ren, Samyam Rajbhandari, Reza Yazdani Aminabadi, Olatunji Ruwase, Shuangyan Yang, Minjia Zhang, Dong Li, and Yuxiong He. {Zero-offload}: Democratizing {billion-scale} model training. In *2021 USENIX Annual Technical Conference (USENIX ATC 21)*, pages 551–564, 2021.

Lawrence Roberts. Picture coding using pseudo-random noise. *IRE Transactions on Information Theory*, 8(2):145–154, 1962.

Bita Darvish Rouhani, Ritchie Zhao, Ankit More, Mathew Hall, Alireza Khodamoradi, Summer Deng, Dhruv Choudhary, Marius Cornea, Eric Dellinger, Kristof Denolf, et al. Microscaling data formats for deep learning. *arXiv preprint arXiv:2310.10537*, 2023a.

Bita Darvish Rouhani et al. Ocp microscaling formats (mx) specification, version 1.0, 2023b.

Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM*, 64(9):99–106, 2021.

L. Schuchman. Dither signals and their effect on quantization noise. *IEEE Transactions on Communication Technology*, 12(4):162–165, 1964. doi: 10.1109/TCOM.1964.1088973.

Noam Shazeer and Mitchell Stern. Adafactor: Adaptive learning rates with sublinear memory cost. In *International Conference on Machine Learning*, pages 4596–4604. PMLR, 2018.

Douglas T Sherwood. Some theorems on quantization and an example using dither. In *Nineteeth Asilomar Conference on Circuits, Systems and Computers, 1985.*, pages 207–212. IEEE, 1985.

John Von Neumann and Herman Heine Goldstine. Numerical inverting of matrices of high order. 1947.

Nikhil Vyas, Depen Morwani, Rosie Zhao, Mujin Kwun, Itai Shapira, David Brandfonbrener, Lucas Janson, and Sham Kakade. Soap: Improving and stabilizing shampoo using adam. *arXiv preprint arXiv:2409.11321*, 2024.

Sike Wang, Pan Zhou, Jia Li, and Hua Huang. 4-bit shampoo for memory-efficient network training. *Advances in Neural Information Processing Systems*, 37:126997–127029, 2024.

Kaiyue Wen, Zhiyuan Li, Jason Wang, David Hall, Percy Liang, and Tengyu Ma. Understanding warmup-stable-decay learning rates: A river valley loss landscape perspective. *arXiv preprint arXiv:2410.05192*, 2024.

Huaijin Wu, Bingrui Li, Yebin Yang, Yi Tu, Zhanpeng Zhou, Jianfei Chen, and Junchi Yan. Achieving low-bit muon through subspace preservation and grid quantization. In *The Fourteenth International Conference on Learning Representations*.

Cong Xu, Wenbin Liang, Mo Yu, Anan Liu, Ke-Yue Zhang, Shunli Wang, Lizhuang Ma, Jianyong Wang, Jun Wang, and Wei Zhang. Pushing the limits of low-bit optimizers: A focus on ema dynamics. *arXiv preprint arXiv:2505.00347*, 2025.

Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T Kwok, Zhenguo Li, Adrian Weller, and Weiyang Liu. Metamath: Bootstrap your own mathematical questions for large language models. *arXiv preprint arXiv:2309.12284*, 2023.

Huizhuo Yuan, Yifeng Liu, Shuang Wu, Xun Zhou, and Quanquan Gu. Mars: Unleashing the power of variance reduction for training large models. *arXiv preprint arXiv:2411.10438*, 2024.

Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. Hellaswag: Can a machine really finish your sentence? *arXiv preprint arXiv:1905.07830*, 2019.

Yushun Zhang, Congliang Chen, Ziniu Li, Tian Ding, Chenwei Wu, Diederik P Kingma, Yinyu Ye, Zhi-Quan Luo, and Ruoyu Sun. Adam-mini: Use fewer learning rates to gain more. *arXiv preprint arXiv:2406.16793*, 2024.

A Pre-training Optimizer Hyperparameters

This appendix reports the complete optimizer hyperparameters used in all experiments. Unless otherwise stated, we keep **all training settings identical** across methods (model architecture, data recipe, training length, batch size, learning-rate schedule, and evaluation), and only vary the **storage precision** of optimizer states.

A.1 Global settings

All runs share the following optimizer-related settings:

- **Gradient clipping:** global norm clip at 1.0 (disabled if set to 0).
- **Epsilon:** $\epsilon = 1 \times 10^{-8}$.
- **LR scheduler:** WSD [Hu et al., 2024, Wen et al., 2024].
- **Warmup ratio:** 0.1 (warmup steps = $0.1\times$ total steps).
- **Min LR ratio:** 0.1 (minimum LR = $0.1\times$ peak LR).
- **Batch Size:** 1024.
- **Sequence length:** 1024.

A.2 Hyperparameters for 120M pretraining

Table 6 summarizes the optimizer hyperparameters used for pretraining the 120M model. Unless otherwise noted, all unspecified entries are not applicable to the corresponding optimizer.

Table 6: Optimizer hyperparameters for the 120M pretraining runs.

	AdamW	Muon	NAdamW	SOAP	MARS
Peak LR	1×10^{-3}	1×10^{-3}	1×10^{-3}	2×10^{-3}	1×10^{-3}
(β_1, β_2)	(0.9, 0.95)	—	(0.9, 0.95)	(0.95, 0.95)	(0.9, 0.99)
Weight decay	0.1	0.1	0.1	0.01	0
Weight decay (1D params)	—	—	—	—	0.01
Muon momentum β	—	0.95	—	—	—
NS steps	—	5	—	—	—
Momentum decay	—	—	0.004	—	—
Shampoo β	—	—	—	0.95	—

A.3 Settings for larger models

Following Liu et al. [2025], we use the same learning rate and weight decay for Muon and AdamW. For the 1.1B and 3.4B models, we use:

- **Peak LR for 1.1B model:** 6×10^{-4}.
- **Peak LR for 3.4B model:** 6×10^{-4}.
- **AdamW betas:** $(\beta_1, \beta_2) = (0.9, 0.95)$.
- **Muon momentum:** $\beta = 0.95$.
- **NS steps:** 5.

B Fine-tuning hyperparameters

Following Yu et al. [2023], Liu et al. [2023], Li et al. [2024], we use the learning rate $2 * 10^{-5}$ with a cosine learning rate decay schedule for 3 epochs. The warmup ratio is 0.03. The batch size is 128.

B.1 Packed low-bit storage.

All reported BReD low-bit runs use physically packed state tensors rather than high-precision emulation. We implement custom Triton kernels that store encoded optimizer states in packed integer tensors, decode them on the fly during the optimizer step, perform the EMA update in higher precision, and requantize and pack the updated states back to memory.

C Experiments results

C.1 Task-wise pretraining results

Table 7: Zero-shot LM evaluation accuracies for the 120M model. Subscripts denote standard deviations over three runs. Parentheses report the difference relative to the full-precision optimizer-state baseline of the same optimizer.

OPTIMIZER	HS	OBQA	WG	ARC_E	ARC_C	BOOLQ	PIQA	AVG.↑
ADAMW	$29.41_{0.4}$	$27.60_{0.7}$	$51.20_{0.4}$	$36.36_{0.1}$	$21.79_{1.1}$	$56.12_{1.9}$	$60.58_{0.6}$	40.4(+0.0)
ADAMW-4BIT+**BRED**	$29.12_{0.2}$	$27.07_{0.6}$	$52.04_{0.7}$	$35.84_{1.0}$	$21.82_{0.7}$	$59.35_{1.0}$	$59.58_{0.5}$	40.7(+0.3)
ADAMW-3BIT+**BRED**	$28.68_{0.2}$	$27.27_{0.5}$	$51.93_{0.2}$	$36.10_{0.9}$	$22.33_{0.6}$	$59.63_{1.3}$	$58.62_{0.6}$	40.6(+0.2)
MUON	$29.85_{0.1}$	$26.73_{0.4}$	$51.30_{0.5}$	$36.60_{0.3}$	$22.50_{0.3}$	$59.16_{2.6}$	$60.19_{0.3}$	40.9(+0.0)
MUON-4BIT+**BRED**	$29.44_{0.2}$	$27.13_{1.5}$	$52.57_{1.1}$	$36.67_{1.3}$	$22.55_{0.3}$	$60.15_{1.7}$	$60.08_{0.9}$	41.2(+0.3)
MUON-3BIT+**BRED**	$29.02_{0.1}$	$27.47_{0.6}$	$51.49_{0.8}$	$36.74_{0.4}$	$22.24_{0.5}$	$59.67_{0.6}$	$59.99_{0.4}$	40.9(+0.0)
NADAMW	$28.91_{0.5}$	$26.47_{0.4}$	$50.85_{0.9}$	$36.08_{0.4}$	$21.96_{0.9}$	$56.71_{2.3}$	$59.39_{0.4}$	40.1(+0.0)
NADAMW-4BIT+**BRED**	$28.98_{0.3}$	$26.93_{1.2}$	$49.59_{1.1}$	$36.00_{0.6}$	$22.04_{0.8}$	$60.44_{2.4}$	$60.01_{0.1}$	40.6(+0.5)
NADAMW-3BIT+**BRED**	$28.91_{0.2}$	$25.93_{1.4}$	$51.12_{1.6}$	$35.64_{0.7}$	$21.96_{0.5}$	$60.20_{2.4}$	$59.67_{0.6}$	40.5(+0.4)
SOAP	$29.69_{0.1}$	$24.93_{0.3}$	$50.12_{1.2}$	$36.60_{0.7}$	$22.47_{0.5}$	$54.08_{8.8}$	$59.43_{1.1}$	39.6(+0.0)
SOAP-4BIT+**BRED**	$28.97_{0.1}$	$27.20_{0.2}$	$50.54_{0.9}$	$36.02_{0.7}$	$21.73_{0.4}$	$58.03_{5.1}$	$59.36_{0.6}$	40.3(+0.7)
SOAP-3BIT+**BRED**	$28.86_{0.1}$	$26.47_{0.8}$	$50.70_{0.5}$	$35.62_{0.8}$	$22.47_{0.8}$	$61.04_{1.7}$	$59.01_{1.2}$	40.6(+1.0)
MARS	$29.03_{0.1}$	$27.20_{0.5}$	$51.65_{1.1}$	$36.00_{1.0}$	$21.56_{0.5}$	$60.89_{2.3}$	$59.25_{1.0}$	40.8(+0.0)
MARS-4BIT+**BRED**	$28.88_{0.1}$	$25.20_{0.7}$	$50.54_{0.3}$	$35.69_{0.3}$	$22.58_{0.7}$	$59.67_{1.0}$	$59.99_{0.3}$	40.4(-0.4)
MARS-3BIT+**BRED**	$28.59_{0.3}$	$27.47_{1.2}$	$51.93_{0.5}$	$35.91_{0.6}$	$22.55_{1.2}$	$61.22_{1.0}$	$59.49_{1.0}$	41.0(+0.2)

Table 8: Zero-shot LM evaluation accuracies and validation PPL on 1.1B and 3.4B models. Subscripts denote standard deviations over three runs. Parentheses report the difference relative to the full-precision optimizer-state baseline of the same optimizer.

#PARAM	OPTIMIZER	HS	OBQA	WG	ARC_E	ARC_C	BOOLQ	PIQA	AVG. ↑	PPL ↓
1.1B	ADAMW	$43.45_{0.1}$	$30.80_{0.4}$	$52.62_{0.5}$	$45.38_{0.4}$	$25.66_{0.7}$	$59.66_{0.1}$	$68.05_{0.7}$	46.5(+0.0)	10.8(+0.0)
	ADAMW-4BIT+**BRED**	$43.05_{0.3}$	$30.67_{1.3}$	$52.12_{0.2}$	$45.50_{2.0}$	$25.51_{0.5}$	$59.63_{1.8}$	$68.03_{0.2}$	46.4(-0.1)	10.9(+0.1)
	ADAMW-3BIT+**BRED**	$42.21_{2.2}$	$30.60_{0.5}$	$52.78_{1.6}$	$43.74_{2.4}$	$24.52_{1.1}$	$59.14_{1.6}$	$67.45_{1.6}$	45.8(-0.7)	10.9(+0.1)
	MUON	$46.01_{0.5}$	$31.07_{0.6}$	$52.75_{1.1}$	$48.23_{1.1}$	$26.37_{0.5}$	$59.57_{2.7}$	$69.10_{0.5}$	47.6(+0.0)	10.3(+0.0)
	MUON-4BIT+**BRED**	$44.46_{0.4}$	$30.87_{1.6}$	$53.57_{0.3}$	$47.28_{1.1}$	$25.57_{0.6}$	$58.96_{4.2}$	$68.93_{0.3}$	47.1(-0.5)	10.6(+0.3)
	MUON-3BIT+**BRED**	$43.08_{0.3}$	$31.73_{0.7}$	$51.73_{0.7}$	$46.09_{1.5}$	$25.48_{0.1}$	$60.58_{0.5}$	$68.19_{0.3}$	46.7(-0.9)	10.9(+0.6)
3.4B	ADAMW	$48.01_{0.3}$	$31.33_{0.3}$	$53.99_{1.8}$	$48.45_{0.4}$	$26.93_{0.4}$	$56.47_{4.7}$	$69.98_{0.6}$	47.9(+0.0)	9.7(+0.0)
	ADAMW-4BIT+**BRED**	$47.08_{0.7}$	$31.47_{0.7}$	$52.80_{1.8}$	$47.53_{1.2}$	$26.08_{0.2}$	$59.25_{1.0}$	$69.24_{0.2}$	47.6(-0.3)	9.8(+0.1)
	MUON	$48.37_{0.2}$	$31.53_{0.3}$	$54.91_{1.3}$	$49.17_{0.4}$	$26.48_{0.1}$	$57.50_{2.8}$	$69.68_{0.5}$	48.2(+0.0)	9.6(+0.0)
	MUON-4BIT+**BRED**	$46.31_{0.0}$	$30.47_{0.3}$	$54.96_{0.9}$	$47.11_{1.2}$	$26.23_{0.4}$	$59.77_{1.0}$	$68.74_{0.4}$	47.7(-0.5)	9.8(+0.2)

C.2 3-bit results

Table 9: Aggressive 3-bit BReD compression. For pretraining, Acc. denotes average zero-shot accuracy; for supervised fine-tuning, Avg. denotes average downstream score. Parentheses denote changes relative to the corresponding full-precision optimizer-state baseline.

Stage	Architecture	Size	Optimizer	Acc. ↑		PPL ↓	
				FP32-state	3-bit BReD	FP32-state	3-bit BReD
Pretrain	LLaMA-style	120M	AdamW	40.4	40.6 (+0.2)	18.3	19.0 (+0.7)
Pretrain	LLaMA-style	120M	Muon	40.9	40.9 (+0.0)	17.8	18.5 (+0.7)
Pretrain	LLaMA-style	120M	NAdamW	40.1	40.5 (+0.4)	18.9	19.0 (+0.1)
Pretrain	LLaMA-style	120M	SOAP	39.6	40.6 (+1.0)	17.6	18.8 (+1.2)
Pretrain	LLaMA-style	120M	MARS	40.8	41.0 (+0.2)	18.7	19.4 (+0.7)
Pretrain	LLaMA-style	1.1B	AdamW	46.5	45.8 (-0.7)	10.8	10.9 (+0.1)
Pretrain	LLaMA-style	1.1B	Muon	47.6	46.7 (-0.9)	10.3	10.9 (+0.6)
SFT	Qwen2.5	7B	AdamW	54.2	55.0 (+0.8)	4.6	4.7 (+0.1)
SFT	Qwen2.5	7B	Muon	55.2	54.4 (-0.8)	4.6	4.5 (-0.1)

C.3 Higher EMA decay

We additionally evaluate 120M AdamW pretraining with $\beta_2 = 0.999$, changing only β_2 from the default training recipe. FP32-state AdamW and 4-bit BReD use the same β_2 within each comparison. Table 10 shows that BReD remains stable at $\beta_2 = 0.999$, with a PPL increase of 0.7. This is a single paired run, rather than a systematic high-β sweep; we do not infer a cross-β_2 trend from these results.

Table 10: Validation PPL for 120M AdamW at two EMA decay factors. The $\beta_2 = 0.95$ results are means over three runs; the $\beta_2 = 0.999$ results are from one paired run. ΔPPL is BReD minus FP32-state at the same β_2.

β_2	FP32-state	4-bit BReD	ΔPPL
0.95	18.3	18.7	+0.4
0.999	19.4	20.1	+0.7

C.4 Additional quantization format

We evaluate MXINT4 in the same 120M AdamW setup, changing only the quantization format used for optimizer-state storage. As shown in Table 11, 4-bit BReD with MXINT4 remains convergent, reaching a validation PPL of 18.9, compared with 18.7 for MXFP4.

Table 11: Validation PPL for 120M AdamW with different state-storage formats. All values are means over three runs.

State storage	Format	Validation PPL
FP32-state	FP32	18.3
4-bit BReD	MXFP4	18.7
4-bit BReD	MXINT4	18.9

C.5 SOLO checks

We evaluate the author-released SOLO implementation in our 120M from-scratch pretraining setup. For each precision setting, we retain SOLO's official quantization configuration and vary the AdamW betas as listed in Table 12. The model, data, learning-rate schedule, weight decay, batch size, and seed are fixed across runs. All SOLO runs use quantile 0.1 and block size 128.

For reference, 3/3-bit BReD reaches validation PPL 19.0 in the same pretraining setup with the original AdamW hyperparameters. These limited checks do not imply that SOLO fails in its reported

Table 12: SOLO checks in our 120M pretraining setup. Precision denotes first-/second-moment bitwidths. Each entry reports one run.

Beta setting	Precision	(β_1, β_2)	Result
Original AdamW	4/2-bit	$(0.9, 0.95)$	Diverged
SOLO from-scratch	4/2-bit	$(0.3, 0.95)$	Diverged
SOLO from-scratch	2/2-bit	$(0.1, 0.95)$	Diverged
Author example	4/2-bit	$(0.8, 0.999)$	Diverged

settings, nor do they identify which differences in quantization or optimizer dynamics caused the observed failures.

C.6 Small-scale Shampoo check

We additionally evaluate 120M Shampoo with 4-bit preconditioner states, keeping the model, data, optimizer update, and hyperparameters unchanged across configurations. As shown in Table 13, BReD completes training with validation PPL 18.5, compared with 18.3 for FP32-state Shampoo, whereas NR and SR diverge.

Table 13: Validation PPL for the 120M Shampoo check. Low-bit variants quantize the preconditioner states to 4 bits. Each configuration reports one run.

Method	Validation PPL
FP32-state Shampoo	18.3
Shampoo + NR	Diverged
Shampoo + SR	Diverged
Shampoo + BReD	18.5

This single-run check does not support a statistical comparison or replace large-scale validation. It also does not establish superiority over dedicated methods such as 4-bit Shampoo [Wang et al., 2024] or 4-bit Preconditioned Optimization [Li et al., 2025], which are not evaluated here.

C.7 memory reduction

Table 14: Optimizer-state memory reduction with 4-bit BReD. Reduction is relative to the corresponding full-precision optimizer-state baseline.

Scale	Optimizer	FP32-state	4-bit BReD	Reduction
120M	AdamW	0.90GB	0.19GB	78.9%
120M	Muon	0.64GB	0.41GB	35.9%
120M	NAdamW	0.90GB	0.19GB	78.9%
120M	SOAP	1.66GB	0.55GB	66.9%
120M	MARS	1.36GB	0.64GB	52.9%
1.1B	AdamW	8.20GB	1.18GB	85.6%
1.1B	Muon	4.59GB	1.47GB	68.0%
3.4B	AdamW	25.53GB	3.58GB	86.0%
3.4B	Muon	13.53GB	3.17GB	76.6%

D Implementation note.

PRNG We use stateless PRNG (e.g., Philox/Threefry-style) so that each u can be generated independently in parallel from its key. This avoids maintaining PRNG states across threads and guarantees bitwise reproducibility under re-ordering.

Block-wise SD. For SD, we similarly reuse $u^{(k)}$ to quantize all $\alpha(x_i)$ within the block and replay the same $u^{(k)}$ during dequantization. This keeps the coordinate-wise error unbiased while greatly reducing PRNG calls.

E Proof

Lemma E.1 (mean induces persistent drift under EMA recursion)**.** *Assume the injected errors satisfy* $\mathbb{E}[\epsilon_t] = b$*. For the recursion* $e_{t+1} = \beta e_t + \epsilon_{t+1}$ *with* $0 < \beta < 1$*, we have*

$$\mathbb{E}[e_t] = \beta^t \mathbb{E}[e_0] + \frac{1-\beta^t}{1-\beta}\, b, \tag{18}$$

and therefore the mean error converges to

$$\lim_{t\to\infty} \mathbb{E}[e_t] = \frac{b}{1-\beta}. \tag{19}$$

In particular, even a small nonzero bias b *yields a persistent drift whose magnitude scales as* $\mathcal{O}\Big(\frac{\|b\|}{1-\beta}\Big)$ *when* $\beta \to 1$.

Proof. Unrolling $e_t = \beta^t e_0 + \sum_{k=1}^{t} \beta^{t-k}\epsilon_k$ and taking expectation gives the result.

Theorem E.2 (Steady-state second moment under quantized EMA)**.** *Assume* $\epsilon_t = b + \xi_t$*, where* $\{\xi_t\}_{t\geq 1}$ *are independent with* $\mathbb{E}[\xi_t] = 0$ *and* $\mathbb{E}\|\xi_t\|_F^2 = \sigma^2$*. Assume also that* e_0 *has finite second moment and is independent of* $\{\xi_t\}_{t\geq 1}$*. For* $0 < \beta < 1$*, the state error satisfies*

$$\lim_{t\to\infty} \mathbb{E}\|e_t\|_F^2 = \underbrace{\frac{\|b\|_F^2}{(1-\beta)^2}}_{\textit{mean (bias) term}} + \underbrace{\frac{\sigma^2}{1-\beta^2}}_{\textit{variance (noise) term}}. \tag{20}$$

Proof. Unrolling Eq. (2) gives $e_t = \beta^t e_0 + \sum_{k=1}^{t} \beta^{t-k}\epsilon_k$. Taking expectations yields $\lim_{t\to\infty} \mathbb{E}[e_t] = b/(1-\beta)$. Writing $\epsilon_k = b + \xi_k$ with $\mathbb{E}[\xi_k] = 0$ and using independence cancels cross terms, leading to Eq. (4).

Proof (SD unbiasedness and variance). Let $v := \{\alpha + u\}$ denote the fractional part of $\alpha + u$. Since $u \sim \mathrm{Unif}(0,1)$, we also have $v \sim \mathrm{Unif}(0,1)$. Moreover,

$$\hat{\alpha} - \alpha = c_{\mathrm{sd}}(\alpha; u) - u + \tfrac{1}{2} - \alpha = \tfrac{1}{2} - \{\alpha + u\} = \tfrac{1}{2} - v.$$

Hence $\hat{\alpha} - \alpha \sim \mathrm{Unif}(-1/2,\, 1/2)$, which implies

$$\mathbb{E}[\hat{\alpha} - \alpha \mid x] = 0, \qquad \mathrm{Var}(\hat{\alpha} - \alpha \mid x) = 1/12.$$

Multiplying by $\Delta(x)$ yields

$$\mathbb{E}[\epsilon_{\mathrm{elem}} \mid x] = 0, \qquad \mathrm{Var}(\epsilon_{\mathrm{elem}} \mid x) = \Delta(x)^2/12.$$

E.1 Rounding-error tails of SR and SD

Under the local non-clipped interval model, let $r = (\hat{x} - x)/\Delta(x)$. For a fixed α, SR gives errors $1-\alpha$ and $-\alpha$ with probabilities α and $1-\alpha$, respectively, whereas SD gives $r_{\mathrm{SD}} \sim \mathrm{Unif}(-1/2, 1/2)$. Under the illustrative assumption $\alpha \sim \mathrm{Unif}(0,1)$,

$$\Pr(|r_{\mathrm{SR}}| > \tau) = (1-\tau)^2 \quad (0 \leq \tau \leq 1), \qquad \Pr(|r_{\mathrm{SD}}| > \tau) = 0 \quad (\tau > 1/2).$$

Thus, unbiased SR can produce larger rounding errors than SD. This offers a possible explanation for the observed SR failures, but does not establish a training-divergence condition under Theorem 3.2.

F Additional In-Training Diagnostics

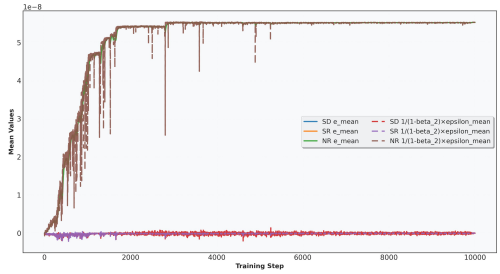

(a) Layer 1.attn.attn, show the $\mathbb{E}[e_t]$ and $\mathbb{E}[\epsilon_t]/(1-\beta_2)$.

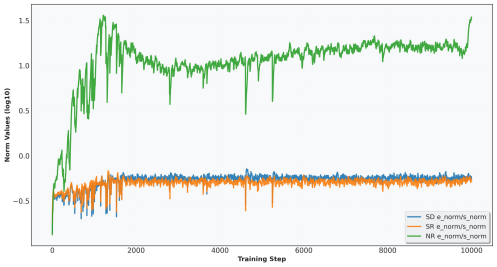

(b) Layer 1.attn.attn, error accumulates quickly.

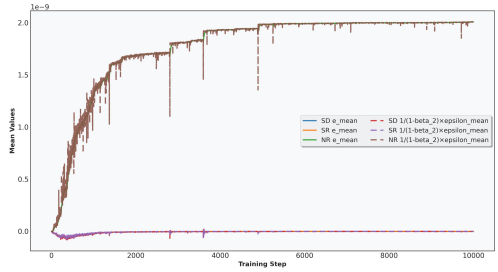

(c) Layer 5.mlp.swiglu.w1, show the $\mathbb{E}[e_t]$ and $\mathbb{E}[\epsilon_t]/(1-\beta_2)$.

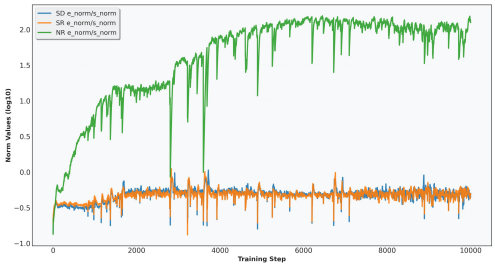

(d) Layer 5.mlp.swiglu.w1, error accumulate quickly.

Figure 3: **Additional quantized-EMA diagnostics for two layers:** Panels (a) and (c) compare the measured mean accumulated error $\mathbb{E}[e_t]$ with the predicted amplified injected error $\mathbb{E}[\epsilon_t]/(1-\beta_2)$. Panels (b) and (d) show the relative error norm over training. Across layers, nearest rounding accumulates large error, whereas stochastic rounding and subtractive dithering maintain bounded relative errors in this diagnostic.