arXiv:2611.24644v1 [cs.LG] 18 Nov 2026

Equal Steps, Unequal Curvature: Where Muon's Update Goes Wrong

Biswa Sengupta **Jinhua Wang** **Léo Maxime Brunswic**
LLM Suite Team, JPMorgan Chase & Co.
{biswa.sengupta, jinhua.wang, leo.brunswic}@jpmorgan.com

Disclaimer: This paper was prepared for informational purposes by the LLM Suite group of JP Morgan Chase and its affiliates ('JPMC') and is not a product of the Research Department of JP Morgan. JP Morgan makes no representation, warranty or undertaking whatsoever and disclaims all liability for the completeness, accuracy or reliability of the information contained herein. This document is not intended as investment research or investment advice, or a recommendation, offer or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction, and shall not constitute a solicitation under any jurisdiction or to any person, if such solicitation under such jurisdiction or to such person would be unlawful.

Abstract

Muon updates each hidden weight matrix along the polar factor of its momentum, the steepest-descent direction under the spectral norm. That choice fixes two allocations without consulting the loss: every singular direction of the momentum receives the same step, and in a non-square matrix each neuron's step is set by its leverage score. Whether these allocations are good is a question about curvature, which steepest-descent arguments do not answer. We measure it. Using held-out gradients and Hessian–vector products, we decompose Muon's step by singular-value band and by neuron group and compute each group's second-order optimal multiplier. In a 124M-parameter decoder trained on FineWeb-Edu, the multipliers rise monotonically from the top singular directions to the tail: Muon over-steps the high-curvature top of its spectrum. We prove in Lean when such a diagnostic licenses an intervention at Muon's own update norm. Applied to the measured curvature, the criterion predicts, and direct measurements of the held-out loss confirm, that moving step length off the top improves the one-step loss severalfold mid-training and at the start of the cooldown, but not at its end. We then test interventions under per-step norm matching, tuned baselines and paired seeds. Orthogonalizing Q, K and V separately rather than as one fused matrix lowers the loss by 0.0250 nats and accounts for all of the gain of block-wise variants. A matmul-only, rotation-equivariant spectral tilt $\sigma^{-1/4}$ applied in the cooldown lowers the loss of this stronger baseline by a further 0.0120 nats, additively with NorMuon (0.0158 nats together, three seeds; the baseline needs about 4.3% more tokens to match), and by 0.0188 nats at 355M parameters (three seeds). Energy-matched Schatten-1 weight decay adds a further 0.0052 nats. Geometrically principled changes that do not act on the measured miscalibration, namely exact polar factors, steepest-descent input whitening and top-heavy decay, do not help.

1 Introduction

Muon (Jordan et al., 2024b) has become the reference optimizer for the hidden matrices of transformers at small and medium scale (Liu et al., 2025; Essential AI et al., 2025; Jordan et al., 2024a). Its theory is clean. For a weight matrix with momentum $N = U\Sigma V^\top$, the update direction $\mathrm{polar}(N) = UV^\top$ is the steepest-descent direction under the spectral norm (Bernstein and Newhouse, 2024; Large et al., 2024; Pethick et al., 2025), and Muon with decoupled weight decay solves a spectral-norm-constrained problem (Chen et al., 2026).

That theory picks the *norm* in which the step is steepest. It is silent on whether the step's *distribution* is right, and $UV^\top$ fixes two distributions without asking. (i) *Across singular directions*, every direction $u_i v_i^\top$ receives a step of the same length, whatever its gradient signal σ_i and whatever the loss curvature along it. (ii) *Across neurons*, a tall polar factor has row norms equal to the square roots of the leverage scores of N's column space. Each output neuron therefore moves by an amount set by the momentum's geometry and by how matrices happen to be fused in the implementation.

Recent variants change both distributions. NorMuon, Muon+ and Aurora rebalance neurons (Li et al., 2026; Zhang et al., 2026a; Dewulf et al., 2026). DynMuon, Freon/Kaon and Spectrally Targeted Muon reshape the spectrum (Wu et al., 2026a; Shumaylov et al., 2026; Goyal and Ramkumar, 2026). MuonBP and HiMuon orthogonalize blocks (Khaled et al., 2026; Tang et al., 2026). Some work also measures why. SAMuon probes held-out loss-optimal step sizes along the momentum's singular directions and builds an optimizer from that profile (Wu et al., 2026b).

Preprint.

© 2026 JP Morgan Chase & Co. All rights reserved

BulkBoost proves when a norm-matched two-band reallocation helps to first order (Chaitanya et al., 2026). Curvature analyses explain Muon's advantage through its actual update (Wang et al., 2026a; Shumaylov et al., 2026; Shulgin et al., 2026). These measurements treat directions or bands one at a time or compare Muon with Adam, and gains against loosely tuned baselines often shrink against well-tuned ones (Wen et al., 2026; Zhao et al., 2025).

This paper measures the joint picture. We treat the flat spectrum and the leverage-weighted neuron allocation as hypotheses and test them against the loss's local geometry. We split Muon's next step into groups of singular directions or of neurons. On held-out data we measure each group's gradient alignment g_b and the full curvature matrix H between groups, including the cross-group terms that isolated probes omit. The ratio g_b/H_{bb} is the step multiplier that would be optimal for group b on its own to second order, and Muon corresponds to multiplier 1 everywhere. A one-step greedy analysis is known to misjudge the overall step size (Wu et al., 2018). We therefore use an exact identity, proved in Lean, that splits Muon's second-order gap into a *scale* part (the overall step size) and a *shape* part (the relative weighting across groups). Our experiments compare steps of equal norm, so we also prove when a same-norm reallocation can beat Muon's step, and we evaluate that criterion on the measured curvature.

What the anatomy shows (Section 4). The optimal multiplier rises monotonically from the top singular directions to the tail. On ranks 513–768 it is 38–190 times its value on ranks 1–8. Muon thus over-steps the high-curvature top of its spectrum relative to the rest. At Muon's own update norm, which is the comparison our experiments make, reallocating step length off the top would change the one-step loss by -0.00589 mid-training and -0.00488 at the start of the cooldown. Muon's own step changes it by -0.00079 and $+0.00026$. By the end of the cooldown the possible gain is only -0.00009. Measuring the actual held-out loss confirms this: at matched norm, halving only the top eight singular directions improves Muon's one-step loss by -4.4 and -3.2 ($\times 10^{-3}$), and not at the end. Per-neuron step sizes vary most in the fused 2304×768 QKV matrix (coefficient of variation 0.43, against 0.18 in the MLP).

What acting on it buys (Section 6). We test interventions under strict controls. Each direction-changing arm is norm-matched to Muon at every step, baselines use tuned (interior-optimum) learning rates, and seeds are paired. Three findings stand out.

1. *The neuron allocation that matters is per projection.* Orthogonalizing Q, K and V separately, a known practice (Jordan et al., 2024a), changes the final loss by -0.0250 nats [-0.0290, -0.0210] at 100.7M tokens. It accounts for *all* of the gain of square-block variants (residual $+0.0011$ [-0.0029, $+0.0051$]). Finer rebalancing, per head or per row, does worse. The fused-QKV baseline would need about 6.8% more tokens to match, which confounds comparisons: NorMuon's gain shrinks from 0.0150 to 0.0038 against the split baseline.
2. *Tilting the spectrum away from its top helps.* A cooldown tilt toward small singular values lowers the stronger baseline's loss by a further 0.0120 nats (-0.0120 [-0.0137, -0.0102]), better in every seed. We implement it with a Chebyshev filter evaluated by Clenshaw's recurrence, using matrix products only, and prove that it keeps Muon's rotational equivariance.
3. *The two corrections add.* NorMuon + tilt gives -0.0158 [-0.0178, -0.0138], against -0.0038 and -0.0120 separately. At 355M parameters the combination gains 0.0188 nats [-0.0202, -0.0175] (three seeds). Energy-matched Schatten-1 weight decay adds a further -0.0052 [-0.0066, -0.0038] at 124M.

The negatives are equally informative. Input whitening is provably $\mathrm{GL}(n)$-equivariant and conditioning-free on least squares, yet it hurts (up to $+0.06$ nats). More accurate polar factors do nothing, and top-heavy decay hurts; none acts on the miscalibration we measured. The diagnostic and the first screens ran concurrently, so the anatomy's implications are consistency checks, not pre-registered predictions; the factorial ablations and the 355M runs came later.

Contributions. (1) A held-out, curvature-resolved diagnostic for a matrix optimizer's step, with an exact scale/shape decomposition and a proved criterion for when a same-norm reallocation can improve the step. **(2)** An anatomy of Muon: monotone over-stepping of the top singular directions, which a same-norm reallocation can exploit until late in the cooldown, and a neuron imbalance concentrated in fused matrices. **(3)** A controlled attribution, including a 2×2 factorial with three seeds and an A/A noise check. **(4)** Machine-checked statements of the geometry used (Section A), and an account of what they did and did not predict (Section 7). Every optimizer ingredient we evaluate has a precedent, and so does probing Muon's spectrum against held-out loss (Section 8). What is new is the joint-curvature diagnostic with its fixed-norm criterion, the neuron-level and QKV attribution, and the controls that test them.

2 Two allocations hidden in Muon's step

For a hidden matrix $W \in \mathbb{R}^{m\times n}$ with gradient G_t, the Muon variant we study forms the sum-form Nesterov input and steps along a quintic Newton–Schulz map:

$$\begin{aligned} B_t &= \beta B_{t-1} + G_t, \qquad N_t = G_t + \beta B_t, \\ W_{t+1} &= (1-\eta_t\lambda)W_t - \eta_t\, s\, \mathrm{NS}_5(N_t), \end{aligned} \tag{1}$$

with shape factor $s = \sqrt{\max(1, m/n)}$. NS_5 normalizes N by its Frobenius norm and applies five iterations of

© 2026 JP Morgan Chase & Co. All rights reserved

$X \mapsto aX + bXX^\top X + c(XX^\top)^2 X$ with $(a, b, c) = (3.4445, -4.775, 2.0315)$ (Jordan et al., 2024b). On singular values it acts as a scalar map ϕ_5 that is close to 1 on most of the spectrum but well below 1 on near-null directions. Hence $\mathrm{NS}_5(N) \approx \mathrm{polar}(N) = UV^\top = \arg\max_{\|D\|_{\mathrm{op}} \le 1} \langle N, D \rangle$.

Allocation across singular directions. Write any spectral filter as $D_f = Uf(\Sigma)V^\top$. SGD is $f(\sigma) = \sigma$, Muon is (approximately) $f \equiv 1$, and power-law tilts are $f(\sigma) = (\sigma/\sigma_{\mathrm{rms}})^p$ (Wu et al., 2026a). Muon's choice is optimal for the worst case over a spectral-norm ball. It is not obviously optimal for a loss whose curvature differs across singular directions.

Allocation across neurons. Every theorem and proposition in this paper is proved in Lean 4; Section A states each in full. Row i of W is the incoming weight vector of output neuron i. Muon moves it by $\|(UV^\top)_{i,:}\|$.

Proposition 1 (Polar row norms are leverage scores; Theorem A.10)**.** *Let $O \in \mathbb{R}^{m \times n}$ with $O^\top O = I_n$ ($m \ge n$). Then $\sum_i \|O_{i,:}\|^2 = n$, but the individual $\|O_{i,:}\|^2$, the leverage scores of* $\mathrm{range}(O)$ *(Drineas et al., 2012), can be arbitrarily uneven: one row can have norm 1 and another 0. If $m = n$ and $OO^\top = I$, every row has unit norm.*

So Muon's per-neuron step depends on the *shape* of the matrix it is applied to. Storing Q, K and V as one fused $3d \times d$ matrix (the GPT-2 layout (Radford et al., 2019)) makes each projection's step depend on the momenta of the other two. The block-wise alternative is the exact steepest step of a natural norm:

Proposition 2 (Block-wise polar steps; Section A.2)**.** *Split $W = [W_1; \ldots; W_k]$ into k blocks and let $\|D\|_{\mathrm{blk}} = \max_b \|D_b\|_{\mathrm{op}}$. Then $[\mathrm{polar}(N_1); \ldots; \mathrm{polar}(N_k)]$ maximizes $\langle N, D \rangle$ over $\|D\|_{\mathrm{blk}} \le 1$, and the joint polar factor can be strictly worse (for $N_1 = N_2 = (1, 0)$ it attains $\sqrt{2}$ against 2). For a tall O split into k exactly orthogonal $n \times n$ row blocks, Muon's shape factor $s = \sqrt{k}$ gives the joint step the same Frobenius energy as the block-wise step.*

Both propositions concern exact polar factors. NS_5 is only approximately orthogonal, so in practice block-wise steps shrink the row spread but do not remove it. The block-wise steepest step for this max-norm is MuonBP's Block-Muon (Khaled et al., 2026). The leverage-score reading of tall polar factors is also given by Dewulf et al. (2026). We use both diagnostically, to locate where Muon's per-neuron steps are uneven, not to show that the unevenness hurts.

3 A curvature-resolved anatomy of the step

Grouped multipliers. We partition Muon's next step into groups $b = 1, \ldots, B$ in one of two ways. *Bands:* singular-value ranks of the momentum, with $D_{\ell,b} = \eta\, s_\ell\, U_{I_b} V_{I_b}^\top$ for every matrix ℓ. This is the exact-polar idealization of Muon's step; NS_5 itself gives near-null directions less than unit length. *Neurons:* sextiles of neurons ranked by the norm of their row (or column) in Muon's actual NS_5 step. Now scale group b by c_b, i.e. $W_\ell \leftarrow W_\ell - \sum_b c_b D_{\ell,b}$. To second order in c, the held-out loss changes by

$$q(c) = -g^\top c + \tfrac{1}{2} c^\top H c, \qquad g_b = \langle \nabla \hat{L}, D_b \rangle, \qquad H_{bb'} = \langle D_b, \nabla^2 \hat{L}\, D_{b'} \rangle, \tag{2}$$

where $\hat{L}$ is the loss on held-out tokens. Thus g and H refer to the held-out loss, not the training minibatch, though they still carry finite-sample noise. Muon is $c = \mathbf{1}$. The stationary point of q is $c^\star = H^{-1} g$. It is the minimizer of q only when $H \succeq 0$, and the unique one when $H \succ 0$ (Section A.3). Every curvature matrix we estimate is positive definite, with smallest eigenvalue at least 4.6×10^{-8}. The per-group multiplier $c_b = g_b / H_{bb}$ minimizes q for group b on its own. Muon's flat filter is second-order optimal up to a common scale exactly when $c^\star \propto \mathbf{1}$. All c_b scale as $1/\eta$, so only their ratios are comparable across learning rates.

Scale versus shape. A greedy one-step model is a poor guide to the overall step size, because it ignores the step's effect on future iterates (Wu et al., 2018; Baydin et al., 2018). It can still say which groups get too much or too little of a given step. The following identity separates the two.

Theorem 3 (Exact scale/shape split; Section A.3)**.** *Let H be symmetric positive definite, $c^\star = H^{-1} g$ the minimizer of q, and $\gamma^\star = \mathbf{1}^\top g / \mathbf{1}^\top H \mathbf{1}$ the minimizer of $q(\gamma \mathbf{1})$ over γ, i.e. the best uniform rescaling of Muon's step. Then*

$$q(\mathbf{1}) - q(c^\star) = \underbrace{\tfrac{1}{2}(1 - \gamma^\star)^2\, \mathbf{1}^\top H \mathbf{1}}_{\textit{step size}} + \underbrace{\tfrac{1}{2}(\gamma^\star \mathbf{1} - c^\star)^\top H (\gamma^\star \mathbf{1} - c^\star)}_{\textit{shape}},$$

and the shape term is strictly positive unless $c^\star$ is flat.

The identity itself needs only symmetry and $Hc^\star = g$. Positive definiteness is what makes $c^\star$ and $\gamma^\star$ optima, and both terms losses relative to them. The step-size term carries the known short-horizon bias. Removing it does not make the shape immune: in an explicit example, a one-step shape preference reverses after one more gradient step (Theorem A.28). We therefore treat the shape as a one-step diagnostic whose implications must be tested in training, not as a prescription.

Fixed norm. Theorem 3 compares Muon with $c^\star$, whose norm generally differs from Muon's. Our experiments instead compare steps of *equal* Frobenius norm. Band

© 2026 JP Morgan Chase & Co. All rights reserved

directions are mutually orthogonal, with squared norms proportional to the band sizes $|I_b|$, so equal norm means $c^\top Nc = \mathbf{1}^\top N\mathbf{1}$ with $N = \mathrm{diag}(|I_b|)$. A positive shape term does not by itself imply that such a step beats Muon:

Proposition 4 (Same-norm reallocation; Section A.3)**.** *Let $N \succ 0$, $r = H\mathbf{1} - g$ the model gradient at Muon's step, and p the component of r orthogonal to $N\mathbf{1}$. (i) If $p \neq 0$, some c with $c^\top Nc = \mathbf{1}^\top N\mathbf{1}$ has $q(c) < q(\mathbf{1})$. (ii) If $r = \lambda N\mathbf{1}$ and $H - \lambda N \succeq 0$ for some λ, no such c exists. (iii) The shape term does not decide between (i) and (ii). For $H = \mathrm{diag}(1, 3)$, $g = (0, 2)$ and $N = I$, the shape term is $\frac{1}{6} > 0$, yet Muon is optimal among all same-norm steps.*

We therefore also compute the best same-norm step, an exactly solvable trust-region problem whose solution is certified globally optimal on the ellipsoid (Theorem A.22), and the model's prediction for the norm-matched filters we train with. Two remarks make this the right comparison. First, every hidden matrix has rank 768 and bands are defined by rank, so each matrix's step has squared norm proportional to $\sum_b |I_b| c_b^2$; matching each matrix's norm separately, as the experiments do, is the same constraint for exact-polar bands. Second, the ratios g_b/H_{bb} are not the fixed-norm criterion: even with diagonal H, the band with the larger ratio can be the one a same-norm step should shrink (Theorem A.27); only p, which uses the full H, decides.

Estimation and validity. We take checkpoints of a Muon training run (Section 6) at steps 1,024 (constant learning rate), 1,536 (start of the cooldown) and 2,048 (end), and form the exact next Nesterov input. Bands are ranks $\{1\text{–}8, 9\text{–}32, 33\text{–}96, 97\text{–}256, 257\text{–}512, 513\text{–}768\}$ of every matrix. g, the Hessian–vector products for H, and realized loss changes use three disjoint held-out sets (393k, 98k and 393k tokens; Section B). Across 12 test steps (Muon and $c^\star$; bands and neurons; three checkpoints), the predicted and realized changes have the same sign in every case and agree within a factor of three (Figure 3). Most importantly, the $c^\star$ step *realizes* a much larger decrease than Muon's step on unseen tokens: -0.00730 against -0.00054 at step 1,024. H is ill-conditioned (condition number $\sim 10^4$), so $c^\star$ is irregular across bands (Section B) and we use it only through realized losses and Theorem 3.

4 What the anatomy shows

F1: Muon over-steps its top singular directions. At every checkpoint the per-band multiplier rises monotonically from the top band to the tail (Figure 1a). The tail-to-top ratio is 48 at step 1,024, 38 at the start of the cooldown and 190 at the end. Within each of the four layer types separately, the top band has the smallest multiplier in all twelve type–checkpoint cases (Section B). The raw quantities show the mechanism. At step 1,024 the top band's curvature H_{11} is $29\times$ that of the next band, but its gradient alignment g_1 is only $9\times$ larger. Curvature falls off faster than signal with rank. Equal step lengths therefore overstep the high-curvature top directions, in which gradients concentrate (Gur-Ari et al., 2018; Ghorbani et al., 2019; Cohen et al., 2021). The ratios are isolated-band quantities, however. At fixed norm, the full-curvature analysis (F2) confirms the top as the place to take step length from but gives it to ranks 97–512, not to the extreme tail.

F2: at Muon's norm, moving step length off the top helps until late in the cooldown. At step 1,024 and at the start of the cooldown, only 17% and 12% of the gap of Theorem 3 is shape. The rest is the greedy model asking for a smaller step ($\gamma^\star = 0.52, 0.49$). Even so, most of the gap is reachable *without* changing the step's norm. The model gradient is almost entirely tangent to the same-norm ellipsoid ($\|p\|/\|r\| = 1.00$ and 1.00), so by Theorem 4 a same-norm step does better. The best one is predicted to change the one-step loss by -0.00589 and -0.00488, against -0.00079 and $+0.00026$ for Muon; that is 87% and 92% of the full gap. It takes step length from ranks 1–96 and gives it to ranks 97–512 (Section B). We confirmed this by measuring the actual held-out loss after one step at matched norm (Figure 1b, Table 4). Relative to Muon, halving ranks 1–8 changes it by -4.4 and -3.2 ($\times 10^{-3}$), and the best same-norm filter by -4.5 and -3.4, which is 77% and 69% of the predicted gain. Along the tilt family σ^p, dL/dp =+0.090 and +0.060 at $p = 0$, so a negative power helps and a positive one hurts. A proved certificate (Theorem A.26) turns a predicted gain into a guaranteed one once the model's errors are bounded; its constants are not yet measured (Section E). As a consistency check, the observed errors, at most 5.4×10^{-4} at 12 calibration steps and $1.4\text{–}1.5 \times 10^{-3}$ for the same-norm gains, are well below the predicted gains here and far above the one at the end. At the end of the cooldown the picture changes. The top band is no longer over-stepped ($c_1 = 0.98$), the model predicts a gain of only -0.00009, and the measured one-step losses of all filters agree with Muon's to within 1.2×10^{-4} ($dL/dp \approx 0$). Theorem 3 still attributes 97% of the gap to shape there, but that gap is measured against a step of very different norm and is out of reach at Muon's norm. The headroom a norm-matched tilt can use is therefore largest early in the cooldown and gone by its end.

F3: neuron imbalance is concentrated in fused QKV. Per-neuron update norms of Muon's step vary most in the fused QKV matrix (CV 0.43, per-layer range 0.14–0.73). The MLP's expand rows and contract columns have CV 0.18, and the square attention-output matrix has CV 0.02 (Figure 1c). Grouping neurons by Muon's update size gives the same monotone picture as the bands: the most-moved sextile has multiplier 1.5 and the least-moved 7.1 (Section B). Only 5% of the neuron-group gap at step 1,024 is shape

© 2026 JP Morgan Chase & Co. All rights reserved

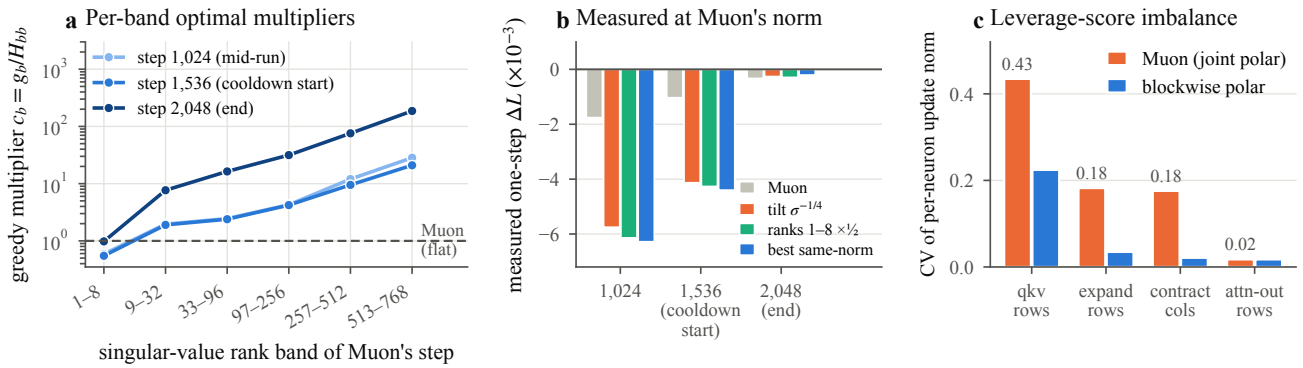


Figure 1: Anatomy of Muon's step (124M decoder, Muon seed 43; screen schedule with cooldown over steps 1,536–2,048). **(a)** Per-band second-order optimal multiplier $c_b = g_b/H_{bb}$ for the exact-polar step restricted to each singular-value band. Muon's flat filter is $c_b = 1$ (dashed). Multipliers scale as $1/\eta$, and η is $10\times$ smaller at step 2,048, so compare profiles rather than levels. The tail-to-top ratio is 48, 38 and 190 at the three checkpoints. These isolated-band ratios are not the fixed-norm criterion (Theorem A.27); panel (b) uses the full curvature. **(b)** Measured change of the held-out loss after one step at Muon's own update norm (393k tokens; Table 4). Bars show Muon, the exact tilt $\sigma^{-1/4}$, halving ranks 1–8, and the best same-norm filter from the quadratic model (Theorem 4). Moving step length off the top helps severalfold mid-training and at the start of the cooldown, and not at its end, where Muon's top band is no longer over-stepped. **(c)** Spread (coefficient of variation) of per-neuron update norms of Muon's step, mean over the 12 matrices of each type, at step 1,024. The fused 2304×768 QKV matrix carries the largest imbalance. Block-wise polar factors halve the QKV spread and nearly remove the MLP spread.

(65% at the end), so any benefit of block-wise steps must come mostly from the new directions they produce, not from reweighting.

Implications. We test three implications in Section 6. **I1:** rebalancing should matter most where the imbalance is largest, in QKV, and little in the MLP. **I2:** at fixed norm, moving step length off the top of the spectrum ($p < 0$) should help, independently of I1, because it acts on different structure. By F2 its per-step benefit is largest from mid-training to the start of the cooldown and gone by the end, because the decaying learning rate itself removes the top's over-stepping. A tilt confined to the cooldown should therefore build an early lead that partly erodes. **I3:** changes that barely alter the allocation should not help. Making NS_5 a more exact polar map is one: it removes NS_5's small deviations from unit length, mostly by lifting near-null directions, and leaves the top's share relative to the bulk essentially unchanged.

5 From measurements to interventions

Neuron allocation: block-wise polar steps. Following Theorem 2, we orthogonalize blocks separately. *Split QKV* uses three square 768×768 blocks, one each for Q, K and V. *Per-head* uses 36 wide 64×768 blocks, whose polar factors have exactly orthonormal rows. *Square blocks* splits QKV and also splits the MLP's 3072×768 and 768×3072 matrices into four square blocks each. Wide matrices' column blocks are scaled by $1/\sqrt{k}$ to keep Muon's update energy. All reported block arms are norm-matched anyway; the factor matters for the deployable version (Section D). The stateful alternative from prior work is NorMuon's row normalization by an EMA of row second moments (Li et al., 2026).

Spectral allocation: a matmul-only cooldown tilt. The exact tilt $U\,\phi_5(\Sigma)(\Sigma/\sigma_{\mathrm{rms}})^p V^\top$ needs an eigendecomposition per matrix and step. We avoid it as follows. Let $O = \mathrm{NS}_5(N) = U\,\mathrm{diag}(f_i)V^\top$ with $f_i = \phi_5(\hat{\sigma}_i)$, and $S = ON^\top = U\,\mathrm{diag}(f_i\sigma_i)U^\top \succeq 0$. A polynomial in S acts in the same singular basis (Theorem 5):

$$D = P(S/\Lambda)\,O = U\,\mathrm{diag}\left(f_i\,P(f_i\sigma_i/\Lambda)\right)V^\top.$$

With $P(x) \approx x^p$ the filter is $f_i^{1+p}\sigma_i^p$ up to scale. On the Newton–Schulz plateau ($f_i \in [0.68, 1.20]$) this is the exact tilt $f_i\sigma_i^p$ up to a factor in $[0.96, 1.10]$; near-null directions are lifted further. We use a degree-12 least-squares Chebyshev fit of x^p on $[10^{-3}, 1]$. Λ is 1.01 times a 20-step power-iteration estimate of $\lambda_{\max}(S)$, which is an estimate rather than a proved bound. We then rescale D to $\|O\|_{\mathrm{F}}$ (Algorithm 1).

Proposition 5 (What the tilt computes; Section A.1)**.** *For any coefficients and depth k and any real polynomial P, let $O = \varphi^k(X)$ and $D = P(OX^\top)\,O$. (i) D has the form $F(XX^\top)X$, so $D(XQ) = D(X)Q$ for every orthogonal Q. (ii) If $X = U\,\mathrm{diag}(d)V^\top$ with orthonormal columns, then $O = U\,\mathrm{diag}(f_i)V^\top$ and $D = U\,\mathrm{diag}(f_iP(f_id_i))V^\top$, with f_i the scalar iterate of d_i.*

Part (ii), not (i), is what keeps the singular vectors: $X \mapsto AX$ is also a Gram map. The normalization, Λ and the final rescaling are scalars. The tilted step therefore changes *only* how singular directions are weighted, so a comparison

© 2026 JP Morgan Chase & Co. All rights reserved

with Muon isolates the spectral allocation. One numerical detail decides whether the method works. Evaluating P in monomials by Horner's rule in TF32 cancels catastrophically (coefficients reach 4×10^4) and flipped the filter's sign on the top direction. With square blocks as the base (-0.0219 alone, against fused-QKV Muon), that version reached -0.0240, against -0.0322 for Clenshaw's recurrence (Clenshaw, 1955; Trefethen, 2019) in FP32 and -0.0332 for the exact tilt; the Clenshaw filter matches the exact tilt to direction cosine 0.998–0.999, so the normalized steps differ by at most 6.3% in Frobenius norm.

6 Experiments

Setup. We train a 124.5M-parameter pre-norm decoder (12 layers, width 768, 12 heads, tied embeddings) on a 167.8M-token FineWeb-Edu sample (Penedo et al., 2024) with the GPT-2 tokenizer, sequence length 512 and 24,576 tokens per step. Muon (1) acts on the 48 hidden matrices with momentum 0.95 and decoupled weight decay 0.1. Its state is in FP32, and Newton–Schulz runs in FP32 with TF32 matrix products. AdamW (Loshchilov and Hutter, 2019) (3×10^{-4}) handles embeddings and norms, and gradients are clipped at norm 1.0. The schedule is warmup–constant–cooldown (Hägele et al., 2024): 128 warmup steps, then linear decay to $0.1\times$ over the last 25% of steps. The *screen* runs 2,048 steps (50.3M tokens). The *confirmation* runs 4,096 steps (100.7M tokens). The metric is final loss on 4.18M held-out tokens. A 355M model (24 layers, width 1024, 16 heads) tests scale.

Controls. **Norm matching:** unless stated otherwise, every direction-changing arm is rescaled, per matrix and per step, to the Frobenius norm of Muon's NS_5 output on the same momentum, so differences reflect direction, not step size. **Tuned baselines:** on the screen, Muon's learning rate was swept over $\{0.0035, 0.005, 0.007, 0.01\}$, split-QKV Muon's over $\{0.0035, 0.005, 0.007\}$, and the 355M baseline's over the same three values. 0.005 is an interior optimum in all three cases, and every new arm uses the baseline's rate. **Paired seeds:** each seed fixes initialization and data order. Confirmations use seeds 42–44 for every arm, and we report paired 95% t-intervals. **A/A check:** before the cooldown the tilt arms run an optimizer mathematically identical to their baselines. Their intermediate evaluations differ by at most 0.009 nats, an empirical bound on numerical non-determinism plus evaluation noise. The screen baseline's seed spread is 0.0032 nats. **Gating:** midway through the study we adopted a rule that a full-length run is launched only if its screen beats the baseline by more than 0.005 nats. The factorial ablations and the 355M seeds 43–44 were released under it; the earlier confirmations were not.

Table 1: Confirmation at 100.7M tokens (124M model, three paired seeds per arm). Mean final held-out loss and mean paired differences against the original fused-QKV Muon and against the stronger split-QKV Muon. Brackets give 95% t-intervals for the latter; intervals against fused-QKV Muon are in the text. Lower is better.

Optimizer	Loss	Δ vs. Muon	Δ vs. split-QKV
Muon (fused QKV)	3.9427	—	+0.0250
NorMuon (fused QKV)	3.9277	−0.0150	+0.0100 [+0.003, +0.017]
Square blocks	3.9187	−0.0239	+0.0011 [−0.003, +0.005]
Split-QKV Muon	3.9176	−0.0250	—
+ cooldown tilt	3.9057	−0.0370	−0.0120 [−0.014, −0.010]
+ NorMuon	3.9139	−0.0288	−0.0038 [−0.006, −0.001]
+ NorMuon + tilt	3.9018	−0.0408	−0.0158 [−0.018, −0.014]
+ NorMuon + tilt + Schatten-1 decay	3.8966	−0.0460	−0.0210 [−0.024, −0.018]
Sq. blocks + NorMuon + tilt	3.9050	−0.0376	−0.0126 [−0.015, −0.010]

I1: the neuron gain is the QKV split. Splitting Q, K and V changes the final loss by -0.0250 nats [-0.0290, -0.0210] relative to fused-QKV Muon (Table 1). It is better in every seed, and fused-QKV Muon would need about 6.8% more tokens to match. Square blocks everywhere add nothing beyond it ($+0.0011$ [-0.0029, $+0.0051$]). Square blocks in the MLP alone, with QKV fused, change the loss by $+0.0039$ (screen). As I1 implies, the large QKV imbalance is the costly one and the moderate MLP imbalance is harmless. The data also locate the effect at the level of projections rather than rows. Per-head blocks equalize rows exactly yet do worse than three blocks (-0.0172 against -0.0255, screen). NorMuon's row normalization on fused QKV is worse than the split ($+0.0100$ [$+0.0031$, $+0.0170$]), and on top of the split it adds only -0.0038 [-0.0062, -0.0013]. One reading is that in a fused matrix Q, K and V share a single spectral budget, so the projection with the strongest momentum claims most of it, whereas splitting gives each projection its own budget. We have not isolated this mechanism. Either way, methods evaluated against fused-QKV Muon can take credit for a gain the split provides at no cost.

I2: the tilt helps, and adds to neuron balance. On the split baseline the cooldown tilt $\sigma^{-1/4}$ changes the loss by -0.0120 nats [-0.0137, -0.0102] and is better in all three seeds (-0.0113, -0.0127, -0.0119). The 2×2 factorial over {NorMuon, tilt} on split-QKV Muon is additive within noise: -0.0038 and -0.0120 separately, -0.0158 [-0.0178, -0.0138] together. Each component's effect is the same with or without the other: the tilt given NorMuon gives -0.0120, and NorMuon given the tilt gives -0.0038. The tilt supplies about three quarters of the combined gain. By the split baseline's own loss-versus-tokens curve (0.260 nats per doubling, a two-point log-linear estimate), split-QKV Muon would need about 4.3% more tokens to match the combination. Other tilts and band filters behave consistently on the screen (Section C); the dose was not optimized. Consistent with F2, timing and dose mat-

© 2026 JP Morgan Chase & Co. All rights reserved

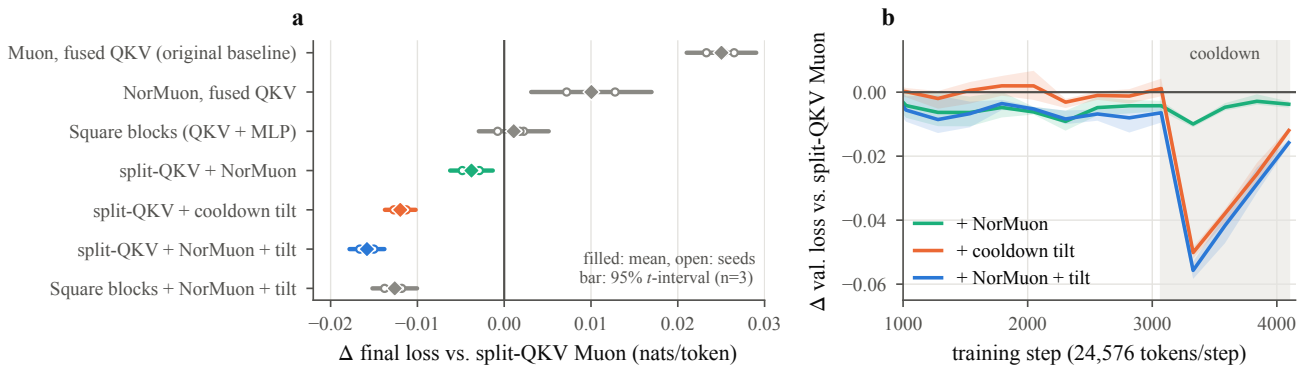


Figure 2: Controlled results at 100.7M tokens (124M, three paired seeds). **(a)** Paired difference in final loss against split-QKV Muon. Splitting Q, K and V explains the whole square-block gain; NorMuon on fused QKV is worse than the split; the tilt and NorMuon add on top of the split. **(b)** Paired difference over training (mean, with the seed range shaded). Intermediate points use 98k held-out tokens; the last point uses 4.18M. Before the cooldown the tilt arm is mathematically identical to its baseline (A/A differences $\leq$0.009). The tilt's lead peaks at -0.050 nats 256 steps into the cooldown and settles at -0.0120.

ter little: on the full stack, starting the tilt at step 1,024 or using $\sigma^{-1/2}$ changes the screen loss by only -0.0007 and $+0.0009$ (one seed each, within the 0.003 seed spread). We keep the tilt in the cooldown, where it costs least.

How the gain develops. Figure 2b shows the tilt's lead peaking at -0.050 nats 256 steps into the cooldown and narrowing to -0.0120 at the end. This is the pattern F2 predicts. At fixed norm the per-step benefit of moving step length off the top is large at the start of the cooldown and gone by its end, because the decaying learning rate itself removes the top's over-stepping, so the baseline recovers part of the early lead. The baseline's cooldown length and final learning rate were not tuned, so a different cooldown might recover part of the final gain; we have not tested this.

Scale. At 355M parameters (GPT-2-medium shape, 96 Muon matrices, same data and tokens), the split-QKV baseline's learning-rate screen again peaks at 0.005 (4.0642/4.0547/4.0647 for 0.0035/0.005/0.007). Against the baseline at the same seeds, NorMuon + tilt changes the final loss by -0.0188 nats [-0.0202, -0.0175] (per seed -0.0188, -0.0183, -0.0194), comparable to the 124M effect; the baseline's 3 seeds span only 0.0008 nats.

I3 and other geometric changes. Table 8 (Section C) collects them. Three extra cubic Newton–Schulz iterations drive every singular value of the NS_5 output to 1 and change the loss by -0.0020, within noise, as I3 implies. *Input whitening*, steepest descent in the data metric $\|DC^{1/2}\|_{\mathrm{op}}$, has the cleanest theory of anything we tried. In its exact form it is $\mathrm{GL}(n)$-equivariant and conditioning-free on least squares (Section A.4); the damped version we run is approximately so. It is harmful: $+0.021$ to $+0.060$ when norm-matched, and $+0.0524$ at its own magnitude. A whitened step concentrates its energy in low-variance input directions, so at a fixed norm it moves the layer's outputs less. The published Toner recipe (Mohamadi et al., 2026) gives -0.0068 (one seed), the sign of its reported gain at smaller size. For *decay geometry*, energy-matched Schatten-4 decay, which mainly shrinks the top of the spectrum, hurts ($+0.0177$). Schatten-1 decay, which removes about the same amount from every singular value and so prunes the bulk, helps (-0.0057 [-0.0100, -0.0014], three seeds). Decay along the polar factor of W is prior art, used with a fixed coefficient for low rank, where with Muon it costs loss (Andrianov et al., 2026), and for grokking (Li, 2026). With our energy-matched coefficient $\|W\|_{\mathrm{F}}^2/\|W\|_*$ (Section A.5) it is additive with our method: -0.0052 [-0.0066, -0.0038] on top of split QKV + NorMuon + tilt at 100.7M tokens (three seeds, Table 1), -0.0210 against split-QKV Muon, at $+26\%$ more wall time. In this regime the useful role of decay is to clean the bulk, not to cap the spectral norm.

Cost. Splitting QKV adds no FLOPs; its $+7\%$ recorded wall time is mostly the norm-matching scaffold. In our unfused FP32 implementation NorMuon + tilt takes $+12\%$ more wall time than split-QKV Muon at 124M and $+12\%$ at 355M, against a token saving of about 4.3%. **The gain is in tokens, not yet in wall clock;** folding the tilt into the Newton–Schulz coefficients, as DynMuon does, may remove most of the overhead.

7 What formal verification did and did not buy

All formal results are stated in Section A. The Newton–Schulz results hold for the iteration Muon actually runs, with any coefficients, depth and normalizer. The block results concern exactly orthogonal blocks. **What it bought.** (i) *Bookkeeping.* The energy identity (Theorem 2), carried over to wide matrices, gives the $1/\sqrt{k}$ factor without which a deployable block step on the MLP contraction is $1.9\times$ too large. (ii) *A symmetry guarantee.* The cheap tilt keeps Muon's right equivariance (Theorem 5). (iii) *A cor-*

© 2026 JP Morgan Chase & Co. All rights reserved

rection to our own reasoning. Theorem 3's stationary points are optima only under positive definite curvature, and its shape term compares Muon with a step of a different norm. Formalizing the question our norm-matched experiments ask produced Theorem 4 and a finite-step certificate, and changed our reading of the diagnostic (F2). (iv) *A change of direction.* Earlier counterexamples (Section A) showed that spatial orientation and curvature, not temporal statistics, decide momentum preferences under NS_5. **What it did not buy** was a prediction of training benefit. Whitened Muon is provably equivariant and conditioning-free yet hurts; the joint polar factor is provably suboptimal for the block norm, yet MLP blocks do not help. Whether a certified property is the bottleneck is an empirical question.

8 Related work

Muon and its geometry. Muon (Jordan et al., 2024b) and its scaling studies (Liu et al., 2025; Essential AI et al., 2025) have been analyzed as steepest descent under operator norms (Bernstein and Newhouse, 2024; Large et al., 2024), as linear-minimization-oracle methods (Pethick et al., 2025; Riabinin et al., 2025), as non-Euclidean trust regions (Kovalev, 2025) and as constrained optimization (Chen et al., 2026; Shen et al., 2026). Prior work also studies polar approximations (Amsel et al., 2026) and the benefit of flat spectra for tail-end memories (Wang et al., 2026b). On matrix quadratics, Shulgin et al. (2026) show that the polar direction beats the gradient direction, each at its best step size, when large singular directions sit in high curvature, and that finite Newton–Schulz acts as a damped polar update.

Spectra, neurons and blocks. DynMuon schedules a spectral power from $+1$ to $-\frac{1}{4}$ (Wu et al., 2026a). Freon/Kaon show that random or inverted spectra can work as well as flat ones (Shumaylov et al., 2026), and Spectrally Targeted Muon, AdaMuon and PolarGrad adapt magnitudes (Goyal and Ramkumar, 2026; Si et al., 2025; Lau et al., 2025). NorMuon and Muon+ normalize rows (Li et al., 2026; Zhang et al., 2026a), Aurora gives the leverage-score argument (Dewulf et al., 2026), and row-sphere and Stiefel constraints are compared by Hu et al. (2026). MuonBP and HiMuon orthogonalize blocks (Khaled et al., 2026; Tang et al., 2026), and separate Q/K/V orthogonalization is standard in modded-nanogpt (Jordan et al., 2024a). Closest to our diagnostic, SAMuon measures held-out optimal step sizes along each singular direction of the momentum, finding a volatile head and a tolerant bulk, and builds an optimizer from the profile (Wu et al., 2026b). BulkBoost proves when a norm-matched bulk/spike reweighting improves an objective to first order (Chaitanya et al., 2026), and Wang et al. (2026a) attribute Muon's advantage over Adam to within- and cross-layer curvature of its update. Our tilt uses DynMuon's dose and late-training sign. We add the full cross-group curvature, neuron groups, scale versus fixed-norm shape with proved criteria, and a controlled attribution: the QKV split accounts for the block-wise gains and most of NorMuon's gain over fused-QKV Muon.

Data metric, curvature and evaluation. Shampoo, SOAP and K-FAC precondition with layer statistics (Gupta et al., 2018; Vyas et al., 2025; Martens and Grosse, 2015), and whitened Muon has close 2026 relatives (Mohamadi et al., 2026; Che, 2026; Zhang et al., 2026b; Du and Su, 2026; Xu et al., 2026). Gradients concentrate in a few top Hessian directions (Gur-Ari et al., 2018) above a bulk (Sagun et al., 2017; Ghorbani et al., 2019; Papyan, 2020), near the edge of stability (Cohen et al., 2021). Greedy step sizes suffer short-horizon bias (Wu et al., 2018; Baydin et al., 2018), and optimizer gains shrink against tuned baselines (Wen et al., 2026; Zhao et al., 2025). Like Wu et al. (2026b) and Wang et al. (2026a), we measure curvature along the optimizer's own step, here with held-out Hessian–vector products (Pearlmutter, 1994). Formal verification of ML systems goes back at least to Selsam et al. (2017).

Novelty. Held-out spectral probing of Muon's step (Wu et al., 2026b) and first-order analysis of norm-matched reallocation (Chaitanya et al., 2026) precede this work. We found no prior work that measures the full cross-group curvature of the step by singular band and by neuron group, separates scale from fixed-norm shape with proved criteria and counterexamples, and uses the result in a controlled attribution of published ingredients.

9 Limitations and conclusion

Scale and regime. One data distribution; 124M and 355M parameters (three seeds each); at most 100.7M tokens, under one token per parameter. Effect sizes are specific to this regime. **The diagnostic is local and noisy.** It is second-order, one step ahead and from a single run, without intervals. g and H come from separate held-out folds, and its joint optimum is ill-conditioned. It perturbs only the hidden matrices along exact-polar band directions, leaving out the AdamW and decay parts of the update. The measured one-step losses use one checkpoint seed and one validation subset, without intervals, and the certificate's constants are not yet measured. We used the diagnostic for direction, not dose; the tilt power comes from prior work. The fixed-norm analysis covers bands only. **Design and statistics.** The anatomy was not pre-registered; intervals are nominal (three paired seeds); most screen arms are single-seed. **Alternatives.** The baseline's cooldown was not tuned; the QKV mechanism is a hypothesis. **Provenance.** Comparisons with fused-QKV or square-block arms cross code versions; split-QKV arms share one (Section E).

Conclusion. Measured against curvature, both of Muon's hidden allocations are miscalibrated. Its spectrum oversteps the top, which a same-norm reallocation can exploit

© 2026 JP Morgan Chase & Co. All rights reserved

until late in the cooldown, and its per-neuron steps are most uneven in fused QKV matrices. Corrections aimed at these help additively; principled changes that miss them do not.

References

Noah Amsel, David Persson, Christopher Musco, and Robert M. Gower. The polar express: Optimal matrix sign methods and their application to the Muon algorithm. In *International Conference on Learning Representations (ICLR)*, 2026. arXiv:2505.16932.

Dmitrii Andriianov, Andrey Veprikov, and Aleksandr Beznosikov. Spectral weight decay: Inducing low-rank structure in neural network weights. *arXiv preprint arXiv:2610.11730*, 2026.

Atilim Gunes Baydin, Robert Cornish, David Martinez Rubio, Mark Schmidt, and Frank Wood. Online learning rate adaptation with hypergradient descent. In *International Conference on Learning Representations (ICLR)*, 2018.

Jeremy Bernstein and Laker Newhouse. Old optimizer, new norm: An anthology. *arXiv preprint arXiv:2409.20325*, 2024.

Meher Chaitanya, Tianyi Zhou, and Aristides Gionis. Does Muon need fine-grained spectral shaping? *arXiv preprint arXiv:2610.07497*, 2026.

Tong Che. Second-order Muon done right: A principled marriage of spectral geometry and curvature. *arXiv preprint arXiv:2608.09763*, 2026.

Lizhang Chen, Jonathan Li, and Qiang Liu. Muon optimizes under spectral norm constraints. *Transactions on Machine Learning Research*, 2026. ISSN 2835-8856. URL https://openreview.net/forum?id=Blz4hjxLwU.

C. W. Clenshaw. A note on the summation of Chebyshev series. *Mathematical Tables and Other Aids to Computation*, 9(51):118–120, 1955. doi: 10.1090/S0025-5718-1955-0071856-0.

Jeremy M. Cohen, Simran Kaur, Yuanzhi Li, J. Zico Kolter, and Ameet Talwalkar. Gradient descent on neural networks typically occurs at the edge of stability. In *International Conference on Learning Representations (ICLR)*, 2021.

Alec Dewulf, Dhruv Pai, Li Yang, Ashley Zhang, and Ben Keigwin. Aurora: A leverage-aware spectral optimizer. *arXiv preprint arXiv:2606.27715*, 2026.

Petros Drineas, Malik Magdon-Ismail, Michael W. Mahoney, and David P. Woodruff. Fast approximation of matrix coherence and statistical leverage. *Journal of Machine Learning Research*, 13(111):3475–3506, 2012.

Zhehang Du and Weijie Su. The Newton-Muon optimizer. *arXiv preprint arXiv:2604.01472*, 2026.

Essential AI, Ishaan Shah, et al. Practical efficiency of Muon for pretraining. *arXiv preprint arXiv:2505.02222*, 2025.

Behrooz Ghorbani, Shankar Krishnan, and Ying Xiao. An investigation into neural net optimization via Hessian eigenvalue density. In *International Conference on Machine Learning (ICML)*, volume 97 of *Proceedings of Machine Learning Research*, pages 2232–2241, 2019.

Vishrut Goyal and Rohan Ramkumar. Spectrally targeted Muon. *arXiv preprint arXiv:2610.10965*, 2026.

Vineet Gupta, Tomer Koren, and Yoram Singer. Shampoo: Preconditioned stochastic tensor optimization. In *International Conference on Machine Learning (ICML)*, volume 80 of *Proceedings of Machine Learning Research*, pages 1842–1850, 2018.

Guy Gur-Ari, Daniel A. Roberts, and Ethan Dyer. Gradient descent happens in a tiny subspace. *arXiv preprint arXiv:1812.04754*, 2018.

Alexander Hägele, Elie Bakouch, Atli Kosson, Loubna Ben Allal, Leandro Von Werra, and Martin Jaggi. Scaling laws and compute-optimal training beyond fixed training durations. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 37, 2024.

Kai Hu, Li Zhang, Weichen Yu, Matt Fredrikson, and Haoqi Hu. Do orthogonality constraints help language-model training beyond sphere optimization? *arXiv preprint arXiv:2611.00593*, 2026.

Keller Jordan et al. modded-nanogpt: Speedrunning the NanoGPT baseline. GitHub repository, https://github.com/KellerJordan/modded-nanogpt, 2024a.

Keller Jordan et al. Muon: An optimizer for hidden layers in neural networks. Blog post, https://kellerjordan.github.io/posts/muon/, December 2024b.

Ahmed Khaled, Kaan Ozkara, Tao Yu, Mingyi Hong, and Youngsuk Park. MuonBP: Faster Muon via block-periodic orthogonalization. In *International Conference on Learning Representations (ICLR)*, 2026. arXiv:2510.16981.

Dmitry Kovalev. Understanding gradient orthogonalization for deep learning via non-Euclidean trust-region optimization. *arXiv preprint arXiv:2503.12645*, 2025.

© 2026 JP Morgan Chase & Co. All rights reserved

Tim Large, Yang Liu, Minyoung Huh, Hyojin Bahng, Phillip Isola, and Jeremy Bernstein. Scalable optimization in the modular norm. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 37, 2024. doi: 10.52202/079017-2338.

Tim Tsz-Kit Lau, Qi Long, and Weijie Su. PolarGrad: A class of matrix-gradient optimizers from a unifying preconditioning perspective. *arXiv preprint arXiv:2505.21799*, 2025.

Mingyu Li. Low-rank decay for grokking in scale-invariant transformers: A spectral-geometric view. *arXiv preprint arXiv:2606.04405*, 2026.

Zichong Li, Liming Liu, Chen Liang, Weizhu Chen, and Tuo Zhao. NorMuon: Making Muon more efficient and scalable. In *International Conference on Machine Learning (ICML)*, 2026. arXiv:2510.05491.

Jingyuan Liu et al. Muon is scalable for LLM training. *arXiv preprint arXiv:2502.16982*, 2025.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations (ICLR)*, 2019.

James Martens and Roger Grosse. Optimizing neural networks with Kronecker-factored approximate curvature. In *International Conference on Machine Learning (ICML)*, volume 37 of *Proceedings of Machine Learning Research*, pages 2408–2417, 2015.

M. Mohamadi, K. Wen, T. Wang, and Z. Li. Toner / MuonT: input-whitened Muon. Blog post and code, November 2026.

Vardan Papyan. Traces of class/cross-class structure pervade deep learning spectra. *Journal of Machine Learning Research*, 21(252):1–64, 2020.

Barak A. Pearlmutter. Fast exact multiplication by the Hessian. *Neural Computation*, 6(1):147–160, 1994. doi: 10.1162/neco.1994.6.1.147.

Guilherme Penedo et al. The FineWeb datasets: Decanting the web for the finest text data at scale. In *Advances in Neural Information Processing Systems (NeurIPS), Datasets and Benchmarks Track*, volume 37, 2024. doi: 10.52202/079017-0970.

Thomas Pethick, Wanyun Xie, Kimon Antonakopoulos, Zhenyu Zhu, Antonio Silveti-Falls, and Volkan Cevher. Training deep learning models with norm-constrained LMOs. In *International Conference on Machine Learning (ICML)*, volume 267 of *Proceedings of Machine Learning Research*, pages 49069–49104, 2025.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. Technical report, OpenAI, 2019.

Artem Riabinin, Egor Shulgin, Kaja Gruntkowska, and Peter Richtárik. Gluon: Making Muon & Scion great again! (Bridging theory and practice of LMO-based optimizers for LLMs). *arXiv preprint arXiv:2505.13416*, 2025.

Levent Sagun, Utku Evci, V. Ugur Guney, Yann Dauphin, and Léon Bottou. Empirical analysis of the Hessian of over-parametrized neural networks. *arXiv preprint arXiv:1706.04454*, 2017.

Daniel Selsam, Percy Liang, and David L. Dill. Developing bug-free machine learning systems with formal mathematics. In *International Conference on Machine Learning (ICML)*, volume 70 of *Proceedings of Machine Learning Research*, pages 3047–3056, 2017.

Wei Shen, Ruichuan Huang, Minhui Huang, Cong Shen, and Jiawei Zhang. On the convergence analysis of Muon. *Transactions on Machine Learning Research*, 2026. ISSN 2835-8856. URL https://openreview.net/forum?id=4nH4CulGaP.

Egor Shulgin, Sam Laing, Antonio Orvieto, and Peter Richtárik. A quadratic lens on Muon's polar update: Benefits, pitfalls, and the role of Newton-Schulz. *arXiv preprint arXiv:2610.30280*, 2026.

Zakhar Shumaylov et al. Muon is not that special: Random or inverted spectra work just as well. *arXiv preprint arXiv:2605.11181*, 2026.

Chongjie Si, Debing Zhang, and Wei Shen. AdaMuon: Adaptive Muon optimizer. *arXiv preprint arXiv:2507.11005*, 2025.

Ziyuan Tang, Tianshi Xu, Yousef Saad, and Yuanzhe Xi. Hierarchical Muon: Tiled Newton-Schulz updates for efficient Muon optimization. *arXiv preprint arXiv:2606.27216*, 2026.

Lloyd N. Trefethen. *Approximation Theory and Approximation Practice, Extended Edition*. Society for Industrial and Applied Mathematics (SIAM), 2019. doi: 10.1137/1.9781611975949.

Nikhil Vyas et al. SOAP: Improving and stabilizing Shampoo using Adam for language modeling. In *International Conference on Learning Representations (ICLR)*, 2025. arXiv:2409.11321.

Shuche Wang, Fengzhuo Zhang, Jiaxiang Li, Dirk Bergemann, and Zhuoran Yang. Why Muon outperforms Adam: A curvature perspective. *arXiv preprint arXiv:2606.04662*, 2026a.

© 2026 JP Morgan Chase & Co. All rights reserved

Shuche Wang et al. Muon outperforms Adam in tail-end associative memory learning. In *International Conference on Learning Representations (ICLR)*, 2026b. arXiv:2509.26030.

Kaiyue Wen, David Hall, Tengyu Ma, and Percy Liang. Fantastic pretraining optimizers and where to find them. In *International Conference on Learning Representations (ICLR)*, 2026. arXiv:2509.02046.

Fangzhou Wu, Rikhav Shah, Sandeep Silwal, and Qiuyi Zhang. DynMuon: A dynamic spectral shaping view of Muon. *arXiv preprint arXiv:2605.17109*, 2026a. ICML 2026 Workshop on High-dimensional Learning Dynamics (HiLD).

Xiaodong Wu, Wenyi Yu, Chao Zhang, and Philip Woodland. Spectral allocation: Why Muon outperforms Adam, and how to improve Muon. *arXiv preprint arXiv:2608.25990*, 2026b.

Yuhuai Wu, Mengye Ren, Renjie Liao, and Roger Grosse. Understanding short-horizon bias in stochastic meta-optimization. In *International Conference on Learning Representations (ICLR)*, 2018.

Chenrui Xu, Wenjing Yan, and Ying-Jun Angela Zhang. FISMO: Fisher-structured momentum-orthogonalized optimizer. *arXiv preprint arXiv:2601.21750*, 2026.

Ruijie Zhang et al. MUON+: Towards more effective Muon via one additional normalization step for LLM pre-training. *arXiv preprint arXiv:2602.21545*, 2026a.

Yechen Zhang et al. Mousse: Rectifying the geometry of Muon with curvature-aware preconditioning. *arXiv preprint arXiv:2603.09697*, 2026b.

Rosie Zhao, Depen Morwani, David Brandfonbrener, Nikhil Vyas, and Sham Kakade. Deconstructing what makes a good optimizer for language models. In *International Conference on Learning Representations (ICLR)*, 2025.

© 2026 JP Morgan Chase & Co. All rights reserved

A Formal results

This appendix states, in ordinary mathematical language, every formal result the paper uses. All of them are proved in Lean 4 with Mathlib, in exact real arithmetic. They certify the algebra and geometry behind each method. They say nothing about floating-point behaviour, convergence or training benefit.

Notation. $\langle A, B\rangle = \operatorname{tr}(A^\top B)$. For real coefficients a, b, c let $\varphi(X) = aX + b\,XX^\top X + c\,(XX^\top)^2 X$, and let φ^k be its k-fold composition. For any function ν of an $m \times m$ matrix, the *Newton–Schulz map* is

$$\mathrm{NS}(X) = \varphi^k\big(\nu(XX^\top)\,X\big).$$

Muon's NS_5 is the case $k = 5$, $(a, b, c) = (3.4445, -4.775, 2.0315)$, $\nu(XX^\top) = 1/\sqrt{\operatorname{tr} XX^\top}$. Results stated for NS hold for *every* choice of a, b, c, k and ν, and so apply to the iteration Muon actually runs, not only to an exact polar factor. A matrix Q is orthogonal if $QQ^\top = I$.

A.1 Newton–Schulz maps and the matmul-only tilt

Definition A.1 (Gram map)**.** A map $G : \mathbb{R}^{m\times n} \to \mathbb{R}^{m\times n}$ is a *Gram map* if there is a function $F : \mathbb{R}^{m\times m} \to \mathbb{R}^{m\times m}$ such that $G(X) = F(XX^\top)\,X$ for every X.

Theorem A.2 (Gram maps commute with right rotations)**.** *If G is a Gram map and $Q \in \mathbb{R}^{n\times n}$ is orthogonal, then $G(XQ) = G(X)\,Q$ for every X.*

Proof idea. $(XQ)(XQ)^\top = XX^\top$, so $G(XQ) = F(XX^\top)XQ$.

Theorem A.3 (Newton–Schulz iterates are Gram maps)**.** *For all a, b, c and k, the map $X \mapsto \varphi^k(X)$ is a Gram map. Consequently $\mathrm{NS}(XQ) = \mathrm{NS}(X)\,Q$ for every orthogonal Q, for any normalizer ν.*

Proof idea. Each step multiplies X on the left by the matrix $aI + b\,XX^\top + c\,(XX^\top)^2$. That matrix depends only on the Gram matrix, and the class of Gram maps is closed under composition.

Theorem A.4 (The polynomial tilt preserves Muon's symmetry)**.** *Let P be any real polynomial and $O = \varphi^k(X)$. The tilt $T_P(X) = P(OX^\top)\,O$ is a Gram map. Hence $T_P(XQ) = T_P(X)\,Q$ for every orthogonal Q.*

Right equivariance alone does not mean singular vectors are kept: $X \mapsto AX$ for a fixed non-scalar A is also a Gram map. The next theorem shows that the Newton–Schulz iterates and the tilt do keep them.

Theorem A.5 (Singular-basis action of Newton–Schulz and of the tilt)**.** *Let $X = U\,\mathrm{diag}(d)V^\top$, where U and V have orthonormal columns ($U^\top U = I$, $V^\top V = I$; any shapes, e.g. a thin SVD). Let s_k be the scalar map $x \mapsto ax + bx^3 + cx^5$ iterated k times, and $f_i = s_k(d_i)$. (i) $\varphi^k(X) = U\,\mathrm{diag}(f_i)V^\top$. (ii) For every real polynomial P, $P(\varphi^k(X)X^\top)\,\varphi^k(X) = U\,\mathrm{diag}\big(f_i\,P(f_i d_i)\big)V^\top$.*

Proof idea. $XX^\top = U\,\mathrm{diag}(d_i^2)U^\top$ and $U^\top U = I$ make every product collapse onto the diagonal. Induct on k for (i); for (ii), induct on the powers of S, then use linearity in P.

In the implementation X is first normalized and S is divided by an estimate Λ of its largest eigenvalue. Both are scalar factors, absorbed into d and P. Both singular frames are therefore kept, and the implemented filter is $f_i\,P(f_i\sigma_i/\Lambda)$, not $f_i\,P(\sigma_i)$. With $P(x) \approx x^p$ it is $f_i^{1+p}\sigma_i^p$ up to scale. On the Newton–Schulz plateau ($f_i \in [0.68, 1.20]$) this differs from the exact tilt $f_i\sigma_i^p$ by the factor f_i^p, which lies in $[0.96, 1.10]$ for $p = -\frac{1}{4}$.

Theorem A.6 (From filter error to step error)**.** *(i) For non-zero $Y, \hat{Y}$ in a real normed space and $\rho \geq 0$, $\big\|\rho\hat{Y}/\|\hat{Y}\| - \rho Y/\|Y\|\big\| \leq 2\rho\,\|\hat{Y} - Y\|/\|Y\|$. (ii) For unit vectors u, v in a real inner-product space, $\|u - v\|^2 = 2(1 - \langle u, v\rangle)$.*

By (i), a raw filter error of $\epsilon\|Y\|$ moves the norm-matched step by at most 2ϵ of its norm. By (ii), the measured direction cosine of 0.998 between the matmul and exact tilts means their normalized steps are within 6.3%.

© 2026 JP Morgan Chase & Co. All rights reserved

A.2 Block geometry and per-neuron step sizes

Definition A.7 (Steepest direction)**.** For a norm-like function $\mathcal{N}$ and a gradient G, a matrix D is a *steepest direction* if $\mathcal{N}(D) \le 1$ and $\langle G, D' \rangle \le \langle G, D \rangle$ for every D' with $\mathcal{N}(D') \le 1$.

Theorem A.8 (Block-wise steepest descent)**.** *Let $G = (G_1, \dots, G_p)$ be a block gradient and $\mathcal{N}_1, \dots, \mathcal{N}_p$ block norms. If each D_i is a steepest direction for $\mathcal{N}_i$ at G_i, then $(D_1, \dots, D_p)$ is a steepest direction at G for the max-norm $\|D\|_{\max} = \max_i \mathcal{N}_i(D_i)$, whose unit ball is the product of the block unit balls. With every $\mathcal{N}_i$ the spectral norm, the block-wise polar factors are therefore steepest for $\|\cdot\|_{\mathrm{blk}}$.*

Proposition A.9 (The joint polar step can be strictly worse)**.** *Take two 1×2 blocks $G_1 = G_2 = (1, 0)$. The polar factor of the stacked matrix has blocks $D_1 = D_2 = (1/\sqrt{2},\, 0)$, a partial isometry with $D_1^\top D_1 + D_2^\top D_2 = e_1 e_1^\top$, and attains $\langle G, D \rangle = \sqrt{2}$. The block-wise step $D_1 = D_2 = (1, 0)$ is feasible for the max-norm and attains $2 > \sqrt{2}$.*

Proposition A.10 (Row norms of polar factors)**.** *(i) If O is square and $OO^\top = I$, every row of O has unit norm. (ii) If $O \in \mathbb{R}^{m \times k}$ has orthonormal columns ($O^\top O = I_k$), its squared row norms sum to k. (iii) They need not be equal: $O = (1;\, 0) \in \mathbb{R}^{2 \times 1}$ satisfies $O^\top O = 1$ and has row norms 1 and 0.*

The squared row norms in (ii) are the leverage scores of $\mathrm{range}(O)$. In Muon's step they are the per-neuron step sizes of a tall matrix.

Proposition A.11 (Energy identity for block-wise steps)**.** *Let $O \in \mathbb{R}^{m \times k}$ with $O^\top O = I_k$, let $B_1, \dots, B_p \in \mathbb{R}^{k \times k}$ be orthogonal, and let $c^2 = p$. Then*

$$\|c\, O\|_{\mathrm{F}}^2 = \sum_{b=1}^{p} \|B_b\|_{\mathrm{F}}^2 \; (= pk).$$

For a tall matrix with $m = pk$ rows, Muon's shape factor $s = \sqrt{m/k} = \sqrt{p}$ therefore gives the joint step the same Frobenius energy as p orthogonal $k \times k$ blocks, and the learning rate transfers.

A.3 Second-order model of a filtered step

For group multipliers $c \in \mathbb{R}^B$, a gradient-alignment vector g and a symmetric curvature matrix H, let $q(c) = -g^\top c + \frac{1}{2} c^\top H c$ (Equation (2)).

Theorem A.12 (Gap formula)**.** *If H is symmetric and $Hc^\star = g$, then for every c,*

$$q(c) - q(c^\star) = \tfrac{1}{2}\,(c - c^\star)^\top H\,(c - c^\star).$$

Theorem A.13 (Scale/shape split)**.** *Let H be symmetric with $Hc^\star = g$, let u be a direction (Muon's step is $u = \mathbf{1}$), and let γ satisfy $u^\top H(\gamma u - c^\star) = 0$. Then*

$$q(u) - q(c^\star) = \tfrac{1}{2}(1 - \gamma)^2\, u^\top H u \;+\; \tfrac{1}{2}\,(\gamma u - c^\star)^\top H\,(\gamma u - c^\star).$$

When H is positive definite (Theorem A.16), the first term is the loss due to the overall step size along u and the second the loss due to the shape of u.

Proof idea. Write $u - c^\star = (1 - \gamma)u + (\gamma u - c^\star)$ and apply Theorem A.12. The cross term vanishes by the condition on γ, so the identity is Pythagoras in the H-inner product.

Lemma A.14 (The best uniform rescaling is an H-projection)**.** *If $u^\top H u \neq 0$, then $\gamma^\star = u^\top g / u^\top H u$ satisfies $u^\top H(\gamma^\star u - c^\star) = 0$. So Theorem A.13 applies with $\gamma = \gamma^\star$, the best scalar multiple of u.*

Lemma A.15 (The shape term is strictly positive)**.** *If H is positive definite and $\gamma u \neq c^\star$, then $\frac{1}{2}(\gamma u - c^\star)^\top H(\gamma u - c^\star) > 0$. Muon's flat step therefore loses something to shape unless the second-order optimum is itself flat.*

When the stationary points are optima. Theorems A.12 and A.13 are identities that hold for any symmetric H. Calling $c^\star$ and $\gamma^\star$ optima, and the two terms of Theorem A.13 losses, needs curvature assumptions:

Theorem A.16 (Stationary points are minimizers under positive curvature)**.** *Let H be symmetric and $Hc^\star = g$. (i) If H is positive semidefinite, then $q(c^\star) \le q(c)$ for every c. (ii) If H is positive definite, then $q(c^\star) < q(c)$ for every $c \neq c^\star$.*

Lemma A.17 (The best uniform rescaling minimizes along the ray)**.** *If $u^\top H u > 0$, then $q(\gamma^\star u) \le q(\gamma u)$ for every γ, where $\gamma^\star = u^\top g / u^\top H u$. Indeed $q(\gamma u) - q(\gamma^\star u) = \frac{1}{2}\, u^\top H u\, (\gamma - \gamma^\star)^2$.*

© 2026 JP Morgan Chase & Co. All rights reserved

Comparison at a fixed update norm. The shape term of Theorem A.13 compares u with $c^\star$, which generally has a different norm. Experiments that norm-match every arm compare steps on the ellipsoid $\{c : c^\top N c = u^\top N u\}$ for a symmetric positive definite N. Write $r = Hu - g$ for the gradient of q at u, $w = Nu$, and

$$p = r - \frac{r^\top w}{w^\top w}\, w$$

for the component of r orthogonal to w.

Theorem A.18 (Lagrange identity on the same-norm ellipsoid)**.** *Let H and N be symmetric. For every c with $c^\top N c = u^\top N u$ and every real λ,*

$$q(c) - q(u) = (r - \lambda N u)^\top (c - u) + \tfrac{1}{2}\,(c - u)^\top (H - \lambda N)(c - u).$$

Proof idea. Expand $q(u+d)-q(u)$, and subtract $\frac{\lambda}{2}$ times the constraint, which in terms of $d = c-u$ reads $2u^\top N d + d^\top N d = 0$.

Corollary A.19 (Certificate that no same-norm step improves)**.** *If $r = \lambda N u$ and $H - \lambda N$ is positive semidefinite for some λ, then $q(u) \le q(c)$ for every c with $c^\top N c = u^\top N u$.*

Theorem A.20 (A same-norm improvement exists when the tangential gradient is non-zero)**.** *Let N be symmetric positive definite, $u \neq 0$, and $p \neq 0$. Then there is c with $c^\top N c = u^\top N u$ and $q(c) < q(u)$.*

Proof idea. The vector $v = -p$ satisfies $w^\top v = 0$ and $r^\top v = -\|p\|^2 < 0$. Rescale it to $e = \kappa v$ with $e^\top N e = u^\top N u$; then $u^\top N e = 0$. Every $c = xu + ye$ with $x^2 + y^2 = 1$ lies on the ellipsoid. Along the rational parametrization $x = (1 - s^2)/(1 + s^2)$, $y = 2s/(1 + s^2)$,

$$q(c) - q(u) = \frac{2s\big[\, r^\top e + s\,\alpha_1 + s^2\alpha_2 + s^3\alpha_3 \,\big]}{(1 + s^2)^2}$$

for constants α_i. Choosing $s = \min\{1, -r^\top e/(2K)\}$ with $K = |\alpha_1| + |\alpha_2| + |\alpha_3| + 1$ makes the bracket at most $r^\top e/2 < 0$.

Proposition A.21 (A positive shape term does not imply a same-norm improvement)**.** *Let $H = \mathrm{diag}(1, 3)$, $g = (0, 2)$, $u = (1, 1)$ and $N = I$. The minimizer is $c^\star = (0, \frac{2}{3})$, the best uniform rescaling is $\gamma^\star = \frac{1}{2}$, and the shape term of Theorem A.13 equals $\frac{1}{6} > 0$. Yet $q(u) \le q(c)$ for every c with $\|c\| = \|u\|$. Indeed, on that circle $q(c) - q(u) = (c_2 - 1)^2$.*

Here $r = (1, 1) = 1 \cdot Nu$ and $H - N = \mathrm{diag}(0, 2) \succeq 0$, so Theorem A.19 applies. Theorems A.19 and A.20 together give the decision rule used in Section 4: compute p, and if it is non-zero, solve for the best same-norm step.

Certifying the best same-norm step and a finite-step gain.

Theorem A.22 (Certificate for the best same-norm step)**.** *Let H and N be symmetric. If c satisfies $c^\top N c = R$, $(H+\lambda N)c = g$ and $H + \lambda N \succeq 0$, then $q(c) \le q(c')$ for every c' with $c'^\top N c' = R$.*

The best same-norm steps reported in Section 4 satisfy these conditions at all three checkpoints: $\lambda > 0$, so $H + \lambda N \succ 0$, and the stationarity residual is below 10^{-15}. They are therefore the global optima of the model on the ellipsoid.

Lemma A.23 (Model-error budget)**.** *For estimates $(\hat g, \hat H)$ of (g, H) and any c, $|\hat q(c) - q(c)| \le |(\hat g - g)^\top c| + \frac{1}{2}|c^\top (\hat H - H) c|$.*

With $\|\hat g - g\| \le \delta_g$ and $\|\hat H - H\|_{\mathrm{op}} \le \delta_H$ this is at most $\delta_g \|c\| + \frac{1}{2}\delta_H \|c\|^2$.

Theorem A.24 (A-posteriori finite-step certificate)**.** *Let $F(x)$ be the true loss change of an implemented step x and $\hat q(x)$ its predicted change, with $|F(x) - \hat q(x)| \le E(x)$ for a candidate x_1 and every comparator x in a set S. If $\hat q(x_1) + E(x_1) < \hat q(x) - E(x)$ for every $x \in S$, then $F(x_1) < F(x)$ for every $x \in S$.*

The certificate needs no inverse of H. Each part of the error E is bounded by a proved result:

© 2026 JP Morgan Chase & Co. All rights reserved

Theorem A.25 (The three error terms)**.** *(i)* Taylor remainder. *Let $\phi(t)$ be the loss after a fraction $t \in [0,1]$ of a step, twice continuously differentiable with $|\phi'''| \le M$ on $(0,1)$. Then $|\phi(1) - \phi(0) - \phi'(0) - \frac{1}{2}\phi''(0)| \le M/6$. (ii)* Estimation. *$|x^\top y| \le \|x\|\,\|y\|$. Hence, if $\|\hat{g} - g\| \le \delta_g$ and $|x^\top(\hat{H} - H)x| \le \delta_H\|x\|^2$ for all x, then $|\hat{q}(c) - q(c)| \le \delta_g\|c\| + \frac{1}{2}\delta_H\|c\|^2$. (iii)* Implementation. *If $\|\nabla L\| \le G$ on a convex set containing the endpoints of the exact and the implemented step, and both steps are norm-matched to radius ρ from raw outputs Y and $\hat{Y}$, then their losses differ by at most $2G\rho\,\|\hat{Y} - Y\|/\|Y\|$.*

In (i), $\phi'(0) = -g^\top c$ and $\phi''(0) = c^\top H c$ are the first and second directional derivatives of the held-out loss along the step, which the diagnostic measures. Part (iii) combines the mean-value inequality with Theorem A.6.

Corollary A.26 (End-to-end certificate)**.** *For an implemented candidate step and a comparator, let I, T and E bound the implementation, Taylor and estimation errors of each. If $\hat{q}_1 + (I_1 + T_1 + E_1) < \hat{q}_0 - (I_0 + T_0 + E_0)$, then the implemented candidate truly lowers the loss more than the comparator.*

Which inputs are established. The theorems are unconditional; their *inputs* differ in status. Established for our measurements: the trust-region conditions of Theorem A.22 ($\lambda > 0$, residual below 10^{-15}), so the reported best same-norm steps are certified optima of the model; the tangential gradient $p \neq 0$ of Theorem A.20 ($\|p\|/\|r\| = 1.00, 1.00, 0.79$), so a same-norm improvement of the model exists at every checkpoint; and the hypotheses of Theorems A.5 and A.6, which hold by construction. *Not* established: the constants M, δ_g, δ_H and G of Theorem A.25, and sampling intervals for the measured one-step losses. Theorem A.26 is therefore not instantiated as a guarantee. The comparison in F2 with observed prediction errors is a consistency check only. Section E specifies the measurements that would instantiate it.

Limits of the diagnostic. Two further facts bound what the measurements can say.

Proposition A.27 (Per-band ratios are not the fixed-norm criterion)**.** *Let $H = \mathrm{diag}(1,3)$, $g = (\frac{1}{2}, 2)$, $u = (1,1)$ and $N = I$. The isolated ratio g_b/H_{bb} is larger for band 2 ($\frac{2}{3}$ against $\frac{1}{2}$). Yet every same-norm step that gives band 2 more weight is strictly worse than u: $\|c\| = \|u\|$ and $c_2 > 1$ imply $q(c) > q(u)$. The same-norm step $c = (\frac{31}{25}, \frac{17}{25})$, which gives band 1 more weight, is strictly better.*

On the circle, $q(c) - q(u) = \frac{1}{2}(1 - c_1) + (c_2 - 1)^2$. No cross-curvature is involved: at fixed norm what matters is the tangential part of $Hu - g$ (Theorem A.20), not the ratio g_b/H_{bb}.

Proposition A.28 (A one-step shape preference can reverse over two steps)**.** *Let $F(x) = \frac{1}{2}x_1^2 + \frac{1}{4}x_2^2$, whose gradient step of size 1 maps x to $(0, \frac{1}{2}x_2)$, and let $V(x) = F(0, \frac{1}{2}x_2)$ be the loss after one more such step. The points $A = (\frac{142}{61}, \frac{98}{61})$ and $B = (\frac{98}{61}, \frac{142}{61})$ have the same norm as $x = (2,2)$. Then $F(B) < F(x) < F(A)$ but $V(A) < V(x) < V(B)$.*

Removing the step-size term therefore does not protect a one-step shape preference from horizon bias. This is why we treat the measured shape as a hypothesis for training experiments rather than as a prescription.

A.4 Input-whitened Muon

For a whitening factor $P \in \mathbb{R}^{n\times n}$, the whitened direction is $\mathcal{W}(G;P) = \mathrm{NS}(GP)\,P^\top$. With $PP^\top = C^{-1}$ it is the steepest direction for the data-dependent norm $\|DC^{1/2}\|_{\mathrm{op}}$.

Theorem A.29 (Only the metric matters)**.** *If $P_1P_1^\top = P_2P_2^\top$ and P_2 is invertible, then $\mathcal{W}(G;P_1) = \mathcal{W}(G;P_2)$. The whitened step depends on C only, not on the square root chosen.*

Proof idea. $P_1 = P_2Q$ with Q orthogonal, and Theorem A.3 moves Q through NS.

Theorem A.30 (Exact $\mathrm{GL}(n)$-equivariance)**.** *Reparametrize the inputs by an invertible M: $x' = Mx$, $W' = WM^{-1}$, $G' = GM^\top$, $C' = MCM^\top$. (i) If $PP^\top = C^{-1}$ and P' is invertible with $P'P'^\top = C'^{-1}$, then $\mathcal{W}(G';P') = \mathcal{W}(G;P)\,M^{-1}$. (ii) The full step with decoupled weight decay is equivariant too:*

$$(1-\eta\lambda)WM^{-1} - \eta\,\mathcal{W}\big(GM^\top; (M^\top)^{-1}P\big) = \big[(1-\eta\lambda)W - \eta\,\mathcal{W}(G;P)\big]M^{-1}.$$

The trained function is therefore independent of the input coordinates.

© 2026 JP Morgan Chase & Co. All rights reserved

Theorem A.31 (Steepest directions transfer through a change of metric)**.** *Let P be invertible and $\mathcal{N}$ any norm-like function. If D is a steepest direction for $\mathcal{N}$ at GP, then $DP^\top$ is a steepest direction at G for the pulled-back function $\Delta \mapsto \mathcal{N}\big(\Delta(P^\top)^{-1}\big)$.*

Theorem A.32 (Whitened Muon is conditioning-free on least squares)**.** *Let $L(W) = \frac{1}{2}\operatorname{tr}\big((W - W_\star)C(W - W_\star)^\top\big)$ with $C = RR^\top$ and R invertible. Run whitened Muon with $P = (R^\top)^{-1}$ and any step sizes $\eta_0, \eta_1, \ldots$. Then the whitened error $E_t = (W_t - W_\star)R$ satisfies*

$$E_{t+1} = E_t - \eta_t \,\mathrm{NS}(E_t), \qquad L(W_t) = \tfrac{1}{2}\|E_t\|_{\mathrm{F}}^2,$$

which is plain Muon on the isotropic loss $\frac{1}{2}\|E\|_{\mathrm{F}}^2$, for every C, however ill-conditioned.

Proposition A.33 (Plain Muon is not $\mathrm{GL}(n)$-equivariant)**.** *For $G = (1, 1)$ and $M = \mathrm{diag}(2, 1)$: if plain Muon's step were the same in both coordinate systems, $\mathrm{NS}(GM^\top)M = \mathrm{NS}(G)$, then $\mathrm{NS}(G) = 0$.*

A.5 Energy-matched decay

Ordinary decoupled decay $W \leftarrow W - \eta\lambda W$ shrinks every singular value by the same factor. Energy-matched Schatten-1 decay replaces λW by $\lambda\,(\|W\|_{\mathrm{F}}^2/\langle W, P\rangle)\,P$ with P the polar factor of W. Since $\langle W, P\rangle = \|W\|_* = \sum_i \sigma_i$, every singular value then loses the same absolute amount $c_1 = \lambda \sum_j \sigma_j^2 / \sum_j \sigma_j$.

Theorem A.34 (Energy match)**.** *(i) If $\sum_j \sigma_j \neq 0$, then $\sum_i \sigma_i\, c_1 = \sum_i \sigma_i\,(\lambda\sigma_i)$: subtracting c_1 from every singular value removes the same first-order energy as ordinary decay. (ii) More generally, for matrices W and D with $\langle W, D\rangle \neq 0$, the rescaled direction $\tilde{D} = (\|W\|_{\mathrm{F}}^2/\langle W, D\rangle)\,D$ satisfies $\langle W, \tilde{D}\rangle = \langle W, W\rangle$ and $\langle W, \tilde{D} - W\rangle = 0$. The energy-matched step is ordinary decay plus a step orthogonal to W.*

Proposition A.35 (Which singular values shrink more)**.** *Let $\lambda > 0$ and $\sigma_i > 0$. Under energy-matched Schatten-1 decay, σ_i shrinks by a larger fraction than under ordinary decay ($c_1/\sigma_i > \lambda$) if and only if $\sigma_i < \sum_j \sigma_j^2 / \sum_j \sigma_j$. If $c_1 > 0$ and $0 < \sigma_i < \sigma_j$, then $c_1/\sigma_j < c_1/\sigma_i$: smaller singular values always lose a larger fraction.*

The decay therefore prunes the bulk of the spectrum and spares its top. That is the opposite of Schatten-4 decay, which hurt in our experiments while Schatten-1 helped.

A.6 Results from earlier formal work cited in Section 7

These results come from the earlier, temporal-filter part of the project. Most are explicit counterexamples or elementary identities, and we state them for what they rule out.

Proposition A.36 (Moments do not determine the momentum preference under NS_5)**.** *(i) Any risk that is affine in the first two moments of a finite noise law takes the same value on two laws with equal mass, mean and second moment. (ii) There are two explicit laws for the amplitude of rank-two gradient noise with the following properties: equal means, equal second moments, identical temporal covariance of the noise streams, and identical linearized risk curves $\beta \mapsto R_{\mathrm{lin}}(\beta)$ for every momentum β. Yet under the actual five-step NS_5 map with stabilized normalization, the first law strictly prefers $\beta = 0.95$ to $\beta = 0.85$ and the second strictly prefers 0.85 to 0.95. (iii) Hence no rule that observes only the linearized risk curve can choose the better momentum for both laws.*

Proposition A.37 (Spatial orientation decides the momentum preference)**.** *In an explicit full-rank construction with two consecutive updates, the signal path, drift, noise energy and time dependence are all matched. The Euclidean block risk is identical for every β whether the noise is tangential or radial. After Newton–Schulz with any finite depth (including five), tangential noise makes $\beta = 0.95$ strictly better than $\beta = 0.85$, and radial noise makes $\beta = 0.85$ strictly better than $\beta = 0.95$.*

Proposition A.38 (First-order alignment can mis-rank finite steps)**.** *Let $L(x, y) = x + 2y + \frac{1}{2}(x^2 + 10y^2)$. The unit steps $A = (-1, 0)$ and $B = (0, -1)$ have equal length. At the origin, first-order analysis prefers B (directional decrease 2 against 1), yet $L(A) - L(B) = -\frac{7}{2}$, so A is the better finite step.*

Proposition A.39 (Curvature decides between the polar and gradient directions)**.** *Let a step d change a diagonal quadratic by $\Delta(d) = -\langle g, d\rangle + \frac{1}{2}\sum_i h_i d_i^2$, with gradient $g = (7, 1)$. The polar direction $d_P = (1, 1)$ and the norm-matched gradient direction $d_G = (7/5, 1/5)$ both have $\|d\|^2 = 2$. (i) With curvature $h = (1, 1)$, d_G is better. (ii) With $h = (6, 1)$, d_P is better. (iii) With $h = (4, 1)$, the norm-matched intermediate direction $(17/13,\ 7/13)$ is strictly better than both.*

© 2026 JP Morgan Chase & Co. All rights reserved

This is the two-direction version of Figure 1a: which spectral allocation is best depends on how curvature compares with gradient signal across directions.

Proposition A.40 (Greedy step-size facts)**.** *For a quadratic loss with gradient $W \mapsto HW - b$ and a step $W - \eta O$: (i) $\langle \nabla L(W - \eta O), O\rangle = \langle \nabla L(W), O\rangle - \eta \langle HO, O\rangle$, which is affine in η. (ii) With $a = \langle \nabla L(W), O\rangle$ and $q = \langle HO, O\rangle > 0$, the one-step loss $L(W - \eta O)$ decreases with η exactly when $\eta < a/q$. (iii) With no aligned signal ($a = 0$), every positive step looks too large to a greedy criterion. (iv) A bias added equally to the marginal value of every group does not change their comparison.*

B Diagnostic details and additional results

Scope. The diagnostic perturbs only the 48 hidden matrices, along exact-polar band directions or the actual NS_5 neuron groups. The AdamW update of embeddings and norms and the weight decay are left out. g and H are estimated on separate held-out folds, so they are estimates of a common local model with separate errors, not derivatives of one empirical loss. The realized-loss checks use a third fold. The inverse-based quantities ($c^\star$ and the shape shares) are sensitive because H is ill-conditioned. The fixed-norm band-filter predictions evaluate q at a given vector and need no inverse.

Procedure. For a checkpoint at step t (Muon, seed 43, screen schedule), we recompute the next training minibatch's gradient with the same clipping, form $N = G + \beta(\beta B + G)$ exactly as the optimizer would, and take the full SVD of N for each of the 48 matrices. Band directions use the exact polar factor $U_{I_b} V_{I_b}^\top$, scaled by the learning rate at step t and Muon's shape factor. Neuron groups use Muon's actual NS_5 step, split into six equal-count sextiles by row norm (tall or square matrices) or column norm (wide matrices). The held-out gradient uses 768 validation sequences (393k tokens). Each Hessian–vector product uses 192 further sequences (98k tokens), computed by double backward with the math attention kernel. Realized loss changes use 768 further sequences. The three sets are disjoint. For $B = 6$ groups and five group sets (all matrices, and each layer type), the diagnostic needs 30 Hessian–vector products per checkpoint.

Table 2: Band diagnostic over all 48 matrices. Top: per-band multiplier $c_b = g_b/H_{bb}$, best uniform rescaling $\gamma^\star$, second-order gap $q(\mathbf{1}) - q(c^\star)$ (nats) and its shape share (Theorem 3). Bottom: the joint optimum $c^\star = H^{-1}g$ and the condition number of H. $c^\star$ is irregular because H is ill-conditioned, but it also puts nearly all its weight on the lower half of the spectrum. All multipliers scale as $1/\eta$; η at step 2,048 is $10\times$ smaller.

Step	1–8	9–32	33–96	97–256	257–512	513–768	$\gamma^\star$	gap	shape share
1,024	0.59	1.96	2.46	4.27	12.0	28.3	0.52	0.00674	17%
1,536	0.55	1.90	2.39	4.21	9.52	21.0	0.49	0.00532	12%
2,048	0.98	7.67	16.4	31.6	75.8	186.3	1.52	0.00137	97%
1,024 ($c^\star$)	0.53	0.18	−0.10	0.31	5.92	16.4	$\mathrm{cond}(H) \approx 1 \times 10^4$		
1,536 ($c^\star$)	0.48	0.34	−0.15	0.88	4.01	11.5	$\mathrm{cond}(H) \approx 8 \times 10^3$		
2,048 ($c^\star$)	0.40	0.12	3.16	8.05	34.8	115.9	$\mathrm{cond}(H) \approx 4 \times 10^3$		

Table 3: Quadratic-model calibration (bands, all matrices): predicted versus realized change in held-out loss (393k disjoint tokens) for Muon's step ($c = \mathbf{1}$) and the model-optimal band shape ($c = c^\star$).

Checkpoint	Muon, predicted	Muon, realized	$c^\star$, predicted	$c^\star$, realized
1,024	−0.00079	−0.00054	−0.00753	−0.00730
1,536	+0.00026	+0.00073	−0.00506	−0.00471
2,048	−0.00028	−0.00046	−0.00165	−0.00183

Measured one-step losses at matched norm. To test the fixed-norm predictions with the actual loss rather than the quadratic model, we start from the same checkpoints and form each matrix's next Nesterov input. We take the exact tilted direction $\phi_5(\sigma)(\sigma/\sigma_{\mathrm{rms}})^p$, rescaled per matrix to the Frobenius norm of Muon's NS_5 step ($p = 0$ is Muon), or a norm-matched band filter. We then measure the held-out loss after one step on a fixed 393k-token set (Table 4). The slope dL/dp at $p = 0$ is a central finite difference. The best same-norm filter is the step-1,536 model optimum. One checkpoint seed and one validation subset are used, without intervals.

© 2026 JP Morgan Chase & Co. All rights reserved

Table 4: Measured change of the held-out loss ($\times 10^{-3}$) after one step at Muon's own update norm: the exact tilt σ^p for several p, halving ranks 1–8, and the best same-norm filter, plus the slope dL/dp at $p = 0$.

Step	$p = -\frac{1}{2}$	$-\frac{1}{4}$	-0.1	0 (Muon)	$+0.1$	$+\frac{1}{4}$	ranks 1–8 $\times\frac{1}{2}$	best same-norm	$dL/dp\|_0$
1,024	-3.07	-5.75	-5.93	-1.76	$+12.08$	$+80.45$	-6.14	-6.28	$+0.090$
1,536	-2.48	-4.13	-3.99	-1.03	$+7.95$	$+48.44$	-4.27	-4.40	$+0.060$
2,048	-0.13	-0.26	-0.32	-0.33	-0.38	-0.22	-0.29	-0.21	-0.000

Fixed-norm analysis. Band directions $D_{\ell,b} = \eta s_\ell U_{I_b} V_{I_b}^\top$ are mutually Frobenius-orthogonal, with $\|D_b\|_F^2 \propto |I_b|$ (every hidden matrix has rank 768). Every matrix's step therefore has the same norm as Muon's exactly when $c^\top N c = \mathbf{1}^\top N \mathbf{1}$ with $N = \mathrm{diag}(8, 24, 64, 160, 256, 256)$. The best same-norm step minimizes q on this ellipsoid. With $\tilde{H} = N^{-1/2} H N^{-1/2}$ and $\tilde{g} = N^{-1/2} g$, it is $c = N^{-1/2}(\tilde{H} + \lambda I)^{-1}\tilde{g}$, with $\lambda > -\lambda_{\min}(\tilde{H})$ chosen by a one-dimensional root search so that the constraint holds. Norm-matched filters are rescaled onto the ellipsoid before evaluation. Table 5 gives the predictions.

Table 5: Second-order predictions at Muon's own update norm (bands, all 48 matrices). Column 2 is Muon's predicted one-step change $q(\mathbf{1})$. Columns 3–5 are $q(c) - q(\mathbf{1})$ for the best same-norm step, the norm-matched band filter used in training ($m_b \propto c_b^{1/4}$ at step 1,024), and halving ranks 1–8 with norm matching. The last column is the relative size of the tangential gradient p (Theorem A.20). Bottom: the best same-norm step's band multipliers.

Checkpoint	Muon's $q(\mathbf{1})$	best same-norm c	band filter	ranks 1–8 halved	$\|p\|/\|r\|$
1,024	-0.00079	-0.00589	-0.00570	-0.00527	1.00
1,536	$+0.00026$	-0.00488	-0.00478	-0.00443	1.00
2,048	-0.00028	-0.00009	$+0.00005$	$+0.00001$	0.79

Best same-norm c at step 1,024: (0.54, 0.22, 0.21, 1.26, 1.26, 0.62)
Best same-norm c at step 1,536: (0.49, 0.35, 0.17, 1.38, 1.17, 0.63)
Best same-norm c at step 2,048: (0.64, 2.92, 2.30, 1.05, 0.39, 0.17)

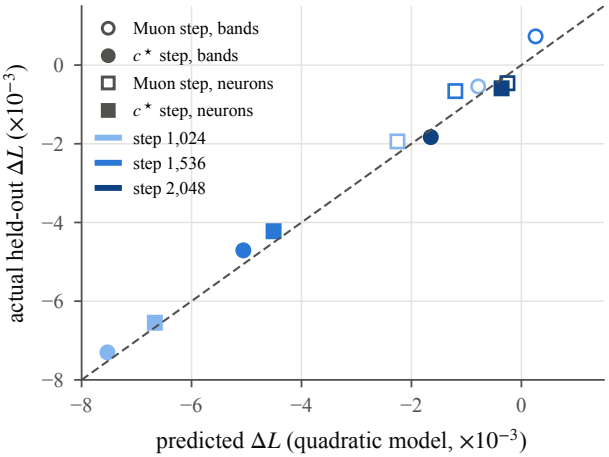


Figure 3: Predicted (quadratic model) versus realized held-out loss change for 12 test steps.

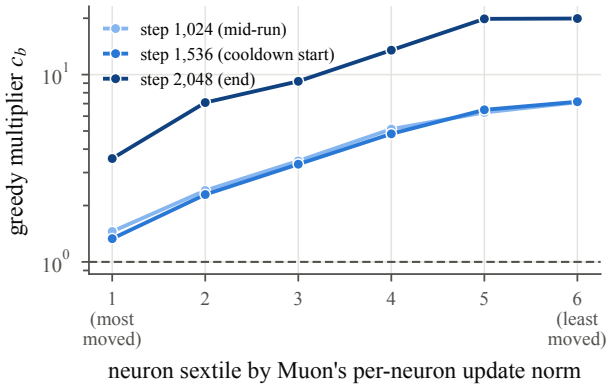


Figure 4: Greedy multiplier per neuron sextile, ranked by Muon's per-neuron update norm. The most-moved neurons sit in higher-curvature directions.

Implied spectral exponent. To relate the band multipliers to a power law $f(\sigma) \propto \sigma^p$, we take each matrix's singular values at ranks 1, 8, 32, 96 and 256, normalize by the value at rank 96, and average them in log space over the 48 matrices. Each of the top four bands is represented by the geometric mean of its two edges, and we regress $\log c_b$ on $\log \sigma_b$. The slopes are -0.8, -0.8 and -1.4 at the three checkpoints. The estimate is crude: the band representatives are coarse, and the tail bands, whose lower edges we did not log, are excluded. It is also a greedy quantity, so it overstates how strong a tilt to use. We report it only to show that the prescribed sign is negative and of order one in magnitude. The deployed power, $-\frac{1}{4}$, is deliberately conservative.

© 2026 JP Morgan Chase & Co. All rights reserved

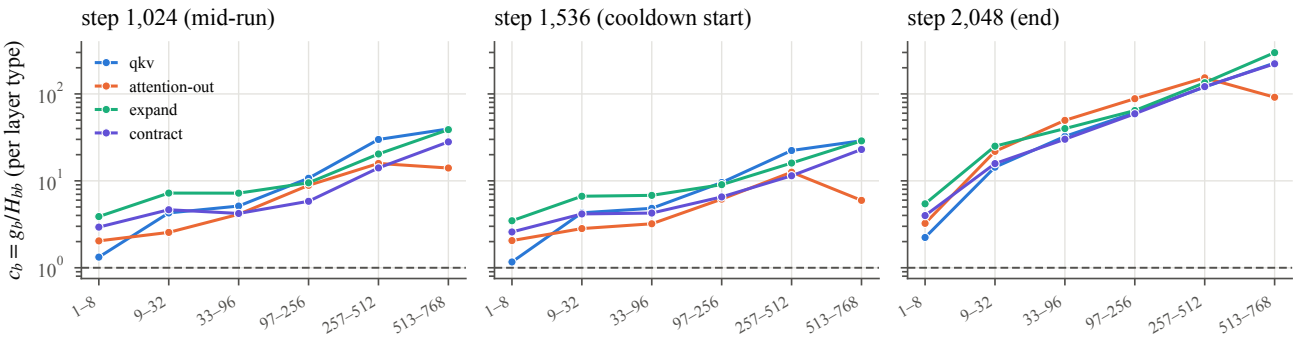


Figure 5: Band multipliers computed separately for each layer type. The top band has the smallest multiplier in every case.

Per-neuron spread. Table 6 reports the coefficient of variation (CV) of per-neuron update norms at the three checkpoints. Each entry is the mean over the 12 matrices of a type, along the axis that indexes neurons: rows for qkv, expand and attention-out, columns for contract. For comparison, the raw momentum's QKV row-norm CV is 0.93. Muon reduces the spread but does not remove it. Block-wise NS_5 steps leave a residual QKV spread (CV 0.22) although exact orthogonal blocks would have none (Theorem A.10). A plausible cause is that NS_5 maps the near-null singular directions of low-rank blocks to values well below 1; we have not verified this.

Table 6: Coefficient of variation of per-neuron update norms of Muon's step and of the block-wise polar step.

Layer (neuron axis)	Muon, step 1,024	Muon, step 2,048	Block-wise, step 1,024	Muon max/median, step 1,024
qkv (rows)	0.43	0.47	0.22	1.84
expand (rows)	0.18	0.18	0.03	1.72
contract (columns)	0.18	0.17	0.02	1.67
attention-out (rows)	0.02	0.02	0.02	1.05

C All screen runs

Table 7 and Figure 6 list every completed screen arm (2,048 steps, 50.3M tokens). Differences are paired against fused-QKV Muon at LR 0.005: the seed-42 difference for every arm, and the mean over seeds 42–44 with a 95% t-interval where three seeds exist. The arms were selected adaptively during the study; no multiplicity correction is applied.

Learning-rate robustness of block-wise steps. With square blocks, raising the learning rate from 0.005 to 0.007 and 0.01 gives final losses of 4.1840 and 4.2045. Muon at the same rates gives 4.2077 and 4.2376 (+0.0034 and +0.0332 relative to its optimum). Doubling the learning rate leaves square-block Muon level with Muon's tuned optimum. MuonBP reported that block-only training needed smaller learning rates at scale (Khaled et al., 2026); in our setting block-wise steps are, if anything, more tolerant of large rates.

D Implementation details

Norm matching. For every matrix and step, the candidate direction D is rescaled to $\|\mathrm{NS}_5(N)\|_{\mathrm{F}}$, with N the same momentum Muon would use, and then enters (1) with Muon's shape factor and learning rate. Norm-matched whitening arms are matched to $\|\mathrm{NS}_5(NP)\|_{\mathrm{F}}$, the norm of their own whitened polar factor. The trainer enables TF32 matrix products globally, so NS_5 runs in FP32 storage with TF32 products; only the Chebyshev filter disables TF32. Matching costs one extra joint NS_5 per matrix for block-wise arms, used only to read its norm. The deployable square-block step omits it: it scales wide-matrix column blocks by $1/\sqrt{k}$ and runs all 144 square blocks in one batched NS_5 call. It gives -0.0200 on the screen, against -0.0219 norm-matched.

NorMuon. Following Li et al. (2026), the orthogonalized update is divided row-wise by $\sqrt{v} + 10^{-8}$, where v is an EMA ($\beta_2 = 0.95$) of the row mean of $D \odot D$. It is then norm-matched. A stateless variant with $\beta_2 = 0$ gives -0.0192, against -0.0189 for the EMA version (screen), so the EMA is not what helps.

© 2026 JP Morgan Chase & Co. All rights reserved

Table 7: Screen results. Final held-out loss (4.18M tokens) at seed 42 and paired differences from fused-QKV Muon.

Family	Arm	Final (s42)	Δ (s42)	Δ, 3 seeds [95% CI]
	Muon, LR 0.005 (seeds 42/43/44: 4.2043, 4.2066, 4.2035)	4.2043	—	—
Learning rate	Muon, LR 0.0035	4.2179	+0.0136	
	Muon, LR 0.007	4.2077	+0.0034	
	Muon, LR 0.01	4.2376	+0.0332	
	Split QKV, LR 0.0035	4.1931	−0.0112	
	Split QKV, LR 0.007	4.1847	−0.0196	
	Square blocks, LR 0.007	4.1840	−0.0203	
	Square blocks, LR 0.01	4.2045	+0.0001	
Neuron balance	Split QKV (Q, K, V orthogonalized separately)	4.1788	−0.0255	−0.0267 [−0.0304, −0.0230]
	Per-head Q/K/V blocks (36 blocks of 64×768)	4.1871	−0.0172	
	Square blocks (QKV + MLP)	4.1824	−0.0219	−0.0238 [−0.0282, −0.0194]
	Square blocks, not norm-matched (deployable)	4.1843	−0.0200	
	Square blocks in the MLP only (QKV fused)	4.2082	+0.0039	
	NorMuon row normalization	4.1854	−0.0189	−0.0188 [−0.0193, −0.0183]
	Row normalization without EMA	4.1852	−0.0192	
Spectral shape	Cooldown tilt $\sigma^{-1/4}$ (exact)	4.1910	−0.0133	
	Cooldown tilt $\sigma^{-1/2}$ (exact)	4.1888	−0.0155	
	Tilt $\sigma^{-1/4}$ from step 1,024	4.1896	−0.0147	
	Cooldown band filter, $m_b \propto c_b^{1/4}$ from step 1,024	4.1927	−0.0116	
	Accurate polar factor (NS5 + 3 cubic steps)	4.2023	−0.0020	
	Square blocks + cooldown tilt (Chebyshev, monomial/TF32; faulty)	4.1804	−0.0240	
	Square blocks + cooldown tilt (Chebyshev, Clenshaw/FP32)	4.1721	−0.0322	
	Square blocks + cooldown tilt (exact)	4.1711	−0.0332	
Decay geometry	Schatten-1 decay (bulk shrinkage)	4.1991	−0.0052	−0.0057 [−0.0100, −0.0014]
	Schatten-4 decay (top-heavy)	4.2220	+0.0177	
Data metric (whitening)	Input whitening $(\frac{1}{4}, \frac{1}{4})$, norm-matched	4.2254	+0.0211	
	Input whitening $(\frac{1}{2}, \frac{1}{4})$, norm-matched	4.2642	+0.0599	
	Input whitening $(\frac{1}{2}, \frac{1}{2})$, δ=0.3, norm-matched	4.2474	+0.0431	
	Input whitening $(\frac{1}{2}, \frac{1}{2})$, data-metric magnitude	4.2567	+0.0524	
	Input whitening $(\frac{1}{2}, \frac{1}{2})$ + shaped decay	4.2638	+0.0595	
	Toner recipe $(\frac{1}{2}, \frac{1}{8})$ + shaped decay	4.1975	−0.0068	
Combinations	Tilt $\sigma^{-1/4}$ + row normalization (no EMA)	4.1737	−0.0306	
	Band filter (step-1,024 multipliers) + row normalization (no EMA)	4.1742	−0.0302	
	Square blocks + NorMuon	4.1780	−0.0263	
	Square blocks + NorMuon + cooldown tilt (exact)	4.1661	−0.0382	
	Square blocks + NorMuon + cooldown tilt	4.1665	−0.0378	
	... + Schatten-1 decay	4.1648	−0.0396	
	Split QKV + cooldown tilt	4.1667	−0.0376	
	Split QKV + NorMuon	4.1722	−0.0321	
	Split QKV + NorMuon + cooldown tilt	4.1632	−0.0411	−0.0436 [−0.0511, −0.0362]
	Split QKV + NorMuon + tilt from step 1,024	4.1625	−0.0418	
	Split QKV + NorMuon + cooldown tilt $\sigma^{-1/2}$	4.1641	−0.0402	
	Split QKV + NorMuon + cooldown tilt + Schatten-1 decay	4.1581	−0.0462	−0.0487 [−0.0555, −0.0420]

Spectral tilt. The exact tilt maps singular value σ_i of N to $\phi_5(\sigma_i/\|N\|_{\mathrm{F}}) \cdot (\sigma_i/\sigma_{\mathrm{rms}})^p$ via an eigendecomposition of $NN^\top$. With $p = 0$ it reproduces NS_5 to 2×10^{-4} relative error. The matmul-only tilt (Algorithm 1) uses a degree-12 Chebyshev least-squares fit of x^p on $[10^{-3}, 1]$ (4,000 Chebyshev nodes) in the variable $t = 2(x - x_{\mathrm{lo}})/(1 - x_{\mathrm{lo}}) - 1$. It is evaluated by Clenshaw's recurrence with TF32 disabled. $S = ON^\top$ is symmetrized and scaled by $\Lambda = 1.01$ times a 20-step power-iteration estimate of $\lambda_{\max}(S)$. That is not a proved upper bound: eigenvalues above Λ, or below $10^{-3}\Lambda$, fall outside the fitted interval, where the Chebyshev polynomial is not controlled. By Theorem A.5, the implemented filter is $f_i^{1+p}\sigma_i^p$ rather than the exact $f_i\sigma_i^p$; on the Newton–Schulz plateau they differ by a factor in $[0.96, 1.10]$. On real momenta the matmul tilt matches the exact tilt to direction cosine 0.998–0.999; enabling TF32 lowers the worst block to 0.992. By Theorem A.6(ii) the normalized steps differ by at most 6.3% in Frobenius norm. Turning that into a loss bound would need a curvature bound along the difference (Theorem A.24), which we do not have. When Q, K and V are split, the three 768×768 blocks of all 12 layers and the square attention-output matrices are filtered in one batched chain. The MLP matrices are filtered whole. With TF32 off the tilt costs about 0.17 s per step on an A6000, during the cooldown only.

Input whitening. C is an EMA of the uncentered input second moment, trace-normalized with damping, $\hat{C} = C/\overline{\lambda(C)} + \delta I$ ($\delta = 0.1$ unless stated). It is refreshed from 4,096-token activation samples. The step is $\mathrm{polar}(N\hat{C}^{-q_{\mathrm{in}}})\hat{C}^{-q_{\mathrm{out}}}$: $(\frac{1}{4}, \frac{1}{4})$ and $(\frac{1}{2}, \frac{1}{2})$ norm-matched, $(\frac{1}{2}, \frac{1}{2})$ at its own magnitude, and the Toner recipe $(\frac{1}{2}, \frac{1}{8})$ with shaped decay $\lambda W\hat{C}^{-1/4}$.

© 2026 JP Morgan Chase & Co. All rights reserved

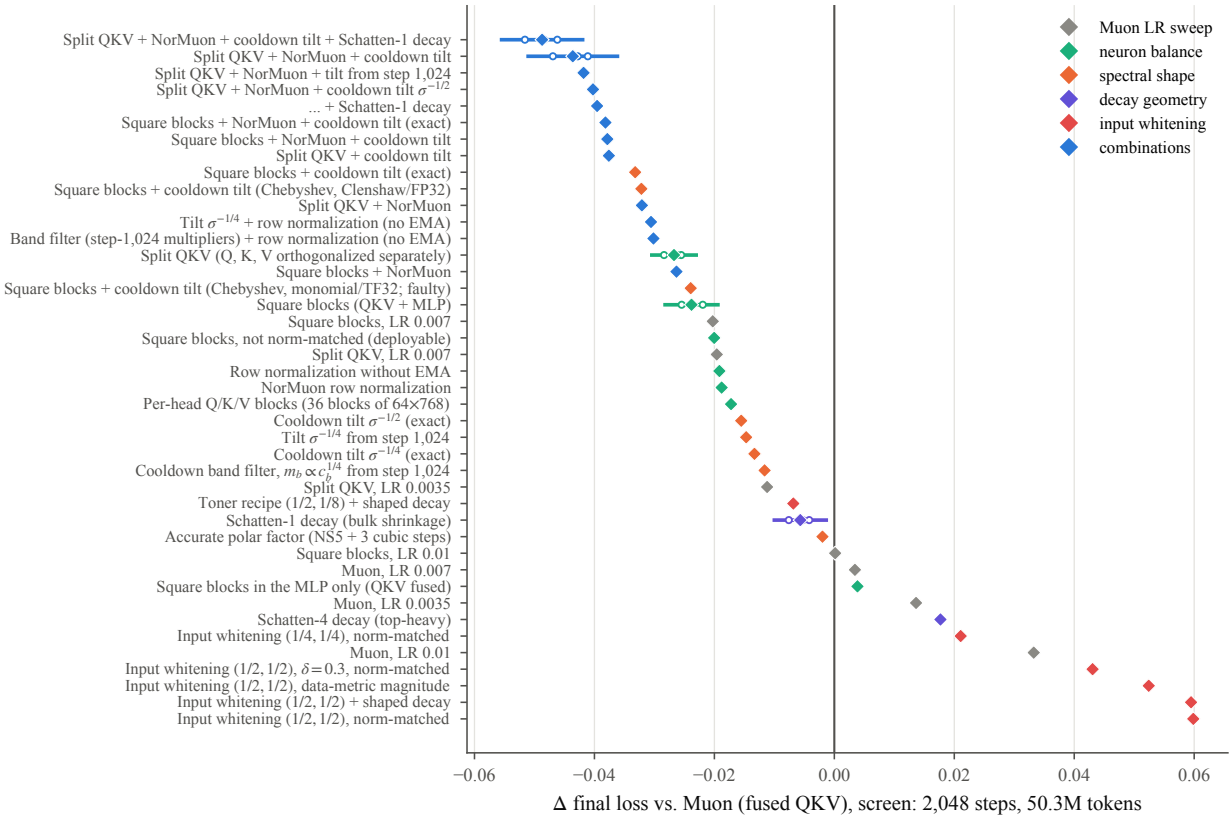


Figure 6: Screen arms sorted by their mean paired difference from fused-QKV Muon (seed 42, or the three-seed mean with 95% t-interval and open circles for seeds where three seeds exist).

Schatten-p decay. We replace λW by $\lambda \frac{\|W\|_F^2}{\langle W, W(W^\top W)^{(p-2)/2}\rangle} W(W^\top W)^{(p-2)/2}$. This removes the same Frobenius energy per step as uniform decay, with relative rate $\lambda \sigma_i^{p-2} \sum \sigma^2 / \sum \sigma^p$ on singular value σ_i. For $p = 1$ we use $\mathrm{NS}_5(W)$ for the polar factor.

Cost. Measured seconds per step on one RTX A6000 (100.7M-token runs): Muon 0.343; NorMuon 0.329; norm-matched square blocks 0.415; square blocks + NorMuon + tilt 0.418 before and 0.588 during the cooldown. Mean wall time per 100.7M-token run, including evaluations: Muon 23.3 min, split-QKV 24.9, split-QKV + tilt 26.4, split-QKV + NorMuon 25.1, split-QKV + NorMuon + tilt 27.8. Three runs shared a host, so these timings are indicative.

E Reproducibility and provenance

Logging. Each run writes an event log with its recipe, optimizer variant, seed, software versions and a hash of the optimizer source. The results tables, figures and run-derived numbers in the results text are generated automatically from these logs. A few descriptive quantities (per-step timings, curvature ratios, some leverage-table entries and filter cosines) were transcribed and checked in an independent pass against the logs. Training uses `torch.compile` and BF16 autocast for the forward pass; the optimizer state is FP32, and Newton–Schulz runs in FP32 with TF32 matrix products. The trainer reproduces an earlier, slower harness's Muon loss (3.9445 against 3.9463 at seed 42).

Order of experiments. The screen arms, the diagnostic and the first confirmations ran in parallel on three GPUs over one day. The anatomy was *not* pre-registered: the square-block, whitening and accurate-polar screens started before the curvature diagnostics, and the neuron diagnostic postdates the first block arms. The square-block arm was motivated by the leverage analysis; the attribution runs later showed that the QKV split carries its gain. Two confirmations, square blocks + NorMuon + tilt and split QKV + NorMuon + tilt, were launched before their own screens had finished. The 355M seed-42 pair was launched after the positive 124M confirmation. The gating rule of Section 6 was adopted after that; it released the full-length factorial ablations (split QKV + tilt; split QKV + NorMuon), whose screens beat the baseline by 0.0121 and

© 2026 JP Morgan Chase & Co. All rights reserved

Table 8: Theory versus training. What the Lean statements certify, whether the anatomy implies a gain, and the measured paired Δ against fused-QKV Muon (screen, one seed, unless marked). Formal properties alone did not predict which changes help.

Change	Certified property	Anat.	Δ
Split QKV	steepest for block max-norm; exact-polar blocks have unit rows; energy identity	yes	$-0.0250^{\dagger}$
MLP square blocks	same	weak	$+0.0039$
Cooldown tilt	keeps Muon's $O(n)$-equivariance	yes	$-0.0120^{\dagger\ddagger}$
Accurate polar	—	no	-0.0020
Input whitening	$\mathrm{GL}(n)$-equivariance; conditioning-free LS (undamped form)	—	$+0.021$–0.060
Toner recipe	(prior art)	—	-0.0068
Schatten-4 decay	—	—	$+0.0177$
Schatten-1 decay	—	—	$-0.0057^{\S}$

†100.7M tokens, $n = 3$. ‡vs. split-QKV Muon. $^{\S}n = 3$.

Algorithm 1 One step for a hidden matrix: split-QKV Muon with NorMuon row normalization and the cooldown tilt (`cooldown` is true in the last 25% of steps).

$B \leftarrow \beta B + G;\ N \leftarrow G + \beta B$
Split N into blocks N_j (Q, K, V for the attention input matrix; one block otherwise)
for each block N_j **do**
  $O_j \leftarrow \mathrm{NS}_5(N_j)$
  if `cooldown` **then**
    $S \leftarrow \mathrm{sym}(O_j N_j^{\top});\ \lambda \leftarrow 1.01\,\lambda_{\max}(S)$ (20 power iters)
    $D_j \leftarrow \mathrm{ClenshawCheb}_{x^{-1/4}}(S/\lambda)\,O_j$ (FP32)
  else
    $D_j \leftarrow O_j$
  end if
end for
$D \leftarrow [D_1; \ldots];\ v \leftarrow \beta_2 v + (1-\beta_2)\,\mathrm{rowmean}(D \odot D);\ D \leftarrow D/(\sqrt{v} + \epsilon)$
$D \leftarrow D \cdot \|\mathrm{NS}_5(N)\|_{\mathrm{F}} / \|D\|_{\mathrm{F}}$ (norm match)
$W \leftarrow (1 - \eta\lambda)W - \eta\, s\, D$

0.0066, and the 355M seeds 43–44.

Source versions. The optimizer file changed during the study. The original fused-QKV Muon, NorMuon and square-block confirmations span four source hashes, and a read-only audit of the raw logs could not certify, from the retained files, that the changes touched only inactive branches. All split-QKV arms (split-QKV Muon, the factorial and the 355M runs) ran on a single source version, so the factorial comparisons are unaffected. Every comparison involving fused-QKV Muon, NorMuon on fused QKV, square blocks or square blocks + NorMuon + tilt crosses source versions; this includes the attribution results "square blocks vs. split QKV" and "NorMuon vs. split QKV". The A/A bound and the reproduction of the earlier harness's loss suggest the effect is small, but we cannot certify it.

Measurements needed to instantiate the certificate. Experiments were paused before these could run. Each uses saved checkpoints and no training: (1) *Sampling intervals*: record per-sequence held-out losses for every step in Table 4 and bootstrap paired differences over sequences, across several checkpoint seeds. (2) *Estimation errors* δ_g, δ_H: recompute g and H on disjoint folds and take the spread of $\hat{g}$ and of $c^{\top}\hat{H}c$ over folds. (3) *Third-derivative bound* M: evaluate the loss at $t \in \{0, \pm\frac{1}{2}, \pm 1\}$ along each step and take third finite differences. (4) *Gradient bound* G: the held-out gradient norm at both endpoints and the midpoint of each step pair.

Statistics. Intervals are two-sided 95% Student-t intervals over three paired seeds (2 degrees of freedom). They are nominal: the same seeds were used in the screen, and no correction is made for the number of arms examined. The A/A bound (0.009 nats) uses 98k-token intermediate evaluations; final evaluations use 4.18M tokens and are less noisy.

© 2026 JP Morgan Chase & Co. All rights reserved