Capped-Grade Traverse Schedules Price as a Learning Rate: Traverse Amplitude Is Free, Slow Legs Are Not

A Preprint

Caitlyn Meeks
AuxiLab Tenerife — research@auxi.cafe

August 11, 2026

registered 2026-08-11 | internal run log #10–#11 | 2 seeds × 2 feature sets, 18 fits each | replicated before publication | promoted 2026-08-11

Abstract

This study extends the preceding study of normalized fixed-stride descent with a capped-grade traverse schedule: a unit step tilted downhill at a fixed grade, traversing orthogonally to the descent direction and reversing that traverse direction every L steps. The question is whether traverse amplitude is a second stability axis, independent of the learning rate. Three results each hold their registered inequalities on both seeds, although the registered headline prediction itself missed, informatively. (1) At matched effective downhill stride η_{eff} = s·grade, small-amplitude traverse schedules are **indistinguishable from straight descent**: all 16 small-amplitude cells (2 seeds × 8) tie within ±0.01 bits — the grade cap prices as a learning rate and nothing else. (2) The registered traverse-cost prediction missed, informatively: the cost attaches to **slow** legs, not to wide ones. Reversing direction every 64 steps costs +0.076–0.090 bits on both feature sets, while stride-built amplitudes as large as those of the penalized cells are free. (3) On the ill-conditioned feature set, wide *fast* traverses **beat** their matched straight descent by 0.010–0.025 bits in all four cells — orthogonal dither recovers part of the large-stride oscillation penalty, at the same nominal step size at which the preceding study observed a divergence that the program's tripwire failed to catch.

Keywords stochastic optimization · normalized gradient methods · constrained descent directions · learning-rate parameterization · gradient noise · pre-registration

---

§1 Motivation

A descent direction need not be the steepest available one. A capped-grade traverse schedule constrains every step to make a fixed angle with the momentum direction, spending the remaining component of the step on a direction orthogonal to it and reversing that orthogonal direction at a fixed period. The preceding study established what normalized fixed-stride steps buy on the program's two feature sets; its notes named, but did not evaluate, the capped-grade variant. This study evaluates it, to answer a precise question: at *matched effective downhill rate*, does the orthogonal component cost anything? Is traverse amplitude a second stability axis — or does the grade cap simply reparameterize the learning rate?

§2 Method

The fitting procedure is the program's reference logistic readout, verbatim (27-way softmax, batch 8192, gradient clip 1.0, exponentially weighted moving-average momentum [1] β=0.9, validation early stopping, and the divergence tripwire — a fit-invalidating rule that voids any run whose validation loss blows up during training). The capped-grade traverse step is a *unit* step tilted toward the descent direction at fixed grade:

```
# m̂ = normalized momentum (the descent direction)
# û = persistent traverse direction, re-projected every step
#     (orthogonalized against m̂, renormalized)
# σ ∈ {+1, −1} reverses every L steps
d = grade·m̂ + √(1 − grade²)·σ·û
W ← W − s·d
# effective downhill stride:  η_eff = s·grade
# grade = sin θ, so cos θ = √(1 − grade²)
# traverse amplitude:         A = s·L·cos θ
```

A *traverse leg* is the run of L steps between two consecutive direction reversals. At grade = 1 the method reduces *bitwise* to the promoted probe-free normalized fixed-stride descent of the preceding study (its `noprobe` arm); this was verified before registration, in that the grade-1 pilot run reproduces the on-disk values to four decimals on both feature sets. Straight-descent anchors at matched η_{eff} therefore serve as exact controls, and every cell is scored as Δ = traverse validation bits per character (bpc) – anchor validation bpc at the same η_{eff}, within feature set and seed.

Feature sets as in the preceding study. The **well-conditioned** set is the program's reference sparse-tanh echo state network (ESN) features [3] (spectral radius ρ=0.8); the **ill-conditioned** set is the addressed-memory ESN's *linear-activation* variant at ρ=0.95 — the worst-conditioned features in the program's inventory (see whitepaper 01). Standard text8 splits [4], 2M training characters. Nine cells per feature set and seed: six matched-η_{eff} (s, grade) pairs spanning amplitudes A ≈ 0.06 to 15.9 at L=16, two leg-length cells (L=4 and L=64 at s=0.1, grade=0.1), and a fresh grade-1 anchor at s=1e-2.

§3 Pre-registered predictions

Five ordinal predictions and two quantitative bands were registered at seed 0, each with a named alternative. As in the preceding study, the finding arrived partly through the misses.

- **S1** (the grade cap as a reparameterized learning rate): ≥7 of the 8 small-amplitude cells (A ≤ 0.7) tie their matched anchor within ±0.01 bits. (Hit, 8/8.)
- **S2** (the traverse cost — the registered headline): all 6 large-amplitude cells (A ≥ 6) pay > +0.01 bits. Named loss: wide traverses tie or win, which would make the traverse schedule a reparameterized learning rate at every amplitude. (Missed — *and the named loss's own reading was falsified too*; see §4.)
- **S3** (amplitude collapse): a stride-built and a leg-built cell of matched amplitude land within a stated band of each other — A is the controlling parameter, not stride or leg length separately.

(Half-falsified, decisively on the well-conditioned set.)

- **S4** (ordering across feature sets): the ill-conditioned set penalizes traversing more heavily at every large-amplitude setting. (Missed, 0/3 — the ordering *reversed*.)
- **S5** (stability): zero divergence-tripwire events among the ten non-large-amplitude cells. (Hit; in fact 0/18 overall.)

The seed-1 rerun then registered the surviving structure as R1–R4 with a promotion rule stated before running: R1 (both stride-built wide cells on the ill-conditioned set improve on their anchor, $\Delta < -0.01$), R2 (both L=64 cells penalized, $\Delta > +0.03$), and R3 (small-amplitude ties, ≥7/8) all holding would promote the finding — that is, enter it in the program's confirmed-findings record; any split result would be recorded without promotion. R4 registered within-η_{eff} monotonicity of the penalty on the well-conditioned set. Bands RQ1–RQ3 placed numeric bounds on the improvement, the penalty, and the fresh anchors.

---

§4 Results

Table 1. Seed 0 grid (internal run log #10). Δ = traverse validation bpc − matched-η_{eff} straight-descent anchor. Zero divergence-tripwire events in 18/18 cells.

cell (s, grade, L)	A	η_{eff}	well-cond. val / Δ	ill-cond. val / Δ
(4e−3, .25, 16)	0.062	1e−3	3.0169 / +0.0004	3.1752 / +0.0003
(1e−2, .1, 16)	0.159	1e−3	3.0191 / +0.0026	3.1769 / +0.0020
(1e−1, .1, 4)	0.40	1e−2	2.7261 / +0.0011	3.1671 / −0.0001
(4e−2, .25, 16)	0.62	1e−2	2.7273 / +0.0023	3.1671 / −0.0001
(1e−1, .1, 16) mid	1.59	1e−2	2.7354 / +0.0104	3.1688 / +0.0016
(4e−1, .25, 16)	**6.20**	**1e−1**	**2.6395 / +0.0011**	**3.2256 / −0.0185**
(1e−1, .1, 64)	**6.37**	**1e−2**	**2.8131 / +0.0881**	**3.2480 / +0.0808**
(1, .1, 16)	**15.92**	**1e−1**	**2.6439 / +0.0055**	**3.2221 / −0.0220**

The bolded rows carry the study. S1 hit 8/8 — every small-amplitude cell ties its anchor, max |Δ| = 0.0026. S2 missed: only the two L=64 cells (of six large-amplitude cells) show the predicted penalty, and its named loss — a reparameterized learning rate at *every* amplitude — is contradicted by those same L=64 penalties. The registered variable A splits *by mechanism*: amplitude built from a larger stride is free or better, while amplitude built from longer legs is penalized. S3's well-conditioned half was decisively falsified: stride-built A=6.20 and leg-built A=6.37 differ by 0.1736 bits against a 0.03 band — a penalty ratio of roughly 80× at matched amplitude. (Its ill-conditioned half fell inside the registered band but was flagged in the run journal as a band-width artifact: the two cells sit on opposite sides of zero.) S4 reversed outright: at every large-amplitude setting the ill-conditioned set penalized traversing *less* than the well-conditioned one, and at both stride-built settings the ill-conditioned Δ was negative — traversing *helped*. That improvement — two independent cells, same sign, well outside the tie band — was the run's surprise, and was held provisional under the program's standing rules. Of the seed-0 bands, the fresh-anchor band hit 2/2; the amplitude power-law band was unjudgeable by its own registered gate (only two qualifying cells, where the gate required three).

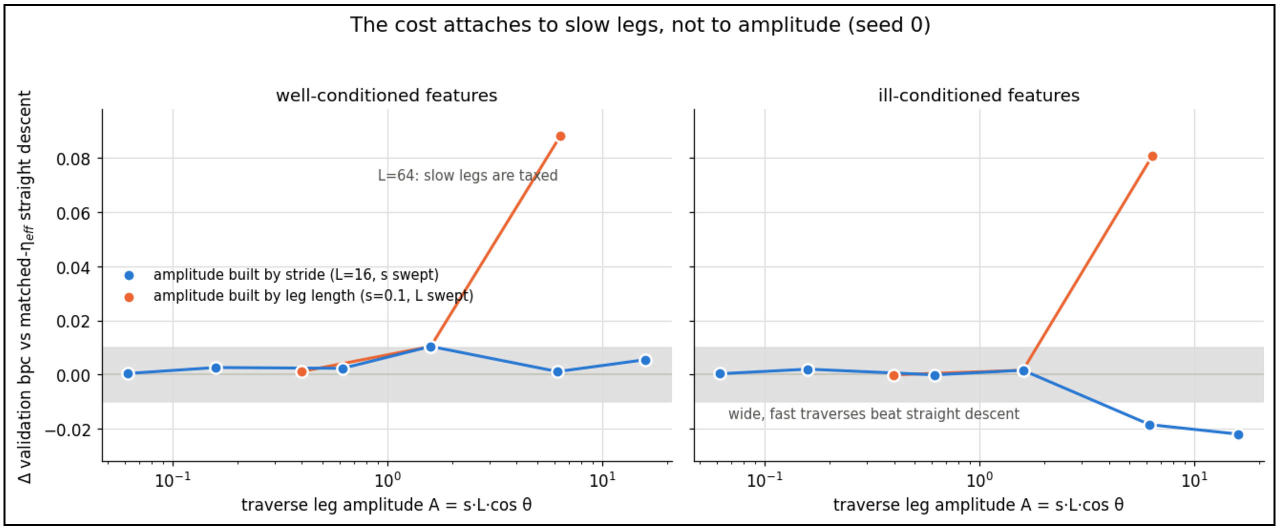

Figure 1. The penalty follows leg length, not amplitude (seed 0, internal run log #10). Both panels: Δ validation bpc relative to the matched-η_{eff} straight-descent anchor, against traverse amplitude A = s·L·cos θ (log scale); shaded: the tie band. Blue: amplitude built by stride (L=16, s swept) — inside the band everywhere on the well-conditioned set, and *negative* (better than straight descent) at the two large-A ill-conditioned cells (−0.0185, −0.0220). Orange: amplitude built by leg length (s=0.1, L swept) — the two L=64 cells pay +0.0881 (well-conditioned) and +0.0808 (ill-conditioned).

Table 2. The registered replication, seed 1 (internal run log #11): 4/4 ordinals, 3/3 bands. Fresh reservoir draws and batch order; 0/18 divergence-tripwire events.

claim	seed 0	seed 1
R1 ill-cond. improvement, stride-built wide cells (Δ)	−0.0185 · −0.0220	−0.0102 · −0.0251
R2 leg-length penalty, L=64 (Δ well-cond. · ill-cond.)	+0.0881 · +0.0808	+0.0897 · +0.0756
R3 small-amplitude ties (max |Δ| of 8)	0.0026	0.0022
R4 well-cond. penalty monotone within η_{eff}=1e-2	+0.0881 > +0.0104 > +0.0017*	+0.0897 > +0.0100 > +0.0015*

(*The last term is the mean of the two smallest-amplitude cells at that η_{eff}, as registered; the mean absorbs the noise floor.) One disclosure, recorded as registered: seed 1's weaker improving cell landed at Δ = −0.0102, only 0.0002 inside the registered −0.01 threshold — a thin margin, flagged in the run journal and judged as written. Promotion followed the pre-stated R1∧R2∧R3 rule on two independent passes.

Two scope notes. First, the improvement is *of the operating point, not of the feature set*: it is not a claim that the traverse schedule is the best method on the ill-conditioned features. The preceding study's best ill-conditioned run (NFSD, 3.1647) still leads that feature set, and the improved cells (3.2221–3.2256) sit far above both it and probe-free normalized fixed-stride descent at η=1e-3 (3.1749). What the wide fast traverse improves is the large-stride operating point specifically (straight descent at η_{eff}=0.1 sits at 3.2441 / 3.2437 on the two seeds). Second, the candidate collapse variable for the leg-length penalty — the leg aspect ratio A/η_{eff} = L·cos θ/grade, ≈637 in the penalized cells against

62–159 in the free wide cells — was registered only as the within-η_{eff} monotonicity above (R4, both seeds). The aspect-ratio law itself, and the decomposition of the orthogonal dither, are named follow-ups, not results.

§5 Discussion

The grade cap prices as a learning rate. At matched effective downhill stride, the grade cap does nothing that the stride had not already done — sixteen of sixteen small-amplitude cells tie their straight-descent controls. Whatever intuition suggests about the stability of a shallow grade, on this convex fitting problem it prices as η_{eff} = s·grade and nothing else.

The cost attaches to leg length, not to amplitude. Where the traverse does cost bits, the cost tracks how *long* a leg runs relative to the descent rate, not how far the iterate is displaced orthogonally: legs of 64 steps pay +0.076–0.090 bits on both feature sets and both seeds, while stride-built displacements of the same amplitude are free. A plausible mechanism, recorded in the run journal as unregistered conjecture: a long leg gives the momentum estimate time to absorb the orthogonal displacement — the restoring gradient produced by that displacement enters the momentum estimate, and part of the descent budget is spent undoing the traverse. A short leg reverses before this contamination accumulates.

On ill-conditioned features, fast wide traverses help. The ill-conditioned feature set was predicted to penalize traversing most heavily; instead, fast wide dither *recovered* 0.010–0.025 bits of the ≈0.07-bit large-stride oscillation penalty, in all four cells across both seeds — at the same operating point at which the preceding study observed a divergence that the tripwire never caught. An orthogonal high-frequency perturbation appears to act as a smoothing term on oscillating dynamics, related to the broader observation that noise deliberately injected during training can act as a regularizer [2]. Both directional expectations were therefore inverted: amplitude is not the destabilizing quantity, leg length is; and on the worst-conditioned features, a step that is partly orthogonal to the descent direction outperforms one that is not.

§6 References

1. B. T. Polyak, "Some methods of speeding up the convergence of iteration methods," *USSR Computational Mathematics and Mathematical Physics* 4(5), 1964.
2. N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *Journal of Machine Learning Research* 15(56), 2014.
3. H. Jaeger, "The 'echo state' approach to analysing and training recurrent neural networks," GMD Report 148, German National Research Center for Information Technology, 2001.
4. M. Mahoney, "Large text compression benchmark" (text8), mattmahoney.net/dc/textdata.

§7 Provenance

- **00764148** — the design brief and main grid; 18 fits plus a 6-fit pilot run, seed 0 (internal run log #10). Registered before running; the grade-1 identity gate and the pilot run's divergence behaviour were disclosed in the registration.

- **9e62e253** — full 18-fit grid on seed 1, fresh reservoir draws and batch order (run internal run log #11). Promotion by the pre-stated R1∧R2∧R3 rule on two independent passes.

Audit status: independently audited 2026-08-12, passed — 30 checks, 0 failures. Every difference, amplitude and leg aspect ratio was recomputed from the raw step-size, grade and leg-length settings rather than read from the stored summaries, and the registered anchor pairing rule was rebuilt from scratch: small-amplitude maximum $|\Delta|$ 0.0026 (16/16), leg tax +0.0881 / +0.0808 / +0.0897 / +0.0756, ill-conditioned rescue −0.0185 / −0.0220 / −0.0102 / −0.0251, aspect ratios 636.8 against 62.0 / 159.2, and the within-η monotonicity holding on both seeds — every promoted number exact. The identity between the grade-1 traverse and the probe-free straight descent, which is what licenses using one script's cells as anchors for the other, was verified line by line *and* by the smoke-budget identity cells agreeing on every entry of the per-epoch validation trace rather than only on the final number. The auditor's main attack was winner's curse: validation bpc is a minimum over epochs on the same split early stopping uses, and the rescue cells' traces oscillate by 0.14 and 3.22 bpc against their anchor's 0.02 — exactly the shape needed to manufacture a 0.02-bit rescue. It fails. On the untouched test split the rescue keeps its sign 4/4, clears −0.01 4/4, and gets *stronger* (−0.0159 / −0.0220 / −0.0143 / −0.0291), with the thinnest cell moving from −0.0102 to −0.0143; the leg tax and the small-amplitude ties also survive on test.

Three disclosures accompany the pass, none of them recorded as concerns. The full-budget grade-1 identity is never cross-checked directly — the fresh grade-1 cells sit at an effective stride where the comparison script has no matching probe-free cell — and so rests on the code reading rather than on a measurement. The "free at amplitudes 62–159" range omits a shorter-leg cell at aspect ratio 39.8, which is also free. And the statement that probe-free descent at η=1e-3 still wins the ill-conditioned terrain is true of the rescue cells reported here, but not of the whole traverse grid, where the cells at effective stride 1e-2 (3.1671 / 3.1668) beat it. No independent re-measurement was attempted for this entry and none is claimed: the auditor's own run showed the large-stride regime does not engage below the full training budget, so a cheap re-measurement would have been uninformative.

All claims are judged strictly against the registered wording, with misses included above; 2 seeds × 2 feature sets; replicated before publication and independently audited afterward. Internal designation: switchback descent.
