RIEMANNIAN NEWTON METHODS WITH STATISTICAL INFERENCE

Anonymous authors
Paper under double-blind review

ABSTRACT

We develop online statistical inference for Polyak-Ruppert averages of inexact Newton iterates on Riemannian manifolds. A local asymptotic representation separates observation noise, solver randomization, and accumulated implementation bias, and yields one-pass covariance estimators and asymptotically valid Wald regions under local regularity and centering conditions. Independent fresh solves add a positive semidefinite term to the score-sandwich covariance; averaging solver replicas reduces this term in inverse proportion to their number. Slowly varying predictable gains instead cancel from the first order law. Product torus simulations show how Newton scaling shortens an ill-conditioned transient and how retaining solver variation restores coverage. Intrinsic means and principal subspaces provide geometric applications, with a streaming data illustration.

1 INTRODUCTION

Means of directional observations, covariance matrices, and principal subspaces are parameters with intrinsic geometry. Riemannian stochastic gradient methods respect that geometry, and Polyak-Ruppert averaging converts their local iterates into regular root-n estimators (Ruppert, 1988; Polyak and Juditsky, 1992; Kushner and Yang, 1993; Pelletier, 2000; Tripuraneni et al., 2018). In an ill-conditioned problem, however, a Newton or scaled gradient direction can be preferable to a raw gradient. A stochastic Newton direction solves a linear system with an estimated Hessian and a stochastic gradient; a randomized sketch or an iterative solver stopped after a fixed number of steps reduces the cost of this solve (Dembo et al., 1982; Pilanci and Wainwright, 2017; Na et al., 2023). Randomized linear algebra supplies coordinate-free approximations for optimization (Nie et al., 2026), but its inferential effect on an averaged estimator is not determined by its optimization accuracy alone.

The obstacle is that computational randomness can enter at two different times. An operator built from past observations is predictable when the current score arrives; if its one step variation is sufficiently small, summation by parts removes it from the first order law even when the operator has no fixed limit (An and Huo, 2026a). A solver drawn with a fresh seed at the current iteration is different: the centered product with the current score is the martingale innovation, and its conditional mean contributes to the root-n bias, so replacing the realized product by the mean solver times the score discards genuine sampling variation. This temporal distinction also separates the present argument from the augmented state analyses required for momentum methods (An and Huo, 2026b).

Polyak-Ruppert averaging makes the distinction transparent. On the target tangent space, let H be the invertible derivative of the population field, S the covariance of its centered target score ξ, and $\widetilde{\mathrm{P}}$ a current solver independent of ξ with invertible mean $\bar{\mathrm{P}}$; a star $*$ denotes the adjoint. Under the local central limit theorem below, the averaged estimator has covariance

$$\mathrm{V} = \mathrm{H}^{-1}\mathrm{S}(\mathrm{H}^{*-1}) + \mathrm{H}^{-1}\bar{\mathrm{P}}^{-1}\mathbb{E}\{(\widetilde{\mathrm{P}} - \bar{\mathrm{P}})\mathrm{S}(\widetilde{\mathrm{P}} - \bar{\mathrm{P}})^{*}\}(\bar{\mathrm{P}}^{*-1})(\mathrm{H}^{*-1}). \tag{1}$$

The additional term is positive semidefinite, records sampling variation, and is divided by m when m independent solver operators are averaged.

We develop this mechanism for estimating equations on finite-dimensional Riemannian manifolds. A common tangent space representation, with direct control of accumulated remainders, isolates the centered solver score product and an optional root-n local bias; it yields (1), the exact $1/m$

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

replication law, and a contrastwise variance ordering. A local quadratic comparison with Riemannian stochastic gradient descent (RSGD) separates this first order inferential law from the optimization role of preconditioning: the best worst direction mean contraction is governed by the condition number of the preconditioned drift, which exact Newton scaling reduces to one. The same representation identifies the feasible estimator, implemented directions in the fresh solver regime but raw scores after predictable cancellation. Product centering justifies ordinary Wald regions, and a consistently estimable local bias is removed by an explicit tangent correction. We identify population derivatives and target scores for intrinsic Fréchet means, affine-invariant covariance centers, and Grassmann principal subspaces, and a fixed-budget rank one sketch and project solver verifies the remaining conditions within an invariant local reporting phase.

1.1 Related work

Averaging is the classical route from a slower stochastic approximation sequence to a root-n limit (Polyak and Juditsky, 1992; Ruppert, 1988; Kushner and Yang, 1993). For Riemannian stochastic gradients, Tripuraneni et al. (2018) establish an averaged central limit theorem under local geometric and state rate conditions; the present analysis retains a generic C^2 retraction and separates the state rate argument from the geometric and averaging remainders. Stochastic Newton and conditional gradient methods have related Euclidean limits (Leluc and Portier, 2023; Boyer and Godichon-Baggioni, 2023).

Randomized linear solvers introduce a second source of variation beyond the observation stream. The sketch and project framework of Gower and Richtárik (2015) provides the rank one solver used below. Last iterate inference and covariance estimation for sketched Newton or sequential quadratic programming are developed by Na and Mahoney (2025); Kuang et al. (2026); Wang et al. (2026). For averaged iterates, Du et al. (2025) prove a solver dependent CLT and random scaling inference for sketched SQP. Concurrent work by Du et al. (2026) treats averaged Newton iterations with accelerated sketch and project and identifies covariance inflation. Our local manifold analysis complements these results by treating tangent space solver score products, complete implementation error, and an implemented direction covariance estimator.

A distinct cancellation mechanism applies to sufficiently slowly varying predictable preconditioners. An and Huo (2026a) give a Euclidean inverse gain decomposition and a sufficient stabilization threshold; the scope of its sharpness is restricted to their stated rate class. Our predictable branch uses the corresponding identity in a coherent local tangent representation. Neither predictability alone nor an optimization convergence guarantee implies the required averaged CLT. Model-specific online inference includes entrywise streaming PCA (Kumar et al., 2025) and online multiplier bootstrap inference for streaming r-PCA (Xu and Li, 2026); these concern different estimators and studentizers from the Newton procedure studied here.

We call $\mathrm{H}^{-1}\mathrm{S}(\mathrm{H}^{*-1})$ the sandwich covariance. It is an efficiency bound only when an information identity or a separate efficiency argument justifies that interpretation.

The rest of the paper is organized as follows. Section 2 further gives the online procedure; Section 3 develops its local limit, covariance, and conditioning results. Section 4 examines small sample inference and solver replication. Section B contains proofs, sufficient conditions, and geometric applications; Section A gives complete numerical comparisons and the streaming principal subspace analysis.

2 Averaged randomized Newton inference

Let $x_\star$ be a selected root of a population tangent field G on a d-dimensional Riemannian manifold $\mathcal{M}$. For optimization, $G = \operatorname{grad} F$ for an objective F. Let $E = T_{x_\star}\mathcal{M}$, let R be a C^2 retraction, and write $u_k = R_{x_\star}^{-1}(x_k)$ for the iterate coordinates. The working neighborhood U has compact closure in a larger common retraction/inverse-retraction domain inside a normal neighborhood. A smooth orthonormal frame gives isometries $\mathfrak{e}_x : E \to T_x\mathcal{M}$, with $\mathfrak{e}_{x_\star} = \mathrm{I}$. Define $\mathcal{T}_{x\to y} = \mathfrak{e}_y\mathfrak{e}_x^*$: it identifies tangent vectors by their frame coordinates and satisfies $\mathcal{T}_{y\to z}\mathcal{T}_{x\to y} = \mathcal{T}_{x\to z}$. It is generally different from Levi-Civita parallel transport. A public anchor frame makes these maps computable without knowing $x_\star$. Stars denote adjoints, and $(v \otimes w)z = v\langle w, z\rangle$. Vector and induced operator norms are denoted by $\|\cdot\|$ and $\|\cdot\|_{\mathrm{op}}$; L^2 norms take the corresponding squared norm in expectation.

Algorithm 1 Online Newton averaging and covariance tracking

Require: x_0, anchor a, horizon n, step sizes, averaging rule, solver budgets, spectral band, and admissible safeguards.

Initialize derivative/solver means and Welford state $(\mu, \mathrm{W}) = (0, 0)$.
for $k = 0, \dots, n-1$ **do**
  Construct P_{k+1} from past derivatives and fresh seeds, or use a predictable gain.
  Obtain the score and derivative; apply the solver and safeguards to update x_{k+1}.
  Record $D_{k+1} = -\eta_{k+1}^{-1} R_{x_k}^{-1}(x_{k+1})$; update the online average and safeguarded derivative mean.
  if fresh solves **then**
    Update $\widehat{\mathrm{P}}^a$; set $z = \mathcal{T}_{x_k \to a} D_{k+1}$.
  else
    Set $z = \mathcal{T}_{x_k \to a} \widehat{G}_{k+1}(x_k)$.
  end if
  Update Welford: $\delta = z - \mu$, $\mu^+ = \mu + \delta/(k+1)$, $\mathrm{W}^+ = \mathrm{W} + \delta \otimes (z - \mu^+)$.
end for
Form the fresh solver covariance in Section 3.2, or the admissible predictable covariance in (19).

Ensure: $\bar{x}_n$, $\mathcal{T}_{a \to \bar{x}_n} \widehat{\mathrm{V}}^a \mathcal{T}_{\bar{x}_n \to a}$; report unavailable covariance if the inverse does not exist.

The actual update is $x_{k+1} = R_{x_k}(-\eta_{k+1} D_{k+1})$, where $D_{k+1} = \mathrm{P}_{k+1} \widehat{G}_{k+1}(x_k) + e_{k+1}$, $P_{k+1} : T_{x_k}\mathcal{M} \to T_{x_k}\mathcal{M}$, and $\eta_k = \eta_0 k^{-\alpha}$, $\alpha \in (1/2, 1)$. The discrepancy e_{k+1} collects inner solver, clipping, and localization errors. Current scores, operators and updated iterates are $\mathcal{F}_{k+1}$-measurable, $\mathcal{F}_k$ contains their history, and all displayed conditional moments exist. The reported estimator is either the anchor average $R_a\{n^{-1} \sum_{k=1}^{n} R_a^{-1}(x_k)\}$ or the recursive average in (16); both have the same first order limit under the conditions below. Algorithm 2 gives the full implementation. Full dimensional Wald inference also requires a positive definite covariance and a common local chart, and a scalar interval a positive contrast variance. Burn-in keeps the global step index but restarts reporting accumulators with the retained count. A step cap controls individual moves, and localization of the actual trajectory requires a separate invariant basin or safeguard argument. A nonsymmetric estimating field uses a consistent Jacobian estimator and a corresponding invertibility safeguard in place of spectral Hessian clipping.

3 Local limit and one-pass inference

3.1 Conditions and influence representation

The expansion rests on a local linear score, control of the implemented solver and its accumulated errors, and a martingale limit for the resulting directions. We summarize them here; their full statements are in Section B.1.4, primitive criteria in Section B.4 (Propositions B.25 to B.27), and proofs in Section B.6.

Local score and solver. The reported phase starts at deterministic $x_0 \in U$, remains in U almost surely, and its reported average lies in the common chart with probability tending to one. In the reference tangent space, $\mathcal{T}_{x \to x_\star} G(x) = \mathrm{H} R_{x_\star}^{-1}(x) + r_G(x)$, where H is invertible, $r_G(x) = o(\|R_{x_\star}^{-1}(x)\|)$, and $\|\mathrm{H} R_{x_\star}^{-1}(x) + r_G(x)\| \le C \|R_{x_\star}^{-1}(x)\|$ on U. The score decomposes as $\mathcal{T}_{x_k \to x_\star} \widehat{G}_{k+1}(x_k) = \mathrm{H} u_k + r_G(x_k) + \xi_{k+1} + q_{k+1}$, where ξ_{k+1} is the target score and q_{k+1} is the centered state dependent score increment with $\mathbb{E}(\xi_{k+1} \mid \mathcal{F}_k) = \mathbb{E}(q_{k+1} \mid \mathcal{F}_k) = 0$. Put $\widetilde{\mathrm{P}}_{k+1} = \mathcal{T}_{x_k \to x_\star} \mathrm{P}_{k+1} \mathcal{T}_{x_\star \to x_k}$ and $\Pi_k = \mathbb{E}(\widetilde{\mathrm{P}}_{k+1} \mid \mathcal{F}_k)$. Define $b_{k+1}^{\xi} = \mathbb{E}(\widetilde{\mathrm{P}}_{k+1} \xi_{k+1} \mid \mathcal{F}_k)$, $b_{k+1}^{q} = \mathbb{E}(\widetilde{\mathrm{P}}_{k+1} q_{k+1} \mid \mathcal{F}_k)$, and centered products $\zeta_{k+1} = \widetilde{\mathrm{P}}_{k+1} \xi_{k+1} - b_{k+1}^{\xi}$, $\chi_{k+1} = \widetilde{\mathrm{P}}_{k+1} q_{k+1} - b_{k+1}^{q}$. Assume a deterministic invertible $\bar{\mathrm{P}}$ with positive stable (all eigenvalues have positive real parts) $\mathrm{A} = \bar{\mathrm{P}}\mathrm{H}$, a uniform conditional second moment bound for $\widetilde{\mathrm{P}}_{k+1}$, and $\mathbb{E}(\|\chi_{k+1}\|^2 \mid \mathcal{F}_k) \le C \|u_k\|^2$.

Accumulated errors. Decompose $\mathcal{T}_{x_k \to x_\star} e_{k+1} = a_{k+1} + \nu_{k+1}$ into its conditional mean and centered part, and put $\beta_{k+1} = b_{k+1}^{\xi} + b_{k+1}^{q} + a_{k+1}$. The actual coordinate recursion is $u_{k+1} =$

$u_k - \eta_{k+1}\mathcal{T}_{x_k \to x_\star} D_{k+1} + \eta_{k+1}(c_{k+1} + d_{k+1})$, where c_{k+1} and d_{k+1} are centered and accumulated geometric remainders, $\mathbb{E}(c_{k+1} \mid \mathcal{F}_k) = 0$. For deterministic $b \in E$, require

$$n^{-1/2} \sum_{k<n} (\Pi_k - \bar{\mathrm{P}})\mathrm{H} u_k \xrightarrow{p} 0, \qquad n^{-1/2} \sum_{k<n} \Pi_k r_G(x_k) \xrightarrow{p} 0,$$
$$n^{-1/2} \sum_{k<n} \beta_{k+1} \xrightarrow{p} b, \qquad n^{-1/2} \sum_{k<n} d_{k+1} \xrightarrow{p} 0,$$
$$n^{-1} \sum_{k<n} \mathbb{E}(\|\nu_{k+1}\|^2 + \|c_{k+1}\|^2 \mid \mathcal{F}_k) \xrightarrow{p} 0.$$

Here and below $\sum_{k<n}$ runs from $k = 0$ to $n - 1$.

State rate and limiting noise. Finally, $\mathbb{E}\,\|u_k\|^2 \le C\eta_k$, $\sqrt{n}\{R_{x_\star}^{-1}(\bar{x}_n) - n^{-1}\sum_{k=1}^n u_k\} \xrightarrow{p} 0$, and for deterministic positive semidefinite Q and every $\epsilon > 0$,

$$n^{-1} \sum_{k<n} \mathbb{E}(\zeta_{k+1} \otimes \zeta_{k+1} \mid \mathcal{F}_k) \xrightarrow{p} \mathrm{Q}, \qquad n^{-1} \sum_{k<n} \mathbb{E}\{\|\zeta_{k+1}\|^2 \mathbf{1}_{\{\|\zeta_{k+1}\| > \epsilon\sqrt{n}\}} \mid \mathcal{F}_k\} \xrightarrow{p} 0.$$

The appendix supplies a Lyapunov state rate criterion (Proposition B.10) and C^2 averaging arguments that avoid a fourth moment state recursion, and shows that restarting the average after a proportional burn-in leaves the limit unchanged (Proposition B.9). Conditional product centering, rather than separate unbiasedness, determines whether $b = 0$ (Corollary B.8).

Theorem 3.1. *Suppose Assumptions B.1 to B.7 hold. Then*

$$\sqrt{n}\, R_{x_\star}^{-1}(\bar{x}_n) = -\mathrm{A}^{-1} \frac{1}{\sqrt{n}} \sum_{k=1}^n \zeta_k - \mathrm{A}^{-1} b + o_p(1), \qquad \mathrm{A} = \bar{\mathrm{P}}\mathrm{H}. \tag{2}$$

Consequently,

$$\sqrt{n}\, R_{x_\star}^{-1}(\bar{x}_n) \xrightarrow{d} N_E(-\mathrm{A}^{-1}b, \mathrm{V}), \qquad \mathrm{V} = \mathrm{A}^{-1}\mathrm{Q}(\mathrm{A}^{*-1}). \tag{3}$$

Here N_E denotes a Gaussian law on E. For every C^1 map $\psi : \mathcal{M} \to \mathbb{R}^p$ with differential $D\psi_{x_\star} : E \to \mathbb{R}^p$,

$$\sqrt{n}\{\psi(\bar{x}_n) - \psi(x_\star)\} \xrightarrow{d} N_p\big(-D\psi_{x_\star}\mathrm{A}^{-1}b, D\psi_{x_\star}\mathrm{V}D\psi_{x_\star}^*\big). \tag{4}$$

The mean drift is A, whereas fluctuations retain the realized solver score product. The shift $-\mathrm{A}^{-1}b$ persists under covariance estimation.

Dividing the coordinate recursion by its step size and summing separates the average state from a summation by parts boundary. The state bound and $\alpha < 1$ make that boundary negligible at the root-n scale. The centered solver acting on the current state and the state dependent score increment also vanish after averaging, whereas the target score product ζ_k remains. Its conditional variance and Lindeberg conditions give the Gaussian limit. Finally, the C^2 chart and averaging bounds transfer the linear expansion back to the manifold. Section B.6.4 gives the proof.

3.2 Covariance estimation and Wald regions

In anchor coordinates, let $\mathrm{H}_a = \mathcal{T}_{x_\star \to a}\mathrm{H}\mathcal{T}_{a \to x_\star}$, $\bar{\mathrm{P}}_a = \mathcal{T}_{x_\star \to a}\bar{\mathrm{P}}\mathcal{T}_{a \to x_\star}$, $\mathrm{Q}_a = \mathcal{T}_{x_\star \to a}\mathrm{Q}\mathcal{T}_{a \to x_\star}$ and $\mathrm{V}_a = (\bar{\mathrm{P}}_a\mathrm{H}_a)^{-1}\mathrm{Q}_a((\bar{\mathrm{P}}_a\mathrm{H}_a)^{*-1})$. Define $Y_{k+1} = \mathcal{T}_{x_k \to a} D_{k+1}$, $\bar{Y}_n = n^{-1}\sum_{k<n} Y_{k+1}$,

$$\widehat{\mathrm{Q}}_n^a = n^{-1} \sum_{k<n} (Y_{k+1} - \bar{Y}_n) \otimes (Y_{k+1} - \bar{Y}_n) = \mathrm{W}_n/n, \qquad \widehat{\mathrm{A}}_n^a = \widehat{\mathrm{P}}_n^a \widehat{\mathrm{H}}_n^a.$$

Set $\widehat{\mathrm{V}}_n^a = (\widehat{\mathrm{A}}_n^a)^{-1}\widehat{\mathrm{Q}}_n^a((\widehat{\mathrm{A}}_n^a)^{*-1})$ when invertible, extend it by zero otherwise for convergence statements, and transport it to obtain $\widehat{\mathrm{V}}_n^{\bar{x}} = \mathcal{T}_{a \to \bar{x}_n}\widehat{\mathrm{V}}_n^a\mathcal{T}_{\bar{x}_n \to a}$.

Assumption 3.2. *The online derivative and mean solver estimators $\widehat{\mathrm{H}}_n^a$ and $\widehat{\mathrm{P}}_n^a$ satisfy*

$$\widehat{\mathrm{H}}_n^a \xrightarrow{p} \mathrm{H}_a, \qquad \widehat{\mathrm{P}}_n^a \xrightarrow{p} \bar{\mathrm{P}}_a, \tag{5}$$

and the total conditional bias $\beta_{k+1} = b_{k+1}^\xi + b_{k+1}^q + a_{k+1}$ obeys $n^{-1}\sum_{k<n} \|\beta_{k+1}\|^2 \xrightarrow{p} 0$.

The squared bias requirement is distinct from the signed root-n condition: alternating errors can cancel in the estimator and remain in its covariance.

Theorem 3.3. *Suppose the assumptions of Theorem 3.1 hold and Assumption 3.2 is satisfied. Then*

$$\widehat{\mathrm{Q}}_n^a \xrightarrow{p} \mathrm{Q}_a, \qquad \widehat{\mathrm{V}}_n^a \xrightarrow{p} \mathrm{V}_a, \qquad \widehat{\mathrm{V}}_n^{\bar{x}} - \mathcal{T}_{x_\star \to \bar{x}_n} \mathrm{V} \mathcal{T}_{\bar{x}_n \to x_\star} \xrightarrow{p} 0. \tag{6}$$

If Q *is positive definite,* $b = 0$*, and* $x_\star$ *and* $\bar{x}_n$ *lie in a common inverse-retraction neighborhood with probability tending to one, fix* $\gamma \in (0,1)$ *and let* $\chi^2_{d,1-\gamma}$ *be the* $(1-\gamma)$ *quantile of the chi-squared law with* d *degrees of freedom. When the required inverse exists and* $\widehat{\mathrm{V}}_n^{\bar{x}}$ *is positive definite, define the uncorrected Wald region*

$$\mathcal{C}_{n,1-\gamma} = \left\{ x \in U : R_{\bar{x}_n}^{-1}(x) \text{ is defined and } n \left\langle R_{\bar{x}_n}^{-1}(x), (\widehat{\mathrm{V}}_n^{\bar{x}})^{-1} R_{\bar{x}_n}^{-1}(x) \right\rangle_{\bar{x}_n} \le \chi^2_{d,1-\gamma} \right\}$$

and otherwise put $\mathcal{C}_{n,1-\gamma} = \varnothing$*. Availability has probability tending to one, and*

$$\mathbb{P}\{x_\star \in \mathcal{C}_{n,1-\gamma}\} \longrightarrow 1 - \gamma. \tag{7}$$

Under the same positive definiteness and neighborhood conditions, if instead $b \neq 0$*, the uncorrected statistic at* $x_\star$ *converges to a noncentral chi-squared law with noncentrality*

$$\lambda_b = \left\langle \mathrm{A}^{-1} b, \mathrm{V}^{-1} \mathrm{A}^{-1} b \right\rangle, \tag{8}$$

and therefore does not in general have nominal coverage.

Bias-corrected inference is feasible when the signed accumulated bias can be estimated. Put $b_a = \mathcal{T}_{x_\star \to a} b$*, and suppose an anchor space estimator satisfies* $\widehat{b}_n^a \xrightarrow{p} b_a$*. Since the limiting shift is* $-\mathrm{A}^{-1} b$*, define*

$$\begin{aligned} \bar{x}_n^{\mathrm{bc}} &= R_{\bar{x}_n}\left[n^{-1/2} \mathcal{T}_{a \to \bar{x}_n} (\widehat{\mathrm{A}}_n^a)^{-1} \widehat{b}_n^a \right], \\ \widehat{\mathrm{V}}_n^{\mathrm{bc}} &= \mathcal{T}_{\bar{x}_n \to \bar{x}_n^{\mathrm{bc}}} \widehat{\mathrm{V}}_n^{\bar{x}} \mathcal{T}_{\bar{x}_n^{\mathrm{bc}} \to \bar{x}_n}. \end{aligned} \tag{9}$$

Then $\sqrt{n} R_{x_\star}^{-1}(\bar{x}_n^{\mathrm{bc}}) \xrightarrow{d} N_E(0, \mathrm{V})$*, and the Wald region obtained by replacing* $\bar{x}_n$ *and* $\widehat{\mathrm{V}}_n^{\bar{x}}$ *in* $\mathcal{C}_{n,1-\gamma}$ *by* $\bar{x}_n^{\mathrm{bc}}$ *and* $\widehat{\mathrm{V}}_n^{\mathrm{bc}}$ *has asymptotic coverage* $1-\gamma$*.*

For a scalar C^1 *functional* ψ *with* $D\psi_{x_\star} \mathrm{V} D\psi_{x_\star}^* > 0$*, a consistent derivative estimate gives the usual two-sided interval when it is centered at* $\bar{x}_n$ *under* $b = 0$*, or at* $\bar{x}_n^{\mathrm{bc}}$ *under the displayed consistent bias correction.*

Section B.6.2 develops a computable bias estimator from recorded implementation discrepancies, with its additional centering conditions.

3.3 Fresh solves and covariance inflation

For independent current solver randomization, the covariance separates into the score sandwich $\mathrm{V}_{\mathrm{sand}} = \mathrm{H}^{-1} \mathrm{S} (\mathrm{H}^{*-1})$ and a solver contribution.

Theorem 3.4. *Suppose that, at* $x_\star$*, a solver operator* $\widetilde{\mathrm{P}}$ *is independent of a centered score* ξ*, with*

$$\mathbb{E}\widetilde{\mathrm{P}} = \bar{\mathrm{P}}, \qquad \mathbb{E}(\xi \otimes \xi) = \mathrm{S}, \qquad \mathbb{E} \left\| \widetilde{\mathrm{P}} \right\|_{\mathrm{op}}^2 < \infty, \qquad \mathrm{H} \text{ and } \bar{\mathrm{P}} \text{ invertible}.$$

Define the target product covariance and the solver variance operator by

$$\mathrm{Q} = \mathbb{E}\{(\widetilde{\mathrm{P}}\xi) \otimes (\widetilde{\mathrm{P}}\xi)\}, \qquad \mathrm{R} = \mathbb{E}\{(\widetilde{\mathrm{P}} - \bar{\mathrm{P}}) \mathrm{S} (\widetilde{\mathrm{P}} - \bar{\mathrm{P}})^*\}. \tag{10}$$

Then

$$\mathrm{Q} = \bar{\mathrm{P}} \mathrm{S} \bar{\mathrm{P}}^* + \mathrm{R}, \qquad \mathrm{R} \succeq 0,$$

and

$$\mathrm{V} = \mathrm{V}_{\mathrm{sand}} + \mathrm{H}^{-1} \bar{\mathrm{P}}^{-1} \mathrm{R} (\bar{\mathrm{P}}^{*-1}) (\mathrm{H}^{*-1}). \tag{11}$$

The second term is positive semidefinite and is zero if and only if

$$(\widetilde{\mathrm{P}} - \bar{\mathrm{P}}) \mathrm{S}^{1/2} = 0 \quad \textit{almost surely}.$$

Averaging m conditionally independent replicas with a common conditional law and mean divides the extra term by m, which orders the limiting covariances in the Loewner sense and contrastwise (Corollaries B.12 and B.13). Sandwich recovery requires the solver error applied to the target score to vanish in conditional mean square; probability convergence of the operator alone is insufficient (Corollaries B.11 and B.14).

For a general fixed-budget randomized sketch, Proposition B.19 states the solver conditions under which Theorems 3.1 and 3.4 apply. For rank one sketch and project, Proposition B.20 verifies these conditions from a common spectral band, a spanning sketch law, and the local score and invariant basin assumptions. Its coordinate sketch specialization underlies Section 4.

3.4 Conditioning and predictable gains

Newton scaling can improve the transient even when averaging ultimately recovers the same sandwich covariance as RSGD; a local quadratic comparison identifies the relevant spectral ratio.

Proposition 3.5. *Let $\mathrm{H} \succ 0$ be self-adjoint and consider the noiseless linearized recursion on E with a fixed self-adjoint preconditioner $\mathrm{P} \succ 0$, $u_{k+1} = (\mathrm{I} - \eta \mathrm{PH})u_k$. Let $\lambda_{\min}$ and $\lambda_{\max}$ be the smallest and largest eigenvalues of $\mathrm{P}^{1/2}\mathrm{H}\mathrm{P}^{1/2}$, and define $\kappa_{\mathrm{P}} = \frac{\lambda_{\max}}{\lambda_{\min}}$. In the P^{-1} energy norm (Section B.6.1),*

$$\inf_{\eta>0} \|\mathrm{I} - \eta \mathrm{PH}\|_{\mathrm{op}} = \frac{\kappa_{\mathrm{P}} - 1}{\kappa_{\mathrm{P}} + 1}, \qquad \eta_\star = \frac{2}{\lambda_{\min} + \lambda_{\max}}.$$

For RSGD, $\mathrm{P} = \mathrm{I}$ and $\kappa_{\mathrm{P}} = \kappa(\mathrm{H})$; for exact Newton conditioning, $\mathrm{P} = \mathrm{H}^{-1}$ and $\kappa_{\mathrm{P}} = 1$, so the unit step removes the linearized deterministic error. More generally, a smaller effective condition number yields a smaller best worst direction contraction factor in the corresponding energy norm.

For a fresh solver with mean P, this contraction governs the linearized mean drift; its realized variation contributes the additional covariance above. A predictable gain has a different inferential effect: multiplication by its inverse removes the gain from the leading score term, provided the resulting summation by parts boundary is negligible.

For a predictable invertible C_k, consider $u_{k+1} = u_k - \eta_{k+1}\{\mathrm{C}_k(\mathrm{H}u_k + \xi_{k+1}) + \varepsilon_{k+1}\}$, where ε_{k+1} contains geometric and other remainders.

Theorem 3.6. *Assume that H is invertible. Let $\mathrm{M}_k = \mathrm{C}_k^{-1}$ and $\Gamma_k = \mathrm{M}_k/\eta_{k+1}$. Suppose*

$$\frac{1}{\sqrt{n}} \left\{ \|\Gamma_0 u_0\| + \|\Gamma_{n-1} u_n\| + \sum_{k=1}^{n-1} \|(\Gamma_k - \Gamma_{k-1})u_k\| \right\} \xrightarrow{p} 0, \tag{12}$$

$$\frac{1}{\sqrt{n}} \sum_{k=0}^{n-1} \mathrm{M}_k \varepsilon_{k+1} \xrightarrow{p} 0, \tag{13}$$

assume also that the reported online average is root-n equivalent to the arithmetic average of u_k and that $n^{-1/2}(u_n - u_0) \xrightarrow{p} 0$. If $n^{-1/2} \sum_{k=0}^{n-1} \xi_{k+1} \xrightarrow{d} N_E(0, \mathrm{S})$, then

$$\sqrt{n}\, R_{x_\star}^{-1}(\bar{x}_n) = -\mathrm{H}^{-1} \frac{1}{\sqrt{n}} \sum_{k=0}^{n-1} \xi_{k+1} + o_p(1) \xrightarrow{d} N_E(0, \mathrm{V}_{\mathrm{sand}}). \tag{14}$$

No deterministic limit of C_k is required.

The inverse gain criterion holds, for example, when $\sup_k \left\|\mathrm{C}_k^{-1}\right\|_{L^2(\mathrm{op})} < \infty$, $\mathbb{E}\,\|u_k\|^2 = O(k^{-\alpha})$, and $\|C_k^{-1} - C_{k-1}^{-1}\|_{L^2(\mathrm{op})} = O(k^{-\beta})$ with $\beta > (1+\alpha)/2$ (Corollary B.15); the weighted remainder condition is still required. Under the raw score approximation and covariance conditions in Corollary B.16, the raw score Welford covariance $\widehat{\mathrm{S}}_n^a$ yields the estimator $(\widehat{\mathrm{H}}_n^a)^{-1}\widehat{\mathrm{S}}_n^a((\widehat{\mathrm{H}}_n^a)^{*-1})$ even if C_k has no limit. Exact predictable solves of a spectrally clipped streaming Hessian satisfy the increment criterion under Corollary B.18, with the averaging rate supplied by Proposition B.17.

When $\mathcal{M} = \mathbb{R}^d$, the retraction is addition and all geometric remainders vanish, so the representation reduces to the Euclidean averaged law: a deterministic or predictable preconditioner returns the classical sandwich covariance of Polyak and Juditsky (1992), and only current solver randomization

adds the replica-averaged inflation term of (11) (Section B.2.8 and corollary B.13). Section B.3 identifies the derivative and target score covariance for intrinsic Fréchet means, the affine-invariant covariance center, and streaming principal subspaces on the Grassmann manifold (Corollaries B.21 to B.24), and Section A.3 applies the Grassmann case to a streaming principal subspace analysis of the Wisconsin Diagnostic Breast Cancer data.

4 NUMERICAL EXPERIMENTS

We compare RSGD, exact Newton, and our randomized Newton inference method (proposed), implemented by coordinate sketch and project to separate the effect of conditioning on initialization bias from the estimation and coverage effects of solver randomization. A product torus model keeps the iterates in an invariant angular chart and gives exact linear benchmarks for both effects.

4.1 MODEL AND IMPLEMENTATION

On $\mathbb{S}^1 \times \mathbb{S}^1$ with the product metric, consider

$$f(\theta, Y) = \frac{1}{2} \sum_{j=1}^{2} h_j \operatorname{ang}(\theta_j - Y_j)^2, \qquad (h_1, h_2) = (1, \kappa),$$

where ang is the shortest signed angular difference. Generate $Y_j = \theta_{\star j} + \varepsilon_j$ with $\theta_\star = (0.20, -0.15)$ and independent $\varepsilon_j \sim \mathrm{Unif}[-\sqrt{3}\sigma, \sqrt{3}\sigma]$, $\sigma = 0.20$. The known loss weights give $\mathrm{H} = \mathrm{diag}(1, \kappa)$; the iteration updates use the observations, not the target or noise covariance. The initial angles are $(-0.35, 0.30)$ and the public anchor is $(0, 0)$. We cross the condition number $\kappa \in \{1, 30, 300\}$ with the number of observations $n \in \{256, 1024, 4096\}$, using $m \in \{1, 4, 16\}$ independent coordinate sketches per observation.

RSGD uses $\mathrm{P} = \mathrm{I}$ and exact Newton uses $\mathrm{P} = \mathrm{H}^{-1}$. For proposed, draw m independent uniform coordinate sketches $s_\ell \in \{e_1, e_2\}$ per observation and use

$$\mathrm{P}^{(m)} = \frac{2}{m} \sum_{\ell=1}^{m} \frac{s_\ell s_\ell^*}{s_\ell^* \mathrm{H} s_\ell}, \qquad \mathbb{E}\mathrm{P}^{(m)} = \mathrm{H}^{-1}.$$

The scalar factor two matches the mean Newton drift. The limiting covariance is $\mathrm{V}_m = (1+1/m)\sigma^2\mathrm{I}$ for proposed and $\mathrm{V} = \sigma^2\mathrm{I}$ for RSGD and exact Newton.

The schedule is $\eta_k = \min\{c(k+10)^{-0.65}, \eta_{\max}\}$, with caps $1/\kappa$, 1, and $1/2$ for RSGD, exact Newton, and proposed, respectively. For each method and configuration, select $c \in \{0.2, 0.4, 0.8, 1.6, 3.2, 6.4\}$ by minimum mean squared geodesic error over 48 pilot replications per candidate, then freeze it before independent validation and the 500-replication evaluation. This simulation tuning uses the known target; the subsequent updates use only the observations and specified loss weights. Methods share observation streams; solver randomness is independent. Discard the first $n/8$ iterates, retaining the global step index, and restart the average and reporting accumulators.

In replication r, the reported estimate $\widehat{\theta}_r$ is the fixed anchor average of the $7n/8$ iterates retained after burn-in. Proposed uses the full observed directions and empirical mean solver for covariance estimation; RSGD and exact Newton use the score sandwich. Joint 95% Wald regions use the retained count $7n/8$ and the χ^2_2 quantile. Estimation error is $\mathrm{RMSE}_g = \{500^{-1} \sum_{r=1}^{500} d^2(\widehat{\theta}_r, \theta_\star)\}^{1/2}$. The Monte Carlo standard error of coverage near 95% is about one percentage point. Sections A.1 and A.2 give the remaining design details, all selected constants, checkpoint errors, and timings.

4.2 CONDITIONING AND SMALL SAMPLE INFERENCE

At $\kappa = 1$, RSGD and exact Newton coincide, while proposed adds randomization error (Table 1). At $\kappa = 30$, Newton scaling improves the shortest horizon comparison, but RSGD approaches the exact Newton error as n increases. At $\kappa = 300$, substantial differences persist: with $n = 256$, proposed with $m = 4$ attains 95.4% coverage, whereas none of the RSGD regions covers the target. Exact Newton is the deterministic preconditioning benchmark and has the smallest errors; Proposed is its

Table 1: Product torus comparison across spectral ratios and sample sizes. Coverage is for joint 95% Wald regions; geodesic RMSE is multiplied by 10^3. Proposed uses $m = 4$ independent coordinate solves. Every entry uses 500 paired replications; solver budgets and checkpoint errors are in the appendix.

		RSGD		Proposed, $m = 4$		Exact Newton	
κ	n	Cov. (%)	RMSE	Cov. (%)	RMSE	Cov. (%)	RMSE
1	256	95.8	18.99	95.2	21.00	95.8	18.99
1	1024	95.2	9.84	94.6	11.11	95.2	9.84
1	4096	96.2	4.63	94.4	5.38	96.2	4.63
30	256	75.4	29.60	96.4	21.60	96.6	18.85
30	1024	96.0	9.56	95.0	10.77	94.8	9.56
30	4096	95.4	4.67	93.2	5.53	95.8	4.66
300	256	0.0	347.57	95.4	20.68	95.8	18.88
300	1024	0.0	114.03	94.8	10.86	95.4	9.65
300	4096	41.2	9.55	96.2	5.25	95.4	4.70

randomized counterpart. The improvement over RSGD reflects conditioning under the stated step schedules, while randomization carries a separate statistical cost.

The diagonal exact solve is inexpensive in this design, and Proposed takes longer per path (Table 7; Section A.2 reports the timing values and the aggregated coordinate implementation), so the gain over RSGD is in performance at a fixed observation budget.

The linear chart makes the source of this transient explicit. For RSGD, the mean error in coordinate j at step k is

$$\mathbb{E}(\theta_{kj} - \theta_{\star j}) = (\theta_{0j} - \theta_{\star j}) \prod_{t=1}^{k} (1 - \eta_t h_j).$$

Exact Newton and mean rescaled proposed replace h_j by one in this expression, using their respective step schedules. The caps make every coordinate update a convex combination of its preceding value and the observation. Since both lie in $(-0.6, 0.6)^2$, all paths stay in this angular chart, where the recursion and transport are exactly linear. The RSGD cap $1/\kappa$ protects the stiff coordinate but slows the removal of initial error in the other coordinate, and tail averaging reduces this error only as the iterates themselves approach the target. Poor coverage can therefore persist despite a decreasing RMSE: the remaining offset is still large relative to the sampling standard error. Newton scaling directly reduces that offset, while the direction covariance accounts for proposed's additional variation.

Error along the averaging window. To distinguish a transient from the eventual sampling error, Table 2 follows the $n = 4096$ runs through three checkpoints. Each method keeps the step constant chosen for this horizon; the intermediate averages are not separately tuned. In the well-conditioned case, error roughly halves when the retained sample count quadruples, as the root-n approximation suggests. The Newton methods show the same pattern at $\kappa = 300$. RSGD instead has a much larger early error and continues to shed its initialization effect throughout the window.

At $\kappa = 30$, the RSGD checkpoint errors are already close to the exact Newton benchmark, between the well-conditioned and persistent transient regimes.

4.3 Covariance correction and solver replication

For proposed with $m = 1$ in this product torus design, retaining solver variation substantially improves the reported coverage (Table 3): direction based coverage ranges from 93.0% to 95.8%, against 74.0%–80.0% for the same estimates using only score covariance, which isolates the inferential effect of the omitted solver term. Increasing m reduces RMSE throughout the grid, consistent with the $1/m$ covariance contribution. Corrected coverage is generally close to nominal across the budgets

Table 2: Geodesic RMSE along the $n = 4096$ averaging window, multiplied by 10^3. The checkpoints contain 896, 1792 and 3584 retained observations after a burn-in of 512. Each row follows the same 500 paths with a fixed step constant; Proposed uses $m = 4$.

κ	Method	25%	50%	100%
1	RSGD	9.35	6.57	4.63
	Exact Newton	9.35	6.57	4.63
	Proposed, $m = 4$	10.86	7.57	5.38
30	RSGD	9.20	6.67	4.67
	Exact Newton	8.80	6.51	4.66
	Proposed, $m = 4$	11.01	7.99	5.53
300	RSGD	32.81	17.77	9.55
	Exact Newton	9.40	6.60	4.70
	Proposed, $m = 4$	9.81	7.25	5.25

Table 3: Solver covariance correction and replication. The two coverage columns use the same proposed estimator with $m = 1$: direction covariance versus the score only counterfactual that omits solver variation. Geodesic RMSEs for $m \in \{1, 4, 16\}$ use 500 replications and are multiplied by 10^3.

		Coverage (%), $m = 1$		RMSE $\times 10^3$		
κ	n	Direction	Score only	$m = 1$	$m = 4$	$m = 16$
1	256	95.8	76.2	27.01	21.00	19.57
1	1024	93.0	75.2	14.01	11.11	10.16
1	4096	95.6	80.0	6.61	5.38	4.85
30	256	95.4	78.4	26.94	21.60	19.21
30	1024	95.4	78.6	13.40	10.77	9.86
30	4096	95.4	77.8	6.62	5.53	4.81
300	256	94.4	79.0	27.19	20.68	19.29
300	1024	94.0	76.6	13.92	10.86	9.94
300	4096	93.4	74.0	6.88	5.25	4.80

but ranges from 93.0% to 98.4%; the 98.4% coverage at $(\kappa, n, m) = (30, 256, 16)$ is a finite-sample departure rather than evidence of exact calibration (Tables 4 to 6 give the complete comparison).

The score only coverage loss also has an explicit benchmark: the limiting covariance is $(1+1/m)\sigma^2\mathrm{I}$, so studentizing by the score sandwich $\sigma^2\mathrm{I}$ gives the limiting Wald statistic $(1+1/m)\chi^2_2$ and coverage

$$\Pr\{(1 + 1/m)\chi^2_2 \le \chi^2_{2,0.95}\} = 1 - 0.05^{\,m/(m+1)},$$

that is 77.6%, 90.9%, and 94.0% for $m \in \{1, 4, 16\}$. The observed score only results follow this progression: replication reduces the omitted variation, but additional observations alone do not remove it at a fixed solver budget, whereas direction based studentization targets the covariance at that budget.

5 Concluding remarks

Our proposed framework enables valid inference for averaged randomized Newton methods without requiring solver randomness to vanish. Implemented direction covariance captures fresh solver variation, while replication reduces this cost. In ill-conditioned settings, the method retains Newton-type conditioning benefits and improves finite-sample estimation and coverage over RSGD.

Slowly varying predictable gains instead permit score based inference. Simulations support these mechanisms; dependent streams and residual adaptive solvers remain for future work.

AI USE STATEMENT

Generative AI tools were used to revise wording and improve the clarity and readability of the manuscript; to identify potentially relevant literature, references, and search directions, and to assist with research execution, including drafting, debugging, and refining code. For mathematical proofs, AI was used primarily to review and check the authors' arguments for potential errors, inconsistencies, or gaps. AI-assisted manuscript revisions, literature-search suggestions, feedback on research ideas and proofs, and generated or modified code were reviewed and verified by the authors before incorporation into the final work. The authors independently assessed the accuracy and appropriateness of these outputs and retain full responsibility for the methodology, proofs, implementation, analysis, and conclusions presented in the paper.

REPRODUCIBILITY STATEMENT

The code and an accompanying `README` file are submitted with the manuscript. We provide a dedicated reproduction script for each figure and table, allowing the reported results to be reproduced individually. The `README` provides instructions for running the code.

REFERENCES

Afsari, B. (2011). Riemannian L^p center of mass: existence, uniqueness, and convexity. *Proceedings of the American Mathematical Society*, **139**, 655–673.

An, S. and Huo, X. (2026a). When does dynamic preconditioning preserve the Polyak–Ruppert CLT? A stabilization threshold. arXiv:2604.23498.

An, S. and Huo, X. (2026b). A Polyak–Ruppert central limit theorem for SA–Adam with momentum and non-convergent adaptive preconditioning. arXiv:2606.17364.

Bhattacharya, R. and Patrangenaru, V. (2005). Large sample theory of intrinsic and extrinsic sample means on manifolds—II. *Annals of Statistics*, **33**, 1225–1259.

Bhatia, R. (2007). *Positive Definite Matrices*. Princeton University Press.

Bonnabel, S. (2013). Stochastic gradient descent on Riemannian manifolds. *IEEE Transactions on Automatic Control*, **58**, 2217–2229.

Boyer, C. and Godichon-Baggioni, A. (2023). On the asymptotic rate of convergence of stochastic Newton algorithms and their weighted averaged versions. *Computational Optimization and Applications*, **84**, 921–972. See also the correction in *Computational Optimization and Applications*, **87** (2024), 705–706.

Dembo, R. S., Eisenstat, S. C., and Steihaug, T. (1982). Inexact Newton methods. *SIAM Journal on Numerical Analysis*, **19**, 400–408.

Du, X., Zhu, W., Wu, W. B., and Na, S. (2025). Online statistical inference of constrained stochastic optimization via random scaling. arXiv:2505.18327v1.

Du, X., Rebrova, E., Dereziński, M., and Na, S. (2026). Inference for Newton methods with accelerated sketch-and-project via random scaling. arXiv:2609.12421v1, 11 September 2026.

Edelman, A., Arias, T. A., and Smith, S. T. (1998). The geometry of algorithms with orthogonality constraints. *SIAM Journal on Matrix Analysis and Applications*, **20**, 303–353.

Hall, P. and Heyde, C. C. (1980). *Martingale Limit Theory and Its Application*. Academic Press.

Gower, R. M. and Richtárik, P. (2015). Randomized iterative methods for linear systems. *SIAM Journal on Matrix Analysis and Applications*, **36**, 1660–1690.

Hotz, T., Le, H., and Wood, A. T. A. (2024). Central limit theorem for intrinsic Fréchet means in smooth compact Riemannian manifolds. *Probability Theory and Related Fields*, **189**, 1219–1246.

Kuang, W., Anitescu, M., and Na, S. (2026). Online covariance matrix estimation in sketched Newton methods. *Information and Inference: A Journal of the IMA*, **15**(2), iaag012.

Kumar, S., Pandey, S., and Sarkar, P. (2025). Beyond sin-squared error: linear-time entrywise uncertainty quantification for streaming PCA. arXiv:2506.12655v2.

Kushner, H. J. and Yang, J. (1993). Stochastic approximation with averaging of the iterates: optimal asymptotic rate of convergence for general processes. *SIAM Journal on Control and Optimization*, **31**, 1045–1062.

Leluc, R. and Portier, F. (2023). Asymptotic analysis of conditioned stochastic gradient descent. *Transactions on Machine Learning Research*, 2023.

Li, C. J. and Jordan, M. I. (2023). Nonconvex stochastic scaled gradient descent and generalized eigenvector problems. In *Proceedings of the Thirty-Ninth Conference on Uncertainty in Artificial Intelligence*, PMLR **216**, 1230–1240.

Na, S., Dereziński, M., and Mahoney, M. W. (2023). Hessian averaging in stochastic Newton methods achieves superlinear convergence. *Mathematical Programming*, **201**, 473–520.

Na, S. and Mahoney, M. W. (2025). Statistical inference of constrained stochastic optimization via sketched sequential quadratic programming. *Journal of Machine Learning Research*, **26**(33), 1–75.

Nie, H., Gao, B., Han, A., Jawanpuria, P., Mishra, B., and Wen, Z. (2026). Nyström approximation on manifolds. arXiv:2605.14933v1.

Pelletier, M. (2000). Asymptotic almost sure efficiency of averaged stochastic algorithms. *SIAM Journal on Control and Optimization*, **39**, 49–72.

Pilanci, M. and Wainwright, M. J. (2017). Newton sketch: a near linear-time optimization algorithm with linear-quadratic convergence. *SIAM Journal on Optimization*, **27**, 205–245.

Polyak, B. T. and Juditsky, A. B. (1992). Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, **30**, 838–855.

Ruppert, D. (1988). Efficient estimations from a slowly convergent Robbins–Monro process. Technical Report 781, School of Operations Research and Industrial Engineering, Cornell University.

Tripuraneni, N., Flammarion, N., Bach, F., and Jordan, M. I. (2018). Averaging stochastic gradient descent on Riemannian manifolds. In *Proceedings of the 31st Conference on Learning Theory*, PMLR **75**, 650–687.

Wang, H., Du, X., and Na, S. (2026). Inference of online Newton methods with Nesterov's accelerated sketching. arXiv:2604.23436v2; accepted at ICML 2026.

Wolberg, W., Mangasarian, O., Street, N., and Street, W. (1993). Breast Cancer Wisconsin (Diagnostic). UCI Machine Learning Repository.

Wolberg, W. H., Street, W. N., and Mangasarian, O. L. (1994). Machine learning techniques to diagnose breast cancer from image-processed nuclear features of fine needle aspirates. *Cancer Letters*, **77**, 163–171.

Xu, H. and Li, H. (2026). Inference and uncertainty quantification for streaming r-PCA. arXiv:2608.18374v1.

Zhou, J. and Said, S. (2019). Fast, asymptotically efficient, recursive estimation in a Riemannian manifold. *Entropy*, **21**, 1021.

A NUMERICAL AND APPLICATION DETAILS

A.1 COORDINATE SKETCH AND PROJECT IMPLEMENTATION

For the model in Section 4, let $B \sim \mathrm{Bin}(m, 1/2)$ count the first-coordinate sketches. Then

$$\mathrm{P}^{(m)}\mathrm{H} = \mathrm{diag}\{2B/m, 2(m-B)/m\}.$$

Sampling this count directly is distributionally equivalent to summing m independent sketches. Its first two moments give mean drift I and limiting covariance $(1 + 1/m)\sigma^2\mathrm{I}$, whereas exact Newton and RSGD have limiting covariance $\sigma^2\mathrm{I}$.

The step caps make each coordinate update a convex combination of its previous value and the current observation. The initial point and all observations lie in $(-0.6, 0.6)^2$, so this rectangle is invariant. Angular differences never meet the cut locus, transport is the identity, and fixed anchor averaging is arithmetic averaging in these coordinates. Thus the experiment realizes the local linear model exactly on the torus, separating conditioning from curvature.

After burn-in, all reporting accumulators restart while the step index is unchanged. Proposed uses $\widehat{\mathrm{A}} = \widehat{\mathrm{P}}\mathrm{H}$ and the centered empirical covariance of implemented directions; RSGD and exact Newton use the score sandwich. Methods share observations within each replication, with independent observation and solver random generators. The following tables give every selected step constant, both covariance calculations, and errors at three fractions of the retained window.

After selecting step constants using 48 pilot replications per candidate, 96 independent validation replications per configuration checked that outputs were finite and paths stayed in the invariant rectangle. This stage did not alter the selected constants. The subsequent 500 evaluation replications used independent seeds.

For proposed with $m = 4$ minus RSGD, the paired mean squared-geodesic-error difference has a 95% Monte Carlo interval $[-4.74, -3.45] \times 10^{-4}$ at $(\kappa, n) = (30, 256)$ and $[-0.121, -0.120]$ at $(\kappa, n) = (300, 256)$. Both comparisons use the same 500 paired replications as the reported RMSEs.

Table 4: Complete product torus results at $\kappa = 1$. Cov. is implemented direction coverage for proposed and score-sandwich coverage for RSGD and exact Newton. Score is the score only counterfactual for the same estimate. All RMSEs are multiplied by 10^3; fractions refer to the retained averaging window. The constant c was selected before evaluation. Each row has 500 replications.

			Coverage (%)		RMSE $\times 10^3$		
n	Method	c	Cov.	Score	25%	50%	100%
256	RSGD	3.2	95.8	95.8	36.35	26.29	18.99
256	Exact Newton	3.2	95.8	95.8	36.35	26.29	18.99
256	Proposed, $m = 1$	0.8	95.8	76.2	46.82	36.69	27.01
256	Proposed, $m = 4$	1.6	95.2	91.6	39.29	29.49	21.00
256	Proposed, $m = 16$	1.6	95.8	94.4	36.85	26.79	19.57
1024	RSGD	6.4	95.2	95.2	19.38	14.04	9.84
1024	Exact Newton	6.4	95.2	95.2	19.38	14.04	9.84
1024	Proposed, $m = 1$	0.4	93.0	75.2	24.47	18.99	14.01
1024	Proposed, $m = 4$	6.4	94.6	89.6	21.50	15.82	11.11
1024	Proposed, $m = 16$	6.4	93.6	92.4	19.92	14.33	10.16
4096	RSGD	6.4	96.2	96.2	9.35	6.57	4.63
4096	Exact Newton	6.4	96.2	96.2	9.35	6.57	4.63
4096	Proposed, $m = 1$	6.4	95.6	80.0	13.58	9.28	6.61
4096	Proposed, $m = 4$	6.4	94.4	90.2	10.86	7.57	5.38
4096	Proposed, $m = 16$	6.4	95.0	94.0	9.59	6.77	4.85

Table 5: Complete product torus results at $\kappa = 30$. Cov. is implemented direction coverage for proposed and score-sandwich coverage for RSGD and exact Newton. Score is the score only counterfactual for the same estimate. All RMSEs are multiplied by 10^3; fractions refer to the retained averaging window. The constant c was selected before evaluation. Each row has 500 replications.

			Coverage (%)		RMSE $\times 10^3$		
n	Method	c	Cov.	Score	25%	50%	100%
256	RSGD	1.6	75.4	75.4	87.75	52.02	29.60
256	Exact Newton	6.4	96.6	96.6	36.18	26.87	18.85
256	Proposed, $m = 1$	0.8	95.4	78.4	46.75	36.03	26.94
256	Proposed, $m = 4$	6.4	96.4	93.6	42.95	31.40	21.60
256	Proposed, $m = 16$	6.4	98.4	97.4	37.41	27.80	19.21
1024	RSGD	0.4	96.0	96.0	19.43	13.78	9.56
1024	Exact Newton	0.4	94.8	94.8	18.16	13.65	9.56
1024	Proposed, $m = 1$	1.6	95.4	78.6	26.23	18.64	13.40
1024	Proposed, $m = 4$	0.4	95.0	91.6	20.12	15.03	10.77
1024	Proposed, $m = 16$	0.4	94.2	93.6	18.64	14.08	9.86
4096	RSGD	0.4	95.4	95.4	9.20	6.67	4.67
4096	Exact Newton	0.4	95.8	95.8	8.80	6.51	4.66
4096	Proposed, $m = 1$	0.4	95.4	77.8	12.24	9.13	6.62
4096	Proposed, $m = 4$	0.2	93.2	89.2	11.01	7.99	5.53
4096	Proposed, $m = 16$	3.2	95.2	93.8	9.91	7.01	4.81

Table 6: Complete product torus results at $\kappa = 300$. Cov. is implemented direction coverage for proposed and score-sandwich coverage for RSGD and exact Newton. Score is the score only counterfactual for the same estimate. All RMSEs are multiplied by 10^3; fractions refer to the retained averaging window. The constant c was selected before evaluation. Each row has 500 replications.

			Coverage (%)		RMSE $\times 10^3$		
n	Method	c	Cov.	Score	25%	50%	100%
256	RSGD	0.2	0.0	0.0	450.73	412.01	347.57
256	Exact Newton	3.2	95.8	95.8	35.90	25.89	18.88
256	Proposed, $m = 1$	0.8	94.4	79.0	49.39	37.28	27.19
256	Proposed, $m = 4$	0.8	95.4	92.6	36.57	27.67	20.68
256	Proposed, $m = 16$	3.2	96.6	95.4	37.68	26.91	19.29
1024	RSGD	0.4	0.0	0.0	252.63	186.01	114.03
1024	Exact Newton	6.4	95.4	95.4	18.53	13.49	9.65
1024	Proposed, $m = 1$	1.6	94.0	76.6	26.39	18.86	13.92
1024	Proposed, $m = 4$	0.4	94.8	90.6	19.26	14.48	10.86
1024	Proposed, $m = 16$	3.2	95.0	94.8	19.01	13.73	9.94
4096	RSGD	0.8	41.2	41.2	32.81	17.77	9.55
4096	Exact Newton	6.4	95.4	95.4	9.40	6.60	4.70
4096	Proposed, $m = 1$	0.2	93.4	74.0	13.13	9.57	6.88
4096	Proposed, $m = 4$	0.4	96.2	92.0	9.81	7.25	5.25
4096	Proposed, $m = 16$	0.4	95.4	94.6	9.18	6.67	4.80

Table 7: Product torus median serial time per complete path, in milliseconds. Medians are over 30 randomized-order blocks of 100 paths per cell. Timing includes observation generation and uses aggregated coordinate sketches.

κ	n	RSGD	Exact Newton	Proposed, $m = 1$	Proposed, $m = 4$	Proposed, $m = 16$
1	256	0.045	0.046	0.068	0.102	0.205
1	1024	0.182	0.180	0.257	0.407	0.799
1	4096	0.711	0.719	1.023	1.581	3.191
30	256	0.045	0.046	0.066	0.105	0.207
30	1024	0.176	0.177	0.262	0.403	0.803
30	4096	0.692	0.708	1.020	1.590	3.181
300	256	0.045	0.045	0.067	0.104	0.206
300	1024	0.182	0.176	0.259	0.401	0.806
300	4096	0.708	0.704	1.020	1.592	3.197

A.2 Computing times

Serial timing favors the inexpensive diagonal exact solve. For example, at $\kappa = 300$ and $n = 256$, RSGD and exact Newton each take about 0.045 ms per complete path, compared with 0.104 ms for proposed with $m = 4$. These measurements include data generation and one-pass inference, using aggregated coordinate counts for proposed. The observed gain is therefore in sample-budget performance relative to RSGD, rather than computation per observation.

Table 7 gives medians across 30 blocks of 100 complete paths per cell, with randomized method order within each block. These times include observation generation, updates, averaging, covariance tracking, and numeric output construction. Coordinate counts are sampled as binomials, so these timings describe an aggregated coordinate solver, not m general matrix solves.

A.3 WDBC analysis details

The Wisconsin Diagnostic Breast Cancer data provide a heterogeneous stream of correlated image measurements. The official UCI release contains 569 aspirate records and 30 continuous features derived from digitized images of cell nuclei, including measures of size, texture, boundary irregularity, smoothness, symmetry, and fractal dimension (Wolberg et al., 1993; 1994). It reports no missing values and distributes the data under CC BY 4.0. The analysis is an unsupervised dimension reduction illustration, not an evaluation of diagnostic accuracy.

We allocate 169 label free records to calibration and stream the remaining 400. Calibration means and standard deviations define the feature scale; the calibration covariance supplies the initial line and a fixed inverse eigengap preconditioner. Diagnosis labels play no role in either step. The stream uses $\eta_k = 0.45k^{-0.65}$, a maximum Grassmann step of 0.30 radians, and a 50-observation burn-in. Conditional on calibration, the target is the leading eigenspace of the second moment matrix about the fixed calibration center. Pointwise Wald intervals use the implemented direction covariance and the delta method under a working i.i.d. model for the analysis records. After burn-in, the average, direction covariance, and analysis-stream Hessian accumulator are restarted; the Wald multiplier is the 350 retained observations. The Hessian estimate is the retained average of the PCA score derivative, distinct from the fixed calibration preconditioner used in the updates.

Individual loading signs are arbitrary, so we summarize the leading rank one projector Π by its mass on five scientifically interpretable feature groups that partition the 30 standardized coordinates. If D_g is the coordinate-selection projector for group g, its mass is $\theta_g = \mathrm{tr}(\mathrm{D}_g\Pi)$, so $\sum_g \theta_g = \mathrm{tr}(\Pi) = 1$. Boundary irregularity has estimated mass 0.455 with interval $[0.390, 0.521]$, while the size group has mass 0.392 with interval $[0.254, 0.530]$ (Table 8). Texture contributes relatively little, with mass 0.023 and interval $[0.005, 0.041]$. These quantities describe second moment structure in this sample; they are not causal effects of cell morphology.

Table 8: WDBC leading-projector mass by morphology-based feature group. Intervals are pointwise 95% delta-method intervals under the conditional working model.

Feature group	Projector mass	95% interval
Size (radius/perimeter/area)	0.392	[0.254, 0.530]
Boundary irregularity	0.455	[0.390, 0.521]
Texture	0.023	[0.005, 0.041]
Smoothness and symmetry	0.092	[0.041, 0.143]
Fractal dimension	0.038	[0.004, 0.071]

B THEORY SUPPLEMENTS AND PROOFS

This appendix develops the local conditions and implementation underlying the main results, then gives solver comparisons and model-specific criteria. Section B.6 contains the proofs and explicit constructions.

B.1 AVERAGED RANDOMIZED NEWTON INFERENCE

The common tangent representation links the randomized recursion to the influence expansion. It also specifies the directions and derivative quantities needed for one-pass inference.

B.1.1 TARGET AND COMMON TANGENT REPRESENTATION

Let $(\mathcal{M}, g)$ be a d-dimensional Riemannian manifold. A population estimating field assigns a tangent vector $G(x) \in T_x\mathcal{M}$ to each x in its domain. We study a selected root $x_\star$ satisfying $G(x_\star) = 0$. For optimization, $G = \operatorname{grad} F$; more generally, G may be a population estimating equation. Its local differentiability and nondegeneracy are specified in Assumption B.1.

Let R be a C^2 retraction and put $E = T_{x_\star}\mathcal{M}$. Fix a sufficiently small working neighborhood U whose closure is contained in a larger common domain of the retractions and local inverse retractions used below. The domain is chosen before the reported stochastic approximation phase. Choose it inside a normal neighborhood of $x_\star$; the coordinate used here is $R_{x_\star}^{-1}(x) \in E$, with $u_k = R_{x_\star}^{-1}(x_k)$; for a generic retraction it need not equal $\operatorname{Log}_{x_\star}(x)$. All local derivative bounds are taken on compact subsets of the larger domain, and all inverse-retraction expressions below are restricted to that common domain.

Fix a smooth orthonormal frame $\mathfrak{e}_x : E \to T_x\mathcal{M}$ with $\mathfrak{e}_{x_\star} = \mathrm{I}$, and define $\mathcal{T}_{x\to y} = \mathfrak{e}_y\mathfrak{e}_x^*$. This identification is an isometry and satisfies $\mathcal{T}_{y\to z}\mathcal{T}_{x\to y} = \mathcal{T}_{x\to z}$; it is not generally direct Levi-Civita parallel transport from x to y. A computable frame may be fixed using the public anchor: identifying its coordinates with E is only an analytical change of basis and does not require knowledge of $x_\star$ to compute $\mathcal{T}_{x\to y}$. For an operator L, L^* denotes its adjoint. For vectors in a common inner-product space, $(v \otimes w)z = v\,\langle w, z\rangle$. We use $\|\cdot\|$ for the tangent norm, $\|\cdot\|_{\mathrm{op}}$ for the induced operator norm, and $\|\mathrm{L}\|_{\mathrm{F}}^2 = \operatorname{tr}(L^*\mathrm{L})$ for the Frobenius norm. For a random operator, $\|\mathrm{L}\|_{L^2(\mathrm{op})} = \{\mathbb{E}\,\|\mathrm{L}\|_{\mathrm{op}}^2\}^{1/2}$, with analogous notation for the Frobenius and vector norms.

B.1.2 RECURSION AND ONLINE AVERAGE

At time $k+1$, let $\widehat{G}_{k+1}(x_k)$ be a stochastic score and let $\mathrm{P}_{k+1} : T_{x_k}\mathcal{M} \to T_{x_k}\mathcal{M}$ be the random linear operator produced by an inner solver. The update is

$$x_{k+1} = R_{x_k}\left[-\eta_{k+1}\{\mathrm{P}_{k+1}\widehat{G}_{k+1}(x_k) + e_{k+1}\}\right], \qquad \eta_k = \eta_0 k^{-\alpha}, \quad \alpha \in (1/2, 1), \tag{15}$$

where e_{k+1} is the discrepancy between the implemented direction and a chosen linear reference. If an inner algorithm is initialized at zero and uses coefficients independent of the current right hand side, a fixed number of linear iterations is exactly represented by P_{k+1} and $e_{k+1} = 0$. Data-dependent stopping or residual-based acceleration can be handled by retaining a linear reference and placing the remaining error in e_{k+1}.

Two online averages are considered. The recursive Riemannian average is

$$\bar{x}_1^{\mathrm{rec}} = x_1, \qquad \bar{x}_k^{\mathrm{rec}} = R_{\bar{x}_{k-1}^{\mathrm{rec}}}\left[\frac{1}{k} R^{-1}_{\bar{x}_{k-1}^{\mathrm{rec}}}(x_k)\right], \qquad k \geq 2. \tag{16}$$

A fixed anchor chart average, which is often simpler for streaming inference, is

$$y_1^a = R_a^{-1}(x_1), \qquad y_k^a = \left(1 - \frac{1}{k}\right) y_{k-1}^a + \frac{1}{k} R_a^{-1}(x_k), \qquad \bar{x}_k^a = R_a(y_k^a). \tag{17}$$

The anchor a is fixed before the reported averaging phase. The main theorem applies once the first order equivalence in Assumption B.6 is verified. For both averages, a common-domain C^2 retraction and the state rate suffice on the fixed sufficiently small working domains in Lemmas B.33 and B.34.

A deterministic burn-in may be omitted when its cumulative state contribution is root-n negligible. In particular, under (34), omitting the first $k_0 = \lfloor n^{1/2} \rfloor$ iterates is first order negligible because $n^{-1/2} \sum_{k \leq k_0} \mathbb{E} \|u_k\| = O(n^{-\alpha/4})$. For a proportional burn-in, restarting the average defines a tail estimator rather than a negligible deletion; Proposition B.9 gives its root-N_n limit in terms of the retained count.

Let $\mathcal{F}_k$ contain the initial point, observations, derivative information, iterates, and solver randomness through time k. The current score and new solver seed are not included in $\mathcal{F}_k$. Their joint conditional law may be dependent unless a sample split is imposed. All current scores, solver operators, inexactness variables, and updated iterates are assumed $\mathcal{F}_{k+1}$-measurable, and every conditional expectation displayed below is well defined and integrable to the stated order.

B.1.3 One-pass covariance tracking

The following implementation uses a fixed public anchor $a \in U$. All running matrices live in $T_a\mathcal{M}$, while actual directions are mapped to $T_{x_k}\mathcal{M}$. The displayed derivative branch is for a gradient field, where the Jacobian is a self-adjoint Hessian and spectral safeguarding is meaningful. For a generic estimating field with a possibly nonsymmetric derivative H, the same statistical theorem applies only after replacing this branch by a consistent Jacobian estimator and an invertibility/stability safeguard appropriate to nonsymmetric operators. The Hessian and score mini-batches may be literal independent splits, cross-fitted streams, or any construction satisfying the product conditions below. The map $\mathcal{K}_J(\mathrm{B}, \Omega)$ sends a right hand side to the output of J zero-start linear-solver steps, with coefficients fixed by the curvature operator B and seed Ω. Right-hand-side-dependent coefficients, including ordinary finite-step conjugate gradients, require an explicit linear reference and an error term. For a self-adjoint operator, $\mathrm{SpecClip}_{[h_-,h_+]}$ clips its eigenvalues to the specified positive interval. The raw Hessian average is stored separately from this safeguarded estimate. For a vector z, the Welford update of a count-k mean μ and centered sum of squares W is

$$\delta = z - \mu, \qquad \mu^+ = \mu + \delta/(k+1), \qquad \mathrm{W}^+ = \mathrm{W} + \delta \otimes (z - \mu^+).$$

The localization map $\mathcal{L}$ is either the identity on an independently verified invariant basin, or a specified map with its own drift and error proof. It must leave the actual direction in the common inverse-retraction domain and satisfy the actual-step and error conditions used by the theorem. A step cap alone does not supply this property. The predictable covariance branch is used only under Theorem 3.6 and corollary B.16.

For $c > 0$, radial clipping is $\mathrm{clip}_c(z) = \min\{1, c/\|z\|\} z$ for $z \neq 0$, and $\mathrm{clip}_c(0) = 0$.

The full discrepancy in (15) is $e_{k+1} = D_{k+1} - \mathrm{P}_{k+1} \widehat{G}_{k+1}(x_k)$, including inner error, clipping, and localization. The displayed Hessian branch concerns a gradient field; a nonsymmetric estimating field requires a consistent Jacobian estimator and an appropriate alternative safeguard. Full dimensional Wald inference additionally requires a positive definite reported covariance and the relevant inverse-retraction chart. A singular covariance can still be returned; a scalar interval requires a positive estimated variance for its contrast. For deterministic burn-in, keep the global step size index and restart reporting accumulators with the retained count. If the optimizer learns its Hessian before burn-in, retain its optimizer state separately from a restarted reporting Hessian average. A full fresh solver covariance requires access to a consistent mean solver estimator, not only the ability to apply the solver to a vector.

Algorithm 2 Riemannian Newton averaging with online inference

Require: x_0, public anchor a, horizon n, common-domain C^2 retraction and coherent frame; averaging rule; step sizes η_k; solver budgets J_k, m_k and specified inner-error rule; ridge λ_k; spectral band $[h_-, h_+]$; safeguarded $\widehat{\mathrm{H}}_0^a$; step cutoff r_0 and justified localization map $\mathcal{L}$; fresh or admissible predictable covariance regime.

Initialize the raw Hessian mean $\mathrm{M}_0^a = 0$, solver mean $\widehat{\mathrm{P}}_0^a = 0$, and Welford state $(\mu_0, \mathrm{W}_0) = (0, 0)$.
for $k = 0, \dots, n-1$ **do**
  Set $\mathrm{B}_k^{(x_k)} = \mathcal{T}_{a \to x_k} \widehat{\mathrm{H}}_k^a \mathcal{T}_{x_k \to a}$.
  if the fresh solver regime is selected **then**
    Draw m_k conditionally independent solver seeds, independent of the current observation.
    Set $\mathrm{P}_{k+1} = m_k^{-1} \sum_{j=1}^{m_k} \mathcal{K}_{J_k}(\mathrm{B}_k^{(x_k)}, \Omega_{k+1}^{(j)})$.
  else choose the specified invertible predictable P_{k+1}, for example $(\mathrm{B}_k^{(x_k)})^{-1}$.
  end if
  Obtain the current score $\widehat{G}_{k+1}(x_k)$ and derivative contribution $\widehat{\mathrm{B}}_{k+1}$.
  Form $D_{k+1}^0 = \mathrm{P}_{k+1} \widehat{G}_{k+1}(x_k) + e_{k+1}^{\mathrm{inner}}$ and $D_{k+1}^{\mathrm{cap}} = \mathrm{clip}_{r_0/\eta_{k+1}}(D_{k+1}^0)$.
  Set $x_{k+1} = \mathcal{L}\big(R_{x_k}[-\eta_{k+1} D_{k+1}^{\mathrm{cap}}]\big)$ and $D_{k+1} = -\eta_{k+1}^{-1} R_{x_k}^{-1}(x_{k+1})$.
  if $k = 0$ **then** initialize the average at x_1; **else** update it by (17) or (16).
  Set $\widehat{\mathrm{B}}_{k+1}^a = \mathcal{T}_{x_k \to a} \widehat{\mathrm{B}}_{k+1} \mathcal{T}_{a \to x_k}$ and $\mathrm{M}_{k+1}^a = (k\mathrm{M}_k^a + \widehat{\mathrm{B}}_{k+1}^a)/(k+1)$.
  Set $\widehat{\mathrm{H}}_{k+1}^a = \mathrm{SpecClip}_{[h_-, h_+]}(\mathrm{M}_{k+1}^a + \lambda_{k+1}\mathrm{I})$.
  if the fresh solver regime is selected **then**
    Set $\widehat{\mathrm{P}}_{k+1}^a = \{k\widehat{\mathrm{P}}_k^a + \mathcal{T}_{x_k \to a} \mathrm{P}_{k+1} \mathcal{T}_{a \to x_k}\}/(k+1)$ and $z_{k+1} = \mathcal{T}_{x_k \to a} D_{k+1}$.
  else set $z_{k+1} = \mathcal{T}_{x_k \to a} \widehat{G}_{k+1}(x_k)$.
  end if
Apply the count-k Welford update toz_{k+1}, obtaining $(\mu_{k+1}, \mathrm{W}_{k+1})$.
end for
if fresh **then** set $\widehat{\mathrm{A}}_n^a = \widehat{\mathrm{P}}_n^a \widehat{\mathrm{H}}_n^a$ and, when invertible, $\widehat{\mathrm{V}}_n^a = (\widehat{\mathrm{A}}_n^a)^{-1}(\mathrm{W}_n/n)((\widehat{\mathrm{A}}_n^a)^{*-1})$.
else set $\widehat{\mathrm{V}}_n^a = (\widehat{\mathrm{H}}_n^a)^{-1}(\mathrm{W}_n/n)((\widehat{\mathrm{H}}_n^a)^{*-1})$.
end if
return $\bar{x}_n$, count n, and $\mathcal{T}_{a \to \bar{x}_n} \widehat{\mathrm{V}}_n^a \mathcal{T}_{\bar{x}_n \to a}$; report unavailable covariance if the required inverse does not exist.

When the update uses a predictable preconditioner that is not assumed to converge, the mean operator branch in the final covariance formula is not appropriate. In that regime, use the raw transported score

$$G^a_{k+1} = \mathcal{T}_{x_k \to a} \widehat{G}_{k+1}(x_k) \tag{18}$$

as input to the Welford accumulator. Denote its covariance by $\widehat{\mathrm{S}}^a_n = \mathrm{W}_n/n$. and report

$$\widehat{\mathrm{V}}^{a,\mathrm{pred}}_n = (\widehat{\mathrm{H}}^a_n)^{-1} \widehat{\mathrm{S}}^a_n ((\widehat{\mathrm{H}}^a_n)^{*-1}). \tag{19}$$

Corollary B.16 shows that this branch remains consistent when the predictable preconditioner has no fixed limit. The direction covariance branch is used when the solver is randomized at the current iteration because it retains the corresponding product noise.

Algorithm 2 stores a constant number of $d \times d$ operators, requiring $O(d^2)$ memory. A matrix-free direction routine does not by itself provide the full mean solver estimator. Structured or contrast-specific tracking requires a consistent estimator of the corresponding covariance; a diagonal approximation is justified only when the target structure or the chosen contrast makes it sufficient.

B.1.4 Local conditions for inference

The local conditions separate the target field, the current solver score product, and the errors that accumulate under averaging. The state rate criterion in Proposition B.10 is verified directly from the actual update. The geometric and averaging propositions then supply the corresponding remainder conditions under their explicit domain and moment hypotheses.

Assumption B.1. *The reported phase starts at a deterministic $x_0 \in U$, and the actual updates remain in the fixed working neighborhood U almost surely. The reported average belongs to the common inverse-retraction domain with probability tending to one. On U,*

$$\mathcal{T}_{x \to x_\star} G(x) = \mathrm{H} R^{-1}_{x_\star}(x) + r_G(x), \qquad \|r_G(x)\| = o(\|R^{-1}_{x_\star}(x)\|) \quad (x \to x_\star), \tag{20}$$

where $\mathrm{H} : E \to E$ is invertible. The working domain satisfies the uniform bound $\|\mathrm{H} R^{-1}_{x_\star}(x) + r_G(x)\| \le C \|R^{-1}_{x_\star}(x)\|$ for $x \in U$.

Differentiability at the root gives the uniform field bound on a sufficiently small neighborhood, which must be selected before imposing localization. Alternatively, continuity on a compact larger domain, together with differentiability at the root, gives this bound on a fixed compact working basin. A step-length cap alone does not prove basin invariance; that property is supplied by a model-specific argument or an explicit localization mechanism. An independent initialization phase may be conditioned on, but global root selection is not part of the local inference theorem. At a nondegenerate local minimum of a twice differentiable objective, H is positive definite; a strict minimum alone does not imply this property.

Assumption B.2. *For $x_k \in U$,*

$$\mathcal{T}_{x_k \to x_\star} \widehat{G}_{k+1}(x_k) = \mathrm{H} u_k + r_G(x_k) + \xi_{k+1} + q_{k+1},$$

where the target score $\xi_{k+1} \in E$ and centered score increment $q_{k+1} \in E$ satisfy

$$\mathbb{E}(\xi_{k+1} \mid \mathcal{F}_k) = 0, \qquad \mathbb{E}(q_{k+1} \mid \mathcal{F}_k) = 0. \tag{21}$$

For fresh i.i.d. observations and a tangent-valued sample field $\mathrm{g}(x, Z)$ with $\mathbb{E}\mathrm{g}(x, Z) = G(x)$, a consistent choice is

$$\xi_{k+1} = \mathrm{g}(x_\star, Z_{k+1}), \qquad q_{k+1} = \mathcal{T}_{x_k \to x_\star} \mathrm{g}(x_k, Z_{k+1}) - \mathrm{g}(x_\star, Z_{k+1}) - \mathcal{T}_{x_k \to x_\star} G(x_k).$$

Every term is in E. An L^2 local derivative envelope for the transported sample field gives $\mathbb{E}(\|q_{k+1}\|^2 \mid \mathcal{F}_k) \le C \|u_k\|^2$. To obtain a bound for its product with the solver, one also needs the operator or joint-product conditions below.

Write

$$\widetilde{\mathrm{P}}_{k+1} = \mathcal{T}_{x_k \to x_\star} \mathrm{P}_{k+1} \mathcal{T}_{x_\star \to x_k}, \qquad \Pi_k = \mathbb{E}(\widetilde{\mathrm{P}}_{k+1} \mid \mathcal{F}_k),$$

and define the two conditional product means and their centered parts by

$$b^\xi_{k+1} = \mathbb{E}(\widetilde{\mathrm{P}}_{k+1} \xi_{k+1} \mid \mathcal{F}_k), \qquad \zeta_{k+1} = \widetilde{\mathrm{P}}_{k+1} \xi_{k+1} - b^\xi_{k+1}, \tag{22}$$

$$b^q_{k+1} = \mathbb{E}(\widetilde{\mathrm{P}}_{k+1} q_{k+1} \mid \mathcal{F}_k), \qquad \chi_{k+1} = \widetilde{\mathrm{P}}_{k+1} q_{k+1} - b^q_{k+1}. \tag{23}$$

Assumption B.3. *There is a deterministic invertible $\bar{\mathrm{P}} : E \to E$ such that $\mathrm{A} = \bar{\mathrm{P}}\mathrm{H}$ is positive stable, meaning that all its eigenvalues have positive real part. For a deterministic constant C,*

$$\sup_k \mathbb{E}(\left\|\widetilde{\mathrm{P}}_{k+1}\right\|_{\mathrm{op}}^2 \mid \mathcal{F}_k) \le C, \tag{24}$$

$$\mathbb{E}(\|\chi_{k+1}\|^2 \mid \mathcal{F}_k) \le C \|u_k\|^2. \tag{25}$$

The population remainders satisfy

$$\frac{1}{\sqrt{n}} \sum_{k=0}^{n-1} (\Pi_k - \bar{\mathrm{P}})\mathrm{H}u_k \xrightarrow{p} 0, \tag{26}$$

$$\frac{1}{\sqrt{n}} \sum_{k=0}^{n-1} \Pi_k r_G(x_k) \xrightarrow{p} 0. \tag{27}$$

Conditional independence of the current solver from the whole current observation is one sufficient route to product centering and moment bounds; it is not required by the theorem. Separate unbiasedness or separate second moments do not, by themselves, establish the corresponding properties of a dependent product. Given the state rate below, $\left\|\Pi_k - \bar{\mathrm{P}}\right\|_{L^2(\mathrm{op})} = O(k^{-\delta})$ with $\delta > (1-\alpha)/2$, and $\|r_G(x)\| \le C \left\|R_{x_\star}^{-1}(x)\right\|^{1+\rho}$ with $\alpha(1+\rho) > 1$, are sufficient for the two accumulated remainders. The direct sum conditions permit weaker routes when additional cancellation is available.

Decompose the complete implementation discrepancy, including all safeguard modifications, as

$$\widetilde{e}_{k+1} = \mathcal{T}_{x_k \to x_\star} e_{k+1} = a_{k+1} + \nu_{k+1}, \quad a_{k+1} = \mathbb{E}(\widetilde{e}_{k+1} \mid \mathcal{F}_k), \quad \mathbb{E}(\nu_{k+1} \mid \mathcal{F}_k) = 0. \tag{28}$$

Put $\beta_{k+1} = b_{k+1}^{\xi} + b_{k+1}^{q} + a_{k+1}$.

Assumption B.4. *For a deterministic $b \in E$,*

$$\frac{1}{\sqrt{n}} \sum_{k=0}^{n-1} \beta_{k+1} \xrightarrow{p} b, \tag{29}$$

and

$$\frac{1}{n} \sum_{k=0}^{n-1} \mathbb{E}(\|\nu_{k+1}\|^2 \mid \mathcal{F}_k) \xrightarrow{p} 0. \tag{30}$$

The limit b determines a possible Gaussian centering shift. A centered inner error with nonvanishing average variance must instead be included in the leading innovation and its covariance; it cannot be discarded merely because its conditional mean is zero.

Assumption B.5. *In the coordinate u, the actual update admits*

$$\begin{aligned} u_{k+1} = u_k - \eta_{k+1}\Big[\widetilde{\mathrm{P}}_{k+1}\{\mathrm{H}u_k + r_G(x_k) + \xi_{k+1} + q_{k+1}\} + a_{k+1} + \nu_{k+1}\Big] \\ + \eta_{k+1} c_{k+1} + \eta_{k+1} d_{k+1}, \end{aligned} \tag{31}$$

where $\mathbb{E}(c_{k+1} \mid \mathcal{F}_k) = 0$ and

$$\frac{1}{n} \sum_{k=0}^{n-1} \mathbb{E}(\|c_{k+1}\|^2 \mid \mathcal{F}_k) \xrightarrow{p} 0, \qquad \frac{1}{\sqrt{n}} \sum_{k=0}^{n-1} d_{k+1} \xrightarrow{p} 0. \tag{32}$$

The remainder d_{k+1} need not be predictable.

These are accumulated error conditions, not a consequence of smoothness alone. A common-domain C^2 retraction and frame, the state rate, actual bounded steps, and the direction and bias moments in Proposition B.27 imply them.

Assumption B.6. *With $\widetilde{u}_n = n^{-1} \sum_{k=1}^{n} u_k$, the chosen average satisfies*

$$\sqrt{n}\{R_{x_\star}^{-1}(\bar{x}_n) - \widetilde{u}_n\} \xrightarrow{p} 0. \tag{33}$$

For the fixed anchor and recursive averages, the state rate and common-domain C^2 regularity on the fixed sufficiently small domains imply this equivalence by Lemmas B.33 and B.34. No locally Lipschitz second derivative or third derivative is required for either rule.

Assumption B.7. *The local state satisfies*

$$\mathbb{E}\,\|u_k\|^2 \le C\eta_k. \tag{34}$$

For a deterministic positive semidefinite operator $\mathrm{Q} : E \to E$*,*

$$\frac{1}{n}\sum_{k=0}^{n-1}\mathbb{E}(\zeta_{k+1}\otimes\zeta_{k+1}\mid\mathcal{F}_k)\xrightarrow{p}\mathrm{Q}, \tag{35}$$

and, for every $\epsilon > 0$*,*

$$\frac{1}{n}\sum_{k=0}^{n-1}\mathbb{E}\Big[\|\zeta_{k+1}\|^2\,\mathbf{1}\{\|\zeta_{k+1}\| > \epsilon\sqrt{n}\}\mid\mathcal{F}_k\Big]\xrightarrow{p}0. \tag{36}$$

The state rate follows from the direct Lyapunov criterion in Proposition B.10; it is not used to prove that criterion's one step drift inequality. The covariance limit and Lindeberg condition are separate noise requirements. A uniform conditional $(2+\tau)$ moment of ζ_{k+1} is sufficient for Lindeberg, but the main theorem does not require a fixed moment above two. Positive semidefinite Q permits a degenerate Gaussian limit; full dimensional Wald inversion requires positive definiteness.

In Euclidean stochastic approximation with $R_x(v) = x + v$ and identity transport, the geometric errors vanish and both averages are arithmetic. The remaining restrictions concern stable drift, martingale noise, accumulated solver/error bias, and feasible operator estimation. On a manifold, common coordinate domains and quadratic geometric remainders are additionally needed. The verification order is state rate, population and geometric remainders, average equivalence, the CLT, and covariance consistency. Example B.28 gives a fully checkable instance; the local theorem does not assert that an arbitrary optimizer selects or remains in the required basin.

B.2 Theoretical results

B.2.1 Influence representation

The local recursion has drift $\mathrm{A} = \bar{\mathrm{P}}\mathrm{H}$, but its leading innovation is not $\bar{\mathrm{P}}\xi_{k+1}$. It is the centered realized product ζ_{k+1} in (22). The representation in Theorem 3.1 connects that temporal distinction to inference.

The main-text statement is Theorem 3.1.

The mean pull toward the root is governed by $\mathrm{A} = \bar{\mathrm{P}}\mathrm{H}$, while the leading fluctuation is the realized product of the solver and target score. Replacing this product by the mean solver applied to the score would remove randomness that the algorithm introduces. The limit b records conditional product or implementation bias that survives at the root-n scale. Its mean shift $-\mathrm{A}^{-1}b$ cannot be removed by covariance estimation alone: centered Wald inference requires $b = 0$, or an available consistent estimate of b and the plus-sign correction in (9).

No Hölder exponent, polynomial solver rate, or moment above order two is primitive to the theorem. Proposition B.26 records convenient rate conditions for the two accumulated remainders, and Proposition B.10 supplies the mean-square state rate without a fourth moment recursion.

Corollary B.8. *Under Theorem 3.1, if*

$$\frac{1}{\sqrt{n}}\sum_{k=0}^{n-1}\{b_{k+1}^{\xi} + b_{k+1}^{q} + a_{k+1}\}\xrightarrow{p}0,$$

then

$$\sqrt{n}\,R_{x_\star}^{-1}(\bar{x}_n)\xrightarrow{d}N_E(0,\mathrm{V}).$$

In particular, this holds under product orthogonality and root-n negligible inexactness bias.

Proposition B.9. *Under the centered conditions of Corollary B.8, let (b_n) be a deterministic integer sequence with $0 \le b_n < n$, $b_n/n \to \tau \in [0,1)$, and $N_n = n - b_n$. Restart either online average at x_{b_n+1} and suppose its tangent coordinate satisfies the tail analogue of* (33)*,*

$$\sqrt{N_n}\left\{R_{x_\star}^{-1}(\bar{x}_{n,b_n}) - \frac{1}{N_n}\sum_{k=b_n+1}^{n} u_k\right\} \xrightarrow{p} 0.$$

Then

$$\sqrt{N_n}\, R_{x_\star}^{-1}(\bar{x}_{n,b_n}) \xrightarrow{d} N_E(0, \mathrm{V}), \qquad \mathrm{V} = \mathrm{A}^{-1}\mathrm{Q}(\mathrm{A}^{*-1}).$$

The displayed tail equivalence holds for the fixed anchor average under the same local C^2 chart conditions used in Lemma B.33; it also holds for the restarted recursive average under the conditions of Lemma B.34.

B.2.2 Mean-square localization

The local rate (34) can be verified without a fourth moment recursion.

Proposition B.10. *Let $\bar{\mathrm{P}}$ be deterministic with $\mathrm{A} = \bar{\mathrm{P}}\mathrm{H}$ positive stable. Suppose there are a deterministic index $k_\star$ and a deterministic $\mathrm{K} \succ 0$ such that, on the localized trajectory, almost surely for every $k \ge k_\star$,*

$$(\Pi_k \mathrm{H})^* \mathrm{K} + \mathrm{K}\Pi_k \mathrm{H} \succeq c_0 \mathrm{I} \tag{37}$$

for a deterministic $c_0 > 0$. Suppose Assumptions B.1 and B.2 and the error decomposition (28) *hold. Without assuming Assumptions B.3 to B.5 and B.7, assume that the retraction and local frame are C^2, $\sup_k \|\Pi_k\|_{\mathrm{op}} \le C$, and the actual direction satisfies $\eta_{k+1}\|D_{k+1}\| \le r_0$ and* (79)*, and $\mathbb{E}\|\beta_{k+1}\|^2 \le C\eta_{k+1}$ for $\beta_{k+1} = b_{k+1}^{\xi} + b_{k+1}^{q} + a_{k+1}$. If U is chosen sufficiently small before imposing the stated localization, then $\mathbb{E}\|u_k\|^2 \le C\eta_k$. One convenient route is to take K as the solution of $\mathrm{A}^*\mathrm{K} + \mathrm{K}\mathrm{A} = \mathrm{I}$ and impose a deterministic eventual uniform operator-norm bound on $\Pi_k - \bar{\mathrm{P}}$. The rank one closure below instead verifies* (37) *directly with $\mathrm{K} = \mathrm{H}$.*

The proposition is deliberately upstream of the geometric reduction: it uses the direct two-variable retraction expansion and the direction second moment bound, not Assumption B.5. Once the state rate is known, Proposition B.27 turns the same local expansion into the two root-n geometric remainder bounds. This order removes any need to assume the conclusion of one result while verifying the other.

A local state rate and eventual localization are logically distinct. The proposition controls the rate once the trajectory is in the stable neighborhood. Global convergence can be established under the usual coercivity, projection, or truncation conditions for stochastic approximation; such conditions are model-dependent and are not needed for the local influence-function calculation. Similar localization is standard in Riemannian stochastic approximation (Bonnabel, 2013; Tripuraneni et al., 2018; Zhou and Said, 2019).

B.2.3 Feasible covariance estimation and Wald inference

The influence representation identifies what must be tracked online: the covariance of the implemented direction in a common frame, together with the mean solver and derivative. This is stronger than storing a score covariance when the current solve is random, but it remains a one-pass calculation.

Let $a \in U$ be the anchor used by Algorithm 2, and define the anchor representations

$$\mathrm{H}_a = \mathcal{T}_{x_\star \to a}\mathrm{H}\mathcal{T}_{a \to x_\star}, \qquad \bar{\mathrm{P}}_a = \mathcal{T}_{x_\star \to a}\bar{\mathrm{P}}\mathcal{T}_{a \to x_\star},$$
$$\mathrm{Q}_a = \mathcal{T}_{x_\star \to a}\mathrm{Q}\mathcal{T}_{a \to x_\star}, \qquad \mathrm{V}_a = (\bar{\mathrm{P}}_a\mathrm{H}_a)^{-1}\mathrm{Q}_a((\bar{\mathrm{P}}_a\mathrm{H}_a)^{*-1}).$$

The main-text statement is Assumption 3.2.

The squared bias condition is additional to the signed root-n bias condition: alternating errors can cancel in an average while remaining visible in an empirical covariance. Under the preceding assumptions, including the uniform field bound on the actual working basin, Lemma B.39 derives the required realized-error and outer-product laws. No fourth moment is needed for this consistency

argument. The streaming Hessian and solver propositions provide sufficient conditions for the two operator estimators.

For the fresh solver branch, define the implemented directions and their sample covariance in the anchor space by

$$Y_{k+1} = \mathcal{T}_{x_k \to a} D_{k+1}, \qquad \bar{Y}_n = \frac{1}{n}\sum_{k=0}^{n-1} Y_{k+1},$$

$$\widehat{\mathrm{Q}}_n^a = \frac{1}{n}\sum_{k=0}^{n-1}(Y_{k+1} - \bar{Y}_n) \otimes (Y_{k+1} - \bar{Y}_n) = \mathrm{W}_n/n, \qquad \widehat{\mathrm{A}}_n^a = \widehat{\mathrm{P}}_n^a \widehat{\mathrm{H}}_n^a.$$

When $\widehat{\mathrm{A}}_n^a$ is invertible, put $\widehat{\mathrm{V}}_n^a = (\widehat{\mathrm{A}}_n^a)^{-1}\widehat{\mathrm{Q}}_n^a((\widehat{\mathrm{A}}_n^a)^{*-1})$ and $\widehat{\mathrm{V}}_n^{\bar{x}} = \mathcal{T}_{a\to\bar{x}_n}\widehat{\mathrm{V}}_n^a\mathcal{T}_{\bar{x}_n\to a}$. For convergence statements, extend these covariance operators by zero when the required inverse is unavailable.

The main-text statement is Theorem 3.3.

The recorded-discrepancy bias construction is given in Section B.6.2.

The direction covariance, rather than a covariance of the iterates, is the natural estimator for the Polyak-Ruppert influence sum. It automatically includes solver randomness but does not, by itself, estimate the signed local bias b. Thus covariance consistency and bias centering are separate requirements. In particular, using only an unpreconditioned score covariance and then reporting $\mathrm{H}^{-1}\widehat{\mathrm{S}}(\mathrm{H}^{*-1})$ is inconsistent for a fixed-budget randomized sketch unless the inflation term vanishes. The estimator in (6) also differs from the weighted iterate-covariance estimator used for the last iterate Lyapunov limit of Wang et al. (2026).

B.2.4 Deterministic and asymptotically deterministic conditioning

Let S denote the deterministic limit, when it exists, in

$$\frac{1}{n}\sum_{k=0}^{n-1}\mathbb{E}(\xi_{k+1}\otimes\xi_{k+1} \mid \mathcal{F}_k) \xrightarrow{p} \mathrm{S}. \tag{38}$$

Corollary B.11. *Suppose the centered conditions of Theorem 3.1 hold with $\bar{\mathrm{P}} = \mathrm{P}_0$, where P_0 is deterministic and invertible and $\mathrm{P}_0\mathrm{H}$ is positive stable. Assume*

$$\frac{1}{n}\sum_{k=0}^{n-1}\mathbb{E}\{\left\|(\widetilde{\mathrm{P}}_{k+1} - \mathrm{P}_0)\xi_{k+1}\right\|^2 \mid \mathcal{F}_k\} \xrightarrow{p} 0.$$

If (38) *holds, then*

$$\mathrm{V} = \mathrm{V}_{\mathrm{sand}} := \mathrm{H}^{-1}\mathrm{S}(\mathrm{H}^{*-1}). \tag{39}$$

This includes $\mathrm{P}_0 = \mathrm{I}$, exact Newton conditioning $\mathrm{P}_0 = \mathrm{H}^{-1}$ in the optimization case, and every fixed deterministic invertible preconditioner for which $\mathrm{P}_0\mathrm{H}$ is positive stable.

In the applications below, "asymptotically deterministic" means that the score-product condition in Corollary B.11 and the covariance limit (38) hold, together with the centered conditions of Theorem 3.1. Convergence $\widetilde{\mathrm{P}}_{k+1} \xrightarrow{p} \mathrm{P}_0$ alone is insufficient.

This is the established preconditioner-invariance mechanism of averaged stochastic approximation (Polyak and Juditsky, 1992; Boyer and Godichon-Baggioni, 2023; An and Huo, 2026a); it locates the additional fresh solver covariance relative to that benchmark.

Under a correctly specified likelihood score, the information identity can make (39) equal to inverse Fisher information. For a generic loss or estimating equation, the sandwich instead describes the covariance of the corresponding M-estimator.

The local contraction comparison is given in Proposition 3.5.

B.2.5 Randomization in the current solve and covariance inflation

The sandwich cancellation above fails only through directions in which the solver varies and the target score has energy. Independence makes this failure an exact covariance decomposition.

The main-text statement is Theorem 3.4.

If S is positive definite, equality requires $\widetilde{\mathrm{P}} = \bar{\mathrm{P}}$ almost surely; for singular S, randomness outside the score's range is harmless. This comparison fixes the target score and uses independent solver seeds. Computational benefits of a cheaper solve depend separately on the finite-budget transient and its implementation cost.

Corollary B.12. *Assume the centered conditions of Theorem 3.1. Suppose that, conditionally on $\mathcal{F}_k$, the target score has the same covariance* S *at every iteration. Let m_k be a positive, $\mathcal{F}_k$-measurable integer and let*

$$\widetilde{\mathrm{P}}_{k+1}^{(m_k)} = \frac{1}{m_k} \sum_{j=1}^{m_k} \widetilde{\mathrm{P}}_{k+1}^{(j)},$$

where the m_k operators are conditionally independent copies from a common time-homogeneous solver law with mean $\bar{\mathrm{P}}$ and variance operator R *as in* (10)*, independent of the current score. If*

$$\frac{1}{n} \sum_{k=1}^{n} \frac{1}{m_k} \xrightarrow{p} \lambda \in [0, 1],$$

then the limiting covariance is

$$\mathrm{V}_{\mathrm{sand}} + \lambda \mathrm{H}^{-1} \bar{\mathrm{P}}^{-1} \mathrm{R}(\bar{\mathrm{P}}^{*-1})(\mathrm{H}^{*-1}).$$

Thus a fixed m gives the factor $1/m$, while $m_k \xrightarrow{p} \infty$ gives $\lambda = 0$ and recovers the sandwich benchmark. The limiting mean operator need not equal H^{-1}; it only needs to be invertible and to satisfy the stability conditions.

Corollary B.13. *Under the conditions of Theorem 3.4 and corollary B.12, define*

$$\Delta = \mathrm{H}^{-1} \bar{\mathrm{P}}^{-1} \mathrm{R}(\bar{\mathrm{P}}^{*-1})(\mathrm{H}^{*-1}) \succeq 0.$$

For a fixed number m of independent solver replicas,

$$\mathrm{V}_m = \mathrm{V}_{\mathrm{sand}} + \frac{1}{m} \Delta. \tag{40}$$

Consequently, if $1 \le m_1 \le m_2$, then

$$\mathrm{V}_{m_1} \succeq \mathrm{V}_{m_2} \succeq \mathrm{V}_{\mathrm{sand}}, \qquad a^* \mathrm{V}_m a = a^* \mathrm{V}_{\mathrm{sand}} a + \frac{1}{m} a^* \Delta a \quad \textit{for every } a \in E.$$

The first inequality is strict in a contrast a exactly when $m_1 < m_2$ and $a^* \Delta a > 0$. Thus solver replication gives a formal variance ordering at equal score information; it does not compare computational budgets.

Corollary B.14. *Suppose the centered conditions of Theorem 3.1 hold with $\bar{\mathrm{P}} = \mathrm{P}_0$, where P_0 is deterministic and invertible and $\mathrm{P}_0\mathrm{H}$ is positive stable. Assume* (38) *and, in addition to the drift condition* (26)*,*

$$\frac{1}{n} \sum_{k=0}^{n-1} \mathbb{E}\{ \| (\Pi_k - \mathrm{P}_0) \xi_{k+1} \|^2 \mid \mathcal{F}_k \} \xrightarrow{p} 0, \tag{41}$$

$$\frac{1}{n} \sum_{k=0}^{n-1} \mathbb{E}\{ \left\| (\widetilde{\mathrm{P}}_{k+1} - \Pi_k) \xi_{k+1} \right\|^2 \mid \mathcal{F}_k \} \xrightarrow{p} 0. \tag{42}$$

Then $\mathrm{Q} = \mathrm{P}_0 \mathrm{S} \mathrm{P}_0^*$ and $\mathrm{V} = \mathrm{V}_{\mathrm{sand}}$. Under conditional solver score independence and a uniformly bounded conditional second moment of the score, the single stronger condition

$$\frac{1}{n} \sum_{k=0}^{n-1} \mathbb{E}(\left\| \widetilde{\mathrm{P}}_{k+1} - \mathrm{P}_0 \right\|_{\mathrm{op}}^2 \mid \mathcal{F}_k) \xrightarrow{p} 0 \tag{43}$$

is sufficient for (41)–(42)*. Thus a growing inner budget recovers the sandwich when both the conditional mean error and the new random error vanish in the directions excited by the score.*

The root-n drift condition (26) alone is not enough for sandwich recovery. It controls how the conditional mean operator acts on the current state, but it does not control its error on the target score or the random part of the solver error on that score. The two conditions above are therefore separate requirements for the covariance conclusion.

B.2.6 PREDICTABLE DYNAMIC PRECONDITIONING

Fresh solver randomness can remain in the leading score direction. For predictable preconditioners, an alternative representation removes the preconditioner when its accumulated inverse gain variation is small. Let $\mathrm{C}_k : E \to E$ be invertible and $\mathcal{F}_k$-measurable, and consider the recursion

$$u_{k+1} = u_k - \eta_{k+1}\{\mathrm{C}_k(\mathrm{H}u_k + \xi_{k+1}) + \varepsilon_{k+1}\}, \tag{44}$$

where geometric, nonlinear, and estimation terms are included in ε_{k+1}.

The main-text statement is Theorem 3.6.

Corollary B.15. *Suppose* $\sup_k \|\mathrm{M}_k\|_{L^2(\mathrm{op})} < \infty$, $\mathbb{E}\,\|u_k\|^2 = O(k^{-\alpha})$, *and*

$$\|\mathrm{M}_k - \mathrm{M}_{k-1}\|_{L^2(\mathrm{op})} = O(k^{-\beta}) \qquad \textit{for some} \quad \beta > \frac{1+\alpha}{2}.$$

Then (12) *holds. In particular, an* $O(k^{-1})$ *one step variation is sufficient for every* $\alpha \in (1/2, 1)$.

The Euclidean stabilization criterion and its polynomial-rate scope are developed by An and Huo (2026a). In a coherent tangent representation, the same Abel cancellation applies after the geometric remainders are controlled. Predictability alone is insufficient: both accumulated inverse gain variation and the weighted remainder condition must hold. Order-one fresh jumps in a predictable operator can violate the criterion. The distinction therefore concerns variation of the inverse preconditioner, not just the time at which its random seed is revealed.

Corollary B.16. *Suppose Theorem 3.6 holds with a martingale target score and*

$$\frac{1}{n}\sum_{k=0}^{n-1} \mathbb{E}(\xi_{k+1} \otimes \xi_{k+1} \mid \mathcal{F}_k) \xrightarrow{p} \mathrm{S}. \tag{45}$$

Let $\mathrm{S}_a = \mathcal{T}_{x_\star \to a}\mathrm{S}\mathcal{T}_{a \to x_\star}$ *and suppose* $\widehat{\mathrm{H}}_n^a \xrightarrow{p} \mathrm{H}_a$. *Assume that the raw anchor score in* (18) *satisfies*

$$G_{k+1}^a = \mathcal{T}_{x_\star \to a}\xi_{k+1} + h_{k+1}^G, \qquad \frac{1}{n}\sum_{k=0}^{n-1} \left\|h_{k+1}^G\right\|^2 \xrightarrow{p} 0, \tag{46}$$

and that the outer products of ξ_{k+1} *obey a martingale weak law. Then the Welford covariance* $\widehat{\mathrm{S}}_n^a$ *of* G_{k+1}^a *and the estimator in* (19) *satisfy*

$$\widehat{\mathrm{S}}_n^a \xrightarrow{p} \mathrm{S}_a, \qquad \widehat{\mathrm{V}}_n^{a,\mathrm{pred}} \xrightarrow{p} \mathrm{H}_a^{-1}\mathrm{S}_a(\mathrm{H}_a^{*-1}). \tag{47}$$

After transport to $T_{\bar{x}_n}\mathcal{M}$, *the corresponding Wald region has asymptotic coverage* $1-\gamma$ *when* S *is positive definite. No limit of* C_k *is required.*

The two feasible estimators serve different regimes. For a solver randomized at the current iteration, the implemented direction covariance estimates the covariance of $\widetilde{\mathrm{P}}_{k+1}\xi_{k+1}$. For a predictable preconditioner without a fixed limit, the mean direction operator need not converge, but the Abel identity removes it from the influence function; the raw score covariance in (19) therefore estimates the correct sandwich law.

B.2.7 HESSIAN AVERAGING AND FIXED-BUDGET SKETCH SOLVERS

A checkable rank one solver. The fixed-budget conditions can be verified for sketch and project without requiring an increasingly accurate solve at every iteration. For a unit sketch s and a positive definite Hessian estimate B, use $\mathcal{K}_{\mathrm{sap}}(\mathrm{B}, s) = (s \otimes s)/\langle s, \mathrm{B}s\rangle$ and average independent replicas. If H, B and the predictable B_k share a spectral band $[h_-, h_+]$, with $h_- > 0$, then

$$\|\mathcal{K}_{\mathrm{sap}}(\mathrm{B}, s)\|_{\mathrm{op}} \le h_-^{-1}, \qquad \|\mathcal{K}_{\mathrm{sap}}(\mathrm{B}, s) - \mathcal{K}_{\mathrm{sap}}(\mathrm{H}, s)\|_{\mathrm{op}} \le h_-^{-2}\,\|\mathrm{B} - \mathrm{H}\|_{\mathrm{op}}\,.$$

A sketch law that spans the tangent space in mean makes the mean solver positive definite, hence its product with H positive stable. The second bound transfers a Hessian rate $\|\mathrm{B}_k - \mathrm{H}\|_{L^2(\mathrm{op})} = O(k^{-\delta})$ to the solver. The state rate then bounds the mean-drift remainder in L^1 by a constant times $n^{-1/2}\sum_{k=1}^n k^{-\delta-\alpha/2} = o(1)$ when $\delta > (1-\alpha)/2$. Under the local score, moment and invariant

basin conditions, Proposition B.20 completes the state and covariance arguments. This supplies a concrete route from a fixed computational budget to centered Wald inference; the experiments below use its coordinate sketch specialization.

Let $\widehat{\mathrm{B}}^a_{k+1} = \mathcal{T}_{x_k \to a} \widehat{\mathrm{B}}_{k+1} \mathcal{T}_{a \to x_k}$ and consider

$$\widehat{\mathrm{H}}^a_n = \lambda_n \mathrm{I} + \frac{1}{n} \sum_{k=0}^{n-1} \widehat{\mathrm{B}}^a_{k+1}.$$

Proposition B.17. *Suppose H_a and the contributions $\widehat{\mathrm{B}}^a_{k+1}$ are self-adjoint, H_a is positive definite, and*

$$\mathbb{E}(\widehat{\mathrm{B}}^a_{k+1} \mid \mathcal{F}_k) = \mathrm{H}_a + \mathrm{R}^B_{k+1}, \qquad \left\|\mathrm{R}^B_{k+1}\right\|_{L^2(\mathrm{op})} \le C \left\|u_k\right\|_{L^2}^{\rho_{\mathrm{H}}}, \tag{48}$$

$$\mathbb{E}\{\left\|\widehat{\mathrm{B}}^a_{k+1} - \mathbb{E}(\widehat{\mathrm{B}}^a_{k+1} \mid \mathcal{F}_k)\right\|_{\mathrm{F}}^2 \mid \mathcal{F}_k\} \le C, \qquad \lambda_n = O(n^{-\delta_\lambda}) \tag{49}$$

for some $\rho_{\mathrm{H}} \in (0,1]$ and $\delta_\lambda > 0$. Under (34)*,*

$$\left\|\widehat{\mathrm{H}}^a_n - \mathrm{H}_a\right\|_{L^2(\mathrm{op})} = O(n^{-\delta_{\mathrm{H}}}), \qquad \delta_{\mathrm{H}} = \min\left\{\frac{1}{2}, \frac{\alpha\rho_{\mathrm{H}}}{2}, \delta_\lambda\right\}. \tag{50}$$

If the spectrum of H_a belongs to a fixed closed interval $[h_-, h_+]$, $h_- > 0$, then spectral projection of the self-adjoint estimate onto that interval preserves the rate in (50)*, and the inverse of the safeguarded estimate obeys the same rate. In particular, for a locally Lipschitz Hessian ($\rho_{\mathrm{H}} = 1$) and regularization satisfying $\delta_\lambda \ge \alpha/2$, $\delta_{\mathrm{H}} = \alpha/2 > (1-\alpha)/2$.*

Corollary B.18. *Let $\widetilde{\mathrm{H}}^a_k = \lambda_k \mathrm{I} + k^{-1} \sum_{j=1}^k \widehat{\mathrm{B}}^a_j$ be the self-adjoint anchor Hessian average. Suppose its individual contributions have uniformly bounded L^2 norm, $\sup_k |\lambda_k| < \infty$, and $|\lambda_k - \lambda_{k-1}| = O(k^{-1})$. Take $\widehat{\mathrm{H}}^a_k$ either as $\widetilde{\mathrm{H}}^a_k$, assumed uniformly positive definite, or as its spectral projection onto a fixed interval $[h_-, h_+]$, $h_- > 0$. Use the exact predictable operator*

$$\mathrm{C}_k = \mathcal{T}_{a \to x_\star} (\widehat{\mathrm{H}}^a_k)^{-1} \mathcal{T}_{x_\star \to a}$$

in the recursion in the reference tangent space. Then $\left\|\mathrm{C}_k^{-1} - \mathrm{C}_{k-1}^{-1}\right\|_{L^2(\mathrm{op})} = O(k^{-1})$, so Corollary B.15 applies. The sandwich benchmark follows even without invoking the convergence rate (50)*, provided the state and remainder conditions hold.*

Proposition B.19. *Let*

$$\widetilde{\mathrm{P}}_{k+1} = \mathcal{K}_J(\mathrm{B}_k, \Omega_{k+1}), \tag{51}$$

where B_k is $\mathcal{F}_k$-measurable in coordinates of the reference tangent space, Ω_{k+1} is conditionally independent of the current target score given $\mathcal{F}_k$, with a fixed deterministic conditional law $\mathcal{L}(\Omega_{k+1} \mid \mathcal{F}_k) = \mu$, and J is fixed. Here $\mathbb{E}_\Omega$ denotes expectation with respect to μ. Put $\Sigma_k = \mathbb{E}(\xi_{k+1} \otimes \xi_{k+1} \mid \mathcal{F}_k)$. Suppose the local field, state, bias, geometric, score increment, centered inexactness, and averaging conditions of Theorem 3.1 hold, including Assumption B.6 and the direct population remainder condition (27)*, and assume*

$$\left\|\mathrm{B}_k - \mathrm{H}\right\|_{L^2(\mathrm{op})} = O(k^{-\delta}), \qquad \delta > \frac{1-\alpha}{2}, \tag{52}$$

$$\sup_k \mathbb{E}(\left\|\xi_{k+1}\right\|^2 \mid \mathcal{F}_k) \le C, \qquad \sup_{\mathrm{B} \in \mathcal{N}_{\mathrm{H}}} \mathbb{E}_\Omega \left\|\mathcal{K}_J(\mathrm{B}, \Omega)\right\|_{\mathrm{op}}^2 \le C,$$

$$\left\{\mathbb{E}_\Omega \left\|\mathcal{K}_J(\mathrm{B}, \Omega) - \mathcal{K}_J(\mathrm{H}, \Omega)\right\|_{\mathrm{op}}^2\right\}^{1/2} \le L \left\|\mathrm{B} - \mathrm{H}\right\|_{\mathrm{op}}, \tag{53}$$

$$\sup_k \left\|\Sigma_k\right\|_{\mathrm{op}} \le C \quad \textit{almost surely},$$

$$\frac{1}{n} \sum_{k=0}^{n-1} \mathbb{E} \left\|\Sigma_k - \mathrm{S}\right\|_{\mathrm{op}} \longrightarrow 0 \tag{54}$$

for a deterministic positive semidefinite S *and an operator neighborhood* $\mathcal{N}_{\mathrm{H}}$ *that contains* B_k *almost surely for every* k *in the reported phase. Assume also the direct product Lindeberg condition: for every* $\epsilon > 0$*,*

$$\frac{1}{n}\sum_{k=0}^{n-1}\mathbb{E}\Big[\|\mathcal{K}_J(\mathrm{B}_k,\Omega_{k+1})\xi_{k+1}\|^2\,\mathbf{1}\{\|\mathcal{K}_J(\mathrm{B}_k,\Omega_{k+1})\xi_{k+1}\|>\epsilon\sqrt{n}\}\mid\mathcal{F}_k\Big]\xrightarrow{p}0. \tag{55}$$

Define

$$\widetilde{\mathrm{P}}_J=\mathcal{K}_J(\mathrm{H},\Omega),\qquad \bar{\mathrm{P}}_J=\mathbb{E}_\Omega(\widetilde{\mathrm{P}}_J).$$

If $\bar{\mathrm{P}}_J\mathrm{H}$ *is positive stable, then the mean operator remainder, conditional Lindeberg condition, and covariance limit in Theorem 3.1 hold with* $\mathrm{Q}_J=\mathbb{E}_\Omega(\widetilde{\mathrm{P}}_J\mathrm{S}\widetilde{\mathrm{P}}_J^*)$*. Consequently, Theorem 3.1 applies with* $\bar{\mathrm{P}}=\bar{\mathrm{P}}_J$ *and* $\mathrm{Q}=\mathrm{Q}_J$*, and Theorem 3.4 gives the independent solver covariance decomposition. For* $J=J_k\to\infty$*, sandwich recovery follows from Corollary B.14 when both solver errors applied to the target score* (41)–(42) *vanish.*

A uniform conditional $(2+\tau)$ moment of $\mathcal{K}_J(\mathrm{B}_k,\Omega_{k+1})\xi_{k+1}$ is a convenient sufficient condition for (55), but is not required by the proposition. A fixed number of Richardson, randomized Kaczmarz, sketch and project, or fixed coefficient accelerated iterations started from zero produces a linear map of the right hand side and is covered by (51). If coefficients or stopping times are selected from the current residual, the map need not be linear; the discrepancy from a reference linear operator must then be included in e_{k+1} and checked through Assumptions B.4 and 3.2.

The generic proposition becomes fully checkable for the following named instance of the sketch and project construction of Gower and Richtárik (2015). It uses one rank one sketch per replica; increasing m averages independent replicas but does not change the one step inner budget.

Proposition B.20. *Suppose* $\mathrm{H}=\mathrm{H}^*\succ 0$*, and let* $\mathrm{B}_k=\mathrm{B}_k^*$ *be* $\mathcal{F}_k$*-measurable. Assume throughout the reported phase, almost surely,*

$$h_-\mathrm{I}\preceq\mathrm{H}\preceq h_+\mathrm{I},\qquad h_-\mathrm{I}\preceq\mathrm{B}_k\preceq h_+\mathrm{I},\qquad 0<h_-<h_+<\infty. \tag{56}$$

For a unit vector s*, define the one step sketch and project operator*

$$\mathcal{K}_{\mathrm{sap}}(\mathrm{B},s)=\frac{s\otimes s}{\langle s,\mathrm{B}s\rangle}. \tag{57}$$

At iteration $k+1$*, draw a fixed number* m *of conditionally independent sketches* $s_{k+1}^{(1)},\dots,s_{k+1}^{(m)}$ *from a time-homogeneous law* μ *on the unit sphere, independently of the current observation, and use*

$$\widetilde{\mathrm{P}}_{k+1}^{(m)}=\frac{1}{m}\sum_{j=1}^{m}\mathcal{K}_{\mathrm{sap}}(\mathrm{B}_k,s_{k+1}^{(j)}).$$

Assume the sketch law spans E *in mean:*

$$\mathbb{E}_\mu(s\otimes s)\succeq\kappa\mathrm{I}\quad\text{for some }\kappa>0. \tag{58}$$

Then, for self-adjoint B,B' *whose spectra lie in* $[h_-,h_+]$ *as in* (56)*,*

$$\|\mathcal{K}_{\mathrm{sap}}(\mathrm{B},s)\|_{\mathrm{op}}\le h_-^{-1},\quad \|\mathcal{K}_{\mathrm{sap}}(\mathrm{B},s)-\mathcal{K}_{\mathrm{sap}}(\mathrm{B}',s)\|_{\mathrm{op}}\le h_-^{-2}\|\mathrm{B}-\mathrm{B}'\|_{\mathrm{op}}, \tag{59}$$

$$\frac{\kappa}{h_+}\mathrm{I}\preceq\bar{\mathrm{P}}_{\mathrm{sap}}:=\mathbb{E}_\mu\{\mathcal{K}_{\mathrm{sap}}(\mathrm{H},s)\}\preceq\frac{1}{h_-}\mathrm{I}. \tag{60}$$

In particular, $\bar{\mathrm{P}}_{\mathrm{sap}}\mathrm{H}$ *is positive stable.*

Assume further that the score decomposition in Assumption B.2 holds and, for some $\delta>(1-\alpha)/2$*,* $\rho\in(0,1]$*, and* $\tau>0$*,*

$$\|\mathrm{B}_k-\mathrm{H}\|_{L^2(\mathrm{op})}=O(k^{-\delta}), \tag{61}$$

$$\|r_G(x)\|\le C\left\|R_{x_\star}^{-1}(x)\right\|^{1+\rho},\qquad \alpha(1+\rho)>1, \tag{62}$$

$$\sup_k\mathbb{E}(\|\xi_{k+1}\|^{2+\tau}\mid\mathcal{F}_k)\le C,\qquad \mathbb{E}(\|q_{k+1}\|^{2+\tau}\mid\mathcal{F}_k)\le C\|u_k\|^{2+\tau}, \tag{63}$$

$$\frac{1}{n}\sum_{k=0}^{n-1}\mathbb{E}\,\|\Sigma_k-\mathrm{S}\|_{\mathrm{op}}\longrightarrow 0,\qquad\qquad \Sigma_k=\mathbb{E}(\xi_{k+1}\otimes\xi_{k+1}\mid\mathcal{F}_k). \tag{64}$$

Suppose the actual clipped recursion, with no additional pre-clipping inexactness or localization modification, remains almost surely in the fixed sufficiently small working basin. The retraction and frame are C^2, and the fixed anchor average is used. Let the derivative estimator be consistent, and let the mean solver and direction covariance estimators be the one-pass averages in Algorithm 2. Then the centered conclusions of Theorems 3.1 and 3.3 hold. With

$$\mathrm{R}_{\mathrm{sap}}=\mathbb{E}_\mu\{(\mathcal{K}_{\mathrm{sap}}(\mathrm{H},s)-\bar{\mathrm{P}}_{\mathrm{sap}})\mathrm{S}(\mathcal{K}_{\mathrm{sap}}(\mathrm{H},s)-\bar{\mathrm{P}}_{\mathrm{sap}})^*\},$$

the limiting covariance is

$$\mathrm{V}_m=\mathrm{V}_{\mathrm{sand}}+\frac{1}{m}\mathrm{H}^{-1}\bar{\mathrm{P}}_{\mathrm{sap}}^{-1}\mathrm{R}_{\mathrm{sap}}(\bar{\mathrm{P}}_{\mathrm{sap}}^{*-1})(\mathrm{H}^{*-1}). \tag{65}$$

For the uniform law on the unit sphere, (58) *holds with $\mathbb{E}_\mu(s\otimes s)=\mathrm{I}/d$.*

Thus the fixed-budget theory covers a named random linear solver on an invariant localized reporting phase. Its common stability inequality is verified before using a Hessian convergence rate. A direction cap alone does not establish basin invariance; any additional localization modification needs its own drift and implementation-error bounds.

The complete invariant model and its verification are in Example B.28.

B.2.8 Relation to established limits

When $\mathcal{M}=\mathbb{R}^d$, $R_x(v)=x+v$, and $\mathcal{T}_{x\to y}=\mathrm{I}$, all geometric terms vanish and Theorem 3.1 becomes

$$\sqrt{n}(\bar{x}_n-x_\star)=-(\bar{\mathrm{P}}\mathrm{H})^{-1}\frac{1}{\sqrt{n}}\sum_{k=1}^{n}\zeta_k-(\bar{\mathrm{P}}\mathrm{H})^{-1}b+o_p(1).$$

For an asymptotically deterministic preconditioner this is the classical sandwich law; for an operator randomized with the current score it gives (11). If instead $\mathrm{P}_{k+1}=\mathrm{I}$ and $e_{k+1}=0$, then

$$\sqrt{n}R_{x_\star}^{-1}(\bar{x}_n)\xrightarrow{d}N_E(0,\mathrm{H}^{-1}\mathrm{S}(\mathrm{H}^{*-1})),$$

which is the local averaged-RSGD covariance of Tripuraneni et al. (2018) under the corresponding centered conditions. Thus deterministic and predictable conditioning lead to the usual averaged stochastic approximation covariance (Polyak and Juditsky, 1992; Boyer and Godichon-Baggioni, 2023; An and Huo, 2026a). Current randomization instead gives the ordered family in Theorem 3.4 and corollary B.13. This comparison is a theorem about first order covariance. Last iterate randomized Newton methods have a different Lyapunov covariance and are not special cases of the present average (Wang et al., 2026).

B.3 Non-Euclidean applications

This section specializes the main theorem to parameters for which the manifold structure is intrinsic to the statistical problem. The purpose is twofold: to identify the derivative H and target score covariance S, and to state primitive conditions under which the abstract remainder assumptions hold. The solver contribution is then obtained from Theorem 3.4 and corollary B.13 without a separate central limit theorem for each model.

B.3.1 Intrinsic Fréchet means

Let Y be manifold-valued and consider

$$F(x)=\frac{1}{2}\mathbb{E}d^2(x,Y),\qquad G(x)=-\mathbb{E}\,\mathrm{Log}_x(Y).$$

Let $x_\star$ be an isolated Fréchet mean and set

$$\mathrm{H}_F=\mathrm{Hess}\,F(x_\star),\qquad \mathrm{S}_F=\mathrm{Cov}\{\mathrm{Log}_{x_\star}(Y)\}=\mathbb{E}\{\mathrm{Log}_{x_\star}(Y)\otimes\mathrm{Log}_{x_\star}(Y)\}.$$

The stochastic score is $\widehat{G}_{k+1}(x)=-\,\mathrm{Log}_x(Y_{k+1})$.

Corollary B.21. *Assume that* $\mathrm{H}_F \succ 0$*, the recursion is localized in a neighborhood on which the relevant logarithm maps are single-valued, and*

$$\left\|\mathcal{T}_{x\to x_\star}\operatorname{Log}_x(Y) - \operatorname{Log}_{x_\star}(Y)\right\| \le L(Y)\left\|R_{x_\star}^{-1}(x)\right\|, \qquad \mathbb{E}L(Y)^2 < \infty. \tag{66}$$

Suppose $\mathbb{E}\left\|\operatorname{Log}_{x_\star}(Y)\right\|^2 < \infty$ *and all remaining conditions of Theorem 3.1 hold with* $b = 0$*. Then*

$$\sqrt{n}\, R_{x_\star}^{-1}(\bar{x}_n) \xrightarrow{d} N_E\big(0, (\bar{\mathrm{P}}\mathrm{H}_F)^{-1}\mathrm{Q}_F((\bar{\mathrm{P}}\mathrm{H}_F)^{*-1})\big), \tag{67}$$

where

$$\mathrm{Q}_F = \lim_{n\to\infty}\frac{1}{n}\sum_{k=0}^{n-1}\operatorname{Cov}\Big(\widetilde{\mathrm{P}}_{k+1}\operatorname{Log}_{x_\star}(Y_{k+1}) \mid \mathcal{F}_k\Big). \tag{68}$$

The covariance, rather than the raw conditional second moment, is required when the current solver and score may be dependent; under exact conditional product centering the two coincide. An asymptotically deterministic solver gives $\mathrm{H}_F^{-1}\mathrm{S}_F\mathrm{H}_F^{-1}$*. An independent nondegenerate current solver adds the positive semidefinite term in* (11)*.*

Condition (66) is a local L^2 condition on the derivative of the logarithm map. A compact support/domain pair uniformly separated from the cut and conjugate loci gives a sufficient condition by uniform boundedness of the transported logarithm derivative. It is not necessary. General central limit theorems for intrinsic means can allow support at the cut locus under more delicate density and geometric conditions (Bhattacharya and Patrangenaru, 2005; Hotz et al., 2024); the present online argument uses the regular local regime in which the score itself has the expansion required by Assumption B.2.

The unit sphere For $\mathcal{M} = \mathbb{S}^q$, write

$$r = d(x_\star, Y), \qquad e = \frac{\operatorname{Log}_{x_\star}(Y)}{r} \quad (r > 0),$$

and interpret the expressions below by continuity at $r = 0$.

Corollary B.22. *Take* $R = \operatorname{Exp}$*. Suppose* $x_\star$ *is the unique intrinsic mean,* $r \le r_0 < \pi/2$ *almost surely, and* $\mathbb{E}\operatorname{Log}_{x_\star}(Y) = 0$*. Then*

$$\mathrm{H}_{\mathbb{S}} = \mathbb{E}[r\cot(r)\mathrm{I} + \{1 - r\cot(r)\}e\otimes e], \tag{69}$$

$$\mathrm{S}_{\mathbb{S}} = \mathbb{E}(r^2 e\otimes e), \tag{70}$$

and $\mathrm{H}_{\mathbb{S}} \succeq r_0\cot(r_0)\mathrm{I}$*. Let* $\mathrm{Q}_{\mathbb{S}}$ *denote the sphere instance of the conditional product-covariance limit* (68)*. Under the algorithmic conditions of Corollary B.21,*

$$\sqrt{n}\,\operatorname{Log}_{x_\star}(\bar{x}_n) \xrightarrow{d} N\big(0, (\bar{\mathrm{P}}\mathrm{H}_{\mathbb{S}})^{-1}\mathrm{Q}_{\mathbb{S}}((\bar{\mathrm{P}}\mathrm{H}_{\mathbb{S}})^{*-1})\big).$$

In particular, an asymptotically deterministic solver has covariance $\mathrm{H}_{\mathbb{S}}^{-1}\mathrm{S}_{\mathbb{S}}\mathrm{H}_{\mathbb{S}}^{-1}$*, while an independent current solver obeys* (11) *and* (40)*.*

The support condition is a transparent sufficient condition for regularity of the logarithm map and a uniformly nonsingular Hessian; uniqueness remains the separate assumption in Corollary B.22. Related convexity radii for Riemannian centers of mass are studied by Afsari (2011). It is stronger than necessary and is imposed here to keep the online score expansion regular.

Affine-invariant covariance centers Let $\mathcal{P}_q$ be the manifold of $q \times q$ symmetric positive definite matrices with affine-invariant metric $\langle \mathrm{U}, \mathrm{V}\rangle_{\mathrm{X}} = \operatorname{tr}(\mathrm{X}^{-1}\mathrm{U}\mathrm{X}^{-1}\mathrm{V})$. Its exponential and logarithm maps are

$$\operatorname{Exp}_{\mathrm{X}}(\mathrm{U}) = \mathrm{X}^{1/2}\exp(\mathrm{X}^{-1/2}\mathrm{U}\mathrm{X}^{-1/2})\mathrm{X}^{1/2}, \tag{71}$$

$$\operatorname{Log}_{\mathrm{X}}(\mathrm{Y}) = \mathrm{X}^{1/2}\log(\mathrm{X}^{-1/2}\mathrm{Y}\mathrm{X}^{-1/2})\mathrm{X}^{1/2}. \tag{72}$$

Corollary B.23. *Let* $\mathrm{X}_\star$ *be the affine-invariant Fréchet mean of a random matrix* $\mathrm{Y} \in \mathcal{P}_q$*. Assume that the support is contained in a bounded geodesic ball,* $\mathrm{H}_{\mathrm{SPD}} = \operatorname{Hess} F(\mathrm{X}_\star) \succ 0$*, and the remaining conditions of Corollary B.21 hold. Let* $\mathrm{Q}_{\mathrm{SPD}}$ *be the SPD instance of* (68) *and define*

$$\mathrm{V}_{\mathrm{SPD}} = (\bar{\mathrm{P}}\mathrm{H}_{\mathrm{SPD}})^{-1}\mathrm{Q}_{\mathrm{SPD}}((\bar{\mathrm{P}}\mathrm{H}_{\mathrm{SPD}})^{*-1}).$$

With

$$\mathrm{S_{SPD}} = \mathbb{E}\{\mathrm{Log}_{\mathrm{X}_\star}(\mathrm{Y}) \otimes \mathrm{Log}_{\mathrm{X}_\star}(\mathrm{Y})\},$$

the limit in (67) *holds with* $\mathrm{H}_F = \mathrm{H_{SPD}}$ *and covariance* $\mathrm{V_{SPD}}$. *In particular, in the asymptotically deterministic-solver regime this reduces to* $\mathrm{H_{SPD}^{-1}S_{SPD}H_{SPD}^{-1}}$. *For* $\psi(\mathrm{X}) = \log\det \mathrm{X}$,

$$\sqrt{n}\{\log\det(\bar{\mathrm{X}}_n) - \log\det(\mathrm{X}_\star)\} \xrightarrow{d} N(0, \ell \mathrm{V_{SPD}} \ell^*), \qquad \ell(\mathrm{U}) = \mathrm{tr}(\mathrm{X}_\star^{-1}\mathrm{U}). \tag{73}$$

The affine-invariant positive definite manifold is complete, simply connected, and nonpositively curved, so it has no cut-locus obstruction to (72); see Bhatia (2007). Bounded support is used only as a convenient route to the derivative envelope and can be replaced by an integrable envelope on the logarithm and its derivative.

B.3.2 Streaming principal subspaces

Let $X \in \mathbb{R}^p$ have finite second moment, and write its uncentered second moment as

$$\mathrm{M} = \mathbb{E}(XX^*) = \mathrm{U}_1\Lambda_1\mathrm{U}_1^* + \mathrm{U}_2\Lambda_2\mathrm{U}_2^*, \tag{74}$$

where $1 \le r < p$, $[\mathrm{U}_1, \mathrm{U}_2]$ is an orthogonal matrix, $\mathrm{U}_1 \in \mathbb{R}^{p\times r}$, and $\mathrm{U}_2 \in \mathbb{R}^{p\times(p-r)}$. Order the eigenvalues as $\lambda_1 \ge \cdots \ge \lambda_p$, set $\Lambda_1 = \mathrm{diag}(\lambda_1, \ldots, \lambda_r)$ and $\Lambda_2 = \mathrm{diag}(\lambda_{r+1}, \ldots, \lambda_p)$, and assume $\lambda_r > \lambda_{r+1}$. Thus U_1 spans the leading eigenspace. The parameter is the point $[\mathrm{U}_1] \in \mathrm{Gr}(r, p)$ minimizing $F([\mathrm{U}]) = -\frac{1}{2}\mathbb{E}\,\|\mathrm{U}^*X\|^2$. Represent a tangent vector by $\mathrm{K} \in \mathbb{R}^{(p-r)\times r}$ through $\mathrm{U}_2\mathrm{K}$, and write $z_1 = \mathrm{U}_1^*X \in \mathbb{R}^r$ and $z_2 = \mathrm{U}_2^*X \in \mathbb{R}^{p-r}$. The manifold dimension is $r(p-r)$. Column-stacking vec identifies tangent matrices isometrically with $\mathbb{R}^{r(p-r)}$; all covariance products below represent the tangent operators in these coordinates. We use the graph/polar retraction $R_{[\mathrm{U}_1]}(\mathrm{U}_2\mathrm{K}) = [(\mathrm{U}_1 + \mathrm{U}_2\mathrm{K})(\mathrm{I} + \mathrm{K}^*\mathrm{K})^{-1/2}]$. For any other C^2 retraction with the same differential, its inverse-retraction coordinate and the graph coordinate differ quadratically; hence their difference at the reported average is $o_p(n^{-1/2})$.

Corollary B.24. *Suppose* $\mathbb{E}\,\|X\|^4 < \infty$, *the eigengap in* (74) *is positive, and the recursion is localized and satisfies the centered conditions of Theorem 3.1. In the graph chart at* $[\mathrm{U}_1]$, *the population derivative and target score are*

$$\mathcal{H}[\mathrm{K}] = \mathrm{K}\Lambda_1 - \Lambda_2\mathrm{K}, \qquad \xi = -z_2 z_1^*.$$

If

$$\mathrm{Q_{PCA}} = \lim_{n\to\infty} \frac{1}{n}\sum_{k=0}^{n-1} \mathrm{Cov}\Big(\mathrm{vec}\{\widetilde{\mathrm{P}}_{k+1}(z_{2,k+1}z_{1,k+1}^*)\} \mid \mathcal{F}_k\Big),$$

then, for the tangent graph coordinate $\bar{\mathrm{K}}_n$ *of the reported average,*

$$\sqrt{n}\,\mathrm{vec}(\bar{\mathrm{K}}_n) \xrightarrow{d} N\big(0, (\bar{\mathrm{P}}\mathcal{H})^{-1}\mathrm{Q_{PCA}}((\bar{\mathrm{P}}\mathcal{H})^{*-1})\big). \tag{75}$$

For an asymptotically deterministic solver, the covariance is $\mathcal{H}^{-1}\mathrm{S_{PCA}}\mathcal{H}^{-1}$, *where*

$$\mathrm{S_{PCA}} = \mathrm{Cov}\{\mathrm{vec}(z_2 z_1^*)\}.$$

In this asymptotically deterministic-solver regime, if X *is centered Gaussian in the eigenbasis, let* Z *denote the limiting Gaussian tangent matrix. Its coordinates are uncorrelated and*

$$\mathrm{Var}(Z_{ij}) = \frac{\lambda_{r+i}\lambda_j}{(\lambda_j - \lambda_{r+i})^2}. \tag{76}$$

For the spectral projector $\Pi_\star = \mathrm{U}_1\mathrm{U}_1^*$,

$$D\Pi_{[\mathrm{U}_1]}[\mathrm{K}] = \mathrm{U}_2\mathrm{K}\mathrm{U}_1^* + \mathrm{U}_1\mathrm{K}^*\mathrm{U}_2^*. \tag{77}$$

For centered X, M in (74) is the covariance. Without centering, the same score and Hessian derivation targets the second moment about the chosen origin; centering is needed only to identify M with the covariance. The eigengap is not a technical convenience: it identifies a unique leading subspace and makes $\mathcal{H}$ invertible. Without it, a regular root-n limit for a fixed subspace is not available without changing the target. The fourth moment of X is a second moment condition for the quadratic score $z_2 z_1^*$ and its local random derivative; it is not a fourth moment condition on the iterate. The local geometry agrees with the standard Grassmann formulation of invariant-subspace optimization (Edelman et al., 1998; Tripuraneni et al., 2018; Li and Jordan, 2023). Entrywise and high-dimensional uncertainty quantification for Oja-type streaming PCA has been studied by Kumar et al. (2025) and Xu and Li (2026). The result here instead concerns a fixed-dimensional Riemannian Polyak-Ruppert average with randomized Newton directions.

B.4 PRIMITIVE CONDITIONS FOR THE HIGH-LEVEL ASSUMPTIONS

This section records three reusable ways to verify the product, remainder, and geometric conditions in the main text. The propositions are sufficient routes, not additional requirements of Theorem 3.1.

The verification order is one-way. Product and local direction bounds first control the raw update; Proposition B.10 then obtains the mean-square state rate directly from the local update map. The accumulated population remainders and the two root-n geometric sums are verified only afterward through Propositions B.26 and B.27. In particular, the state rate argument assumes neither conclusion.

Proposition B.25. *Suppose, conditionally on $\mathcal{F}_k$, that $\widetilde{\mathrm{P}}_{k+1}$ is independent of (ξ_{k+1}, q_{k+1}). If*

$$\sup_k \mathbb{E}(\left\|\widetilde{\mathrm{P}}_{k+1}\right\|_{\mathrm{op}}^2 \mid \mathcal{F}_k) \le C, \qquad \sup_k \mathbb{E}(\|\xi_{k+1}\|^2 \mid \mathcal{F}_k) \le C,$$

then $b_{k+1}^{\xi} = 0$ and $\mathbb{E}(\left\|\widetilde{\mathrm{P}}_{k+1}\xi_{k+1}\right\|^2 \mid \mathcal{F}_k) \le C$. If, in addition, $\mathbb{E}(\|q_{k+1}\|^2 \mid \mathcal{F}_k) \le C\,\|u_k\|^2$ and the conditional second moment of $\widetilde{\mathrm{P}}_{k+1}$ is bounded as above, then $b_{k+1}^{q} = 0$ and (25) *holds. A sufficient tail condition for* (36) *is*

$$\sup_k \mathbb{E}(\left\|\widetilde{\mathrm{P}}_{k+1}\xi_{k+1}\right\|^{2+\tau} \mid \mathcal{F}_k) \le C \quad \text{for some } \tau > 0. \tag{78}$$

Under conditional independence, (78) *follows from separate $(2+\tau_{\mathrm{P}})$ and $(2+\tau_{\xi})$ moments whenever $\tau \le \min(\tau_{\mathrm{P}}, \tau_{\xi})$. Without independence, conditional Hölder still yields the product moment, but not product centering, if*

$$\frac{2+\tau}{2+\tau_{\mathrm{P}}} + \frac{2+\tau}{2+\tau_{\xi}} \le 1.$$

Conditional independence is sufficient, not necessary. A model-specific identity may give $b_{k+1}^{\xi} = 0$ when the score and Hessian use the same observation. Conversely, separate conditional unbiasedness of the solver and score does not imply $\mathbb{E}(\widetilde{\mathrm{P}}_{k+1}\xi_{k+1} \mid \mathcal{F}_k) = 0$.

Proposition B.26. *Assume* (34) *and* (24)*.*

(i) If $\{\mathbb{E}\left\|\Pi_k - \bar{\mathrm{P}}\right\|_{\mathrm{op}}^2\}^{1/2} = O(k^{-\delta})$ with $\delta > (1-\alpha)/2$, then (26) *holds.*

(ii) If, for some $\rho \in (0,1]$, $\|r_G(x)\| \le L\left\|R_{x_\star}^{-1}(x)\right\|^{1+\rho}$ locally and $\alpha(1+\rho) > 1$, then (27) *holds.*

Both conclusions may instead be verified directly when neither a polynomial rate nor a Hölder modulus is available.

Proposition B.27. *Assume Assumptions B.1 to B.4 and* (34)*. Suppose R and the local frame are C^2 on a relatively compact neighborhood and use the safeguard $\eta_{k+1}\|D_{k+1}\| \le r_0$, where $D_{k+1} = \mathrm{P}_{k+1}\widehat{G}_{k+1}(x_k) + e_{k+1}$. Suppose*

$$\mathbb{E}(\|\mathcal{T}_{x_k \to x_\star} D_{k+1}\|^2 \mid \mathcal{F}_k) \le C\{1 + \|u_k\|^2 + \|\beta_{k+1}\|^2\}, \qquad \beta_{k+1} = b_{k+1}^{\xi} + b_{k+1}^{q} + a_{k+1}, \tag{79}$$

with $\mathbb{E}\|\beta_{k+1}\|^2 \le C\eta_{k+1}$. Then Assumption B.5 holds for the actual safeguarded recursion. A uniform conditional $(2+\tau)$ moment of its raw pre-clipping direction is sufficient for the clipping discrepancy to have negligible root-n mean and vanishing average conditional variance along that recursion. An uncapped method requires its own common-domain expansion and localization argument; clipping-tail bounds alone do not transfer a state rate result between different trajectories.

B.5 VERIFICATION OF THE HIGH-LEVEL CONDITIONS

The main theorem is stated in terms of the accumulated quantities that enter the root-n expansion. This section records common primitive routes to those conditions without making them part of every application.

Population linearization and solver mean. Differentiation under the expectation, justified by an integrable local derivative envelope, gives the derivative H. If

$$\|r_G(x)\| \le C \left\|R_{x_\star}^{-1}(x)\right\|^{1+\rho}, \qquad \alpha(1+\rho) > 1,$$

then Proposition B.26 gives (27). Likewise, $\left\|\Pi_k - \bar{\mathrm{P}}\right\|_{L^2} = O(k^{-\delta})$ with $\delta > (1-\alpha)/2$ is sufficient for (26). Neither polynomial rate is necessary if the corresponding sum can be bounded directly.

Scores, products, and inexactness. For fresh i.i.d. observations with tangent-valued estimating function $\mathrm{g}(x, Z)$, $\mathbb{E}\mathrm{g}(x, Z) = G(x)$, take

$$\xi_{k+1} = \mathrm{g}(x_\star, Z_{k+1}), \qquad q_{k+1} = \mathcal{T}_{x_k \to x_\star}\mathrm{g}(x_k, Z_{k+1}) - \mathrm{g}(x_\star, Z_{k+1}) - \mathcal{T}_{x_k \to x_\star} G(x_k).$$

A local L^2 Lipschitz envelope for the transported score gives $\mathbb{E}(\|q_{k+1}\|^2 \mid \mathcal{F}_k) \le C\|u_k\|^2$. Together with a bounded solver operator, suitable joint-product bounds, or the conditional independence and moment assumptions of Proposition B.25, this verifies (25).

Without independence, one must verify the cumulative conditional product mean in (29). A deterministic bias bound $O(k^{-\kappa})$ with $\kappa > 1/2$ gives $b = 0$, whereas an order $k^{-1/2}$ contribution may produce the local shift in Theorem 3.1. The centered inner error is negligible when its average conditional second moment tends to zero.

State and geometry. A local Lyapunov inequality for $\Pi_k \mathrm{H}$, bounded conditional second moments of the implemented direction, and sufficiently small geometric remainder give $\mathbb{E}\|u_k\|^2 = O(\eta_k)$ by Proposition B.10. A C^2 retraction and frame then yield Assumption B.5 after a local direction safeguard. For the fixed anchor average, C^2 chart transitions and $\alpha > 1/2$ make the quadratic averaging error root-n negligible. The recursive average uses the same C^2 regularity on its fixed small coordinate ball; (88) supplies the quadratic error and proves invariance of that average's ball.

Covariance limits. A direct predictable covariance limit and conditional uniform integrability verify (35)–(36). A bounded conditional $(2+\tau)$ moment of $\widetilde{\mathrm{P}}_{k+1}\xi_{k+1}$ is one simple sufficient condition. Consistency of the one-pass covariance estimator additionally requires consistent derivative and mean solver estimates and vanishing Cesàro squared conditional bias. Lemma B.39 derives the realized-direction and outer-product laws from these and the preceding conditions; no additional fourth moment is required.

Recovery of the sandwich covariance. For increasing inner accuracy, convergence of the solver mean on the state is not enough. The two score-direction conditions (41)–(42) control, respectively, the conditional-mean error and fresh random error. For a predictable nonconvergent preconditioner, the alternative route is the one step variation condition in Corollary B.15; the covariance is then estimated from raw scores as in Corollary B.16.

B.6 Unified proofs

B.6.1 Local mean contraction

Only in Proposition 3.5 and its proof, $\|\cdot\|_{\mathrm{op}}$ denotes the operator norm induced by the P^{-1} energy norm: for $v \in E$ and a linear operator $\mathrm{L} : E \to E$,

$$\|v\|_{\mathrm{P}^{-1}}^2 = \left\langle v, \mathrm{P}^{-1}v \right\rangle, \qquad \|\mathrm{L}\|_{\mathrm{op}} := \sup_{v \ne 0} \frac{\|\mathrm{L}v\|_{\mathrm{P}^{-1}}}{\|v\|_{\mathrm{P}^{-1}}}.$$

Proof of Proposition 3.5. Conjugation by $\mathrm{P}^{1/2}$ gives

$$\mathrm{P}^{-1/2}(\mathrm{I} - \eta \mathrm{PH})\mathrm{P}^{1/2} = \mathrm{I} - \eta \mathrm{P}^{1/2}\mathrm{H}\mathrm{P}^{1/2}.$$

The spectral theorem therefore yields $\|\mathrm{I} - \eta\mathrm{PH}\|_{\mathrm{op}} = \max_{\lambda \in [\lambda_{\min}, \lambda_{\max}]} |1 - \eta\lambda|$. The minimum is attained when the errors at the two endpoints have equal magnitude, which gives the displayed step and contraction factor. The RSGD and exact Newton cases follow by substitution. □

B.6.2 RECORDED-DISCREPANCY BIAS ESTIMATION

One sufficient route to bias estimation is available when $b^{\xi}_{k+1} = b^{q}_{k+1} = 0$ and the full discrepancy e_{k+1} is recorded. Then $\widehat{b}^{a}_{n} = n^{-1/2}\sum_{k=0}^{n-1} \mathcal{T}_{x_k \to a} e_{k+1}$ converges in probability to b_a: coherent transport identifies its predictable part with the signed bias sum, while its centered part vanishes by the martingale variance condition in Assumption B.4. This construction requires one additional running vector; it does not estimate an unknown solver score product bias.

B.6.3 AN INVARIANT BOUNDED NOISE CONSTRUCTION

Example B.28. *On $\mathbb{R}^d$, let $G(x) = \mathrm{H}x$ with $\mathrm{H} = \mathrm{diag}(h_1, \dots, h_d) \succ 0$ and $x_\star = 0$. Let the observations supply centered bounded i.i.d. noise ξ with $|\xi_i| \le h_i r_i$, and take x_0 in the box $|x_i| \le r_i$. Use $\mathrm{B}_k = \mathrm{H}$, sketches $s = e_i$ with probabilities $p_i > 0$, one sketch per observation, and $0 < \eta_k \le 1$. When coordinate i is selected, its update is*

$$x_{k+1,i} = (1 - \eta_{k+1})x_{k,i} + \eta_{k+1}(-\xi_{k+1,i}/h_i),$$

while all other coordinates remain unchanged. Convexity makes the box invariant. The solver is bounded with positive definite mean, its noise has all moments, and the population and geometric remainders vanish. The common Lyapunov inequality above proves the state rate without assuming it. A constant Hessian estimator and the empirical mean solver satisfy the remaining covariance requirements. A step cutoff large enough for this bounded box is inactive. This example verifies a complete primitive route; a nonlinear manifold application must separately establish its working basin and model-specific moment bounds.

B.6.4 REDUCTION TO THE REFERENCE TANGENT SPACE AND PROOF OF THE MAIN THEOREM

Throughout the appendices, all norms on the finite-dimensional space E are equivalent. Constants denoted by C may change from line to line. All local estimates use the fixed working domains chosen before the reporting phase. A sufficiently-small-radius requirement is part of the respective localization hypothesis; containment in a larger basin does not imply containment in a smaller one.

A reduced stochastic approximation recursion Set

$$\beta_{k+1} = b^{\xi}_{k+1} + b^{q}_{k+1} + a_{k+1}, \tag{80}$$

$$\begin{aligned}
\omega_{k+1} &= (\widetilde{\mathrm{P}}_{k+1} - \Pi_k)\{\mathrm{H}u_k + r_G(x_k)\},\\
m_{k+1} &= \omega_{k+1} + \chi_{k+1} + \nu_{k+1} - c_{k+1},\\
r_{k+1} &= (\Pi_k - \bar{\mathrm{P}})\mathrm{H}u_k + \Pi_k r_G(x_k) - d_{k+1}.
\end{aligned} \tag{81}$$

Lemma B.29. *Under Assumptions B.2 and B.5,*

$$u_{k+1} = u_k - \eta_{k+1}\{\mathrm{A}u_k + \zeta_{k+1} + m_{k+1} + r_{k+1} + \beta_{k+1}\}, \qquad \mathrm{A} = \bar{\mathrm{P}}\mathrm{H}. \tag{82}$$

Moreover, $\mathbb{E}(m_{k+1} \mid \mathcal{F}_k) = 0$.

Proof. The transported solver contribution to the population part is

$$\begin{aligned}
\widetilde{\mathrm{P}}_{k+1}\{\mathrm{H}u_k + r_G(x_k)\} &= \Pi_k\{\mathrm{H}u_k + r_G(x_k)\} + (\widetilde{\mathrm{P}}_{k+1} - \Pi_k)\{\mathrm{H}u_k + r_G(x_k)\}\\
&= \mathrm{A}u_k + (\Pi_k - \bar{\mathrm{P}})\mathrm{H}u_k + \Pi_k r_G(x_k) + \omega_{k+1}.
\end{aligned}$$

By (22)–(23),

$$\widetilde{\mathrm{P}}_{k+1}\xi_{k+1} = \zeta_{k+1} + b^{\xi}_{k+1}, \qquad \widetilde{\mathrm{P}}_{k+1}q_{k+1} = \chi_{k+1} + b^{q}_{k+1}.$$

Insert these identities and (28) into (31), move the two positive geometric terms into the bracket with negative signs, and collect terms according to (80)–(81). This gives (82) exactly.

The vector $\mathrm{H}u_k + r_G(x_k)$ is $\mathcal{F}_k$-measurable, so the definition of Π_k gives $\mathbb{E}(\omega_{k+1} \mid \mathcal{F}_k) = 0$. The other three terms in m_{k+1} are conditionally centered by (23), (28), and Assumption B.5. Hence $\mathbb{E}(m_{k+1} \mid \mathcal{F}_k) = 0$. □

Lemma B.30. *Under Assumptions B.1 and B.3 to B.5 and the state bound* (34)*,*

$$\frac{1}{\sqrt{n}}\sum_{k=0}^{n-1} m_{k+1} \xrightarrow{p} 0. \tag{83}$$

Proof. The uniform field bound in Assumption B.1 gives $\|\mathrm{H}u_k + r_G(x_k)\| \le C\,\|u_k\|$. Conditional Jensen's inequality and (24) give a uniform conditional second moment bound for $\widetilde{\mathrm{P}}_{k+1} - \Pi_k$. Therefore

$$\mathbb{E}(\|\omega_{k+1}\|^2 \mid \mathcal{F}_k) \le C\,\|u_k\|^2 .$$

By (25), the same bound holds for χ_{k+1}. Using $\|z_1 + \cdots + z_4\|^2 \le 4\sum_{j=1}^4 \|z_j\|^2$, we obtain

$$\begin{aligned}\frac{1}{n}\sum_{k=0}^{n-1}\mathbb{E}(\|m_{k+1}\|^2 \mid \mathcal{F}_k) \le \frac{C}{n}\sum_{k=0}^{n-1}\|u_k\|^2 + \frac{C}{n}\sum_{k=0}^{n-1}\mathbb{E}(\|\nu_{k+1}\|^2 \mid \mathcal{F}_k)\\ + \frac{C}{n}\sum_{k=0}^{n-1}\mathbb{E}(\|c_{k+1}\|^2 \mid \mathcal{F}_k).\end{aligned}$$

The expectation of the first term is bounded by $Cn^{-1}\sum_{k=1}^n k^{-\alpha} = O(n^{-\alpha})$ by (34); hence it converges to zero in probability. The last two terms converge to zero in probability by (30) and (32). Thus the predictable quadratic variation of the finite-dimensional martingale $n^{-1/2}\sum m_{k+1}$ converges to zero in probability. Applying Lenglart's inequality coordinatewise, or the martingale isometry after a standard localization at a bounded quadratic variation, yields (83). □

Lemma B.31. *Under the assumptions of Theorem 3.1,*

$$\frac{1}{\sqrt{n}}\sum_{k=0}^{n-1} r_{k+1} \xrightarrow{p} 0, \qquad \frac{1}{\sqrt{n}}\sum_{k=0}^{n-1}\beta_{k+1} \xrightarrow{p} b. \tag{84}$$

Proof. The second conclusion is (29). From (81),

$$\frac{1}{\sqrt{n}}\sum_{k=0}^{n-1} r_{k+1} = \frac{1}{\sqrt{n}}\sum_{k=0}^{n-1}(\Pi_k - \bar{\mathrm{P}})\mathrm{H}u_k + \frac{1}{\sqrt{n}}\sum_{k=0}^{n-1}\Pi_k r_G(x_k) - \frac{1}{\sqrt{n}}\sum_{k=0}^{n-1} d_{k+1}.$$

The first two terms converge to zero by (26)–(27); the last converges to zero by (32). This proves (84). □

Abel boundary and online averaging

Lemma B.32. *If $\eta_k = \eta_0 k^{-\alpha}$ with $\alpha \in (1/2, 1)$ and* (34) *holds, then*

$$\frac{1}{\sqrt{n}}\sum_{k=0}^{n-1}\frac{u_k - u_{k+1}}{\eta_{k+1}} \xrightarrow{p} 0.$$

Proof. A direct summation by parts gives the exact identity

$$\sum_{k=0}^{n-1}\frac{u_k - u_{k+1}}{\eta_{k+1}} = \frac{u_0}{\eta_1} - \frac{u_n}{\eta_n} + \sum_{k=1}^{n-1}\left(\frac{1}{\eta_{k+1}} - \frac{1}{\eta_k}\right)u_k.$$

The first term divided by $\sqrt{n}$ converges to zero because u_0 is fixed after localization. For the terminal term,

$$\frac{1}{\sqrt{n}}\left\|\frac{u_n}{\eta_n}\right\|_{L^2} \le Cn^{-1/2}n^{\alpha}n^{-\alpha/2} = Cn^{(\alpha-1)/2} \longrightarrow 0.$$

By the mean-value theorem, $\eta_{k+1}^{-1} - \eta_k^{-1} \le Ck^{\alpha-1}$. Hence

$$\begin{aligned}\frac{1}{\sqrt{n}}\sum_{k=1}^{n-1}\left(\frac{1}{\eta_{k+1}} - \frac{1}{\eta_k}\right)\|u_k\|_{L^2} &\le \frac{C}{\sqrt{n}}\sum_{k=1}^{n-1} k^{\alpha/2-1}\\ &= O\{n^{(\alpha-1)/2}\} = o(1).\end{aligned}$$

The L^1 bound and Markov's inequality complete the proof. □

For the averaging argument, put $\widetilde{u}_n = n^{-1}\sum_{k=1}^{n} u_k$.

Lemma B.33. *Let $a \in U$ be fixed, suppose the chart transition $R_a^{-1} \circ R_{x_\star}$ and its local inverse are C^2, and let $\bar{x}_n^a$ be defined by* (17)*. If $\mathbb{E}\,\|u_k\|^2 \le Ck^{-\alpha}$ with $\alpha > 1/2$, then*

$$\sqrt{n}\{R_{x_\star}^{-1}(\bar{x}_n^a) - \widetilde{u}_n\} \xrightarrow{p} 0. \tag{85}$$

Proof. Define the local chart transition

$$f(u) = R_a^{-1}\{R_{x_\star}(u)\}, \qquad y_\star = f(0), \qquad \mathrm{J} = Df(0).$$

Because f is a diffeomorphism between local coordinate neighborhoods, J is invertible. A second-order Taylor expansion on a compact local set gives

$$f(u) = y_\star + \mathrm{J}u + r_f(u), \qquad \|r_f(u)\| \le C\,\|u\|^2. \tag{86}$$

The recursion (17) is the ordinary arithmetic mean in $T_a\mathcal{M}$, hence

$$y_n^a - y_\star = \mathrm{J}\widetilde{u}_n + \bar{r}_n, \qquad \bar{r}_n = \frac{1}{n}\sum_{k=1}^{n} r_f(u_k). \tag{87}$$

By (86) and the state rate,

$$\sqrt{n}\,\mathbb{E}\,\|\bar{r}_n\| \le \frac{C}{\sqrt{n}}\sum_{k=1}^{n} k^{-\alpha} = O(n^{1/2-\alpha}) \longrightarrow 0.$$

Thus $\sqrt{n}\,\bar{r}_n \xrightarrow{p} 0$. Moreover, Minkowski's inequality gives

$$\|\widetilde{u}_n\|_{L^2} \le \frac{1}{n}\sum_{k=1}^{n} \|u_k\|_{L^2} = O(n^{-\alpha/2}),$$

so (87) implies $y_n^a - y_\star = O_p(n^{-\alpha/2})$.

Let $g = R_{x_\star}^{-1} \circ R_a = f^{-1}$. Then $Dg(y_\star) = \mathrm{J}^{-1}$, and another second-order Taylor expansion gives

$$g(y_\star + h) = \mathrm{J}^{-1}h + r_g(h), \qquad \|r_g(h)\| \le C\,\|h\|^2.$$

Using (87),

$$R_{x_\star}^{-1}(\bar{x}_n^a) - \widetilde{u}_n = \mathrm{J}^{-1}\bar{r}_n + r_g(y_n^a - y_\star).$$

The first term is $o_p(n^{-1/2})$, while $\sqrt{n}\,\|r_g(y_n^a - y_\star)\| = O_p(n^{1/2-\alpha}) = o_p(1)$. This proves (85). □

For the recursive rule, write $\bar{u}_k^{\mathrm{rec}} = R_{x_\star}^{-1}(\bar{x}_k^{\mathrm{rec}})$.

Lemma B.34. *Suppose the retraction and relevant local inverse retractions are C^2 on a common domain. There is a sufficiently small coordinate ball, chosen before the recursion, such that the following holds. If all u_k lie in that ball and $\mathbb{E}\,\|u_k\|^2 \le Ck^{-\alpha}$ with $1/2 < \alpha < 1$, then the recursive average in* (16) *is well defined, remains in the ball, and*

$$\sqrt{n}\{\bar{u}_n^{\mathrm{rec}} - \widetilde{u}_n\} \xrightarrow{p} 0.$$

Moreover, $\mathbb{E}\,\|\bar{u}_k^{\mathrm{rec}}\|^2 = O(k^{-\alpha})$.

Proof. For $x(v) = R_{x_\star}(v)$ define

$$f_v(w) = R_{x_\star}^{-1}\big(R_{x(v)}[\mathcal{T}_{x_\star \to x(v)} w]\big).$$

The inverse function theorem, uniformly on a compact neighborhood, gives $w_{v,u}$ with $f_v(w_{v,u}) = u$ and $\|w_{v,u}\| \le C\,\|u - v\|$. The averaging map is $\mathcal{A}(v,u,t) = f_v(tw_{v,u})$. Its values at $t = 0, 1$ are v, u, and bounded second derivatives of f_v give

$$\sup_{0\le t\le 1} \|\partial_{tt}\mathcal{A}(v,u,t)\| \le C\,\|u - v\|^2.$$

For a twice differentiable vector-valued function g on $[0,1]$, subtract its linear endpoint interpolant. Integrating g'' twice, with zero boundary values for the difference, gives the Green-kernel bound $\|g(t) - (1-t)g(0) - tg(1)\| \le \frac{1}{2}t(1-t)\sup_s \|g''(s)\|$. Therefore

$$\mathcal{A}(v,u,t) = (1-t)v + tu + \varrho(v,u,t), \qquad \|\varrho(v,u,t)\| \le Ct(1-t)\|u-v\|^2. \tag{88}$$

Only C^2 regularity has been used.

Let $\|u\|, \|v\| \le r$, put $w = (1-t)v + tu$ and $D = t(1-t)\|u-v\|^2 \le r^2$. Since

$$\|w\|^2 = (1-t)\|v\|^2 + t\|u\|^2 - D,$$

we have

$$\|\mathcal{A}(v,u,t)\|^2 \le (1-t)\|v\|^2 + t\|u\|^2 - \{1 - 2Cr - C^2r^2\}D.$$

Choose r within the common domain and with $2Cr + C^2r^2 \le 1$. This proves invariance of the coordinate ball for the average, and with $t = 1/k$ gives

$$\mathbb{E}\,\|\bar{u}_k^{\mathrm{rec}}\|^2 \le \frac{1}{k}\sum_{j=1}^{k}\mathbb{E}\,\|u_j\|^2 = O(k^{-\alpha}).$$

The recursive identity is

$$\bar{u}_k^{\mathrm{rec}} = \left(1 - \frac{1}{k}\right)\bar{u}_{k-1}^{\mathrm{rec}} + \frac{1}{k}u_k + \varrho_k, \qquad \|\varrho_k\| \le \frac{C}{k}\left\|u_k - \bar{u}_{k-1}^{\mathrm{rec}}\right\|^2.$$

Subtracting the arithmetic-mean recursion and using the common initialization at $k = 1$ yields

$$n(\bar{u}_n^{\mathrm{rec}} - \widetilde{u}_n) = \sum_{k=2}^{n} k\varrho_k.$$

Consequently

$$\mathbb{E}\left\|\sqrt{n}(\bar{u}_n^{\mathrm{rec}} - \widetilde{u}_n)\right\| \le \frac{C}{\sqrt{n}}\sum_{k=2}^{n}\left\{\mathbb{E}\,\|u_k\|^2 + \mathbb{E}\left\|\bar{u}_{k-1}^{\mathrm{rec}}\right\|^2\right\} = O(n^{1/2-\alpha}) \longrightarrow 0.$$

Markov's inequality proves the claim. □

Proof of the main theorem

Proof of Theorem 3.1. Divide (82) by η_{k+1} and rearrange:

$$\mathrm{A}u_k = \frac{u_k - u_{k+1}}{\eta_{k+1}} - \zeta_{k+1} - m_{k+1} - r_{k+1} - \beta_{k+1}.$$

Sum this identity from $k = 0$ to $n-1$ and divide by $\sqrt{n}$. By Lemmas B.30 to B.32,

$$\mathrm{A}\frac{1}{\sqrt{n}}\sum_{k=0}^{n-1}u_k = -\frac{1}{\sqrt{n}}\sum_{k=0}^{n-1}\zeta_{k+1} - b + o_p(1). \tag{89}$$

The arithmetic averages based on $u_0, \ldots, u_{n-1}$ and $u_1, \ldots, u_n$ differ, after multiplication by $\sqrt{n}$, by $(u_n - u_0)/\sqrt{n} = o_p(1)$. By (33), the latter average is root-n equivalent to $\bar{u}_n = R_{x_\star}^{-1}(\bar{x}_n)$. Since A is invertible, (89) gives (2).

The martingale central limit theorem applies to $n^{-1/2}\sum_{k=0}^{n-1}\zeta_{k+1}$. Its predictable covariance converges to Q by (35), and the conditional Lindeberg condition is assumed. Alternatively, under (78), for every $\epsilon > 0$,

$$\begin{aligned}\frac{1}{n}\sum_{k=0}^{n-1}\mathbb{E}\{\|\zeta_{k+1}\|^2\,\mathbf{1}(\|\zeta_{k+1}\| > \epsilon\sqrt{n}) \mid \mathcal{F}_k\} &\le \frac{1}{\epsilon^\tau n^{1+\tau/2}}\sum_{k=0}^{n-1}\mathbb{E}(\|\zeta_{k+1}\|^{2+\tau} \mid \mathcal{F}_k)\\ &\le Cn^{-\tau/2} \longrightarrow 0.\end{aligned}$$

The finite-dimensional martingale CLT (Hall and Heyde, 1980) therefore yields $n^{-1/2}\sum \zeta_{k+1} \xrightarrow{d} N_E(0, \mathrm{Q})$. Slutsky's theorem proves (3).

Finally, in the retraction chart,

$$\psi\{R_{x_\star}(u)\} = \psi(x_\star) + D\psi_{x_\star}u + o(\|u\|).$$

The established central limit theorem implies $\bar{u}_n = O_p(n^{-1/2})$. Multiplication by $\sqrt{n}$ and another application of Slutsky's theorem give (4). □

Proof of Proposition B.9. Sum the rearranged recursion in the proof of Theorem 3.1 from $k = b_n$ to $n-1$ and divide by $\sqrt{N_n}$. The two endpoint terms are $o_p(1)$ by Lemma B.32, applied at n and at b_n; this is unchanged when $\tau = 0$ because the initial endpoint is then root-N_n negligible. The secondary martingale, remainder, and centered bias sums over the tail are differences of the corresponding prefix sums, so Lemmas B.30 and B.31 and the centered condition give

$$\mathrm{A}\frac{1}{\sqrt{N_n}}\sum_{k=b_n}^{n-1} u_k = -\frac{1}{\sqrt{N_n}}\sum_{k=b_n}^{n-1}\zeta_{k+1} + o_p(1).$$

Prefix convergence of the predictable covariance averages implies

$$\frac{1}{N_n}\sum_{k=b_n}^{n-1}\mathbb{E}(\zeta_{k+1}\zeta_{k+1}^* \mid \mathcal{F}_k) \xrightarrow{p} \mathrm{Q},$$

and the prefix Lindeberg condition implies its tail version. The martingale central limit theorem therefore gives the stated Gaussian limit. Replacing the arithmetic tail average by the reported average uses the assumed tail equivalence. For the two concrete averages, restarting their recursions at $b_n + 1$ and repeating the proofs of Lemmas B.33 and B.34 with N_n in place of n establishes that equivalence. □

Proof of Corollary B.8. Take $b = 0$ in Theorem 3.1. □

B.6.5 Product and geometric conditions

Proofs of the product and remainder criteria

Proof of Proposition B.25. Conditional independence and (21) give

$$b_{k+1}^{\xi} = \mathbb{E}(\widetilde{\mathrm{P}}_{k+1} \mid \mathcal{F}_k)\mathbb{E}(\xi_{k+1} \mid \mathcal{F}_k) = 0.$$

Moreover,

$$\mathbb{E}(\left\|\widetilde{\mathrm{P}}_{k+1}\xi_{k+1}\right\|^2 \mid \mathcal{F}_k) \le \mathbb{E}(\left\|\widetilde{\mathrm{P}}_{k+1}\right\|_{\mathrm{op}}^2 \mid \mathcal{F}_k)\mathbb{E}(\|\xi_{k+1}\|^2 \mid \mathcal{F}_k) \le C.$$

The same factorization gives $b_{k+1}^{q} = 0$ and

$$\mathbb{E}(\|\chi_{k+1}\|^2 \mid \mathcal{F}_k) = \mathbb{E}(\left\|\widetilde{\mathrm{P}}_{k+1}q_{k+1}\right\|^2 \mid \mathcal{F}_k) \le C\|u_k\|^2.$$

If (78) holds, conditional centering and Jensen's inequality give a uniform conditional $(2+\tau)$ moment of ζ_{k+1}. Hence, for every $\epsilon > 0$,

$$\begin{aligned}\frac{1}{n}\sum_{k=0}^{n-1}\mathbb{E}\Big[\|\zeta_{k+1}\|^2\,\mathbf{1}\{\|\zeta_{k+1}\| > \epsilon\sqrt{n}\} \mid \mathcal{F}_k\Big] \\ \le \frac{C}{\epsilon^{\tau}n^{1+\tau/2}}\sum_{k=0}^{n-1}\mathbb{E}(\|\zeta_{k+1}\|^{2+\tau} \mid \mathcal{F}_k) \le Cn^{-\tau/2} \longrightarrow 0.\end{aligned}$$

Under conditional independence, the separate moments factorize and imply (78). Without independence, conditional Hölder with exponents determined by the reciprocal-sum condition gives the same product moment. That argument does not factor the first moment and therefore does not establish $b_{k+1}^{\xi} = 0$. □

Proof of Proposition B.26. For part (i), Cauchy-Schwarz and (34) give

$$\mathbb{E}\left\|(\Pi_k - \bar{\mathrm{P}})\mathrm{H}u_k\right\| \le \|\mathrm{H}\|_{\mathrm{op}} \{\mathbb{E}\left\|\Pi_k - \bar{\mathrm{P}}\right\|_{\mathrm{op}}^2\}^{1/2} \{\mathbb{E}\,\|u_k\|^2\}^{1/2} \le Ck^{-\delta-\alpha/2}.$$

Consequently,

$$\mathbb{E}\left\{\frac{1}{\sqrt{n}}\sum_{k=1}^{n-1}\left\|(\Pi_k - \bar{\mathrm{P}})\mathrm{H}u_k\right\|\right\} \le \frac{C}{\sqrt{n}}\sum_{k=1}^{n-1} k^{-\delta-\alpha/2} = o(1)$$

when $\delta > (1-\alpha)/2$. Markov's inequality proves (26).

For part (ii), conditional Jensen and (24) give a uniform bound on $\|\Pi_k\|_{\mathrm{op}}$. Since $1+\rho \le 2$, Lyapunov's inequality and (34) imply

$$\mathbb{E}\,\|u_k\|^{1+\rho} \le \{\mathbb{E}\,\|u_k\|^2\}^{(1+\rho)/2} \le Ck^{-\alpha(1+\rho)/2}.$$

It follows that

$$\mathbb{E}\left\{\frac{1}{\sqrt{n}}\sum_{k=1}^{n-1}\|\Pi_k r_G(x_k)\|\right\} \le \frac{C}{\sqrt{n}}\sum_{k=1}^{n-1} k^{-\alpha(1+\rho)/2} = o(1)$$

when $\alpha(1+\rho) > 1$. Markov's inequality proves (27). □

A two-variable retraction expansion For $u, z \in E$ and $t \ge 0$, define

$$\Phi(u,z,t) = R_{x_\star}^{-1}\left[R_{R_{x_\star}(u)}\{-t\mathcal{T}_{x_\star \to R_{x_\star}(u)} z\}\right]. \tag{90}$$

Lemma B.35. *On a sufficiently small relatively compact coordinate neighborhood, there is a C^1 operator map $\mathrm{L}(u) : E \to E$ and a remainder $\mathcal{R}_2(u,z,t)$ such that, whenever $t\,\|z\| \le r_0$,*

$$\Phi(u,z,t) = u - tz - t\mathrm{L}(u)z + t^2\mathcal{R}_2(u,z,t), \tag{91}$$

with

$$\mathrm{L}(0) = 0, \qquad \|\mathrm{L}(u)\|_{\mathrm{op}} \le C\,\|u\|, \qquad \|\mathcal{R}_2(u,z,t)\| \le C\,\|z\|^2. \tag{92}$$

Proof. The map in (90) is C^2 because R, its local inverse, and the frame transport are C^2. It satisfies $\Phi(u,z,0) = u$. Its t-derivative at zero is linear in z; write $\partial_t\Phi(u,z,0) = -\mathrm{B}(u)z$. At $u = 0$, the retraction identities $R_{x_\star}(0) = x_\star$ and $DR_{x_\star}(0) = \mathrm{I}$ imply $\mathrm{B}(0) = \mathrm{I}$. Since B is C^1 on a compact local set, $\mathrm{L}(u) := \mathrm{B}(u) - \mathrm{I}$ satisfies $\mathrm{L}(0) = 0$ and $\|\mathrm{L}(u)\|_{\mathrm{op}} \le C\,\|u\|$ by the mean-value theorem.

Taylor's theorem in t, with integral remainder, gives

$$\Phi(u,z,t) = u - t\mathrm{B}(u)z + t^2\int_0^1 (1-s)\partial_{tt}\Phi(u,z,st)\,ds.$$

The second derivative is bilinear in z up to bounded chart coefficients. Compactness and the safeguard $t\,\|z\| \le r_0$ therefore give $\|\partial_{tt}\Phi(u,z,st)\| \le C\,\|z\|^2$. Substituting $\mathrm{B}(u) = \mathrm{I} + \mathrm{L}(u)$ proves (91)–(92). □

Proof of the geometric proposition

Proof of Proposition B.27. Put

$$Z_{k+1} = \mathcal{T}_{x_k \to x_\star} D_{k+1}, \qquad \mu_{k+1} = \mathbb{E}(Z_{k+1} \mid \mathcal{F}_k).$$

Coherence of the transport gives $D_{k+1} = \mathcal{T}_{x_\star \to x_k} Z_{k+1}$. Applying Lemma B.35 with $t = \eta_{k+1}$ yields

$$u_{k+1} = u_k - \eta_{k+1}Z_{k+1} - \eta_{k+1}L(u_k)Z_{k+1} + \eta_{k+1}^2\mathcal{R}_{2,k+1}. \tag{93}$$

Define

$$c_{k+1} = -L(u_k)\{Z_{k+1} - \mu_{k+1}\},$$
$$d_{k+1} = -L(u_k)\mu_{k+1} + \eta_{k+1}\mathcal{R}_{2,k+1}.$$

Then (93) has exactly the form (31), and $\mathbb{E}(c_{k+1} \mid \mathcal{F}_k) = 0$.

We next verify the two summability statements. By (92) and conditional variance domination,

$$\begin{aligned}\mathbb{E}(\|c_{k+1}\|^2 \mid \mathcal{F}_k) &\le C\,\|u_k\|^2\,\mathbb{E}(\|Z_{k+1} - \mu_{k+1}\|^2 \mid \mathcal{F}_k)\\ &\le C\,\|u_k\|^2\,\{1 + \|u_k\|^2 + \|\beta_{k+1}\|^2\}.\end{aligned}$$

On the localized neighborhood, $\|u_k\|$ is uniformly small and bounded. Using (34) and $\mathbb{E}\,\|\beta_{k+1}\|^2 \le C\eta_{k+1}$, we obtain

$$\mathbb{E}\left[\frac{1}{n}\sum_{k=0}^{n-1}\mathbb{E}(\|c_{k+1}\|^2 \mid \mathcal{F}_k)\right] \le \frac{C}{n}\sum_{k=0}^{n-1}\eta_{k+1} \longrightarrow 0.$$

Thus the first statement of (32) holds by Markov's inequality.

From the score and error decompositions,

$$\mu_{k+1} = \Pi_k\{\mathrm{H}u_k + r_G(x_k)\} + \beta_{k+1}. \tag{94}$$

Conditional Jensen and (24) bound $\|\Pi_k\|_{\mathrm{op}}$ uniformly. The uniform field bound in Assumption B.1 therefore gives

$$\|\mu_{k+1}\| \le C\,\|u_k\| + \|\beta_{k+1}\|\,.$$

Equations (92), (79), and Cauchy-Schwarz now imply

$$\begin{aligned}\mathbb{E}\,\|d_{k+1}\| &\le C\mathbb{E}\{\|u_k\|\,\|\mu_{k+1}\|\} + C\eta_{k+1}\mathbb{E}\,\|Z_{k+1}\|^2\\ &\le C\mathbb{E}\,\|u_k\|^2 + C\{\mathbb{E}\,\|u_k\|^2\}^{1/2}\{\mathbb{E}\,\|\beta_{k+1}\|^2\}^{1/2} + C\eta_{k+1}\\ &\le C\eta_{k+1}.\end{aligned}$$

Because $\alpha > 1/2$,

$$\mathbb{E}\left\|\frac{1}{\sqrt{n}}\sum_{k=0}^{n-1} d_{k+1}\right\| \le \frac{C}{\sqrt{n}}\sum_{k=1}^{n} k^{-\alpha} = O(n^{1/2-\alpha}) \longrightarrow 0.$$

This proves the second statement of (32).

It remains to justify clipping. Let Z^0_{k+1} be the raw pre-clipping direction evaluated along the actual safeguarded trajectory, and let Z^{c}_{k+1} be that direction clipped at radius r_0/η_{k+1}. With $\Delta_{k+1} = Z^0_{k+1} - Z^{\mathrm{c}}_{k+1}$ and a uniform conditional $(2+\tau)$ moment,

$$\begin{aligned}\mathbb{E}(\|\Delta_{k+1}\| \mid \mathcal{F}_k) &\le \mathbb{E}\{\left\|Z^0_{k+1}\right\|\mathbf{1}(\eta_{k+1}\left\|Z^0_{k+1}\right\| > r_0) \mid \mathcal{F}_k\}\\ &\le r_0^{-(1+\tau)}\eta_{k+1}^{1+\tau}\mathbb{E}(\left\|Z^0_{k+1}\right\|^{2+\tau} \mid \mathcal{F}_k)\\ &\le C\eta_{k+1}^{1+\tau}.\end{aligned}$$

Moreover,

$$\mathbb{E}(\|\Delta_{k+1}\|^2 \mid \mathcal{F}_k) \le r_0^{-\tau}\eta_{k+1}^{\tau}\mathbb{E}(\left\|Z^0_{k+1}\right\|^{2+\tau} \mid \mathcal{F}_k) \le C\eta_{k+1}^{\tau}.$$

Hence

$$\mathbb{E}\left\{\frac{1}{\sqrt{n}}\sum_{k=0}^{n-1}\|\Delta_{k+1}\|\right\} \le \frac{C}{\sqrt{n}}\sum_{k=1}^{n} k^{-\alpha(1+\tau)} = o(1),$$

because $\alpha(1+\tau) > 1/2$, while $\frac{1}{n}\sum_k \mathbb{E}(\|\Delta_{k+1}\|^2 \mid \mathcal{F}_k) \xrightarrow{p} 0$. Writing Δ_{k+1} as its conditional mean plus a centered part also gives the average squared bias condition by conditional Jensen. Thus the clipping discrepancy satisfies all three inexactness requirements used by the safeguarded expansion. This proves the claim without asserting that the safeguarded and unsafeguarded state paths are asymptotically coupled. □

B.6.6 LYAPUNOV PROOF OF THE STATE RATE

Lemma B.36. *Under the assumptions of Proposition B.10, there exist constants $c, C > 0$ such that, for all sufficiently large k,*

$$\mathbb{E}\{\langle u_{k+1}, \mathrm{K}u_{k+1}\rangle\} \le (1 - c\eta_{k+1})\mathbb{E}\{\langle u_k, \mathrm{K}u_k\rangle\} + C\eta_{k+1}^2. \tag{95}$$

Proof. Let $V(u) = \langle u, \mathrm{K}u\rangle$. The local retraction expansion in Lemma B.35 and the direction moment assumptions imply a one step representation

$$u_{k+1} = u_k - \eta_{k+1} Z_{k+1}^{\circ}, \tag{96}$$

where, conditionally on $\mathcal{F}_k$,

$$\begin{aligned} \mathbb{E}(Z_{k+1}^{\circ} \mid \mathcal{F}_k) &= \Pi_k \mathrm{H} u_k + \beta_{k+1} + s_{k+1}, \\ \mathbb{E}(\left\|Z_{k+1}^{\circ}\right\|^2 \mid \mathcal{F}_k) &\le C\{1 + \|u_k\|^2 + \|\beta_{k+1}\|^2\}, \end{aligned} \tag{97}$$

with a geometric/nonlinear mean remainder satisfying, for a fixed $\varepsilon > 0$ when U is chosen sufficiently small before localization,

$$\|s_{k+1}\| \le \varepsilon \|u_k\| + C\{\|u_k\|^2 + \eta_{k+1} + \|u_k\| \|\beta_{k+1}\| + \eta_{k+1} \|\beta_{k+1}\|^2\}. \tag{98}$$

To justify the second moment bound without a fourth moment, set $Z_{k+1} = \mathcal{T}_{x_k \to x_\star} D_{k+1}$ and $\eta = \eta_{k+1}$. The local update expansion gives

$$Z_{k+1}^{\circ} = (\mathrm{I} + L(u_k)) Z_{k+1} - \eta \mathcal{R}_{2,k+1}, \qquad \|\mathcal{R}_{2,k+1}\| \le C \|Z_{k+1}\|^2 .$$

Since the actual step satisfies $\eta \|Z_{k+1}\| \le r_0$,

$$\left\|Z_{k+1}^{\circ}\right\| \le (1 + C\|u_k\| + Cr_0) \|Z_{k+1}\| .$$

Thus its conditional second moment is bounded by a constant multiple of that of the actual direction. The domain and its radius are fixed before the reporting phase; all sufficiently-small-domain requirements in this proof refer to that same choice.

To bound the mean remainder, use (94): the population nonlinearity contributes $\Pi_k r_G(x_k)$, the first chart correction contributes $L(u_k)\mu_{k+1}$, and the second chart derivative contributes $\eta_{k+1}\mathbb{E}(\mathcal{R}_{2,k+1} \mid \mathcal{F}_k)$. Bounds (20), (92), and (79) give (98); the last displayed term is the contribution of the conditional second-order chart remainder. Centered score, solver, inexactness, and chart parts give (97).

Since V is quadratic, (96) gives exactly

$$\begin{aligned} \mathbb{E}\{V(u_{k+1}) \mid \mathcal{F}_k\} = V(u_k) &- 2\eta_{k+1} \langle \mathrm{K}u_k, \Pi_k \mathrm{H} u_k + \beta_{k+1} + s_{k+1}\rangle \\ &+ \eta_{k+1}^2 \mathbb{E}\{\langle Z_{k+1}^{\circ}, \mathrm{K} Z_{k+1}^{\circ}\rangle \mid \mathcal{F}_k\}. \end{aligned} \tag{99}$$

By (37),

$$2\langle \mathrm{K}u_k, \Pi_k \mathrm{H} u_k\rangle = \langle u_k, \{\mathrm{K}\Pi_k \mathrm{H} + (\Pi_k \mathrm{H})^* \mathrm{K}\} u_k\rangle \ge c_0 \|u_k\|^2 .$$

Young's inequality gives, for a sufficiently small fixed $\epsilon > 0$,

$$2 \|\mathrm{K}\|_{\mathrm{op}} \|u_k\| \|\beta_{k+1}\| \le \epsilon \|u_k\|^2 + C_\epsilon \|\beta_{k+1}\|^2 .$$

The terms in (98) are handled similarly. Choose the coefficient ε there small relative to the stable drift. The product $\|u_k\|^3$ is at most a small constant times $\|u_k\|^2$ on this preselected sufficiently small U. The term $\|u_k\|^2 \|\beta_{k+1}\|$ is bounded by the same Young inequality because $\|u_k\|$ is uniformly small. The remaining $\eta_{k+1}\|u_k\|$ term is bounded by $\epsilon \|u_k\|^2 + C\eta_{k+1}^2$. Because the safeguarded reporting phase keeps u_k uniformly bounded, the unconditional contribution of $\eta_{k+1}\|\beta_{k+1}\|^2$ to the drift in (99) is $O(\eta_{k+1}^3)$ under $\mathbb{E}\|\beta_{k+1}\|^2 \le C\eta_{k+1}$. Finally, (97) bounds the quadratic increment in (99).

Choose U and the Young constants so that all $\|u_k\|^2$ remainders consume at most $c_0/2$ of the stable drift. Taking expectations, using $\mathbb{E}\|\beta_{k+1}\|^2 \le C\eta_{k+1}$, and absorbing the harmless η_{k+1}^3 terms gives

$$\mathbb{E}V(u_{k+1}) \le \mathbb{E}V(u_k) - c\eta_{k+1}\mathbb{E}\|u_k\|^2 + C\eta_{k+1}^2.$$

The eigenvalues of K are bounded above and below by positive constants, so this is equivalent, after changing c and C, to (95). □

Proof of Proposition B.10. Put $y_k = \mathbb{E}V(u_k)$. By Lemma B.36, $y_{k+1} \le (1 - c\eta_{k+1})y_k + C\eta_{k+1}^2$. We compare this recursion with $B\eta_k$. Since $\alpha < 1$,

$$\frac{\eta_k - \eta_{k+1}}{\eta_{k+1}^2} = O(k^{\alpha-1}) \longrightarrow 0. \tag{100}$$

Choose $B > 2C/c$. For all sufficiently large k, (100) implies

$$B(\eta_k - \eta_{k+1}) \le \frac{cB - C}{2}\eta_{k+1}^2.$$

If $y_k \le B\eta_k$, then, using $\eta_k/\eta_{k+1} \to 1$,

$$\begin{aligned} y_{k+1} &\le B\eta_k - cB\eta_{k+1}\eta_k + C\eta_{k+1}^2 \\ &\le B\eta_{k+1} + B(\eta_k - \eta_{k+1}) - (cB - C + o(1))\eta_{k+1}^2 \\ &\le B\eta_{k+1}. \end{aligned}$$

Increasing B handles the finite initial segment. Thus $y_k = O(\eta_k)$. Equivalence of $V(u)$ and $\|u\|^2$ proves (34). □

B.6.7 Proof of online covariance consistency

Lemma B.37. *Let $X_{n,k} \ge 0$ be an integrable adapted triangular array and let $\mathrm{A}_n = \sum_k \mathbb{E}(X_{n,k} \mid \mathcal{F}_{k-1})$. If $\mathrm{A}_n \xrightarrow{p} 0$, then $\sum_k X_{n,k} \xrightarrow{p} 0$.*

Proof. Fix $a > 0$ and retain term k only if the predictable cumulative compensator, including term k, is at most a. This is an $\mathcal{F}_{k-1}$-measurable decision. Nonnegativity implies that the expected sum of retained terms is at most a, while every term is retained on $\{\mathrm{A}_n \le a\}$. Hence, for $\epsilon > 0$,

$$\Pr\left(\sum_k X_{n,k} > \epsilon\right) \le \Pr(\mathrm{A}_n > a) + a/\epsilon.$$

Let $n \to \infty$ and then $a \downarrow 0$. □

Lemma B.38. *Let (ζ_k) be a square-integrable martingale difference sequence in a fixed finite-dimensional inner-product space. If its predictable covariance averages converge in probability to a finite deterministic operator Q and it satisfies the conditional Lindeberg condition at scale $\sqrt{n}$, then*

$$\frac{1}{n}\sum_{k=1}^{n} \{\zeta_k \otimes \zeta_k - \mathbb{E}(\zeta_k \otimes \zeta_k \mid \mathcal{F}_{k-1})\} \xrightarrow{p} 0. \tag{101}$$

Proof. Fix $\epsilon > 0$ and split each outer product according to $\|\zeta_k\| > \epsilon\sqrt{n}$. The sum of conditional norms of the normalized tail terms tends to zero by Lindeberg. Its realized sum also tends to zero by Lemma B.37.

For the remaining terms put

$$\mathrm{B}_{n,k} = \frac{1}{n}\left\{\zeta_k \otimes \zeta_k \mathbf{1}\{\|\zeta_k\| \le \epsilon\sqrt{n}\} - \mathbb{E}\big(\zeta_k \otimes \zeta_k \mathbf{1}\{\|\zeta_k\| \le \epsilon\sqrt{n}\} \mid \mathcal{F}_{k-1}\big)\right\}.$$

They are matrix-valued martingale differences, and

$$\mathbb{E}(\|\mathrm{B}_{n,k}\|_{\mathrm{F}}^2 \mid \mathcal{F}_{k-1}) \le \frac{\epsilon^2}{n}\mathbb{E}(\|\zeta_k\|^2 \mid \mathcal{F}_{k-1}).$$

The total predictable quantity $T_n = n^{-1}\sum_k \mathbb{E}(\|\zeta_k\|^2 \mid \mathcal{F}_{k-1})$ is tight, since it converges to $\operatorname{tr}\mathrm{Q}$. Retain a term only when the predictable cumulative quantity defining T_n, including that term, is at most K. The retained martingale has total expected squared Frobenius norm at most $\epsilon^2 K$. Thus, for $\delta > 0$,

$$\Pr\left(\left\|\sum_k \mathrm{B}_{n,k}\right\|_{\mathrm{F}} > \delta\right) \le \Pr(T_n > K) + \epsilon^2 K/\delta^2.$$

For fixed K, first take $n \to \infty$ and then $\epsilon \downarrow 0$; finally let $K \to \infty$. Combining this bound with the tail argument proves the weak law. □

Lemma B.39. *Under Assumptions B.1 to B.7 and 3.2,*

$$\frac{1}{n}\sum_{k=0}^{n-1}\|\nu_{k+1}\|^2 \xrightarrow{p} 0. \tag{102}$$

The actual anchor space directions satisfy

$$Y_{k+1} = \mathcal{T}_{x_\star\to a}\zeta_{k+1} + h_{k+1}, \qquad \frac{1}{n}\sum_{k=0}^{n-1}\|h_{k+1}\|^2 \xrightarrow{p} 0, \tag{103}$$

and (101) *holds.*

Proof. Apply Lemma B.37 with $X_{n,k+1} = \|\nu_{k+1}\|^2/n$ and use (30) to obtain (102). Coherence of the frame and the score and inexactness decompositions give the exact identity

$$h_{k+1} = \mathcal{T}_{x_\star\to a}\left[\widetilde{\mathrm{P}}_{k+1}\{\mathrm{H}u_k + r_G(x_k)\} + \chi_{k+1} + \beta_{k+1} + \nu_{k+1}\right].$$

By the uniform field bound in Assumption B.1, (24), and the state rate,

$$\mathbb{E}\left\|\widetilde{\mathrm{P}}_{k+1}\{\mathrm{H}u_k + r_G(x_k)\}\right\|^2 \le C\mathbb{E}\,\|u_k\|^2 = O(\eta_k).$$

The same estimate holds for χ_{k+1} by (25). Their Cesàro squared norms vanish in probability by Markov's inequality. The squared bias condition and (102), together with the finite-sum squared-norm inequality, prove (103). Finally Lemma B.38 applies by (35) and (36). □

Lemma B.40. *Under the assumptions of Lemma B.39,*

$$\frac{1}{n}\sum_{k=0}^{n-1} Y_{k+1}\otimes Y_{k+1} \xrightarrow{p} \mathrm{Q}_a, \qquad \bar{Y}_n = \frac{1}{n}\sum_{k=0}^{n-1} Y_{k+1} \xrightarrow{p} 0.$$

Proof. Write $Z_{k+1} = \mathcal{T}_{x_\star\to a}\zeta_{k+1}$, so $Y_{k+1} = Z_{k+1} + h_{k+1}$. The weak law in Lemma B.38 and (35) imply

$$\frac{1}{n}\sum_{k=0}^{n-1} Z_{k+1}\otimes Z_{k+1} \xrightarrow{p} \mathrm{Q}_a. \tag{104}$$

In particular, $n^{-1}\sum\|Z_{k+1}\|^2 = O_p(1)$. The tensor inequality $\|v\otimes w\|_{\mathrm{op}} \le \|v\|\,\|w\|$ and Cauchy-Schwarz give

$$\begin{aligned}&\left\|\frac{1}{n}\sum_{k=0}^{n-1}\{Y_{k+1}\otimes Y_{k+1} - Z_{k+1}\otimes Z_{k+1}\}\right\|_{\mathrm{op}}\\ &\quad\le 2\left(\frac{1}{n}\sum\|Z_{k+1}\|^2\right)^{1/2}\left(\frac{1}{n}\sum\|h_{k+1}\|^2\right)^{1/2} + \frac{1}{n}\sum\|h_{k+1}\|^2 = o_p(1).\end{aligned}$$

Together with (104), this proves the first statement.

For the mean, (35) makes the predictable quadratic variation of $n^{-1/2}\sum Z_{k+1}$ bounded in probability. Lenglart's inequality therefore gives $n^{-1}\sum Z_{k+1} = O_p(n^{-1/2})$. Also

$$\left\|\frac{1}{n}\sum h_{k+1}\right\| \le \left(\frac{1}{n}\sum\|h_{k+1}\|^2\right)^{1/2} = o_p(1).$$

Hence $\bar{Y}_n \xrightarrow{p} 0$. □

Proof of Theorem 3.3. Welford's algebraic identity is

$$\widehat{\mathrm{Q}}_n^a = \frac{1}{n}\sum_{k=0}^{n-1}(Y_{k+1}-\bar{Y}_n)\otimes(Y_{k+1}-\bar{Y}_n) = \frac{1}{n}\sum_{k=0}^{n-1} Y_{k+1}\otimes Y_{k+1} - \bar{Y}_n\otimes\bar{Y}_n.$$

Thus Lemma B.40 proves $\widehat{\mathrm{Q}}_n^a \xrightarrow{p} \mathrm{Q}_a$. By (5), $\widehat{\mathrm{A}}_n^a = \widehat{\mathrm{P}}_n^a \widehat{\mathrm{H}}_n^a \xrightarrow{p} \bar{\mathrm{P}}_a \mathrm{H}_a$. The limit is invertible. Matrix inversion and multiplication are continuous on the set of invertible operators, so

$$\widehat{\mathrm{V}}_n^a = (\widehat{\mathrm{A}}_n^a)^{-1} \widehat{\mathrm{Q}}_n^a ((\widehat{\mathrm{A}}_n^a)^{*-1}) \xrightarrow{p} \mathrm{V}_a.$$

Coherence of the frame transport gives

$$\mathcal{T}_{a \to \bar{x}_n} \mathrm{V}_a \mathcal{T}_{\bar{x}_n \to a} = \mathcal{T}_{x_\star \to \bar{x}_n} \mathrm{V} \mathcal{T}_{\bar{x}_n \to x_\star}.$$

Transport is isometric, so transporting the preceding convergence proves the third statement in (6).

We now identify the centering boundary for coverage. Let $\bar{u}_n = R_{x_\star}^{-1}(\bar{x}_n)$ and $s_n = R_{\bar{x}_n}^{-1}(x_\star)$. A C^2 expansion of the two-point inverse retraction map at $(x_\star, x_\star)$ yields

$$\mathcal{T}_{\bar{x}_n \to x_\star} s_n = -\bar{u}_n + O_p(\|\bar{u}_n\|^2). \tag{105}$$

Indeed, the map vanishes on the diagonal and its two first derivatives at the diagonal are $-\mathrm{I}$ and I, respectively. The main theorem gives $\bar{u}_n = O_p(n^{-1/2})$, so the remainder in (105) is $o_p(n^{-1/2})$.

Use the isometry of transport to write the Wald statistic at the true parameter as

$$\begin{aligned} W_n &= n \left\langle s_n, (\widehat{\mathrm{V}}_n^{\bar{x}})^{-1} s_n \right\rangle_{\bar{x}_n} \\ &= n \left\langle \mathcal{T}_{\bar{x}_n \to x_\star} s_n, \{\mathcal{T}_{\bar{x}_n \to x_\star} \widehat{\mathrm{V}}_n^{\bar{x}} \mathcal{T}_{x_\star \to \bar{x}_n}\}^{-1} \mathcal{T}_{\bar{x}_n \to x_\star} s_n \right\rangle_{x_\star}. \end{aligned}$$

The covariance inside braces converges to V. By (105) and Theorem 3.1,

$$\sqrt{n} \mathcal{T}_{\bar{x}_n \to x_\star} s_n \xrightarrow{d} N_E(\mathrm{A}^{-1} b, \mathrm{V}).$$

If Q is positive definite, so is V. Slutsky's theorem therefore gives $W_n \xrightarrow{d} \chi_d^2(\lambda_b)$ with (8). When $b = 0$, this is the central chi-squared law and proves (7); when $b \neq 0$, it proves the stated noncentral boundary.

It remains to verify the correction in (9). Set $\delta_n^{\bar{x}} = \mathcal{T}_{a \to \bar{x}_n} (\widehat{\mathrm{A}}_n^a)^{-1} \widehat{b}_n^a \in T_{\bar{x}_n} \mathcal{M}$. Coherence, consistency of $\widehat{\mathrm{A}}_n^a$, and $\widehat{b}_n^a \xrightarrow{p} b_a$ imply

$$\mathcal{T}_{\bar{x}_n \to x_\star} \delta_n^{\bar{x}} = \mathcal{T}_{a \to x_\star} (\widehat{\mathrm{A}}_n^a)^{-1} \widehat{b}_n^a \xrightarrow{p} \mathrm{A}^{-1} b.$$

A C^2 expansion of the two-point retraction map gives

$$R_{x_\star}^{-1}(\bar{x}_n^{\mathrm{bc}}) = \bar{u}_n + n^{-1/2} \mathcal{T}_{\bar{x}_n \to x_\star} \delta_n^{\bar{x}} + o_p(n^{-1/2}).$$

Combining this identity with (3) yields the centered limit asserted for $\bar{x}_n^{\mathrm{bc}}$. Transport continuity also gives consistency of $\widehat{\mathrm{V}}_n^{\mathrm{bc}}$. Repeating the preceding inverse-retraction argument at $\bar{x}_n^{\mathrm{bc}}$ proves coverage for the corrected Wald region.

For a scalar functional, apply the delta method in Theorem 3.1 to the centered or corrected estimator, replace the limiting derivative and covariance by their consistent transported estimators, and use the standard normal quantile. The resulting studentized statistic converges to $N(0, 1)$. □

B.6.8 Randomized solvers, predictable preconditioners, and Hessian averaging

Deterministic and randomized-solver covariance formulas

Proof of Corollary B.11. Let $\Delta_{k+1} = (\widetilde{\mathrm{P}}_{k+1} - \mathrm{P}_0)\xi_{k+1}$. Since $\mathbb{E}(\xi_{k+1} \mid \mathcal{F}_k) = 0$,

$$\zeta_{k+1} = \mathrm{P}_0 \xi_{k+1} + \Delta_{k+1} - \mathbb{E}(\Delta_{k+1} \mid \mathcal{F}_k).$$

The conditional second moment of the last two terms is bounded by $4\mathbb{E}(\|\Delta_{k+1}\|^2 \mid \mathcal{F}_k)$. The Cesàro average of this bound converges to zero by assumption. Conditional Cauchy-Schwarz then shows that the cross-covariances with $\mathrm{P}_0 \xi_{k+1}$ also vanish in Cesàro mean. Therefore

$$\mathrm{Q} = \mathrm{P}_0 \mathrm{S} \mathrm{P}_0^*.$$

Since $\mathrm{A} = \mathrm{P}_0 \mathrm{H}$,

$$\mathrm{A}^{-1} \mathrm{Q} (\mathrm{A}^{*-1}) = \mathrm{H}^{-1} \mathrm{P}_0^{-1} (\mathrm{P}_0 \mathrm{S} \mathrm{P}_0^*)(\mathrm{P}_0^{*-1})(\mathrm{H}^{*-1}) = \mathrm{H}^{-1} \mathrm{S} (\mathrm{H}^{*-1}).$$

□

Proof of Theorem 3.4. Independence and centering imply $\mathbb{E}(\widetilde{\mathrm{P}}\xi) = 0$, so the leading innovation is $\widetilde{\mathrm{P}}\xi$. Conditional on $\widetilde{\mathrm{P}}$, $\mathbb{E}\{(\widetilde{\mathrm{P}}\xi) \otimes (\widetilde{\mathrm{P}}\xi) \mid \widetilde{\mathrm{P}}\} = \widetilde{\mathrm{P}}\mathrm{S}\widetilde{\mathrm{P}}^*$. Write $\widetilde{\mathrm{P}} = \bar{\mathrm{P}} + (\widetilde{\mathrm{P}} - \bar{\mathrm{P}})$ and take expectations. Both cross terms vanish because $\mathbb{E}(\widetilde{\mathrm{P}} - \bar{\mathrm{P}}) = 0$, giving

$$\mathrm{Q} = \bar{\mathrm{P}}\mathrm{S}\bar{\mathrm{P}}^* + \mathbb{E}\{(\widetilde{\mathrm{P}} - \bar{\mathrm{P}})\mathrm{S}(\widetilde{\mathrm{P}} - \bar{\mathrm{P}})^*\} = \bar{\mathrm{P}}\mathrm{S}\bar{\mathrm{P}}^* + \mathrm{R}.$$

For every $v \in E$,

$$\langle v, \mathrm{R}v\rangle = \mathbb{E}\left\|\mathrm{S}^{1/2}(\widetilde{\mathrm{P}} - \bar{\mathrm{P}})^* v\right\|^2 \geq 0,$$

so $\mathrm{R} \succeq 0$. Moreover, $\mathrm{R} = 0$ if and only if $\mathrm{S}^{1/2}(\widetilde{\mathrm{P}} - \bar{\mathrm{P}})^* v = 0$ almost surely for every vector in a finite basis, which is equivalent to $(\widetilde{\mathrm{P}} - \bar{\mathrm{P}})\mathrm{S}^{1/2} = 0$ almost surely.

Finally, $\mathrm{A} = \bar{\mathrm{P}}\mathrm{H}$, and hence

$$\begin{aligned}\mathrm{V} &= (\bar{\mathrm{P}}\mathrm{H})^{-1}(\bar{\mathrm{P}}\mathrm{S}\bar{\mathrm{P}}^* + \mathrm{R})((\bar{\mathrm{P}}\mathrm{H})^{*-1}) \\ &= \mathrm{H}^{-1}\mathrm{S}(\mathrm{H}^{*-1}) + \mathrm{H}^{-1}\bar{\mathrm{P}}^{-1}\mathrm{R}(\bar{\mathrm{P}}^{*-1})(\mathrm{H}^{*-1}).\end{aligned}$$

Congruence preserves positive semidefiniteness, proving every claim. □

Proof of Corollary B.12. The average operator has mean $\bar{\mathrm{P}}$. Conditional independence of its replicas gives

$$\begin{aligned}&\mathbb{E}\left[\{\widetilde{\mathrm{P}}_{k+1}^{(m_k)} - \bar{\mathrm{P}}\}\mathrm{S}\{\widetilde{\mathrm{P}}_{k+1}^{(m_k)} - \bar{\mathrm{P}}\}^* \mid \mathcal{F}_k\right] \\ &\quad = \frac{1}{m_k^2}\sum_{j=1}^{m_k}\mathbb{E}\{(\widetilde{\mathrm{P}}_{k+1}^{(j)} - \bar{\mathrm{P}})\mathrm{S}(\widetilde{\mathrm{P}}_{k+1}^{(j)} - \bar{\mathrm{P}})^* \mid \mathcal{F}_k\} = \frac{\mathrm{R}}{m_k};\end{aligned}$$

the cross terms are zero. The predictable covariance average in the martingale CLT is therefore

$$\bar{\mathrm{P}}\mathrm{S}\bar{\mathrm{P}}^* + \left(\frac{1}{n}\sum_{k=1}^{n}\frac{1}{m_k}\right)\mathrm{R} \xrightarrow{p} \bar{\mathrm{P}}\mathrm{S}\bar{\mathrm{P}}^* + \lambda\mathrm{R}.$$

Apply the same congruence calculation as in Theorem 3.4. If $m_k \xrightarrow{p} \infty$, boundedness of m_k^{-1} and Cesàro averaging give $\lambda = 0$ in probability. □

Proof of Corollary B.13. For fixed m, Corollary B.12 gives $\mathrm{V}_m = \mathrm{V}_{\mathrm{sand}} + m^{-1}\Delta$. If $m_1 \leq m_2$, then

$$\mathrm{V}_{m_1} - \mathrm{V}_{m_2} = \left(\frac{1}{m_1} - \frac{1}{m_2}\right)\Delta \succeq 0, \qquad \mathrm{V}_{m_2} - \mathrm{V}_{\mathrm{sand}} = \frac{1}{m_2}\Delta \succeq 0.$$

Taking the quadratic form with a gives the contrastwise identity and its strictness criterion. □

Proof of Corollary B.14. Write

$$\begin{aligned}D_{k+1}^{(1)} &= (\Pi_k - \mathrm{P}_0)\xi_{k+1}, \\ D_{k+1}^{(2)} &= (\widetilde{\mathrm{P}}_{k+1} - \Pi_k)\xi_{k+1}.\end{aligned}$$

Because Π_k is $\mathcal{F}_k$-measurable and $\mathbb{E}(\xi_{k+1} \mid \mathcal{F}_k) = 0$, $\mathbb{E}(D_{k+1}^{(1)} \mid \mathcal{F}_k) = 0$. Therefore

$$\zeta_{k+1} = \mathrm{P}_0\xi_{k+1} + D_{k+1}^{(1)} + D_{k+1}^{(2)} - \mathbb{E}(D_{k+1}^{(2)} \mid \mathcal{F}_k).$$

The Cesàro conditional second moments of $D_{k+1}^{(1)}$ vanish by (41). Conditional variance domination and (42) give the same conclusion for the centered second term. Conditional Cauchy-Schwarz shows that all cross covariances involving either error have vanishing Cesàro norm. Hence

$$\frac{1}{n}\sum_{k=0}^{n-1}\mathbb{E}(\zeta_{k+1} \otimes \zeta_{k+1} \mid \mathcal{F}_k) - \frac{1}{n}\sum_{k=0}^{n-1}\mathrm{P}_0\mathbb{E}(\xi_{k+1} \otimes \xi_{k+1} \mid \mathcal{F}_k)\mathrm{P}_0^* \xrightarrow{p} 0.$$

By (38), $\mathrm{Q} = \mathrm{P}_0\mathrm{S}\mathrm{P}_0^*$. Since $\mathrm{A} = \mathrm{P}_0\mathrm{H}$, the cancellation in Corollary B.11 yields $\mathrm{V} = \mathrm{V}_{\mathrm{sand}}$.

Under conditional solver score independence,

$$\mathbb{E}\{\|(\Pi_k - \mathrm{P}_0)\xi_{k+1}\|^2 \mid \mathcal{F}_k\} \le \|\Pi_k - \mathrm{P}_0\|_{\mathrm{op}}^2 \,\mathbb{E}(\|\xi_{k+1}\|^2 \mid \mathcal{F}_k),$$

$$\mathbb{E}\{\left\|(\widetilde{\mathrm{P}}_{k+1} - \Pi_k)\xi_{k+1}\right\|^2 \mid \mathcal{F}_k\} \le \mathbb{E}(\left\|\widetilde{\mathrm{P}}_{k+1} - \Pi_k\right\|_{\mathrm{op}}^2 \mid \mathcal{F}_k)\mathbb{E}(\|\xi_{k+1}\|^2 \mid \mathcal{F}_k).$$

Conditional Jensen gives $\|\Pi_k - \mathrm{P}_0\|_{\mathrm{op}}^2 \le \mathbb{E}(\left\|\widetilde{\mathrm{P}}_{k+1} - \mathrm{P}_0\right\|_{\mathrm{op}}^2 \mid \mathcal{F}_k)$, while

$$\mathbb{E}(\left\|\widetilde{\mathrm{P}}_{k+1} - \Pi_k\right\|_{\mathrm{op}}^2 \mid \mathcal{F}_k) \le 4\mathbb{E}(\left\|\widetilde{\mathrm{P}}_{k+1} - \mathrm{P}_0\right\|_{\mathrm{op}}^2 \mid \mathcal{F}_k).$$

Thus (43) and the bounded score second moment imply both conditions on solver errors applied to the target score. □

The predictable Abel identity

Proof of Theorem 3.6. Multiply (44) by $\mathrm{M}_k = \mathrm{C}_k^{-1}$ and divide by η_{k+1}:

$$\mathrm{H}u_k = \Gamma_k(u_k - u_{k+1}) - \xi_{k+1} - \mathrm{M}_k\varepsilon_{k+1}, \qquad \Gamma_k = \frac{\mathrm{M}_k}{\eta_{k+1}}. \tag{106}$$

Summing the first term on the right gives the exact vector Abel identity

$$\sum_{k=0}^{n-1} \Gamma_k(u_k - u_{k+1}) = \Gamma_0 u_0 - \Gamma_{n-1}u_n + \sum_{k=1}^{n-1}(\Gamma_k - \Gamma_{k-1})u_k.$$

By (12), the right hand side divided by $\sqrt{n}$ is $o_p(1)$. Sum (106), divide by $\sqrt{n}$, and use (13) to obtain

$$\mathrm{H}\frac{1}{\sqrt{n}}\sum_{k=0}^{n-1} u_k = -\frac{1}{\sqrt{n}}\sum_{k=0}^{n-1}\xi_{k+1} + o_p(1).$$

The assumed root-n equivalence of the recursive and arithmetic averages, and the negligible endpoint change from $u_0, \dots, u_{n-1}$ to $u_1, \dots, u_n$, give (14). The stated Gaussian limit follows from the score CLT and Slutsky's theorem. □

Proof of Corollary B.15. The endpoint terms are negligible. Indeed,

$$\begin{aligned}\frac{1}{\sqrt{n}}\mathbb{E}\,\|\Gamma_{n-1}u_n\| &\le \frac{1}{\sqrt{n}\eta_n}\,\|\mathrm{M}_{n-1}\|_{L^2(\mathrm{op})}\,\|u_n\|_{L^2} \\ &\le Cn^{(\alpha-1)/2} \longrightarrow 0,\end{aligned}$$

and $\Gamma_0 u_0/\sqrt{n} \to 0$.

Decompose the one step difference exactly as

$$\Gamma_k - \Gamma_{k-1} = \frac{\mathrm{M}_k - \mathrm{M}_{k-1}}{\eta_{k+1}} + \mathrm{M}_{k-1}\left(\frac{1}{\eta_{k+1}} - \frac{1}{\eta_k}\right). \tag{107}$$

For the first part, Cauchy-Schwarz and the stated rates yield

$$\begin{aligned}\frac{1}{\sqrt{n}}\sum_{k=1}^{n-1}\mathbb{E}\left\|\frac{(\mathrm{M}_k - \mathrm{M}_{k-1})u_k}{\eta_{k+1}}\right\| &\le \frac{C}{\sqrt{n}}\sum_{k=1}^{n-1} k^{\alpha}k^{-\beta}k^{-\alpha/2} \\ &\le \begin{cases} O\{n^{(1+\alpha)/2-\beta}\}, & \beta < 1+\alpha/2, \\ O(n^{-1/2}\log n), & \beta = 1+\alpha/2, \\ O(n^{-1/2}), & \beta > 1+\alpha/2, \end{cases} = o(1).\end{aligned}$$

For the second part, the mean-value theorem gives $\eta_{k+1}^{-1} - \eta_k^{-1} = O(k^{\alpha-1})$, so

$$\begin{aligned}\frac{1}{\sqrt{n}}\sum_{k=1}^{n-1}\mathbb{E}\left\|\mathrm{M}_{k-1}\left(\frac{1}{\eta_{k+1}} - \frac{1}{\eta_k}\right)u_k\right\| &\le \frac{C}{\sqrt{n}}\sum_{k=1}^{n-1} k^{\alpha/2-1} \\ &= O\{n^{(\alpha-1)/2}\} = o(1).\end{aligned}$$

Markov's inequality and (107) prove (12). □

Proof of Corollary B.16. Put $Z_{k+1}^a = \mathcal{T}_{x_\star \to a}\xi_{k+1}$. The assumed martingale weak law and (45) give

$$\frac{1}{n}\sum_{k=0}^{n-1} Z_{k+1}^a \otimes Z_{k+1}^a \xrightarrow{p} \mathrm{S}_a, \qquad \frac{1}{n}\sum_{k=0}^{n-1} Z_{k+1}^a \xrightarrow{p} 0.$$

By (46) and Cauchy-Schwarz,

$$\begin{aligned}&\left\|\frac{1}{n}\sum_{k=0}^{n-1} G_{k+1}^a \otimes G_{k+1}^a - \frac{1}{n}\sum_{k=0}^{n-1} Z_{k+1}^a \otimes Z_{k+1}^a\right\|_{\mathrm{op}} \\ &\quad\le 2\left(\frac{1}{n}\sum \left\|Z_{k+1}^a\right\|^2\right)^{1/2}\left(\frac{1}{n}\sum \left\|h_{k+1}^G\right\|^2\right)^{1/2} + \frac{1}{n}\sum \left\|h_{k+1}^G\right\|^2 = o_p(1).\end{aligned}$$

The sample mean of G_{k+1}^a is also $o_p(1)$. Welford's identity therefore gives $\widehat{\mathrm{S}}_n^a \xrightarrow{p} \mathrm{S}_a$. Together with $\widehat{\mathrm{H}}_n^a \xrightarrow{p} \mathrm{H}_a$ and continuity of matrix inversion, this yields (47).

The influence representation in Theorem 3.6 has covariance $\mathrm{H}_a^{-1}\mathrm{S}_a(\mathrm{H}_a^{*-1})$ after transport to the anchor. The same two-point inverse-retraction expansion used in the proof of Theorem 3.3 then gives convergence of the Wald statistic to χ_d^2. No step uses a limit of C_k. □

Streaming Hessian estimates and fixed-budget sketches

Proof of Proposition B.17. Let

$$\widetilde{\mathrm{H}}_n^a = \lambda_n \mathrm{I} + \frac{1}{n}\sum_{k=0}^{n-1} \widehat{\mathrm{B}}_{k+1}^a$$

be the unsafeguarded estimate. Decompose

$$\begin{aligned}\widetilde{\mathrm{H}}_n^a - \mathrm{H}_a &= \lambda_n \mathrm{I} + \frac{1}{n}\sum_{k=0}^{n-1} \mathrm{R}_{k+1}^B + \frac{1}{n}\sum_{k=0}^{n-1} \mathrm{M}_{k+1}^B, \qquad (108)\\ \mathrm{M}_{k+1}^B &= \widehat{\mathrm{B}}_{k+1}^a - \mathbb{E}(\widehat{\mathrm{B}}_{k+1}^a \mid \mathcal{F}_k).\end{aligned}$$

The final sum is a matrix-valued martingale. In fixed dimension, the Frobenius martingale isometry and (49) give

$$\left\|\frac{1}{n}\sum_{k=0}^{n-1} \mathrm{M}_{k+1}^B\right\|_{L^2(\mathrm{F})}^2 = \frac{1}{n^2}\sum_{k=0}^{n-1} \mathbb{E}\left\|\mathrm{M}_{k+1}^B\right\|_{\mathrm{F}}^2 \le \frac{C}{n}.$$

For the bias term, Minkowski's inequality, (48), and the state rate imply

$$\begin{aligned}\left\|\frac{1}{n}\sum_{k=0}^{n-1} \mathrm{R}_{k+1}^B\right\|_{L^2(\mathrm{op})} &\le \frac{C}{n}\sum_{k=1}^{n} \|u_{k-1}\|_{L^2}^{\rho_{\mathrm{H}}} \\ &\le \frac{C}{n}\sum_{k=1}^{n} k^{-\alpha\rho_{\mathrm{H}}/2} = O(n^{-\alpha\rho_{\mathrm{H}}/2}).\end{aligned}$$

The regularization term is $O(n^{-\delta_\lambda})$. Combining the three bounds in (108), and using $\|\mathrm{B}\|_{\mathrm{op}} \le \|\mathrm{B}\|_{\mathrm{F}}$, proves the stated rate for $\widetilde{\mathrm{H}}_n^a$. Without safeguarding, set $\widehat{\mathrm{H}}_n^a = \widetilde{\mathrm{H}}_n^a$. With safeguarding, let $\widehat{\mathrm{H}}_n^a$ be its spectral projection onto the closed convex set of self-adjoint operators with spectrum in $[h_-, h_+]$. Since the spectrum of H_a belongs to this interval, H_a is fixed by the projection. Nonexpansiveness of metric projection in Frobenius norm gives $\left\|\widehat{\mathrm{H}}_n^a - \mathrm{H}_a\right\|_{L^2(\mathrm{op})} \le \left\|\widetilde{\mathrm{H}}_n^a - \mathrm{H}_a\right\|_{L^2(\mathrm{F})}$, so the safeguarded estimate has the same rate (in fixed dimension). Also $\left\|(\widehat{\mathrm{H}}_n^a)^{-1}\right\|_{\mathrm{op}} \le h_-^{-1}$ and $\left\|\mathrm{H}_a^{-1}\right\|_{\mathrm{op}} < \infty$. The resolvent identity

$$(\widehat{\mathrm{H}}_n^a)^{-1} - \mathrm{H}_a^{-1} = (\widehat{\mathrm{H}}_n^a)^{-1}(\mathrm{H}_a - \widehat{\mathrm{H}}_n^a)\mathrm{H}_a^{-1}$$

then shows that the inverse has the same rate. □

Proof of Corollary B.18. Write $\overline{\mathrm{B}}_k^a = k^{-1}\sum_{j=1}^k \widehat{\mathrm{B}}_j^a$, so $\widetilde{\mathrm{H}}_k^a = \lambda_k \mathrm{I} + \overline{\mathrm{B}}_k^a$. The exact recursion

$$\overline{\mathrm{B}}_k^a - \overline{\mathrm{B}}_{k-1}^a = \frac{1}{k}\{\widehat{\mathrm{B}}_k^a - \overline{\mathrm{B}}_{k-1}^a\}$$

and the bounded L^2 norms of the individual contributions imply $\left\|\overline{\mathrm{B}}_k^a - \overline{\mathrm{B}}_{k-1}^a\right\|_{L^2(\mathrm{op})} = O(k^{-1})$. The regularization increment has the same order by assumption. In fixed dimension this also gives an $O(k^{-1})$ increment in $L^2(\mathrm{F})$. For the safeguarded branch, nonexpansiveness of spectral projection in Frobenius norm yields $\left\|\widehat{\mathrm{H}}_k^a - \widehat{\mathrm{H}}_{k-1}^a\right\|_{L^2(\mathrm{op})} \le \left\|\widetilde{\mathrm{H}}_k^a - \widetilde{\mathrm{H}}_{k-1}^a\right\|_{L^2(\mathrm{F})} = O(k^{-1})$. Coherent transport is isometric and independent of k between the fixed spaces $T_a\mathcal{M}$ and E, so

$$\left\|\mathrm{C}_k^{-1} - \mathrm{C}_{k-1}^{-1}\right\|_{L^2(\mathrm{op})} = \left\|\widehat{\mathrm{H}}_k^a - \widehat{\mathrm{H}}_{k-1}^a\right\|_{L^2(\mathrm{op})} = O(k^{-1}).$$

Minkowski's inequality and the bounded contribution and ridge norms give the uniform L^2 bound on C_k^{-1} in the raw branch; spectral clipping gives that bound directly in the safeguarded branch. Uniform positive definiteness ensures that C_k exists. Apply Corollary B.15. □

Proof of Proposition B.19. Write

$$\mathrm{K}_k(\Omega) = \mathcal{K}_J(\mathrm{B}_k, \Omega), \qquad \mathrm{K}_\star(\Omega) = \mathcal{K}_J(\mathrm{H}, \Omega).$$

By the fixed conditional seed law, the conditional mean operator is $\Pi_k = \mathbb{E}_\Omega \mathrm{K}_k(\Omega)$. Minkowski's inequality and (53) give

$$\left\|\Pi_k - \bar{\mathrm{P}}_J\right\|_{L^2(\mathrm{op})} \le L \left\|\mathrm{B}_k - \mathrm{H}\right\|_{L^2(\mathrm{op})} = O(k^{-\delta}).$$

Thus Proposition B.26(i) verifies (26). The second moment assumptions and conditional independence of seed and score give (24), $b_{k+1}^\xi = 0$, and a bounded conditional second moment of the product by Proposition B.25. The required tail condition is imposed directly in (55). Conditional independence and $\mathbb{E}(\xi_{k+1} \mid \mathcal{F}_k) = 0$ give $\mathbb{E}\{\mathrm{K}_k(\Omega_{k+1})\xi_{k+1} \mid \mathcal{F}_k\} = 0$, so the product itself is ζ_{k+1} and (36) follows without an additional centering argument.

It remains to identify the covariance. Conditional independence and centering give

$$\mathrm{Q}_k := \mathbb{E}(\zeta_{k+1} \otimes \zeta_{k+1} \mid \mathcal{F}_k) = \mathbb{E}_\Omega\{\mathrm{K}_k(\Omega)\Sigma_k \mathrm{K}_k(\Omega)^*\}.$$

Add and subtract $\mathbb{E}_\Omega\{\mathrm{K}_\star(\Omega)\Sigma_k \mathrm{K}_\star(\Omega)^*\}$. The operator inequality $\|\mathrm{UVW}^*\|_{\mathrm{op}} \le \|\mathrm{U}\|_{\mathrm{op}} \|\mathrm{V}\|_{\mathrm{op}} \|\mathrm{W}\|_{\mathrm{op}}$ and Cauchy-Schwarz in Ω yield

$$\begin{aligned}
&\|\mathrm{Q}_k - \mathbb{E}_\Omega\{\mathrm{K}_\star(\Omega)\Sigma_k \mathrm{K}_\star(\Omega)^*\}\|_{\mathrm{op}} \\
&\quad \le C\left\{\mathbb{E}_\Omega \|\mathrm{K}_k(\Omega) - \mathrm{K}_\star(\Omega)\|_{\mathrm{op}}^2\right\}^{1/2}\left[\{\mathbb{E}_\Omega \|\mathrm{K}_k(\Omega)\|_{\mathrm{op}}^2\}^{1/2} + \{\mathbb{E}_\Omega \|\mathrm{K}_\star(\Omega)\|_{\mathrm{op}}^2\}^{1/2}\right] \\
&\quad \le C \|\mathrm{B}_k - \mathrm{H}\|_{\mathrm{op}}.
\end{aligned}$$

Also,

$$\|\mathbb{E}_\Omega\{\mathrm{K}_\star(\Omega)(\Sigma_k - \mathrm{S})\mathrm{K}_\star(\Omega)^*\}\|_{\mathrm{op}} \le C \|\Sigma_k - \mathrm{S}\|_{\mathrm{op}}.$$

Taking Cesàro averages, using the L^2 rate in (52) and (54), gives $n^{-1}\sum_{k<n} \mathrm{Q}_k \xrightarrow{p} \mathrm{Q}_J$. Thus (35) holds. Positive stability of $\bar{\mathrm{P}}_J\mathrm{H}$ and the remaining assumptions complete the application of Theorem 3.1. Because $\mathrm{K}_\star(\Omega)$ is independent of a target score with covariance S, Theorem 3.4 gives the covariance decomposition. The growing-J_k statement follows from Corollary B.14 under its two conditions on solver errors applied to the target score. □

Proof of Proposition B.20. For a fixed positive definite B and right hand side g, one sketch and project step from zero is the solution of

$$\min_v \ \frac{1}{2}\langle v, \mathrm{B}v\rangle \quad \text{subject to} \quad \langle s, \mathrm{B}v\rangle = \langle s, g\rangle.$$

The Lagrange equations give $v = cs$, and the constraint gives $c = \langle s, g\rangle \,/\, \langle s, \mathrm{B}s\rangle$. Thus the update is precisely $\mathcal{K}_{\mathrm{sap}}(\mathrm{B}, s)g$.

Since $\|s\| = 1$ and $\langle s, \mathrm{B}s\rangle \geq h_-$, the rank one operator in (57) has norm at most h_-^{-1}. Moreover,

$$\begin{aligned}&\|\mathcal{K}_{\mathrm{sap}}(\mathrm{B}, s) - \mathcal{K}_{\mathrm{sap}}(\mathrm{B}', s)\|_{\mathrm{op}} \\ &\quad = \left| \frac{\langle s, (\mathrm{B}' - \mathrm{B})s\rangle}{\langle s, \mathrm{B}s\rangle \langle s, \mathrm{B}'s\rangle} \right| \leq h_-^{-2} \|\mathrm{B} - \mathrm{B}'\|_{\mathrm{op}} .\end{aligned}$$

For every $v \in E$,

$$\frac{1}{h_+} \langle v, \mathbb{E}_\mu(s \otimes s)v\rangle \leq \left\langle v, \bar{\mathrm{P}}_{\mathrm{sap}} v\right\rangle = \mathbb{E}_\mu \frac{\langle v, s\rangle^2}{\langle s, \mathrm{H}s\rangle} \leq \frac{1}{h_-} \langle v, \mathbb{E}_\mu(s \otimes s)v\rangle .$$

Because $\|s\| = 1$, also $\mathbb{E}_\mu(s \otimes s) \preceq \mathrm{I}$. Together with (58), this proves (59)–(60). In particular, $\bar{\mathrm{P}}_{\mathrm{sap}}$ is positive definite. The product $\bar{\mathrm{P}}_{\mathrm{sap}}\mathrm{H}$ is similar to the positive definite self-adjoint operator $\mathrm{H}^{1/2}\bar{\mathrm{P}}_{\mathrm{sap}}\mathrm{H}^{1/2}$ and is therefore positive stable. Averaging a fixed number of replicas preserves the mean and all uniform and Lipschitz bounds.

We next verify the stochastic approximation conditions without assuming the state rate or the geometric reduction. Let

$$\Pi_k = \mathbb{E}_\mu\{\mathcal{K}_{\mathrm{sap}}(\mathrm{B}_k, s)\}.$$

The pointwise Lipschitz bound and (61) give

$$\left\|\Pi_k - \bar{\mathrm{P}}_{\mathrm{sap}}\right\|_{L^2(\mathrm{op})} \leq h_-^{-2} \|\mathrm{B}_k - \mathrm{H}\|_{L^2(\mathrm{op})} = O(k^{-\delta}).$$

The spectral band and the spanning condition give, for every k,

$$\Pi_k \succeq \frac{\kappa}{h_+}\mathrm{I}.$$

Taking $\mathrm{K} = \mathrm{H}$ therefore yields the deterministic bound

$$(\Pi_k\mathrm{H})^*\mathrm{H} + \mathrm{H}\Pi_k\mathrm{H} = 2\mathrm{H}\Pi_k\mathrm{H} \succeq \frac{2\kappa}{h_+}\mathrm{H}^2 \succeq \frac{2\kappa h_-^2}{h_+}\mathrm{I}.$$

Thus (37) holds uniformly before the state rate is invoked. The Lipschitz solver bound and the L^2 rate of B_k give $\Pi_k\mathrm{H} \to \bar{\mathrm{P}}_{\mathrm{sap}}\mathrm{H}$ in L^2 for the subsequent mean operator and covariance limits.

Conditional independence of solver and observation and Assumption B.2 give exact centering of the untruncated solver products. The uniform operator bound together with (63) gives the direction second moment bound (79) and a uniform conditional $(2+\tau)$ moment for the raw direction on the localized neighborhood. If Z_{k+1} is the raw transported direction and Δ_{k+1} is its discrepancy from clipping at r_0/η_{k+1}, conditional Markov bounds give directly

$$\mathbb{E}(\|\Delta_{k+1}\| \mid \mathcal{F}_k) \leq C\eta_{k+1}^{1+\tau}, \qquad \mathbb{E}(\|\Delta_{k+1}\|^2 \mid \mathcal{F}_k) \leq C\eta_{k+1}^{\tau}.$$

Thus the discrepancy has root-n negligible conditional mean, vanishing average conditional variance, and vanishing average squared conditional mean. In particular, its predictable contribution satisfies $\mathbb{E}\|\beta_{k+1}\|^2 \leq C\eta_{k+1}$ and accumulates to $b = 0$. Consequently, Proposition B.10 applies before Proposition B.27; it yields $\mathbb{E}\|u_k\|^2 = O(\eta_k)$, after which the latter proposition supplies Assumption B.5. The fixed anchor equivalence follows from Lemma B.33.

The mean operator remainder follows from the displayed L^2 rate and Proposition B.26(i). The field remainder follows from (62) and Proposition B.26(ii). The local score increment moment in (63) gives (25). Boundedness of the sketch operator, the $(2+\tau)$ score moment, and (64) verify the product Lindeberg condition and, by the same covariance comparison used in the proof of Proposition B.19, the limit

$$\mathrm{Q}_m = \mathbb{E}\{\widetilde{\mathrm{P}}_{\mathrm{sap}}^{(m)}\mathrm{S}(\widetilde{\mathrm{P}}_{\mathrm{sap}}^{(m)})^*\}, \qquad \widetilde{\mathrm{P}}_{\mathrm{sap}}^{(m)} = \frac{1}{m}\sum_{j=1}^{m} \mathcal{K}_{\mathrm{sap}}(\mathrm{H}, s^{(j)}).$$

This proves the centered conclusion of Theorem 3.1.

The same primitives also verify the feasible covariance conditions. Decompose the empirical solver mean into

$$\frac{1}{n}\sum_{k=0}^{n-1}(\widetilde{\mathrm{P}}_{k+1}^{(m)} - \Pi_k) + \frac{1}{n}\sum_{k=0}^{n-1}(\Pi_k - \bar{\mathrm{P}}_{\mathrm{sap}}).$$

The first term is a bounded martingale average and the second tends to zero by (61); hence the one-pass mean solver estimate is consistent. The derivative estimate is consistent by assumption. After transport to the anchor, the implemented direction equals the centered target product plus population, state dependent-score, and clipping terms. Their Cesàro squared norm tends to zero by the state rate, (63), the local field bound, and the clipping bounds. Finally, the uniform $(2+\tau)$ product moment yields the outer-product martingale weak law. These observations verify Assumption 3.2, so Theorem 3.3 applies.

Independent replica centering gives

$$\mathrm{Q}_m = \bar{\mathrm{P}}_{\mathrm{sap}} \mathrm{S} \bar{\mathrm{P}}_{\mathrm{sap}}^{*} + \frac{1}{m} \mathrm{R}_{\mathrm{sap}}.$$

Substitution into $(\bar{\mathrm{P}}_{\mathrm{sap}}\mathrm{H})^{-1}\mathrm{Q}_m((\bar{\mathrm{P}}_{\mathrm{sap}}\mathrm{H})^{*-1})$ proves (65). The uniform-sphere identity $\mathbb{E}(s \otimes s) = \mathrm{I}/d$ follows from rotational invariance and unit trace. □

B.6.9 Verification for the applications

Proof of Corollary B.21. For $\widehat{G}(x, Y) = -\operatorname{Log}_x(Y)$, set

$$\xi = -\operatorname{Log}_{x_\star}(Y)$$

and let $q(x, Y)$ be the centered version of $-\mathcal{T}_{x \to x_\star} \operatorname{Log}_x(Y) + \operatorname{Log}_{x_\star}(Y)$. The first order condition gives $\mathbb{E} \operatorname{Log}_{x_\star}(Y) = 0$ and hence $\mathbb{E}\xi = 0$ and $\mathrm{S}_F = \operatorname{Cov}\{\operatorname{Log}_{x_\star}(Y)\}$. The term $q(x, Y)$ is centered by construction. From (66),

$$\mathbb{E} \left\| q(x, Y) \right\|^2 \le \mathbb{E} \left\| -\mathcal{T}_{x \to x_\star} \operatorname{Log}_x(Y) + \operatorname{Log}_{x_\star}(Y) \right\|^2 \le \mathbb{E} L(Y)^2 \left\| R_{x_\star}^{-1}(x) \right\|^2 .$$

Thus the score decomposition and the local second moment bound for the state dependent score follow. Differentiability of G identifies its derivative with H_F in the chosen tangent frame. The remaining direct remainder, solver, geometry, averaging, and martingale conditions are those of Theorem 3.1. The outer product of the leading score is unchanged by its sign, and its conditional product covariance is (68). Applying Corollary B.8 proves (67); the two solver reductions follow from Corollary B.11 and theorem 3.4. □

Proof of Corollary B.22. For a fixed $y \in \mathbb{S}^q$ outside the antipode of $x_\star$, let $r = d(x_\star, y)$ and $e = \operatorname{Log}_{x_\star}(y)/r$. The Hessian of $\frac{1}{2} d^2(x, y)$ at $x_\star$ is

$$r \cot(r) \mathrm{I} + \{1 - r \cot(r)\} e \otimes e.$$

Indeed, the radial direction e has eigenvalue 1, while every direction orthogonal to e has eigenvalue $r \cot r$. Taking expectations gives (69); the target score is $-re$, which gives (70). On $0 \le r \le r_0 < \pi/2$, $r \cot r \ge r_0 \cot r_0 > 0$, so every sample Hessian is bounded below by $r_0 \cot(r_0)\mathrm{I}$ and the same is true of $\mathrm{H}_{\mathbb{S}}$. The support is away from the cut locus and all derivatives used in (66) are uniformly bounded. The conclusion is therefore Corollary B.21 with the displayed $\mathrm{H}_{\mathbb{S}}$ and $\mathrm{S}_{\mathbb{S}}$. □

Proof of Corollary B.23. The affine-invariant positive definite manifold has the globally defined maps (71)–(72). On a bounded geodesic ball, the logarithm and its derivative are bounded in any fixed smooth orthonormal frame, which verifies (66). The first assertion then follows directly from Corollary B.21. Jacobi's formula gives $D \log\det(\mathrm{X})[\mathrm{U}] = \operatorname{tr}(\mathrm{X}^{-1}\mathrm{U})$, so (73) follows from the delta method in (4). □

Proof of Corollary B.24. Use the local graph representative

$$\mathrm{U}(\mathrm{K}) = (\mathrm{U}_1 + \mathrm{U}_2\mathrm{K})(\mathrm{I} + \mathrm{K}^*\mathrm{K})^{-1/2}. \tag{109}$$

At $\mathrm{K} = 0$, a tangent direction is $\mathrm{U}_2\mathrm{K}$. The Riemannian gradient of $-\frac{1}{2} \left\| \mathrm{U}^* X \right\|^2$ is

$$-(\mathrm{I} - \mathrm{U}\mathrm{U}^*) X X^* \mathrm{U}.$$

Its U_2 coordinate at U_1 is $-z_2 z_1^*$. Replacing XX^* by M and differentiating at $\mathrm{K} = 0$ gives $\mathcal{H}[\mathrm{K}] = \mathrm{K}\Lambda_1 - \Lambda_2\mathrm{K}$. For a matrix unit E_{ij},

$$\mathcal{H}[\mathrm{E}_{ij}] = (\lambda_j - \lambda_{r+i})\mathrm{E}_{ij},$$

so the eigengap makes $\mathcal{H}$ self-adjoint positive definite.

On the fixed sufficiently small graph neighborhood, the derivative of the coherently transported sample gradient is bounded by $C\,\|X\|^2$. Conditional on $\mathcal{F}_k$, the current coordinate K_k is fixed and the observation is fresh. Hence

$$\mathbb{E}(\|q_{k+1}\|_{\mathrm{F}}^2 \mid \mathcal{F}_k) \le C\mathbb{E}\,\|X\|^4\,\|\mathrm{K}_k\|_{\mathrm{F}}^2\,,$$

and the same moment controls the second moment of $z_2 z_1^*$. A sufficient $(2+\tau)$ score moment instead follows from $\mathbb{E}\,\|X\|^{4+2\tau} < \infty$. Smoothness of the chart gives a quadratic population remainder. The main theorem now gives (75).

If X is Gaussian in the eigenbasis, z_1 and z_2 are independent and

$$\mathrm{Cov}\{(z_2 z_1^*)_{ij}, (z_2 z_1^*)_{i'j'}\} = \mathbf{1}(i=i')\mathbf{1}(j=j')\lambda_{r+i}\lambda_j.$$

Applying $\mathcal{H}^{-1}$ proves (76). Finally, $D\mathrm{U}(0)[\mathrm{K}] = \mathrm{U}_2\mathrm{K}$ in (109); differentiating $\mathrm{U}(\mathrm{K})\mathrm{U}(\mathrm{K})^*$ yields (77). □